

# Harmonic Triangulations

MARC ALEXA, TU Berlin, Germany

We introduce the notion of harmonic triangulations: a harmonic triangulation simultaneously minimizes the Dirichlet energy of all piecewise linear functions. By a famous result of Rippa, Delaunay triangulations are the harmonic triangulations of planar point sets. We prove by explicit counterexample that in 3D a harmonic triangulation does not exist in general. However, we show that bistellar flips are harmonic: if they decrease Dirichlet energy for one set of function values, they do so for all. This observation gives rise to the notion of locally harmonic triangulations. We demonstrate that locally harmonic triangulations can be efficiently computed, and efficiently reduce sliver tetrahedra. The notion of harmonic triangulation also gives rise to a scalar measure of the quality of a triangulation, which can be used to prioritize flips and optimize the position of vertices. Tetrahedral meshes generated by optimizing this function generally show better quality than Delaunay-based optimization techniques.

CCS Concepts: • **Mathematics of computing** → **Mesh generation**; Discrete optimization; • **Computing methodologies** → *Mesh geometry models*; • **Theory of computation** → Computational geometry.

Additional Key Words and Phrases: Dirichlet energy, simplicial meshes, tetrahedral mesh optimization

## ACM Reference Format:

Marc Alexa. 2019. Harmonic Triangulations. *ACM Trans. Graph.* 38, 4, Article 54 (July 2019), 14 pages. <https://doi.org/10.1145/3306346.3322986>

## 1 INTRODUCTION

The Delaunay triangulation of a point set is a fundamental tool in geometry processing and mesh generation. The reason for this is that it can be computed efficiently in any dimension and the geometric properties of the elements are well defined.

In the plane, Delaunay triangulations satisfy several optimality properties, for example, the minimal interior angle is maximized. A remarkable property is the connection to Dirichlet energy: for the given points in the plane attach a function value to each point. A triangulation of the points gives rise to a piecewise linear function (we present a representation of this function in Section 3). We may ask: which triangulation minimizes the Dirichlet energy of the PL function? The surprising result by Rippa [1990] is that the Delaunay triangulation minimizes Dirichlet energy *independent of the attached function values*.

This property, just as many other optimality properties of planar Delaunay triangulations, fails to generalize to three or more dimensions [Musin 1997]. The starting point of this work was the question: 'Are the good geometric properties of Delaunay triangulations really

due to the classic definitions (see Section 2), or could small Dirichlet energy be the reason?' Our central idea is, consequently, to characterize and define harmonic triangulations as minimizers of Dirichlet energy – see Definition 3.

Our main results are that (1) minimization of Dirichlet energy independent of function values induces a partial order on the triangulations, which is given by the partial order of the positive semi-definite quadratic forms of the Dirichlet energy; and (2) bistellar flips are consistent with this order, in the sense that the effect of a flip on the Dirichlet energy is independent of the function values. This means that flipping can be used to generate local minima of Dirichlet energy in the flip-graph of triangulations. As minimizers of Dirichlet energy we suggest to call such triangulations *harmonic*. The details and additional geometric insights are discussed in Section 4

An important (and to our knowledge open) question is if a globally harmonic triangulation exists. In two dimensions this is the case, since the Delaunay triangulations minimize Dirichlet energy independent of function values. In Section 6.1 we give an explicit counterexample for three dimensions. This means flipping towards locally harmonic triangulations is the best we can do, as the different locally harmonic triangulations are usually not comparable, meaning the comparison of their Dirichlet energies depends on the choice of function values.

There are many functions mapping from the quadratic form of Dirichlet energy to a scalar that are consistent with the partial order. Such functions can be used to optimize not just the combinatorics but also the vertex positions – similar in spirit to optimal Delaunay triangulations [Chen and Xu 2004]. We derive the gradients of a natural function (see Section 5) and use it to optimize vertex positions in a descent scheme (Section 7.3).

Based on our observations we develop basic algorithms (see Section 7) for the generation of locally harmonic triangulations by flipping and potentially moving the vertices. These algorithms are easy to implement given frameworks for robust geometric computations and data structures that support bistellar flips on simplicial complexes. In a series of experiments we demonstrate the favorable properties of these algorithms: harmonic flipping to a locally harmonic triangulation is orders of magnitude faster than sliver exudation [Cheng et al. 2000]. The geometric properties of locally harmonic triangulations are similar or better than techniques based on Delaunay triangulations.

We conclude that harmonic triangulations are a useful new tool for triangulating and optimizing point sets in three dimensions. There are multiple applications for this tool and important follow-up explorations, some of which are discussed in Section 9.

## 2 BACKGROUND AND RELATED WORK

Given a point set  $\{\mathbf{x}_i\}$  in  $\mathbb{R}^d$ , the Delaunay triangulation can be characterized in several seemingly different but intimately related ways. All of the following characterizations are independent of

Author's address: Marc Alexa, TU Berlin, Sekretariat MAR 6-6, Marchstr. 23, 10587 Berlin, Germany, [marc.alex@tu-berlin.de](mailto:marc.alex@tu-berlin.de).

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM. 0730-0301/2019/7-ART54 \$15.00

<https://doi.org/10.1145/3306346.3322986>

the dimension  $d$ . For simplicity, we assume the input points are in *general position* meaning there is only one Delaunay triangulation of the point set.

- A  $d$ -simplex is part of the Delaunay triangulation if its circumsphere contains only the incident vertices.
- Consider the Voronoi diagram of the points, i.e., attach to each point the region of closest points. The Delaunay triangulation is the dual of the Voronoi diagram.
- Embed the points into  $\mathbb{R}^{d+1}$  and lift them to the paraboloid, i.e., the additional coordinate is  $\|\mathbf{x}_i\|^2$ . The lower convex hull, i.e. the faces of the convex hull that can be seen from  $(0, \dots, 0, -\infty)$ , is the Delaunay triangulation.

The literature on Delaunay triangulations is vast – for an in depth discussion of the characterizations above together with proofs of their equivalence we suggest the ‘book in progress’ by Gallier and Quaintance [2017, Ch. 12]; for a general account on triangulations see De Loera et al. [2010].

Delaunay triangulations are an important tool in mesh generation. In many applications, the *quality* of the mesh is important. The notion of quality depends on the particular application, however, in general one wants to avoid simplices that are nearly degenerate. Measures of degeneration are based on angles, sizes of elements (lengths, areas, volumes), radii of insphere and circumsphere. Shewchuk [2002a] provides a detailed, application-dependent discussion of quality measures for dimensions two and three. Bern et al. [Bern et al. 1995] argue that avoiding small dihedral angles is a good measure in arbitrary dimension. Indeed, many works in mesh generation report the smallest dihedral angle. On the other hand, from the perspective of finite element methods is important that large dihedral angles are uniformly bounded away from  $\pi$  as this guarantees convergence. This is true for planar triangulations [Babuška and Aziz 1976] as well as triangulations in 3D [Křížek 1992].

In the plane, the Delaunay triangulation generates high quality elements. Apart from other optimality properties it maximizes the smallest interior angle [Sibson 1978]. Interestingly, in more than two dimensions Delaunay triangulations commonly contain badly shaped elements, so-called *slivers*. One way of removing the slivers is by generalizing to *weighted* Delaunay [de Goes et al. 2014] or *regular* triangulations [De Loera et al. 2010], which are the duals of *power diagrams* [Aurenhammer 1987]. Using this extra degree of freedom of the weights, Cheng et al. [2000] give an algorithm that adjusts the weights so that slivers are removed. The algorithm comes with certain guarantees. We find that in practice the approach we present here performs better (see Section 8.3).

For better quality triangulations based on the Delaunay triangulation it is necessary to adjust the vertex positions. The third of the three notions of Delaunay triangulation can be used to define it as a minimizer of a scalar functional, namely the volume under the triangulation of the points lifted to the paraboloid. Taking the convex hull of the points clearly minimizes the volume (see also [Musin 1997]). This view allows adjusting the point positions to improve the triangulation in a way that is consistent with the Delaunay triangulation. The idea is to fix the vertices at the boundary (the convex hull of the points) and then minimize the volume under the

triangulation of the points lifted to the paraboloid, subject to the combinatorics of the triangulation and the positions of the interior vertices. This approach has been coined *optimal Delaunay triangulation* (ODT) [Chen and Xu 2004]. It has been used to generate high quality triangulations in 3D [Alliez et al. 2005; Chen et al. 2014] and recently been extended to curved meshes [Feng et al. 2018].

A number of algorithms are based on Delaunay triangulations and add additional points to improve the quality of the tetrahedra [Shewchuk 2002b]. Delaunay refinement has been successfully combined with ODT for high-quality mesh generation [Tournois et al. 2009]. If the vertex positions are not important and the mesh only needs to approximate the boundary of a given shape there are also many other ways to generate triangulations, commonly based on starting from a triangulation of the regular lattice and then adjusting the elements close to the boundary [Labelle and Shewchuk 2007; Molino et al. 2003]. These methods generate high quality meshes. We see our work more in relation to Delaunay triangulations.

Harmonic triangulations connect properties of the triangulation with the discrete Laplace-Beltrami operator. One can construct the discrete operator matrix for any given triangulation from the Dirichlet energy [Pinkall and Polthier 1993] (similar to the derivation we provide in Section 3). This construction ensures that the operator will satisfy several pertinent properties of the smooth operator [Wardetzky et al. 2007]. However, the operator generally lacks a maximum principle. In two dimensions, the Delaunay triangulation leads to a maximum principle. This has to do with the Delaunay property guaranteeing that all off-diagonal entries (corresponding to interior edges) of the operator matrix have the same sign. Similar guarantees cannot be deduced from the properties of Delaunay triangulations in higher dimensions (we are unaware of a reference for this statement but it is easy to construct examples). Harmonic triangulations, as we will see, are related to the magnitude of the coefficients of the operator matrix and, therefore, might also be related to the maximum principle.

### 3 PRELIMINARIES

We begin by deriving a description of a piecewise linear function over a triangulation in  $\mathbb{R}^d$ . Then we use this description to compute the Dirichlet energy of the function as a bilinear form. This will lead to a characterization of triangulations based on the symmetric matrices representing the bilinear form.

#### 3.1 Triangulations and piecewise linear functions

Consider  $n$  points  $\mathbf{x}_i \in \mathbb{R}^d$ . We use the notation  $\mathbf{X}_S$  to denote the matrix resulting from considering the points in the index set  $S$ , i.e.

$$\mathbf{X}_S = (\mathbf{x}_{s_0}, \mathbf{x}_{s_1}, \dots) \in \mathbb{R}^{d \times |S|}, \quad S = (s_0, s_1, \dots) \in \mathbb{Z}_n^{|S|} \quad (1)$$

and we write  $\mathbf{X}$  for the matrix consisting of all points. The *convex hull* of the points indexed in  $S$  is

$$\mathbb{R}^d \supset C(\mathbf{X}_S) = \{\mathbf{x} = \mathbf{X}_S \mathbf{a}, \mathbf{a} \in \mathbb{R}_{\geq 0}^{|S|}\}. \quad (2)$$

A  $k$ -simplex is represented by  $k+1$  points  $\Delta^k = (i_0, \dots, i_k) \in \mathbb{Z}_n^{k+1}$  and identifies the region  $C(\mathbf{X}_{\Delta^k})$ . We commonly use  $t = \Delta^d$  to denote a  $d$ -simplex (i.e. *triangle* for  $d = 2$ , *tetrahedron* for  $d = 3$ ).

We are interested in *embedded* triangulations of the convex hull of  $\mathbf{X}$  with all points in  $\mathbf{X}$  as vertices, meaning that each point  $\mathbf{x} \in C(\mathbf{X})$  can be associated uniquely with the relative interior of exactly one simplex  $\Delta^k$ :

$$\mathbf{x} = \mathbf{X}_{\Delta^k} \mathbf{a}, \quad \mathbf{a} \in \mathbb{R}_{>0}^k, \|\mathbf{a}\|_1 = 1. \quad (3)$$

In an embedded triangulation all  $d$ -simplices are non-degenerate. We assume that each  $d$ -simplex  $t$  is represented so that its *signed volume*

$$v_t = \frac{1}{d!} \det(\mathbf{x}_1 - \mathbf{x}_0, \mathbf{x}_2 - \mathbf{x}_0, \dots) = \frac{1}{d!} \det(\mathbf{X}_t \mathbf{M}) \quad (4)$$

is strictly positive. Here  $\mathbf{X}_t \in \mathbb{R}^{d \times (d+1)}$  represents the points spanning the simplex (i.e., the vertex positions) and  $\mathbf{M}$  is a matrix that transforms values associated to the simplex so that they are represented relative to the values of the first vertex. Let  $\mathbf{1} \in \mathbb{R}^d$  be a vector of  $d$  ones and  $\mathbf{I}_d \in \mathbb{R}^d$  the  $d \times d$  identity matrix, then

$$\mathbf{M} = \begin{pmatrix} -\mathbf{1}^T \\ d \\ \mathbf{I}_d \end{pmatrix} \in \mathbb{R}^{(d+1) \times d}. \quad (5)$$

Note that  $\mathbf{X}_t \mathbf{M}$  is a square matrix and based on our assumptions on non-degeneracy it is invertible.

Now consider function values  $\mathbf{f} \in \mathbb{R}^n$  attached to the points. An embedded triangulation gives rise to a piecewise linear function  $f : C(\mathbf{X}) \mapsto \mathbb{R}$  dependent on  $\mathbf{f}$ . The function results from linearly interpolating the function values in each simplex. Note that it suffices to consider  $d$ -simplices, because the linear interpolation in a  $d$ -simplex for a point that is associated to a  $k$ -simplex with  $k < d$  (e.g., an edge) is identical for all  $d$ -simplices incident on that  $k$ -simplex. We perform this interpolation by translating the  $d$ -simplex so that vertex  $i_0$  coincides with the origin, which can be done by multiplication with  $\mathbf{M}$ . Then we solve the linear system

$$\mathbf{X}_t \mathbf{M} \mathbf{b} = \mathbf{x} - \mathbf{x}_{i_0}, \quad (6)$$

providing the coefficients  $\mathbf{b}$  for the linear interpolation of the function values (relative to the function value at vertex  $i_0$ ):

$$f(\mathbf{x}) = f_{i_0} + (f_{i_1} - f_{i_0}, \dots, f_{i_d} - f_{i_0}) \mathbf{b} \quad (7)$$

Putting everything together we get

$$f(\mathbf{x}) = f_{i_0} + \mathbf{f}_t^T \mathbf{M} (\mathbf{X}_t \mathbf{M})^{-1} (\mathbf{x} - \mathbf{x}_{i_0}), \quad C(\mathbf{X}_t) \ni \mathbf{x}, \quad (8)$$

Note that the association  $C(\mathbf{X}_t) \ni \mathbf{x}$  is not unique whenever  $\mathbf{x}$  is contained in a sub-simplex, yet the function  $f$  is identical for all possible assignments.

### 3.2 Dirichlet energy

The Dirichlet energy of a function  $g : \Omega \mapsto \mathbb{R}$  is

$$E(g) = \frac{1}{2} \int_{\Omega} \|\nabla g(\mathbf{x})\|^2 d\mathbf{x}. \quad (9)$$

The gradient of the PL function  $f$  is constant in the relative interior of the  $d$ -simplices and undefined on the sub-simplices. The measure of the sub-simplices is zero, so we can integrate over the relative interiors only. For a simplex, we find the gradient

$$\nabla_{\mathbf{x}} f = \mathbf{f}_t^T \mathbf{M} (\mathbf{X}_t \mathbf{M})^{-1} \quad (10)$$

and since it is constant inside each  $d$ -simplex the Dirichlet energy can be expressed as the sum

$$E(f) = \frac{1}{2} \sum_{t \in \mathcal{T}} v_t \|\mathbf{f}_t^T \mathbf{M} (\mathbf{X}_t \mathbf{M})^{-1}\|^2 \quad (11)$$

$$= \frac{1}{2} \sum_{t \in \mathcal{T}} \mathbf{f}_t^T v_t \mathbf{M} (\mathbf{X}_t \mathbf{M})^{-1} (\mathbf{X}_t \mathbf{M})^{-T} \mathbf{M}^T \mathbf{f}_t \quad (12)$$

$$= \frac{1}{2} \sum_{t \in \mathcal{T}} \mathbf{f}_t^T \mathbf{L}_t \mathbf{f}_t = \frac{1}{2} \mathbf{f}^T \mathbf{L}_{\mathcal{T}} \mathbf{f}. \quad (13)$$

The matrix  $\mathbf{L}_{\mathcal{T}} \in \mathbb{R}^{n \times n}$  is a *discrete Laplace-Beltrami* operator. Since it is derived from a piecewise-linear discretization of Dirichlet energy it is identical to the familiar cotan-operator in two dimensions [Pinkall and Polthier 1993]. It is assembled from the per-triangle Laplace-Beltrami operators

$$\mathbf{L}_t = v_t \mathbf{M} (\mathbf{X}_t \mathbf{M})^{-1} (\mathbf{X}_t \mathbf{M})^{-T} \mathbf{M}^T \quad (14)$$

as submatrices (cf. [Alexa and Wardetzky 2011]).

We believe it is instructive to derive the geometry behind the expression for  $\mathbf{L}_t$ . This will rediscover the more common expressions of Dirichlet energy or the discrete Laplace operator in terms of cotangent weights or dot products of area vectors resp. gradients [Duffin 1959; Dziuk 1988; Meyer et al. 2003; Pinkall and Polthier 1993; Xu and Zikatanov 1999]. The expressions are also useful in the subsequent discussion as well as for implementation.

Consider a simplex with indices  $(0, 1, \dots, d)$  and translate it so that  $\mathbf{x}_0 = \mathbf{0}$ . This means  $\mathbf{X}_t \mathbf{M} = (\mathbf{x}_1, \dots, \mathbf{x}_d)$ . Denote the  $i$ -th row of  $(\mathbf{X}_t \mathbf{M})^{-1}$  as  $\mathbf{y}_i^T$ . By construction of the inverse,  $\mathbf{y}_i$  is orthogonal to all  $\mathbf{x}_j$  for  $j \neq i$ . This means  $\mathbf{y}_i$  is a normal of the *facet* opposite of vertex  $i$ . Let  $\mathbf{n}_i$  be the unit *outward* pointing facet normal. We know  $\mathbf{y}_i^T \mathbf{x}_i = 1$  and can compute the height  $h_i$  of vertex  $i$  over the face opposite to it as  $h_i = -\mathbf{n}_i^T \mathbf{x}_i$ . This implies  $\mathbf{y}_i = -\frac{1}{h_i} \mathbf{n}_i$  and we can write

$$(\mathbf{X}_t \mathbf{M})^{-T} = -(h_1^{-1} \mathbf{n}_1, \dots, h_d^{-1} \mathbf{n}_d). \quad (15)$$

The area  $a_i$  of the facet opposite of vertex  $i$  is connected to the height via the volume:  $a_i h_i / d = v_t$ . This means we can express the inverse also in terms of the facet areas (and the volume of the simplex) as

$$(\mathbf{X}_t \mathbf{M})^{-T} = -\frac{1}{dv_t} (a_1 \mathbf{n}_1, \dots, a_d \mathbf{n}_d). \quad (16)$$

Multiplying this representation from the right with  $\mathbf{M}^T$  yields

$$(\mathbf{X}_t \mathbf{M})^{-T} \mathbf{M}^T = -\frac{1}{dv_t} \left( -\sum_{j=1}^d a_j \mathbf{n}_j, a_1 \mathbf{n}_1, \dots, a_d \mathbf{n}_d \right) \quad (17)$$

We know that the area facet normals sum up to zero (for example [Alexandrov 2005, 2.2.3]), i.e.  $\sum_{j=0}^d a_j \mathbf{n}_j = \mathbf{0}$ , so the sum in the first entry of the vector is simply  $a_0 \mathbf{n}_0$  and we get

$$(\mathbf{X}_t \mathbf{M})^{-T} \mathbf{M}^T = -\frac{1}{dv_t} (a_0 \mathbf{n}_0, a_1 \mathbf{n}_1, \dots, a_d \mathbf{n}_d). \quad (18)$$

With this we can express the elements of  $\mathbf{L}_t$  as

$$(\mathbf{L}_t)_{ij} = \frac{1}{d^2 v_t} a_i a_j \mathbf{n}_i^T \mathbf{n}_j, \quad (19)$$

which is one form of expressing the common *cotan*-weights. Noting that entries of  $\mathbf{L}_{\mathcal{T}}$  result from summing up over  $d$ -simplices, we

find that the  $i$ -th diagonal entry of  $\mathbf{L}_{\mathcal{T}}$  is governed by the volumes of the  $d$ -simplices in the star and the  $d - 1$  simplices in the link. Let  $\mathbf{n}_{t \setminus i}$  be the normal of the facet opposite vertex  $i$  in  $t$ , and likewise  $a_{t \setminus i}$  its area and  $h_{t \setminus i}$  the height of vertex  $i$  over this facet. Then we have

$$(\mathbf{L}_{\mathcal{T}})_{ii} = \frac{1}{d^2} \sum_{t \ni i} \frac{a_{t \setminus i}^2}{|v_t|} = \frac{1}{d^2} \sum_{t \ni i} \frac{a_{t \setminus i}}{h_{t \setminus i}} = \frac{1}{d^2} \sum_{t \ni i} \frac{|v_t|}{h_{t \setminus i}^2}. \quad (20)$$

As we will use the trace of  $\mathbf{L}_{\mathcal{T}}$  to deduce the order of triangulations we note that

$$\text{tr}(\mathbf{L}_{\mathcal{T}}) = \frac{1}{d^2} \sum_i \sum_{t \ni i} \frac{a_{t \setminus i}^2}{|v_t|} = \frac{1}{d^2} \sum_t \frac{\sum_{i \in t} a_{t \setminus i}^2}{|v_t|}. \quad (21)$$

The trace is related to the *harmonic index* [Bobenko and Springborn 2007; Chen et al. 2010; Musin 1997] as a characterization of Delaunay triangulations for  $d = 2$ . The quantity

$$\eta_t = \frac{\sum_{i \in t} a_{t \setminus i}^2}{|v_t|} \quad (22)$$

will be useful for computations that aim at optimizing the trace. For  $d = 3$  a similar (but not identical) term has been suggested by Shewchuk [2002a]<sup>1</sup> as a smooth, size and shape dependent measure for the condition of the stiffness matrix generated by using the tetrahedra as a piecewise linear discretization. We see that for dimension  $d$ ,  $\eta$  and thus  $\text{tr}(\mathbf{L}_{\mathcal{T}})$  scales with exponent  $d - 2$ .

#### 4 HARMONIC ORDER

We start by defining what we mean by harmonic triangulation.

**DEFINITION 1 (HARMONIC TRIANGULATION).** *An embedded triangulation  $\mathcal{T}$  of a set of points  $\mathbf{X} \in \mathbb{R}^{d \times n}$  is harmonic if it spans the convex hull of  $\mathbf{X}$  and satisfies  $\mathbf{f}^T \mathbf{L}_{\mathcal{T}} \mathbf{f} \leq \mathbf{f}^T \mathbf{L}_{\mathcal{T}'} \mathbf{f}$  for any embedded triangulation  $\mathcal{T}'$  of the points  $\mathbf{X}$  spanning their convex hull and any vector of function values  $\mathbf{f} \in \mathbb{R}^n$ .*

By Rippa's theorem [Rippa 1990] the harmonic triangulation for  $d = 2$  exists for any point configuration and coincides with the Delaunay triangulation. For dimension  $d > 2$  every set of  $d + 2$  points has two triangulations and at least one of them is harmonic [Musin 1997]. For more than  $d + 2$  points the question has been open. As we will show in Section 6.1, harmonic triangulations may not exist. This suggests considering a weaker version, where we only ask that a single flip cannot decrease Dirichlet energy (we will make precise what we mean by a flip later).

**DEFINITION 2 (LOCALLY HARMONIC TRIANGULATION).** *A triangulation  $\mathcal{T}$  of set of points  $\mathbf{X} \in \mathbb{R}^{d \times n}$  is locally harmonic if it satisfies the conditions in Definition 1 for all triangulations  $\mathcal{T}'$  that can be reached from  $\mathcal{T}$  with one bistellar flip.*

We now explain how these definitions lead to partial orders and how they are connected to the matrices  $\mathbf{L}_{\mathcal{T}}$ .

<sup>1</sup>this measure only appears in the extended unpublished preprint but not the shortened conference version of the paper




Fig. 1. The two triangulations of  $d + 2$  points for  $d = 3$ . The interior faces are illustrated in blue.

##### 4.1 Partial order

The matrices  $\mathbf{L}_t$  are symmetric PSD by construction (see Eq. 14). It follows that  $\mathbf{L}_{\mathcal{T}}$  is a symmetric PSD matrix. We write this as  $\mathbf{L}_{\mathcal{T}} \geq 0$  (this could be thought of  $\mathbf{f}^T \mathbf{L}_{\mathcal{T}} \mathbf{f} \geq 0$  for any vector  $\mathbf{f} \in \mathbb{R}^n$ ).

We define a partial order on the set of triangulations consisting of all points and covering the convex hull.

**DEFINITION 3 (HARMONIC ORDER).**

$$\mathcal{T}_0 > \mathcal{T}_1 \iff \mathbf{L}_{\mathcal{T}_0} - \mathbf{L}_{\mathcal{T}_1} \geq 0 \wedge \mathbf{L}_{\mathcal{T}_0} \neq \mathbf{L}_{\mathcal{T}_1} \quad (23)$$

In words, two triangulations are ordered, if their Laplace operator matrices are different and the difference is semi-definite. If two triangulations are ordered, their Dirichlet energies are similarly ordered, independent of the function values  $\mathbf{f}$ :

$$\mathcal{T}_0 > \mathcal{T}_1 \implies \mathbf{f}^T \mathbf{L}_{\mathcal{T}_0} \mathbf{f} \geq \mathbf{f}^T \mathbf{L}_{\mathcal{T}_1} \mathbf{f}. \quad (24)$$

If there exists a triangulation  $\mathcal{T}^*$  so that  $\mathcal{T}^* < \mathcal{T}$  for all admissible  $\mathcal{T}$  of  $\mathbf{X}$  it is *harmonic*.

##### 4.2 Bistellar flips are ordered

A set of  $d + 2$  points in  $\mathbb{R}^d$  in convex position has exactly two triangulations (of the convex hull of the points). In the plane ( $d = 2$ ), these are the two triangulations of a (convex) quadrilateral based on choosing one of the two diagonals. In 3-space, the convex hull of 5 points in convex positions always has two vertices with degree 3, and 3 vertices with degree 4 (this is the only configuration admissible by Euler-Poincaré). The five points can be either split into 2 tetrahedra by the triangle formed from the degree 4 vertices, or into three tetrahedra around an additional diagonal between the two degree 3 vertices (see Figure 1).

Replacing the triangulation of a subset of  $d + 2$  points in a triangulation of  $n \geq d + 2$  points by the other possible configuration is called a (bistellar) flip. The *flip-graph* is a graph over the embedded triangulations as nodes and the flips as edges. Two embedded triangulations of a point set are connected by a sequence of flips if they are in the same connected component of the flip graph. Whether the flip-graph is connected depends on the dimension. For  $d = 2$  it is known to be connected and, consequently, any triangulation can be reached from any other triangulation through a sequence of flips [Lawson 1972]. For  $d \geq 5$  the flip graph is known to be disconnected [De Loera et al. 2010]; for  $d = 3, 4$  the general question is still open. We use flips nonetheless as the main tool to explore harmonic triangulations.



DEFINITION 4 (HARMONIC FLIP). A flip from a triangulation  $\mathcal{T}$  to a triangulation  $\mathcal{T}'$  is harmonic if  $\mathcal{T} > \mathcal{T}'$ .

An important insight is that except for flips that have no effect on the PL surface represented by the two triangulations, the two triangulations are always ordered and that this order can be computed efficiently. This is what we show in the following.

OBSERVATION 1. If a flip from a triangulation  $\mathcal{T}$  to a triangulation  $\mathcal{T}'$  changes the Laplace matrix, i.e.,  $\mathbf{L}_{\mathcal{T}} \neq \mathbf{L}_{\mathcal{T}'}$ , then we have either  $\mathcal{T} < \mathcal{T}'$  or  $\mathcal{T} > \mathcal{T}'$ .

This can be shown by generalizing an argument by Chen et al. [2010]. Denote the difference matrix

$$\Delta\mathbf{L} = \mathbf{L}_{\mathcal{T}} - \mathbf{L}_{\mathcal{T}'}, \quad (25)$$

which is symmetric by construction. Consider the quadratic form  $\mathbf{f}^\top \Delta\mathbf{L} \mathbf{f}$ . It is affected only by the  $d + 2$  elements of  $\mathbf{f}$  that correspond to the flip, because all other entries of  $\Delta\mathbf{L}$  are zero. Moreover, it vanishes if the  $d + 2$  values define a linear function (i.e., the points  $\{\mathbf{x}, f_i\}$  lie in a common plane) as both triangulations  $\mathcal{T}$  and  $\mathcal{T}'$  are identical on a plane. Of the  $d + 2$  values of  $\mathbf{f}$  that correspond to the  $d + 2$  vertices, any  $d + 1$  define a linear function in  $\mathbb{R}^d$ . The single remaining value can be chosen in the plane of  $d + 1$  points so that  $\mathbf{f}^\top \Delta\mathbf{L} \mathbf{f}$  vanishes. This means the kernel of  $\Delta\mathbf{L}$  has dimension (at least)  $n - 1$  and the rank of  $\Delta\mathbf{L}$  is (at most) one. A rank-one matrix is necessarily semi-definite, implying the statement above.

Moreover, any rank-one matrix can be written  $\Delta\mathbf{L} = \lambda \mathbf{m} \mathbf{m}^\top$ . From this we see that all diagonal elements of  $\Delta\mathbf{L}$  have the same sign. Consequently, harmonic flips from  $\mathcal{T}$  to  $\mathcal{T}'$  satisfy  $(\mathbf{L}_{\mathcal{T}'} )_{ii} \leq (\mathbf{L}_{\mathcal{T}} )_{ii}$  for all  $i \in [0, n - 1]$  and there exists at least one  $j \in [0, n - 1]$  so that  $(\mathbf{L}_{\mathcal{T}'} )_{ii} < (\mathbf{L}_{\mathcal{T}} )_{ii}$ . This characterization can be conveniently summarized as follows:

OBSERVATION 2 (TRACE CONDITION FOR FLIPS).

$$\mathcal{T}' < \mathcal{T} \iff \text{tr}(\mathbf{L}_{\mathcal{T}'}) < \text{tr}(\mathbf{L}_{\mathcal{T}}). \quad (26)$$

Note that while the implication ' $\implies$ ' follows from the definition, the direction ' $\impliedby$ ' is based on the argument above and only true if  $\mathcal{T}'$  is the result of a flip performed on  $\mathcal{T}$ . Since the trace function is scalar, it introduces an ordering on the graph of flips for edges  $(\mathcal{T}, \mathcal{T}')$  that satisfy  $\mathbf{L}_{\mathcal{T}} \neq \mathbf{L}_{\mathcal{T}'}$ . As pointed out above, the case  $\mathbf{L}_{\mathcal{T}} = \mathbf{L}_{\mathcal{T}'} \iff \Delta\mathbf{L} = 0$  implies that the  $d + 2$  vertices are co-planar. Clearly, the triangulation of a planar patch has no effect on the geometric properties, so this restriction is quite natural.

Since the graph of flips is thus directed (for non-degenerate configurations), it contains sources and sinks. A sink is a locally harmonic triangulation. A sink can be found by starting from an arbitrary triangulation and then following the directed edges, i.e., performing a sequence of harmonic flips.

### 4.3 Geometry of triangulations on $d + 2$ points

Delaunay triangulations can be completely characterized by considering only  $d + 2$  points and asking for the geometric configurations that have no preferred triangulation. Concretely, consider  $d + 1$  points  $\mathbf{X}_t$  forming a  $d$ -simplex. For the  $d + 2$ -nd point we may ask: what are the positions  $\mathbf{x}$  so that both triangulations of the  $d + 2$  points are Delaunay? It turns out that this set is the circumsphere

of  $t$ . Moreover, if  $\mathbf{x}$  lies outside the circumsphere, the triangulation contains  $t$ ; if it lies inside it does not. In other words, every circumsphere of a  $d$ -simplex in the Delaunay triangulation has no vertices in its relative interior. This statement characterizes Delaunay triangulations in every dimension and is commonly known as the *empty sphere property*.

We similarly investigate geometric configurations of  $d + 2$  points for harmonic triangulations. We start with a  $d$ -simplex  $\mathbf{X}_t$  and ask for positions  $\mathbf{x}$  of an additional points so that the difference  $\Delta\mathbf{L}$  of the two triangulations vanishes. It is convenient to focus on a single diagonal element of  $\Delta\mathbf{L}$ , namely the diagonal element of the additional point at  $\mathbf{x}$ . The two possible triangulations can be described as follows: the triangulation containing  $t$  is constructed by connecting  $\mathbf{x}$  to all facets of  $t$  that are *visible* to  $\mathbf{x}$ . To make this precise consider the facet normals  $\mathbf{n}_i$  of  $t$  and denote the height of  $\mathbf{x}$  over the facet opposite to vertex  $i$  as

$$h_i(\mathbf{x}) = \mathbf{n}_i^\top \mathbf{x} - d_i. \quad (27)$$

A facet is visible, if  $h_i(\mathbf{x}) > 0$ . The triangulation not containing  $t$  is composed of all  $d$ -simplices resulting from connecting the vertex at  $\mathbf{x}$  to the remaining facets in  $t$ , i.e., those for which  $h_i(\mathbf{x}) < 0$ . Recall the representation of diagonal entries of the matrix  $\mathbf{L}_{\mathcal{T}}$  from Eq. 20. The fact that the two triangulations are distinguished by the signs of the heights means the difference for the diagonal elements corresponding to  $\mathbf{x}$  is

$$\delta(\mathbf{x}) = \sum_i \frac{a_i}{h_i(\mathbf{x})} \quad (28)$$

and the set of points for which both triangulations are harmonic is characterized by  $\delta(\mathbf{x}) = 0$ . It is important to note that the singularities of this function are the facet planes and that the function changes its sign across the facet planes. Multiplying with all denominators leads to:

$$0 = \sum_i a_i \prod_{j \neq i} h_j(\mathbf{x}) =: c_d(\mathbf{x}) \quad (29)$$

and now the singularities are solutions. The polynomial  $c_d$  has order  $d$ . Since we know by Rippa's theorem [Rippa 1990] that Delaunay and harmonic triangulation coincide in the plane, the quadratic polynomial for  $d = 2$  is a curious way to represent the circumsphere of a triangle. In higher dimension, however, the polynomial degree is larger than 2. We leave the discussion of the case  $d = 3$  for Section 6.2; for now we note that the algebraic surface represented by  $c(\mathbf{x}) = 0$  is not a sphere for  $d > 2$ .

## 5 VERTEX POSITIONS

Recent techniques for generating high quality triangulations in 3D are based on the idea of optimizing a function over the space of all triangulations as well as the positions of (some) vertices. In particular, optimal Delaunay triangulations (ODT) minimize the volume under the PL function resulting from lifting the vertices to a paraboloid [Chen and Xu 2004]. Minimizing the volume among triangulations generates the convex hull, which is known to be the Delaunay triangulation for the vertices lifted to the paraboloid [Musin 1997]. Optimizing the positions of (interior) vertices generates isotropic triangulations [Alliez et al. 2005]. Several techniques can be used to further optimize the distribution of vertices

and the quality of the resulting triangulation [Chen et al. 2014; Feng et al. 2018; Tournois et al. 2009].

The definition of harmonic triangulations is based on being a minimizer for the Dirichlet energy independent on the function values. Yet, minimizing the Dirichlet energy w.r.t. the vertex position does depend on the function values – this situation is an instance of *mesh smoothing* based on a posteriori error estimates [Bank and Smith 1997]. Any fixed choice of function values could be used to define an energy, and this energy would be consistent with the definition of harmonic triangulations. However, fixing a vector of function values is not the only way: every function  $g$  mapping from the space of symmetric PSD matrices to a scalar is admissible if it satisfies  $\mathcal{T}_0 < \mathcal{T}_1 \Rightarrow g(\mathbf{L}_{\mathcal{T}_0}) < g(\mathbf{L}_{\mathcal{T}_1})$ .

Our representation of the matrix  $\mathbf{L}_{\mathcal{T}}$  as a sum of the matrices  $\mathbf{L}_t$  provides an easy way to generate the gradient of any matrix function w.r.t. the vertex positions  $\mathbf{X}$  for functions  $g$  that are *linear* in  $\mathbf{L}_{\mathcal{T}}$ . In these cases we can compute the gradient as the sum of the gradients of  $\mathbf{L}_t$ , i.e.

$$\frac{\partial g(\mathbf{L}_{\mathcal{T}})}{\partial \mathbf{X}} = \sum_{t \in \mathcal{T}} \frac{\partial g(\mathbf{L}_t)}{\partial \mathbf{X}_t}. \quad (30)$$

We focus on the trace as the scalar function to be minimized. Its gradient can be computed as

$$\frac{\partial \text{tr}(\mathbf{L}_t)}{\partial \mathbf{X}_t} = \frac{1}{d!} (\mathbf{X}_t \mathbf{M})^{-\top} \mathbf{M}^{\top} (\text{tr}(\mathbf{L}_t) \mathbf{I} - 2\mathbf{L}_t). \quad (31)$$

Note that  $(\mathbf{X}_t \mathbf{M})^{-\top} \mathbf{M}^{\top}$  is the vector of area-weighted facet normals. This means the entire expression could be evaluated based on the areas and normal vectors of the facets and the volumes of the  $d$ -simplices. We leave the derivation of specific rules for updating the vertices on the boundary for future work.

Since it is not clear what setting the gradient in Eq. 31 to zero means geometrically, we look at a related situation. Consider the case of optimizing a single vertex  $i$  w.r.t. the function  $\mathbf{e}_i^{\top} \mathbf{L}_{\mathcal{T}} \mathbf{e}_i$ , where  $\mathbf{e}_i$  is the  $i$ -th canonical unit vector, i.e. the  $i$ -th row or column of the identity matrix  $\mathbf{I}_n$ . According to the Eq 20 this means we have the following minimization problem:

$$\arg \min_{\mathbf{x}_i} \sum_{t \ni i} \frac{a_t}{h_t(\mathbf{x}_i)} \quad (32)$$

Note that only the heights depend on the vertex position  $\mathbf{x}_i$  while the areas of the facets in the link of vertex  $i$  are not affected by varying  $\mathbf{x}_i$ .

This optimization problem has an elegant solution if the degree of vertex  $i$  is minimal, i.e. its link polytope is a simplex. In this case we can use the heights  $h_t$  as variables, which only have to satisfy the condition

$$\sum_{t \ni i} h_t a_t = d v_i. \quad (33)$$

We can now minimize using the Lagrangian dependent on the heights

$$L(\{h_t\}, \lambda) = \sum_{t \ni i} \frac{a_t}{h_t} + \lambda \sum_{t \ni i} h_t a_t. \quad (34)$$

Setting the partial derivative w.r.t.  $h_t$  to zero

$$\frac{\partial L}{\partial h_t} = -\frac{a_t}{h_t^2} + \lambda a_t = 0 \quad \Rightarrow \quad h_t^{-2} = \lambda \quad (35)$$




Fig. 2. Triangulations of a distorted octahedron. The triangulation with 5 tetrahedra at the top is Delaunay (the green lines indicate the Delaunay ball). There are two harmonic flips leading to triangulations with a single diagonal and 4 tetrahedra. The configuration has one additional triangulation with one diagonal that is not shown. The triangulations based on one diagonal are mutually incomparable, showing that there generally are no (globally) harmonic triangulations in 3D.

reveals that all heights need to be identical. In other words, the optimal placement for a vertex with  $d + 1$  neighbors is at the *incenter* of the link polytope. This result immediately extends to link polytopes that are circumscribed to a common sphere.

This result is interesting in light of the optimal placement of a single vertex in ODT, where the optimal location of a single vertex is at the *circumcenter* if the vertices in the link are co-spherical [Alliez et al. 2005; Chen 2004; Chen et al. 2014]. A potential problem with the circumcenter is that it might be outside the link polytope. In contrast, the incenter is always inside a tangential polytope.

## 6 OBSERVATIONS IN 3D

While the combinatoric of harmonic triangulations is understood in 2D, it is open if they exist in 3 (or more) dimensions [Musin 1997]. We give an example for 6 vertices that settles this issue by showing that there is generally no triangulation that minimizes Dirichlet energy for all function values  $\mathbf{f}$ : The example also illustrates how harmonic flips remove slivers in Delaunay triangulations. Since there is no globally harmonic triangulation in 3D we focus on locally harmonic triangulations, which can be found by flipping. We investigate how harmonic flips relate to Delaunay triangulations and find that harmonic flipping strictly reduces the number of tetrahedra. This suggests computing locally harmonic triangulations by starting from the Delaunay triangulation and then performing only 3/2-flips.

### 6.1 Examples on 6 vertices

Consider 6 points  $\mathbf{x}_0, \dots, \mathbf{x}_5$  in strictly convex position. The convex hull consists of 8 faces, each vertex is incident on 4 of them. This means each vertex is connected to all but one other vertex by an edge. Assume the vertices are numbered so that edges  $(0, 1), (2, 3), (4, 5)$  are *not* part of the convex hull.

In general position the points have 6 different embedded triangulations:

- 3 triangulations arise from inserting one of the three diagonals. Consider the case of inserting the diagonal  $(0, 1)$ . Then there are 4 tetrahedra, all containing the edge  $(0, 1)$  and a pair of the remaining 4 vertices. Two of the three possible triangulations based on one diagonal are illustrated in Figure 2 (lower part).
- 3 triangulations arise from inserting two of the three diagonals. Then the two inserted diagonals form one tetrahedron, and each of the remaining two vertices connects to two faces of this tetrahedron. Thus, these triangulations have a total of 5 tetrahedra. Any of these triangulations may be degenerate if the two diagonals intersect. Figure 2 (top) shows one of the three cases, while the other two in this specific example are degenerate.

Now consider the vertex position of a slightly perturbed octahedron:

$$\mathbf{X} = \begin{pmatrix} -1-\delta & 1+\delta & -\epsilon & -\epsilon & \epsilon & \epsilon \\ 0 & 0 & -1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & -1 & 1 \end{pmatrix}, \quad \delta \geq \epsilon > 0. \quad (36)$$

The situation is visualized for  $\delta = .5, \epsilon = .15$  in Figure 2. The distance of the vertices to the origin is  $1 + \delta^2 + 2\delta > 1 + \epsilon^2$ . This means the sphere around  $(2, 3, 4, 5)$  is empty, so the points form a Delaunay tetrahedron. Consequently, the Delaunay triangulation  $\mathcal{T}_D$  is given by inserting the diagonals  $(2, 3), (4, 5)$  and consists of 5 tetrahedra. There are 3 more embedded triangulations in this configuration, based on the three 4-tetrahedra configurations resulting from inserting one diagonal, which we denote  $\mathcal{T}_0, \mathcal{T}_2, \mathcal{T}_4$ . The remaining two configurations resulting from two diagonals are degenerate because the diagonals intersect. The Delaunay configuration can be flipped into the configurations based on only one of the two diagonals.

In the following we assume that  $\delta$  and  $\epsilon$  are small compared to the lengths of the edges. We first explain why the flip from the Delaunay triangulation is harmonic, i.e.,  $\mathcal{T}_2, \mathcal{T}_4 < \mathcal{T}_D$ , implying that all three triangulations  $\mathcal{T}_0, \mathcal{T}_2, \mathcal{T}_4$  are locally harmonic. Then we show that no two of them are comparable, meaning the vertex set  $\mathbf{X}$  has no harmonic triangulation.

For the flip, recall that it suffices to consider the traces of  $\mathbf{L}_{\mathcal{T}}$ . The  $i$ -th diagonal element can be computed from the ratio of squared face areas and volumes of incident tetrahedra (cf. Eq. 20). Only the Delaunay triangulation contains the sliver  $t = (2, 3, 4, 5)$  with volume  $O(\epsilon)$ . All other tetrahedra in the triangulations have volume  $O(1)$  and the face areas are likewise in  $O(1)$ . This means  $\text{tr}(\mathbf{L}_{\mathcal{T}_D}) = O(\epsilon^{-1})$  where as  $\text{tr}(\mathbf{L}_{\mathcal{T}_i}) = O(1)$ , so the flip is harmonic.

Recall that the for two triangulations to be comparable, all diagonal elements of  $\mathbf{L}_{\mathcal{T}}$  need to be ordered similarly. We only need to consider the triangulations  $\mathcal{T}_0, \mathcal{T}_2, \mathcal{T}_4$  as candidates for a (globally) harmonic triangulation because  $\mathcal{T}_D > \mathcal{T}_2, \mathcal{T}_4$ . Each of these triangulations consists of four tetrahedra incident on the main diagonal. The faces on the convex hull all have area  $\sqrt{3}/2 + O(\delta)$  (here and in the following we make use of  $O(\epsilon) \in O(\delta)$ ). The interior faces are spanned by the diagonal and another vertex, so they have area  $1 + O(\delta)$ . The volumes of the tetrahedra are then  $1/3 + O(\delta)$ .

To compute the diagonal entry  $i$  we need to go over all incident tetrahedra. For the vertices incident on the diagonal there are 4 tetrahedra, and the face not incident on the vertex is on the convex hull; the other vertices are incident on 2 tetrahedra and the face not




Fig. 3. The Cayley nodal cubic is the set of points on which a fifth point leads to identical Laplace operators for the two possible triangulations. It is analogous to the circumsphere in Delaunay triangulations. The color codings shows the different regions induced by the half space of the face planes. In other words, all points with the same color are in the positive half space of the same faces. The left image shows the surface for a regular tetrahedron, the right image is based on an irregular one.

incident on the vertex is interior. This means the diagonal entries are

$$(\mathbf{L}_{\mathcal{T}})_{ii} = \begin{cases} 4 \frac{3/4 + O(\delta)}{1/3 + O(\delta)} = 9 + O(\delta) & i = j, i = j + 1 \\ 2 \frac{1 + O(\delta)}{1/3 + O(\delta)} = 6 + O(\delta) & \text{else} \end{cases} \quad (37)$$

For a wide range of values  $\epsilon, \delta$  we have  $6 + O(\delta) < 9 + O(\delta)$  (for example, the concrete values chosen for Figure 2). In this case, the triangulations are not ordered. To make an example, the choice  $\mathbf{f} = \mathbf{e}_j$  will lead to  $\mathbf{e}_j^T \mathbf{L}_{\mathcal{T}_2} \mathbf{e}_j > \mathbf{e}_j^T \mathbf{L}_{\mathcal{T}_4} \mathbf{e}_j, k \neq j$ . This means we can maximize Dirichlet energy for any choice of triangulation among  $\mathcal{T}_0, \mathcal{T}_2, \mathcal{T}_4$ .

Lastly, note that the two degenerate triangulations missing in this example are not changing this result. Slightly perturbing the vertices would lead to two more triangulations with sliver tetrahedra, which are not harmonic either. Since such configurations of 6 vertices can be part of any point set with more than 6 points we make the following conclusion:

**OBSERVATION 3.** For  $d = 3$  and any  $n \geq 6$  there exist point sets  $\mathbf{X} \in \mathbb{R}^{d \times n}$  without a harmonic triangulation.

While the example on 6 point in  $\mathbb{R}^3$  has no direct analogy in higher dimension, we believe similar situations exist. We make this concrete as follows:

**CONJECTURE 1.** For every dimension  $d > 3$  there exists a configuration of  $n \geq d + 3$  points  $\mathbf{X} \in \mathbb{R}^{d \times n}$  for which none of the embedded triangulations of the convex hull is harmonic.

## 6.2 Delaunay vs. harmonic flip

A flip in 3D is governed by 5 points. We take the view of Section 4.3 and fix 4 points forming a tetrahedron  $t$ . Then we explore the dependence of the triangulation on the position  $\mathbf{x}$  of the fifth vertex. We expect that there is a set of points in space (in fact, a surface) for which both triangulations are harmonic.

As mentioned, for Delaunay triangulations this surface is the circumsphere through the four points of  $t$ . For the harmonic triangulation, the surface is given by Eq. 29. For  $d = 3$  it is a cubic

surface. The gradient of the cubic polynomial is

$$\nabla_{\mathbf{x}} c_3 = \sum_{i=0}^3 a_i \sum_{j=0, j \neq i}^3 \mathbf{n}_j^\top \prod_{k=0, k \neq j, i}^3 h_k(\mathbf{x}). \quad (38)$$

If the position  $\mathbf{x}$  is identical to one of the vertices  $\mathbf{x}_i$ , then three of the heights  $h_j(\mathbf{x}_i)$  are identical zero, while only  $h_i(\mathbf{x}_i) < 0$  (we defined the normals to point outwards). The products  $\prod_{k \neq i, j} h_k(\mathbf{x})$  consist of two terms, of which at least one is zero. This means the gradient vanishes in the vertices  $\mathbf{x}_i$  so the surface has singular points there. Moreover, the Hessian contains only linear terms and is non-zero in the vertices. This means the vertices are *ordinary double points* of the cubic surface. Cayley has characterized cubic surfaces [Cayley 1869] and found that the 4 double points uniquely determine the cubic surface (in projective space). Consequently, the surface for which both triangulations are harmonic is a specific Euclidean instance of *Cayley's nodal cubic*, with the nodal points being the vertices. Figure 3 shows renderings of the surface for a regular tetrahedron as well as one with different edge lengths.

In order to understand the relation of the Delaunay and the harmonic triangulation it is useful to distinguish among the possible locations of  $\mathbf{x}$ . We classify the position based on how many faces of the tetrahedron are visible from  $\mathbf{x}$ , i.e. the number of positive heights  $h_i(\mathbf{x})$ .

- (0) The point is inside the tetrahedron  $t$  and there is only one possible triangulation consisting of 4 tetrahedra.
- (1) The point is in the positive half space of exactly one face. Connecting the point to this face preserves  $t$  and results in 2 tetrahedra. Removing  $t$  requires connecting  $\mathbf{x}$  to the three hidden faces and results in 3 tetrahedra.
- (2) The point is in the positive half space of two faces. Connecting the point to the two faces (i.e. preserving the existing tetrahedron  $t$ ) generates a triangulation with 3 tetrahedra. The triangulation without  $t$  results from connecting  $\mathbf{x}$  to the hidden faces and consists of 2 tetrahedra.
- (3) The point is in the positive half space of three faces. The vertex incident on the three faces is inside the tetrahedron formed by  $\mathbf{x}$  and the other three vertices. There is only one possible triangulation in this case.
- (4) The point cannot be in the positive half spaces of all 4 faces if the tetrahedron has positive volume.

The color coding in Figure 3 illustrates the classification. We now compare the Delaunay and harmonic triangulation. The triangulations are identical for cases 0 and 3, so we only need to consider 1 and 2.

We start with the case that  $\mathbf{x}$  is in the positive half space of only one face. It appears that in this area the Cayley cubic is contained in the circumsphere. This is demonstrate in a cut-away view in Figure 4(left) – note how the Cayley cubic tightly encloses the faces of the tetrahedron, while the circumsphere is far away. This configuration implies that if  $\mathbf{x}$  is outside the circumsphere  $t$  is preserved for both the Delaunay and harmonic triangulation and the triangulation consists of 2 tetrahedra. Between the circumsphere and the Cayley cubic the Delaunay triangulation consists of 3 tetrahedra, while the harmonic triangulation has only 2. Inside the Cayley cubic both Delaunay and harmonic triangulation consist of 3 tetrahedra.




Fig. 4. Cut views of the circumsphere (gray) and the Cayley cubic (blue) for an anisotropic tetrahedron. The circumsphere encloses the Cayley cubic in the positive half space of only one face plane (left), while the Cayley cubic is outside of the circumsphere in the positive half space of two face planes (right).

In the case that  $\mathbf{x}$  is in the positive half space of two faces it appears the circumsphere is inside the Cayley cubic. In Figure 4(right) this is clearly visible on the lower edge. This means if  $\mathbf{x}$  lies outside the Cayley cubic,  $t$  is preserved and both triangulations consist of 3 tetrahedra. Between the Cayley cubic and the circumsphere, the harmonic triangulation removes  $t$  and consists of 2 tetrahedra, while the Delaunay triangulation still consists of 3. Inside the circumsphere both triangulations consist of 2 tetrahedra.

We summarize this observation as follows:

**CONJECTURE 2.** *Given 5 points in  $\mathbb{R}^3$  so that their Delaunay and Harmonic triangulation are unique. Then the two triangulations are either identical or the Harmonic triangulations consists of 2 tetrahedra while the Delaunay triangulation consists of 3.*

We have checked the conjecture by generating random configurations of 5 points and were unable to find a counterexample.

The important practical implication is that harmonic flipping from the Delaunay triangulation decreases the number of tetrahedra unless the Delaunay triangulation coincides with the harmonic triangulation. It also simplifies the implementation of such an algorithm, because only edges need to be checked and possibly flipped. In the following description of algorithms we exploit this property.

## 7 ALGORITHMS

We suggest algorithms for generating locally harmonic triangulations by flipping as well as optimizing the vertex positions.

### 7.1 The basic flipping algorithm

The basic algorithm performs elementary flips. A flip is harmonic if it decreases the trace of  $\mathbf{L}_{\mathcal{T}}$ . A flip either removes two tetrahedra adjacent to a face and replaces it by three tetrahedra; or it removes three tetrahedra incident on an edge and replaces it by two tetrahedra. In both cases we evaluate the change of  $\text{tr}(\mathbf{L}_{\mathcal{T}})$  by considering the values  $\eta_t$  for the 5 tetrahedra  $t$  being removed or created by the flip (since no other tetrahedra are affected). Throughout the procedure we compute  $\eta_t$  for each tetrahedron only once and store

it for later use. This is more efficient than using Eq. 29 that defines the Cayley surface.

When evaluating a possible flip, we first check that the new 2, resp. 3 tetrahedra have positive volume. If this is the case, the flip would lead to another embedded triangulation. Also, this check preserves the convex boundary of the triangulation. If the flip is geometrically valid, we compute the squared areas of the new faces and then  $\eta_t$  for the new tetrahedra. For flipping an edge, let  $t_0, t_1, t_2$  be the current tetrahedra, and  $t_l, t_r$  the new ones. The flip is harmonic if  $\eta_{t_l} + \eta_{t_r} < \eta_{t_0} + \eta_{t_1} + \eta_{t_2}$ . Likewise, let  $t_l, t_r$  be the tetrahedra incident on an existing face and  $t_0, t_1, t_2$  be the tetrahedra resulting from flipping the face. The flip is harmonic if  $\eta_{t_0} + \eta_{t_1} + \eta_{t_2} < \eta_{t_l} + \eta_{t_r}$ .

Initially we iterate over all edges and faces. If a harmonic flip is found it is inserted into a queue. Also, we mark all current tetrahedra as valid. Then we take elements from the queue. If all tetrahedra incident on the face or edge to be flipped are valid, the flip is performed (otherwise it is simply discarded). Performing a flip entails invalidating the tetrahedra incident on the edge or face, removing them from the triangulation. The new tetrahedra are set to be valid and inserted in the triangulation. Their values  $\eta$  have already been computed for checking if the flip is harmonic. All edges and faces on the boundary in the convex hull of the 5 points are checked for possible new flips, similar to the initial stage. If new harmonic flips are found they are inserted into the queue. The process terminates when the queue is empty.

It remains to decide on a way to prioritize the flips. As we will show in Section 8 the order in which flips are performed does affect the outcome of the procedure, however, the effect on the properties of the resulting triangulations is small. Prioritizing flips that lead to a large decrease in trace seems to lead to generally good results. Also, the necessary values are readily available, because we have to compute the sums  $\eta_{t_0} + \eta_{t_1} + \eta_{t_2}$  and  $\eta_{t_l} + \eta_{t_r}$  anyways, and the value for prioritization is simply the difference of these two sums. Note also, that this value only depends on the tetrahedra being affected by a flip. This means the value of a flip never changes – a flip can only become invalid if one of the tetrahedra is being invalidated by another flip. This means the implementation only requires a simple priority queue.

## 7.2 Flipping from the Delaunay triangulation

It seems plausible to start with the Delaunay triangulation. This has the advantage that robust and fast algorithms for generating the starting triangulation are readily available; and that the above algorithm can be simplified because only edges (of degree three) need to be evaluated for possible flips.

A flip induces new possible flips. It is useful and therefore noteworthy that not all of the edges in the two newly generated tetrahedra need to be checked. Let  $x_l, x_r$  be incident on the edge that is flipped so that  $x_0, x_1, x_2$  are forming the new triangle, which is incident on the two new tetrahedra. Since the 5 points are in convex position, all edges are incident on at least two tetrahedra outside of the convex hull (the external dihedral angle is larger than  $\pi$  so it cannot be an internal dihedral angle of a single tetrahedron). Since the edges (1, 2), (2, 3), (3, 1) are incident on the two new tetrahedra on the inside of the convex hull they have at least degree 4 and cannot

be flippable. Conversely, the remaining 6 edges connected to  $x_l, x_r$  used to have degree at least 4 for the same reason, but after the flip they are incident to only one tetrahedron on the inside of the convex hull so they may become flippable and need to be checked.

Perhaps more interestingly, the faces of the two new tetrahedra also never need to be flipped: The inner face is the result of the flip and was created by a harmonic flip. The reverse operation cannot be harmonic. The remaining faces are part of the Delaunay triangulation. This means there exists an empty Delaunay ball around them, suggesting also the Cayley cubic contains no other points (recall that face flips are more restrictive for harmonic flips than Delaunay, Section 6.2), so the face flips would not be harmonic. It seems this argument can be extended to any faces encountered during the flipping algorithm and, indeed, we observe in practice that in the whole sequence no harmonic face flips become available.

## 7.3 Vertex position update

For updating the vertex position we consider the trace function  $\text{tr}(\mathcal{L}_{\mathcal{T}}(\mathbf{X}))$ . This means we try to minimize the trace in all steps of the algorithm. We start by computing the Delaunay triangulation and flip to minimize the trace. Then the vertex positions are updated by gradient descent using line search. After this step, all edges and faces are considered for possible flips to decrease the trace. The two steps of gradient descent with line search and flipping are iterated until no flips are necessary and the gradient descent step is smaller than some threshold.

For line search we use Brent's method [Press et al. 1992, Ch. 10.3]. It is important, however, to choose the interval so that the triangulation is embedded, negative volumes could appear to reduce the trace function. We first approximate the smallest value  $\lambda$  so that the triangulation with coordinates  $\mathbf{X} - \lambda \nabla \mathbf{X}$  is embedded by an interval search. Then we use the interval  $[0, \lambda]$  as the starting point for Brent's method.

For good results it is crucial to also move the vertices on the boundary. Since we have no explicit boundary rule we resort to projecting the vertices back to the surface. In the special case of spherical boundaries it suffices to normalize the vertex positions. In the case of boundaries given as triangle meshes, the closest point on the mesh is taken as the projection. This process has potential for improvement, which we leave for future work.

## 7.4 Implementation

We have implemented the above algorithms using CGAL's 3D triangulation framework [Jamin et al. 2018] and Eigen [Guennebaud et al. 2010]. An implementation of the basic flipping algorithm is available online.

## 8 EXPERIMENTS

In the following experiments we evaluate the geometric properties of harmonic triangulations experimentally. We are interested in computation time relative to the number of input vertices  $n$  and, when applicable, the number of tetrahedra  $|\mathcal{T}|$ . For the properties of the resulting triangulation we look at the number of tetrahedra after optimization and the shape of the tetrahedra.



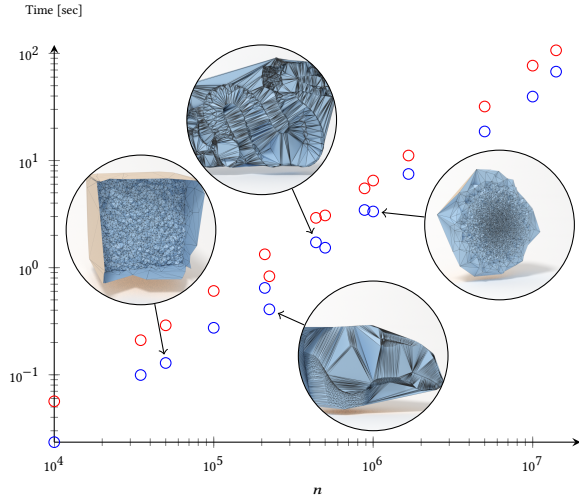


Fig. 5. Time required to generate the Delaunay triangulation (red) and to flip from the Delaunay triangulation to the locally harmonic triangulation (blue). Note that the flipping procedure requires a roughly constant fraction of the time required to build the Delaunay triangulation. The insets illustrate the resulting triangulations.

There is no single measure for the quality of the triangulation that applies to all possible use cases. We focus on dihedral angles  $\phi_{tij}$ , defined in our notation by

$$\cos \phi_{tij} = -\mathbf{n}_{t \setminus i}^T \mathbf{n}_{t \setminus j}. \quad (39)$$

A commonly reported quantity is the smallest dihedral angle:

$$\underline{\phi} = \min_{t \in \mathcal{T}, i, j \in t} \phi_{tij}. \quad (40)$$

While the Delaunay flip in two dimensions (and thus the harmonic flip) selects the configuration that maximizes the smallest interior angle, this is not true for Delaunay flips in three dimensions and dihedral angles. One may wonder if the harmonic flip is a generalization of the two-dimensional situation that does extend this behavior to higher dimension. A test based on sampling 5 uniformly distributed points reveals that the harmonic triangulation maximizes  $\underline{\phi}$  in more than 95% of the cases, but not always. Interestingly, the Delaunay triangulation appears to be uncorrelated with the minimal dihedral angle in three dimensions.

Because of the mentioned importance for the convergence of finite element methods we also report the largest dihedral angle

$$\bar{\phi} = \max_{t \in \mathcal{T}, i, j \in t} \phi_{tij}. \quad (41)$$

For some point sets (or shapes) the smallest or largest dihedral angle is affected only by the boundary, and the different triangulation approaches may have no effect on the elements on the boundary. As an indication on the distribution of small or large dihedral angles we also report the largest angle  $\phi_{5\%}$  among the 5% smallest angles as well as the smallest angle  $\phi_{5\%}^{5\%}$  among the 5% largest angles, i.e. the 5% and 95% rank in the distribution of dihedral angles. In Tables 2 and 3 large angles are given as the difference to  $180^\circ$  so that large numbers are always better.

Table 1. Basic results for Delaunay triangulations and harmonic flipping on point sets. Starting from the Delaunay triangulation, harmonic flipping requires less time than the generation of the Delaunay triangulation. There is a notable reduction in the number of tetrahedra.

| Input    | $n$  | Delaunay |                 | Harmonic |                 |
|----------|------|----------|-----------------|----------|-----------------|
|          |      | time     | $ \mathcal{T} $ | time     | $ \mathcal{T} $ |
| Gaussian | 10K  | 0.06s    | 67K             | 0.02s    | 60K             |
| Bunny    | 35K  | 0.21s    | 238K            | 0.10s    | 204K            |
| Uniform  | 50K  | 0.28s    | 335L            | 0.13s    | 298K            |
| Gaussian | 100K | 0.61s    | 675K            | 0.27s    | 601K            |
| Julius   | 209K | 1.33s    | 1.44M           | 0.64s    | 1.27M           |
| Camel    | 223K | 0.85s    | 858K            | 0.41s    | 727K            |
| Dragon   | 438K | 2.92s    | 3.13M           | 1.71s    | 2.63M           |
| Uniform  | 500K | 3.07s    | 3.37M           | 1.53s    | 2.99M           |
| Blade    | 883K | 5.50s    | 6.07M           | 3.45s    | 5.19M           |
| Gaussian | 1.0M | 6.50s    | 6.77M           | 3.34s    | 6.02M           |
| Hand     | 1.7M | 11.2s    | 11.9M           | 7.48s    | 10.1M           |
| Uniform  | 5.0M | 31.9s    | 33.8M           | 18.6s    | 30.0M           |
| Gaussian | 10M  | 76.9s    | 67.7M           | 39.5s    | 60.2M           |
| Lucy     | 14M  | 107s     | 96.9M           | 67.6s    | 81.8M           |

### 8.1 Statistics for Delaunay and harmonic flipping

We start from various point sets  $X$  and compute the Delaunay triangulation. Some point sets are the vertex sets of common triangle meshes, i.e. they are distributed on a surface embedded in  $\mathbb{R}^3$ . We also use points sampled randomly in space, both based on the uniform distribution in  $[0, 1]^3$  as well as based on a Gaussian distribution with unit variance.

We then take the resulting Delaunay triangulation and perform harmonic flips. The flips are prioritized based on the decrease in trace, i.e. taken from a priority queue with  $\text{tr}(\Delta L)$  as value. Table 1 provides the detailed results of this experiment. We believe the absolute computation times are not relevant, but it is important that harmonic flipping requires less time than computing the initial Delaunay triangulation. Figure 5 shows that the time required for flipping is a constant fraction across a wide range of input sizes. We also see that harmonic flipping reduces the number of tetrahedra significantly.

### 8.2 Dependence on order of flips

Recall that generally more than one primitive can be flipped, and each flip would reduce the Dirichlet energy regardless of function values  $f$ . The decrease in energy depends on the function values  $f$ . Performing one flip may invalidate another flip and in general the order of flips leads to different local minima in the graph of flips.

Here we want to analyze how the different local minima vary and, if there is variation, how prioritizing flips based on  $\text{tr}(\Delta L)$  compares. For this we start with three small tetrahedra meshes. The reason for taken small meshes is that the number of triangulations over a set of points is very large, and we wish to sample a representative part of this space.

Starting from the Delaunay triangulation, we randomly choose elements from the available flips until no harmonic flips are available.

Table 2. Comparison between sliver exudation and harmonic flipping. Starting from a Delaunay triangulation, sliver exudation and harmonic flipping are performed. The results are compared w.r.t. the number of remaining tetrahedra  $|\mathcal{T}|$  as well as small and large dihedral angles. Large dihedral angles are reported as the difference to  $180^\circ$ .

| Input mesh |                 |        |              |              |              | Sliver exudation |                 |             |              |              |              | Harmonic flipping |                 |             |              |              |              |
|------------|-----------------|--------|--------------|--------------|--------------|------------------|-----------------|-------------|--------------|--------------|--------------|-------------------|-----------------|-------------|--------------|--------------|--------------|
| Type       | $ \mathcal{T} $ | $\phi$ | $\phi_{5\%}$ | $\phi^{5\%}$ | $\bar{\phi}$ | time             | $ \mathcal{T} $ | $\phi$      | $\phi_{5\%}$ | $\phi^{5\%}$ | $\bar{\phi}$ | time              | $ \mathcal{T} $ | $\phi$      | $\phi_{5\%}$ | $\phi^{5\%}$ | $\bar{\phi}$ |
| Sphere     | 1.02K           | 0.2    | 35.1         | 70.6         | 0.4          | 0.09s            | 1.00K           | <b>7.9</b>  | 37.6         | 71.7         | <b>31.7</b>  | 0.00s             | 1.00K           | 5.1         | <b>37.8</b>  | <b>72.1</b>  | 23.0         |
| Sphere     | 2.29K           | 0.7    | 33.2         | 61.3         | 1.1          | 0.19s            | 2.20K           | 17.1        | 38.2         | 67.4         | 25.0         | 0.00s             | 2.15K           | <b>19.9</b> | <b>41.2</b>  | <b>71.3</b>  | <b>30.7</b>  |
| Horse      | 2.61K           | 0.6    | 15.7         | 36.4         | 1.6          | 0.23s            | 2.46K           | 3.6         | 20.5         | 44.6         | <b>6.8</b>   | 0.00s             | 2.30K           | <b>3.6</b>  | <b>23.8</b>  | <b>52.1</b>  | 6.7          |
| Sphere     | 18.5K           | 0.3    | 32.1         | 61.2         | 0.6          | 1.18s            | 17.8K           | <b>11.2</b> | 38.1         | 67.8         | <b>19.1</b>  | 0.01s             | 17.4K           | 8.8         | <b>41.3</b>  | <b>72.0</b>  | 13.7         |
| Homer      | 99.3K           | 0.1    | 18.8         | 53.6         | 0.2          | 10.0s            | 95.4K           | <b>0.9</b>  | 25.1         | 59.7         | <b>1.7</b>   | 0.07s             | 92.2K           | 0.9         | <b>29.3</b>  | <b>64.0</b>  | 1.6          |
| Sphere     | 149K            | 0.6    | 31.7         | 60.3         | 0.9          | 11.3s            | 143K            | <b>9.5</b>  | 38.2         | 67.7         | <b>15.3</b>  | 0.09s             | 140K            | 5.0         | <b>41.4</b>  | <b>72.1</b>  | 7.5          |
| Vase       | 295K            | 0.0    | 17.5         | 45.8         | 0.0          | 26.8s            | 282K            | <b>1.3</b>  | 22.0         | 52.7         | <b>2.3</b>   | 0.28s             | 269K            | 1.2         | <b>25.7</b>  | <b>58.5</b>  | 2.0          |
| Sphere     | 1.12M           | 0.0    | 17.8         | 42.1         | 0.1          | 100s             | 1.08M           | <b>1.5</b>  | 21.6         | 48.2         | <b>2.7</b>   | 1.28s             | 1.03M           | 0.7         | <b>24.6</b>  | <b>53.6</b>  | 1.2          |
| Sphere     | 18.6M           | 0.0    | 31.7         | 60.3         | 0.0          | 1733s            | 17.9M           | <b>3.5</b>  | 38.5         | 68.0         | <b>5.7</b>   | 23.4s             | 17.5M           | 3.0         | <b>41.7</b>  | <b>72.5</b>  | 4.8          |

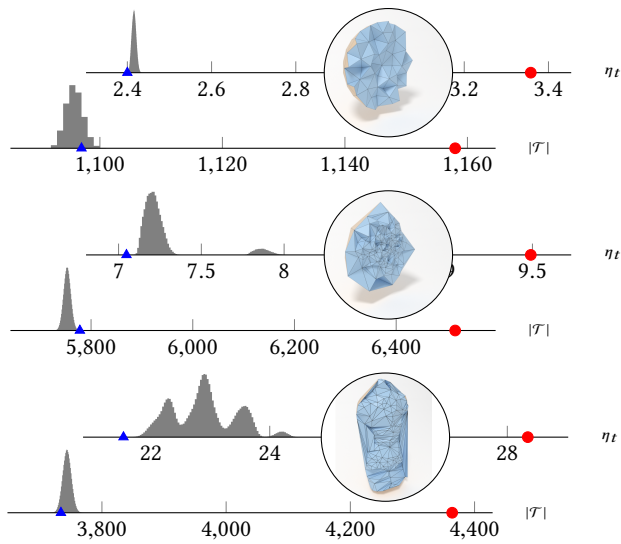


Fig. 6. Variation among different locally harmonic triangulations. Starting from the Delaunay triangulation (red circles) choose harmonic flips randomly until no more harmonic flips are available. The upper graphs show  $\bar{\eta}_t$ , the lower graphs show the histogram of the number of tetrahedra in the resulting mesh. The blue triangles mark the result of prioritizing by  $\text{tr}(\Delta L)$ .

We repeat this process one million times. For a comparison of the resulting tetrahedral meshes we measure  $|\mathcal{T}|$  and, as a measure of the quality of the tetrahedra,  $\bar{\eta}_t$ . We include the initial Delaunay mesh and the result from prioritizing by  $\text{tr}(\Delta L)$ . Figure 6 shows the result of this experiment. We find that the properties of local minima are rather similar relative to the initial Delaunay triangulation. For example, starting with 1158 Delaunay tetrahedra inside a unit sphere, the one million random local harmonic minima cover the range [1092, 1101], which are just 10 different values. The quality measure shows several modes for some of the shapes. Nonetheless, there is a clear difference between the quality of the local harmonic minimum and the Delaunay triangulation.




Fig. 7. Time required by sliver exudation and harmonic flipping relative to the number of tetrahedra  $|\mathcal{T}|$ . The black line indicates linear growth at  $10^5$  tetrahedra per second. Harmonic flipping is two orders of magnitude faster than sliver exudation.

Prioritizing by  $\text{tr}(\Delta L)$  leads to very good shape quality. However, we did find random samples with smaller  $\bar{\eta}_t$  meaning that greedily minimizing the trace provides no guarantee to find the triangulations with smallest trace.

### 8.3 Comparison to sliver exudation

It has become clear through the previous experiment that the quality of elements is improved by harmonic flipping. We now ask how it performs in comparison to techniques that are targeted at eliminating slivers. We chose *sliver exudation* [Cheng et al. 2000] as it appears to be the most successful technique for removing slivers without altering the position of the vertices.

As a starting point we take Delaunay meshes generated by CGAL and then apply either sliver exudation or harmonic flipping. Table 2 shows the result of this comparison for a number of different inputs. The most apparent difference is in running time: harmonic flipping is about two orders of magnitude faster. Figure 7 illustrates that this effect is stable across a wide range of mesh sizes. Another clear trend is that harmonic flipping leads to fewer elements.

The comparison of element shape is more subtle. The extreme dihedral angles show mixed results. While the locally harmonic triangulation has similar or slightly better extreme dihedral angles in a few cases, sliver exudation notably improves the angle bound



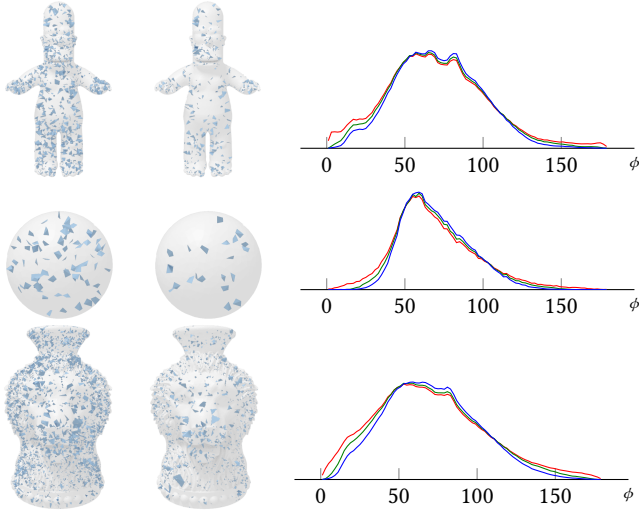


Fig. 8. Comparison of the triangulations resulting from sliver exudation and harmonic flipping started from a triangulation generated by Delaunay refinement (as implemented in CGAL). Left: Tetrahedra with dihedral angles smaller than  $10^\circ$ , resp.  $20^\circ$  for the sphere) after sliver exudation. Middle: the same illustration for harmonic flipping. Right: histograms over dihedral angles. The red curve shows the distribution of dihedral angles for the initial Delaunay triangulation, green for sliver exudation, blue for harmonic flipping.

in particular for the spheres. Looking at the distribution of dihedral angles, harmonic flipping generates consistently better results. This is evident from the angles for the 5% smallest and largest angles. Figure 8 shows examples of the distribution for three of the inputs, including the sphere with largest difference in extremal dihedral angles. In all cases, the histogram of dihedral angles starts and ends flatter for harmonic triangulations, meaning there are fewer elements with extreme dihedral angles. This is also visualized in the left side of Figure 8: locally harmonic triangulations have significantly fewer elements with dihedral angles smaller than  $10^\circ$ , resp. or  $20^\circ$  for the sphere.

#### 8.4 Optimizing vertex positions

The gradient descent scheme described in Section 7.3 can be used to move the vertices for further improvement of the quality of the tetrahedra. We call this process *harmonic optimization*. It is similar in spirit to ODT, which also minimizes a scalar functional across the combinatorics and the vector of vertex positions.

The iterative optimization procedure decreases the quantity  $\text{tr}(\mathbf{L}_{\mathcal{T}})$  in each step. It is interesting to see how the minimal dihedral angle develops over the iterations. The graph in the inset is typical for most examples. Note how the angle oscillates while the trace function monotonically decreases.




Fig. 9. Comparison of runtime performance of ODT (red) and harmonic optimization (blue). The graphs show the smallest and largest dihedral angles relative to the computation time. Both plots are based on tetrahedral meshes with spherical boundary on 500 (left) and 25K (right) input vertices.

This behavior makes it difficult to define a stopping criterion based on  $\text{tr}(\mathbf{L}_{\mathcal{T}})$  when the goal is to optimize extreme dihedral angles.

It seems natural to compare the result of this procedure to ODT. For this we use the implementation of ODT that is part of the Mesh3D package in CGAL. A problem for comparing the runtimes of the iterative procedures is that a large part of the computation time is spent on the boundary projection step. A comparison of dihedral angles versus computation time is best performed on spheres, for which the boundary treatment is trivial. Figure 9 shows the result for two spheres, one with 500 and the other one with 25K input vertices. The oscillating behavior of extramal dihedral angles can be seen for both methods. Harmonic optimization provides much better results much faster. This is quite natural, as ODT is based inherently on a functional that is insensitive to slivers.

It has been noted that the quality of the tetrahedral mesh can be drastically improved by adding sliver perturbation, a process that directly tries to improve element shape by perturbing the vertices incident on slivers. Lastly, we have already seen that harmonic flipping has an effect similar to sliver exudation. For a complete picture, we include sliver perturbation and, additionally, sliver exudation as post-processing steps to ODT.

As mentioned, it is difficult to define a stopping criterion for harmonic optimization. For a fair comparison, first the CGAL implementation of the optimization has been run without time limits. Then harmonic optimization was stopped after roughly the same time. This has led to computation times less than a second for the smaller examples, up to a minute for the larger ones.

Table 3 shows the results of this experiment. Overall, the results are similar to the situation without vertex movement. While the extremal dihedral angles show results in favor of the explicit shape optimization, harmonic optimization exhibits better distribution of dihedral angles. Figure 10 shows the results for 4 cases with inferior extremal dihedral angles. In all cases, the distribution of angles is better at the extreme ends of the spectrum. The illustrations based on tetrahedra with angles smaller than  $15^\circ$  resp.  $25^\circ$  show that harmonic optimization generates significantly less elements with small angles.

## 9 DISCUSSION

The concepts of harmonic flipping and harmonic optimization directly lead to triangulations that improve Delaunay-based methods.

Table 3. Comparison between the optimization of tetrahedral meshes based on optimal Delaunay triangulations, vertex perturbation and sliver exudation as available in CGAL and harmonic optimization. The numbers provided are the same as in Table 2.

| Type             | Input mesh      |        |              |              |              | ODT + perturbation + exudation |             |              |              |              | Harmonic optimization |             |              |              |              |
|------------------|-----------------|--------|--------------|--------------|--------------|--------------------------------|-------------|--------------|--------------|--------------|-----------------------|-------------|--------------|--------------|--------------|
|                  | $ \mathcal{T} $ | $\phi$ | $\phi_{5\%}$ | $\phi^{5\%}$ | $\bar{\phi}$ | $ \mathcal{T} $                | $\phi$      | $\phi_{5\%}$ | $\phi^{5\%}$ | $\bar{\phi}$ | $ \mathcal{T} $       | $\phi$      | $\phi_{5\%}$ | $\phi^{5\%}$ | $\bar{\phi}$ |
| Sphere (uniform) | 2.29K           | 0.6    | 33.2         | 61.3         | 1.1          | <b>2.15K</b>                   | 12.9        | 42.9         | 74.9         | 22.2         | 2.15K                 | <b>30.7</b> | <b>44.9</b>  | <b>75.5</b>  | <b>47.7</b>  |
| Sphere (graded)  | 9.04K           | 0.0    | 25.3         | 60.1         | 0.1          | 8.86K                          | <b>12.0</b> | <b>36.1</b>  | <b>71.6</b>  | <b>26.9</b>  | <b>8.56K</b>          | 6.8         | 34.5         | 66.8         | 5.4          |
| Sphere (uniform) | 18.5K           | 0.3    | 32.1         | 61.2         | 0.6          | <b>17.4K</b>                   | 15.8        | 43.8         | 75.8         | 29.9         | 17.5K                 | <b>23.6</b> | <b>44.5</b>  | <b>75.9</b>  | <b>38.8</b>  |
| Horse            | 40.8K           | 0.2    | 16.2         | 48.0         | 0.3          | 38.7K                          | <b>5.4</b>  | 23.9         | <b>64.8</b>  | <b>8.8</b>   | <b>37.3K</b>          | 3.9         | <b>38.2</b>  | 60.4         | 5.9          |
| Bunny            | 54.5K           | 0.2    | 38.7         | 59.8         | 0.4          | 51.2K                          | <b>12.2</b> | 43.5         | <b>76.3</b>  | <b>20.8</b>  | <b>51.2K</b>          | 11.0        | <b>43.8</b>  | 75.4         | 15.5         |
| Armadillo        | 135K            | 0.1    | 16.3         | 47.1         | 0.2          | 127K                           | <b>6.0</b>  | 23.6         | <b>64.5</b>  | <b>8.7</b>   | <b>123K</b>           | 4.8         | <b>29.7</b>  | 58.5         | 6.8          |
| Sphere (uniform) | 149K            | 0.6    | 31.7         | 60.3         | 0.9          | <b>140K</b>                    | 13.6        | <b>44.9</b>  | <b>77.1</b>  | 24.6         | 140K                  | <b>21.5</b> | 44.7         | 76.2         | <b>36.6</b>  |

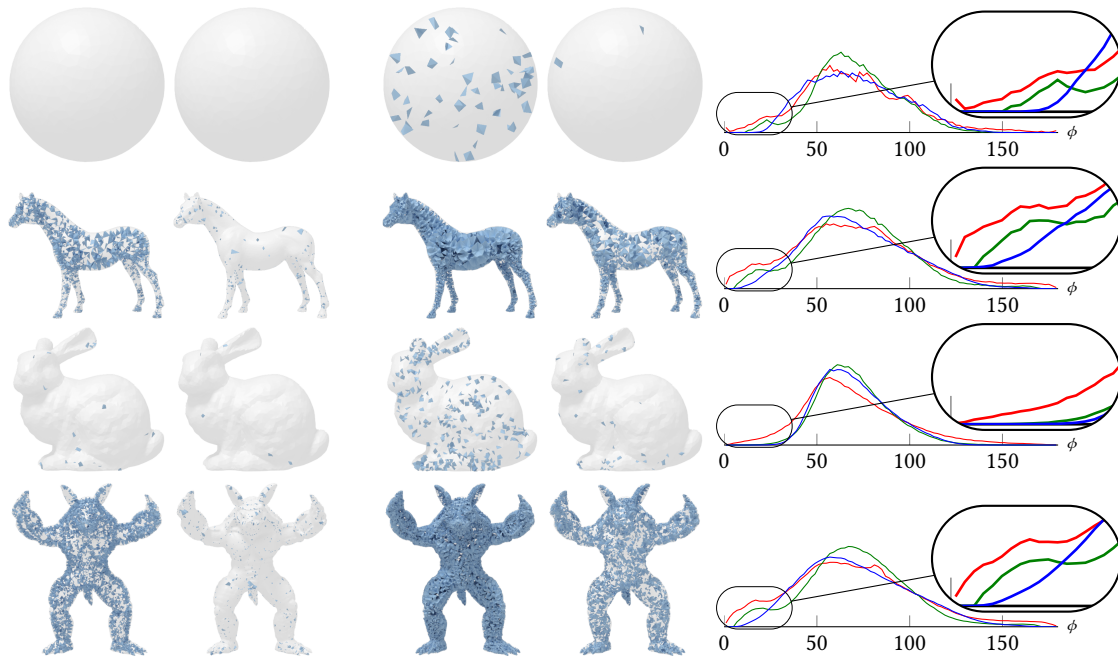


Fig. 10. Comparison of triangulations resulting from ODT plus sliver perturbation and exudation with harmonic optimization. Left: Tetrahedra with dihedral angle smaller than  $15^\circ$  after ODT optimization, sliver perturbation, and sliver exudation (as implemented in CGAL) compared to harmonic optimization. Middle: the same comparison for  $25^\circ$ . Right: histograms of dihedral angles. Red is the initial triangulation, generated by Delaunay refinement, green curves are based on ODT plus sliver perturbation, and additional exudation; blue shows the result of harmonic optimization.

This is notable, as the generation of high quality tetrahedral meshes is usually also a major engineering effort. The implementation of harmonic flipping and harmonic optimization appear to be considerably simpler than alternatives. Our straightforward implementation already handles meshes with several million input vertices and hundreds of millions output elements.

We believe harmonic flipping has a large range of immediate applications. Every process that uses Delaunay triangulations as a way of generating an initial triangulation would likely benefit from harmonic flipping, because it reduces the number of elements to be handled at small additional cost. In addition, harmonic flipping removes exactly those elements that potentially cause numerical issues. For example, TetWild [Hu et al. 2018] starts with a Delaunay

triangulation of the point set and then processes the elements to include the input boundary elements. It falls back to more complex number types in case the elements are near-degenerate. This should happen much less if the tetrahedral mesh is based on harmonic triangulations.

The main drawback of harmonic triangulations is clearly that there is no global optimum in 3 (and probably higher) dimensions. We believe this is an inherent property of triangulations, which can be seen from the example in Section 6.1. There simply is no (globally) preferable choice for removing the sliver tetrahedra. Note that sliver exudation likewise provides no unique way to exude the slivers.

We made a number of claims regarding the properties of harmonic triangulations, in particular relative to the Delaunay triangulation. We believe formally proving these properties is important future work. Of particular relevance is the connection of locally harmonic and regular triangulations. It is currently unclear if locally harmonic triangulations are regular.

We also see interesting avenues for alternative constructions of harmonic triangulations. Note that the common way to generate Delaunay triangulations in three and higher dimensions is incremental [Edelsbrunner and Shah 1992; Joe 1989]. One way to interpret this is that each vertex is inserted into an existing tetrahedron by connecting it to the vertices. Then a sequence of flips restores the Delaunay property. It would be interesting to replace the Delaunay flips by harmonic flips. From a practical perspective it would be important to understand the dependence of the resulting triangulation on the insertion order and also how the incremental construction affects the quality.

Harmonic flipping appears to be easy to implement also in higher dimension. This could immediately improve triangulations in higher dimension.

## ACKNOWLEDGMENTS

Many ideas in this work are the result of discussions with Boris Springborn, who also helped improving the exposition. Philipp Herholz created the renderings of the meshes and the Cayley cubic.

## REFERENCES

- Marc Alexa and Max Wardetzky. 2011. Discrete Laplacians on General Polygonal Meshes. *ACM Trans. Graph.* 30, 4, Article 102 (July 2011), 10 pages. <https://doi.org/10.1145/2010324.1964997>
- Aleksandr D Alexandrov. 2005. *Convex polyhedra*. Springer, Berlin, Heidelberg. <https://doi.org/10.1007/b137434>
- Pierre Alliez, David Cohen-Steiner, Mariette Yvinec, and Mathieu Desbrun. 2005. Variational Tetrahedral Meshing. *ACM Trans. Graph.* 24, 3 (July 2005), 617–625. <https://doi.org/10.1145/1073204.1073238>
- Franz Aurenhammer. 1987. Power diagrams: properties, algorithms and applications. *SIAM J. Comput.* 16, 1 (1987), 78–96.
- Ivo M. Babuška and A. K. Aziz. 1976. On the Angle Condition in the Finite Element Method. *SIAM J. Numer. Anal.* 13, 2 (1976), 214–226. <https://doi.org/10.1137/0713021>
- Randolph E. Bank and R. Kent Smith. 1997. Mesh Smoothing Using A Posteriori Error Estimates. *SIAM J. Numer. Anal.* 34, 3 (June 1997), 979–997. <https://doi.org/10.1137/S0036142994265292>
- Marshall Bern, Paul Chew, David Eppstein, and Jim Ruppert. 1995. Dihedral Bounds for Mesh Generation in High Dimensions. In *Proceedings of the Sixth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA '95)*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 189–196.
- Alexander I. Bobenko and Boris A. Springborn. 2007. A discrete Laplace-Beltrami operator for simplicial surfaces. *Discrete Comput. Geom.* 38, 4 (2007), 740–756. <https://doi.org/10.1007/s00454-007-9006-1>
- Arthur Cayley. 1869. A Memoir on Cubic Surfaces. *Philosophical Transactions of the Royal Society of London* 159 (1869), 231–326. <https://doi.org/10.1098/rstl.1869.0010>
- Long Chen. 2004. Mesh Smoothing Schemes Based on Optimal Delaunay Triangulations. In *Proceedings, 13th International Meshing Roundtable*. Sandia National Laboratories, Williamsburg, VA, USA, 109–120.
- Long Chen and Jinchao Xu. 2004. Optimal Delaunay Triangulations. *Journal of Computational Mathematics* 22, 2 (2004), 299–308.
- Renjie Chen, Yin Xu, Craig Gotsman, and Ligang Liu. 2010. A spectral characterization of the Delaunay triangulation. *Computer Aided Geometric Design* 27, 4 (2010), 295–300. <https://doi.org/10.1016/j.cagd.2010.02.002>
- Zhonggui Chen, Wenping Wang, Bruno Lévy, Ligang Liu, and Feng Sun. 2014. Revisiting Optimal Delaunay Triangulation for 3D Graded Mesh Generation. *SIAM Journal on Scientific Computing* 36, 3 (2014), A930–A954. <https://doi.org/10.1137/120875132>
- Siu-Wing Cheng, Tamal K. Dey, Herbert Edelsbrunner, Michael A. Facello, and Shang-Hua Teng. 2000. Sliver Exudation. *J. ACM* 47, 5 (Sept. 2000), 883–904. <https://doi.org/10.1145/355483.355487>
- Fernando de Goes, Pooran Memari, Patrick Mullen, and Mathieu Desbrun. 2014. Weighted Triangulations for Geometry Processing. *ACM Trans. Graph.* 33, 3 (2014), 28:1–28:13. <https://doi.org/10.1145/2602143>
- Jesus A. De Loera, Jorg Rambau, and Francisco Santos. 2010. *Triangulations: Structures for Algorithms and Applications* (1st ed.). Springer, Berlin, Heidelberg.
- Richard J Duffin. 1959. Distributed and lumped networks. *Journal of Mathematics and Mechanics* 8, 5 (1959), 793–826.
- Gerhard Dziuk. 1988. Finite Elements for the Beltrami operator on arbitrary surfaces. In *Partial Differential Equations and Calculus of Variations*, Stefan Hildebrandt and Rolf Leis (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 142–155. <https://doi.org/10.1007/BFb0082865>
- Herbert Edelsbrunner and Nimish R. Shah. 1992. Incremental Topological Flipping Works for Regular Triangulations. In *Proceedings of the Eighth Annual Symposium on Computational Geometry (SCG '92)*. ACM, New York, NY, USA, 43–52. <https://doi.org/10.1145/142675.142688>
- Leman Feng, Pierre Alliez, Laurent Busé, Hervé Delingette, and Mathieu Desbrun. 2018. Curved Optimal Delaunay Triangulation. *ACM Trans. Graph.* 37, 4, Article 61 (July 2018), 16 pages. <https://doi.org/10.1145/3197517.3201358>
- Jean Gallier and Jocelyn Quaintance. 2017. Aspects of Convex Geometry Polyhedra, Linear Programming, Shellings, Voronoi Diagrams, Delaunay Triangulations. Book in progress, earlier version available as arXiv:0805.0292.
- Gaël Guennebaud, Benoît Jacob, et al. 2010. Eigen v3. <http://eigen.tuxfamily.org>.
- Yixin Hu, Qingnan Zhou, Xifeng Gao, Alec Jacobson, Denis Zorin, and Daniele Panozzo. 2018. Tetrahedral Meshing in the Wild. *ACM Trans. Graph.* 37, 4, Article 60 (July 2018), 14 pages. <https://doi.org/10.1145/3197517.3201353>
- Clément Jamin, Sylvain Pion, and Monique Teillaud. 2018. 3D Triangulations. In *CGAL User and Reference Manual* (4.13 ed.). CGAL Editorial Board. <https://doc.cgal.org/4.13/Manual/packages.html#PkgTriangulation3Summary>
- Barry Joe. 1989. Three-Dimensional Triangulations from Local Transformations. *SIAM J. Sci. Stat. Comput.* 10, 4 (July 1989), 718–741. <https://doi.org/10.1137/0910044>
- Michal Krížek. 1992. On the Maximum Angle Condition for Linear Tetrahedral Elements. *SIAM J. Numer. Anal.* 29, 2 (April 1992), 513–520. <https://doi.org/10.1137/0729031>
- François Labelle and Jonathan Richard Shewchuk. 2007. Isosurface Stuffing: Fast Tetrahedral Meshes with Good Dihedral Angles. *ACM Trans. Graph.* 26, 3, Article 57 (July 2007), 10 pages. <https://doi.org/10.1145/1276377.1276448>
- Charles L. Lawson. 1972. Transforming triangulations. *Discrete Mathematics* 3, 4 (1972), 365–372. [https://doi.org/10.1016/0012-365X\(72\)90093-3](https://doi.org/10.1016/0012-365X(72)90093-3)
- Mark Meyer, Mathieu Desbrun, Peter Schröder, and Alan H. Barr. 2003. Discrete Differential-Geometry Operators for Triangulated 2-Manifolds. In *Visualization and Mathematics III*, Hans-Christian Hege and Konrad Polthier (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 35–57.
- Neil Molino, Robert Bridson, Joseph Teran, and Ronald Fedkiw. 2003. A crystalline, red green strategy for meshing highly deformable objects with tetrahedra. In *Proceedings, 12th International Meshing Roundtable*. Sandia National Laboratories, Santa Fe, NM, USA, 103–114.
- Oleg R. Musin. 1997. Properties of the Delaunay Triangulation. In *Proceedings of the Thirteenth Annual Symposium on Computational Geometry (SCG '97)*. ACM, New York, NY, USA, 424–426. <https://doi.org/10.1145/262839.263061>
- Ulrich Pinkall and Konrad Polthier. 1993. Computing discrete minimal surfaces and their conjugates. *Experim. Math.* 2 (1993), 15–36.
- William H. Press, Saul A. Teukolsky, William T. Vetterling, and Brian P. Flannery. 1992. *Numerical Recipes in C: The Art of Scientific Computing* (second ed.). Cambridge University Press.
- Samuel Rippa. 1990. Minimal Roughness Property of the Delaunay Triangulation. *Comput. Aided Geom. Des.* 7, 6 (Oct. 1990), 489–497. [https://doi.org/10.1016/0167-8396\(90\)90011-F](https://doi.org/10.1016/0167-8396(90)90011-F)
- Jonathan Shewchuk. 2002a. What is a good linear finite element? interpolation, conditioning, anisotropy, and quality measures (preprint). <https://people.eecs.berkeley.edu/~jrs/papers/elemj.pdf>
- Jonathan Richard Shewchuk. 2002b. Delaunay Refinement Algorithms for Triangular Mesh Generation. *Comput. Geom. Theory Appl.* 22, 1-3 (May 2002), 21–74. [https://doi.org/10.1016/S0925-7721\(01\)00047-5](https://doi.org/10.1016/S0925-7721(01)00047-5)
- R. Sibson. 1978. Locally equiangular triangulations. *Comput. J.* 21, 3 (1978), 243–245. <https://doi.org/10.1093/comjnl/21.3.243>
- Jane Tournais, Camille Wormser, Pierre Alliez, and Mathieu Desbrun. 2009. Interleaving Delaunay Refinement and Optimization for Practical Isotropic Tetrahedron Mesh Generation. *ACM Trans. Graph.* 28, 3, Article 75 (July 2009), 9 pages. <https://doi.org/10.1145/1531326.1531381>
- Max Wardetzky, Saurabh Mathur, Felix Kälberer, and Eitan Grinspun. 2007. Discrete Laplace Operators: No Free Lunch. In *Proceedings of the Fifth Eurographics Symposium on Geometry Processing (SGP '07)*. Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, 33–37.
- Jinchao Xu and Ludmil Zikatanov. 1999. A Monotone Finite Element Scheme for Convection-diffusion Equations. *Math. Comput.* 68, 228 (Oct. 1999), 1429–1446. <https://doi.org/10.1090/S0025-5718-99-01148-5>