DOI: 10.1111/cgf.14068
Eurographics Symposium on Geometry Processing 2020
Q. Huang and A. Jacobson
(Guest Editors)

Volume 39 (2020), Number 5

Properties of Laplace Operators for Tetrahedral Meshes

Marc Alexa[1] Philipp Herholz[2] Maximilian Kohlbrenner[1] Olga Sorkine-Hornung[2]

[1] TU Berlin, Germany [2] ETH Zurich, Switzerland

Abstract

Discrete Laplacians for triangle meshes are a fundamental tool in geometry processing. The so-called cotan Laplacian is widely used since it preserves several important properties of its smooth counterpart. It can be derived from different principles: either considering the piecewise linear nature of the primal elements or associating values to the dual vertices. Both approaches lead to the same operator in the two-dimensional setting. In contrast, for tetrahedral meshes, only the primal construction is reminiscent of the cotan weights, involving dihedral angles. We provide explicit formulas for the lesser-known dual construction. In both cases, the weights can be computed by adding the contributions of individual tetrahedra to an edge. The resulting two different discrete Laplacians for tetrahedral meshes only retain some of the properties of their two-dimensional counterpart. In particular, while both constructions have linear precision, only the primal construction is positive semi-definite and only the dual construction generates positive weights and provides a maximum principle for Delaunay meshes. We perform a range of numerical experiments that highlight the benefits and limitations of the two constructions for different problems and meshes.

CCS Concepts

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

1. Introduction

The discrete Laplacian is crucially important in geometry processing. It is commonly defined as a linear operator taking values defined at a discrete set of n points to values on the same n points, i.e., it can be written as a square matrix $\mathbf{L} \in \mathbb{R}^{n \times n}$. For triangle meshes, in the plane or immersed in 3D, Wardetzky et al. [WMKG07] discuss several desirable properties of the discrete operator, such as symmetry, locality, linear precision, maximum principle, and PSD (positive semidefiniteness). De facto, the so-called *cotan* discretization is used almost exclusively, likely because it satisfies all the mentioned desirable properties if the mesh has the Delaunay property, and only lacks the maximum principle if the mesh is not Delaunay.

There are several ways to derive the *cotan* Laplacian on simplicial meshes, and we review them in Section 4. The various approaches fall into two fundamentally different categories:

1. Approaches building on the piecewise linear nature of the triangles and tetrahedra. These include the construction based on FEM with linear basis functions, as well as first order analysis of how the mesh changes when one vertex is moved.
2. Approaches that assume the existence of an orthogonal dual and consider average values assigned to the cells dual to the vertices. In geometry processing this is the basis of discrete exterior calculus (DEC), and in various engineering disciplines deriving discrete differential operators in this way it is known as the finite volume (FV) method.

The essential difference between the two constructions is that the first considers functions over the (simplicial) elements, while the second associates functions to vertices, requiring the definition of cells per vertex. In other words, the first approach is based on primal elements, and the second approach is based on dual elements.

The primal and the circumcentric dual approach yield the same operator for triangle meshes in the plane, yet their direct extensions to simplicial complexes of higher dimensions, notably tetrahedral meshes, turn out to be different. This fact appears to be not widely known. In this paper, we focus on analyzing the different properties of the primal and dual construction for the Laplace operator for tetrahedral meshes. Our contribution lies in showing the following properties:

- The primal and the circumcentric dual construction lead to *different* discrete Laplacians for tetrahedral meshes, i.e., the discrete operator matrices are different (§4), and different sets of properties are satisfied by the operators (§7).
- The edge weights can be constructed by summing up contributions from the tetrahedra incident on an edge (§5), leading to a strong locality property we introduce (§7.1).
- *Both* constructions have *linear precision* (§7.2).
- *Only the primal* construction always yields a *PSD* operator (§7.4).

- *Only the dual* construction guarantees positive edge weights and, thus, a maximum principle for *Delaunay* tetrahedralizations (§7.3).
- This means the dual construction leads to "perfect" Laplacians [WMKG07] for Delaunay tetrahedralizations (similar to the *cotan* for triangulations) while the primal construction requires that all angles are acute.

Our numerical experiments (§8) suggest that accuracy and convergence of discrete Laplacians for tetrahedral meshes more strongly hinge on the Delaunay property than the *cotan* Laplacian in the plane. On Delaunay meshes, the dual construction shows favorable results.

2. Background

The discretization of differential operators has been an active field of research in numerical analysis, computational engineering and computer graphics for decades. Many problems in computer graphics and geometry processing require the discretization of the Laplacian acting on (quantities associated to) the vertices of a triangle mesh [Sor06, SCV14].

Depending on the problem, different discretizations are appropriate and lead to discrete operators with different properties. In order to provide a high level discussion independent of the particular application scenario, Wardetzky et al. [WMKG07] categorize several important properties of the smooth Laplacian and analyze the connection of different constructions of the discrete operator to the continuous counterparts of these properties. Our work is inspired by this approach and offers observations for discrete Laplacians derived for simplicial meshes in dimension higher than 2, particularly tetrahedral meshes.

Our focus is on techniques that are fundamentally first order approximations or based on the linearization of derivatives for an underlying simplicial mesh. These techniques lead to discrete operator matrices with non-zero entries only for vertices and edges of the mesh. We classify them based on the elements used for discretization (primal or dual elements) and provide references along with the derivations presented in Section 4.

When circumcenters of triangles are chosen as the positions for dual vertices, the dual construction of the Laplacian coincides with the primal construction [WMKG07, dGMMD14, SCV14]. This coincidence is commonly exploited in the discussion of operator properties or handling boundary conditions. As we show, this coincidence is a peculiarity of the triangle – we attribute it to the inscribed angle theorem for triangles (§4.3).

The dual construction, however, works for any choice of orthogonal dual [Hir03]. A common approach is to consider weighted Delaunay triangulations and their dual power diagrams [Aur87]. Weighting generates a range of different combinatorial meshes, but also different orthogonal duals for the same combinatorial mesh [Gli07, dGMMD14]. The common primal FEM discretization is generally different from the dual construction in this case. Mullen et al. [MMdGD11] use this space for the optimization of the dual mesh, as the measures of the dual elements affect the properties of differential operators. While they explicitly mention the construction of the Laplacian based on DEC and the connection to *cotan* weights for circumcentric duals, the experiments and discussion focus on triangle meshes. We leave an extension of our analysis to weighted Delaunay tetrahedralizations for future work.

First order discretizations may have drawbacks that could be overcome by using higher order elements in the primal construction. For triangle meshes this has been pointed out for example by Reuter et al. [RWP06]. For tetrahedral meshes, higher order finite elements are often used in simulation and animation [RGTC98, MTPS08, WKS*11, BC14].

It is also possible to generalize the construction of discrete Laplacians to non-simplicial elements, such as polygons instead of triangles [AW11, HKA15, BHKB20]. In 3D, hexahedral elements are a popular choice for FEM; however, it has been recently questioned whether their popularity is justified from the perspective of numerical accuracy [SHD*18]. Extensions to more general polyhedral elements have also been discussed [MKB*08, KBT17, SDG*19].

There are no simple analogies of higher order polynomial functions or non-simplicial primal elements for the dual construction.

3. Setup and overview

Our aim is to construct the matrix $\mathbf{L}$, representing the action of the Laplacian on a function $\mathbf{f} \in \mathbb{R}^n$, defined on the n vertices of the mesh. It is common [WMKG07] to define the operator based on differences along edges:

$$(\mathbf{Lf})_i = \sum_{(i,j)\in\mathcal{M}} w_{ij}(f_j - f_i). \tag{1}$$

This construction implies that the entry w_{ij} is non-zero only if vertices i and j are connected by an edge in the mesh. It also means that the Laplacian vanishes on constants, or, in other words, the rows of $\mathbf{L}$ sum to zero.

As a consequence of (1) we have

$$\mathbf{L}_{ij} = \begin{cases} w_{ij} & \text{for } i \neq j \\ -\sum_{(i,j)\in\mathcal{M}} w_{ij} & \text{else.} \end{cases}$$

For this we assume a triangle or tetrahedral mesh $\mathcal{M}$ is given. We refer to a vertex just by its index, e.g. i, to an edge by the pair (i,j), to a triangle by the triple (i,j,k), and to a tetrahedron by the four vertices (i,j,k,l). When used as subscripts, we drop the commas and parentheses for brevity. Whenever we refer to an element, we assume it is part of the mesh, e.g., (i,j) implies $(i,j) \in \mathcal{M}$.

Vertex positions of the mesh are given as $\mathbf{X} \in \mathbb{R}^{n\times d}$, so that the columns are n-vectors of the components of all vertices, and the i-th row is the position of vertex i in d-dimensional Euclidean space. Note that vertex positions are row vectors, and we refer to them as $\mathbf{x}_i \in \mathbb{R}^{1\times d}$. The dimension d is at least two for triangles and at least three for tetrahedra, but may be higher, in which case we assume the mesh to be immersed.

The *measure* of an element is denoted by $\mathrm{Vol}(\cdot)$, e.g. $\mathrm{Vol}(i,j) = \|\mathbf{x}_i - \mathbf{x}_j\|$ is the length of the edge between vertices i and j, and $\mathrm{Vol}(i,j,k)$ is the area of the triangle (i,j,k). We assume these measures are always positive because the mesh is properly immersed.

The *dual* of an element is denoted as $\star\cdot$, e.g. $\star(i,j)$ is the edge dual to edge (i,j) in a triangle mesh, resp. the face dual to that edge in a tetrahedral mesh. The measures for dual elements are only positive if the dual mesh is immersed. They may become negative if we use the dual mesh even though it contains fold-overs. For instance, if we use the circumcentric dual, this happens whenever the mesh is not Delaunay, as we discuss in more detail in Section 5.

The *mass matrix* $\mathbf{M}$, often required in the discretization of PDEs, accounts for the volume associated with mesh elements. Here we only consider diagonal (lumped) mass matrices that contain volumes for each vertex. A common choice is to split the volume of a full simplex equally among its vertices. For triangle meshes we have

$$\mathbf{M}_{ii} = \frac{1}{3} \sum_{(i,j,k)\in\mathcal{M}} \mathrm{Vol}(i,j,k) \tag{2}$$

and

$$\mathbf{M}_{ii} = \frac{1}{4} \sum_{(i,j,k,l)\in\mathcal{M}} \mathrm{Vol}(i,j,k,l) \tag{3}$$

for tetrahedral meshes. Another common choice is the volume of the dual cell associated with vertex i, i.e.,

$$\mathbf{M}_{ii} = \mathrm{Vol}(\star i). \tag{4}$$

In the following, we first recall the main constructions for the *cotan* Laplacian in the plane and point out how these constructions generalize to tetrahedra (Section 4). We then show how to compute the weights for the dual construction of the tetrahedral Laplacian on a per-cell basis, carefully introducing the necessary signed measures for the dual mesh (Section 5). We believe this is a useful contribution, because it not only leads to a simple and straightforward algorithm[†], but also clearly highlights that the dual construction leads to a discrete operator that is different from the *cotan* generalization. Having established that the primal and dual discrete Laplacians for tetrahedral meshes are different, we analyze their properties in the spirit of Wardetzky et al. [WMKG07] (Section 7), as well as based on numerical experiments (Section 8).

Most of our results for tetrahedra directly generalize to higher dimensional simplical meshes, but for clarity of notation we focus the exposition on the cases of triangles and tetrahedra.

4. Primal and dual construction of the discrete Laplacian

In the following, we recall the primal and dual construction of the Laplacian on the example of triangle meshes and provide the natural generalization to tetrahedral meshes. For the dual construction we first limit ourselves to meshes with an immersed orthogonal dual, in particular, Delaunay meshes. We then explain the well-established fact that the two constructions lead to the same operator for triangles, and how this suggests to apply the dual construction also to non-Delaunay meshes.

[†] We provide the source code for this algorithm in `libigl`-style as supplemental material.

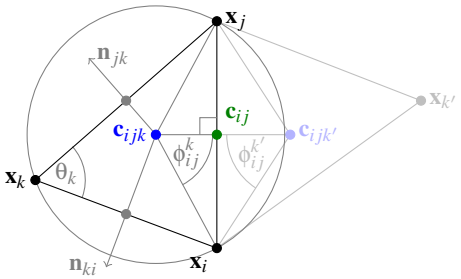

Figure 1: *Illustration of our notations for the various geometric quantities that are used in the contribution of a triangle (i,j,k) to the Laplacian weight w_{ij}. The other triangle incident on edge (i,j), namely (i,j,k'), is also depicted.*

4.1. Primal construction

The two common primal constructions either consider functions defined on vertices or the piecewise linear nature of the immersion of the mesh. We only briefly mention the following different setups of this nature, as all of them are very common in the geometry processing literature. We ask readers to consult the cited work for detailed derivations.

- Construct piecewise linear basis functions over triangles/tetrahedra in FEM. Then use the weak formulation or the Galerkin method for a Poisson problem [Dzi88, XZ99].
- Compute the exact Dirichlet energy of a piecewise linear function [PP93, CPS15, Ale19].
- Consider the mesh immersed into a higher dimension. Then use the fact that the gradient of the area of the triangles or the volumes of the tetrahedra is proportional to $\mathbf{L}$ acting on the vertex positions [DMSB99, MDSB03, Cra19].

The result in all cases is that the coefficient w_{ij} is the sum over contributions from all incident triangles, resp. tetrahedra:

$$\text{triangles: } w_{ij} = \sum_{(i,j,k)} w_{ijk}, \quad \text{tetrahedra: } w_{ij} = \sum_{(i,j,k,l)} w_{ijkl}. \tag{5}$$

This allows computing the weights by looping over the triangles/tetrahedra and summing up their contributions to the edge weights. The contribution of a triangle to the edge (i,j) is

$$w_{ijk} = -\frac{1}{4} \frac{\mathrm{Vol}(j,k)\,\mathrm{Vol}(k,i)\,\mathbf{n}_{jk}^{\mathsf{T}}\mathbf{n}_{ki}}{\mathrm{Vol}(i,j,k)} = \frac{1}{2}\cot\theta_k. \tag{6}$$

See Figure 1 for an illustration. The notation $\mathbf{n}_{ij}$ refers to the unit vector in the plane of the triangle (i,j,k) orthogonal to the edge (i,j), oriented so that it points in the direction opposite of vertex k. The pattern on the left hand side directly generalizes to higher dimension [Ale19, CPS15], the pattern on the right hand side more clearly shows the intrinsic nature of the weights, but its generalization is more subtle [Cra19]. For tetrahedra the formula is:

$$w_{ijkl} = -\frac{1}{9} \frac{\mathrm{Vol}(j,k,l)\,\mathrm{Vol}(k,l,i)\,\mathbf{n}_{jkl}^{\mathsf{T}}\mathbf{n}_{kli}}{\mathrm{Vol}(i,j,k,l)} = \frac{1}{6}\mathrm{Vol}(k,l)\cot\theta_{kl}. \tag{7}$$

Again, $\mathbf{n}_{ijk}$ is the unit vector in the 3-space spanned by the tetrahedron (i,j,k,l), that is orthogonal to the triangle (i,j,k) and oriented in such a way that it points in the direction opposite of vertex

l. For computation of the necessary quantities, in particular if the tetrahedral mesh is immersed in a higher dimensional space, see Appendix A.

4.2. Dual construction

In the dual approach, we start from the observation that $(\mathbf{Lf})_i$ provides the *integrated* Laplacian, i.e., the integral of the Laplacian over a region associated with vertex i [WBH*07]. The common starting point in discrete exterior calculus (DEC) [Hir03] and the two-point flux approximation (TPFA) in finite volume (FV) methods [EGH00] is to assign the region associated with vertex i to be its dual cell $\star i$ and assume the dual graph is *orthogonal* and *immersed*. The Voronoi diagram of the mesh vertices is an orthogonal dual if the mesh has the Delaunay property. The vertices of the Voronoi diagram are the circumcenters of the simplices. For now, assume the mesh is Delaunay – we lift this restriction and work with the circumcentric dual even if it contains fold-overs in Section 5.

We provide the essential derivation of the weights based on these assumptions. Applying Stokes theorem to the integrated Laplacian yields

$$(\mathbf{Lf})_i = \int_{\star i} \Delta f = \int_{\partial \star i} \mathbf{n}^\mathsf{T} \nabla f. \tag{8}$$

We can now decompose the boundary of $\star i$ into the faces dual to the edges (i,j) emanating from vertex i. The normal of each face $\star(i,j)$ is just the normalized edge vector:

$$\int_{\partial \star i} \mathbf{n}^\mathsf{T} \nabla f = \sum_{(i,j)} \int_{\star(i,j)} \frac{\mathbf{x}_j - \mathbf{x}_i}{\mathrm{Vol}(i,j)} \nabla f. \tag{9}$$

Note that the derivation is still exact for any (integrable) function f. In order to compute ∇f on the face $\star(i,j)$ we find two lines of argument with the same result in the literature. Either we assume the unknown functions f are piecewise linear approximations to an unknown exact solution. Or we assume f is the exact solution and we now consider an approximation of the gradient based on linearization along the edge. In either case we have $(\mathbf{x}_j - \mathbf{x}_i)\nabla f = f_j - f_i$, allowing us to compute the integral as

$$\sum_{(i,j)} \int_{\star(i,j)} \frac{\mathbf{x}_j - \mathbf{x}_i}{\mathrm{Vol}(i,j)} \nabla f = \sum_{(i,j)} \mathrm{Vol}(\star(i,j)) \frac{1}{\mathrm{Vol}(i,j)} (f_j - f_i). \tag{10}$$

Comparing to Eq. 1 implies that the coefficients of the matrix $\mathbf{L}$ are

$$w_{ij} = \frac{\mathrm{Vol}(\star(i,j))}{\mathrm{Vol}(i,j)}, \tag{11}$$

or, in words, the measure of the cell dual to edge (i,j) divided by the length of this edge. Note that this derivation is independent of the dimension of the simplicial complex and applies similarly to triangles and tetrahedra. For triangles, $\mathrm{Vol}(\star(i,j))$ is the length of the edge dual to (i,j) and for tetrahedra, $\mathrm{Vol}(\star(i,j))$ is the area of the face dual to this edge.

4.3. Equivalence in the plane

It turns out that the primal and the *circumcentric* dual constructions lead to the same coefficients for planar triangles under our assumptions.

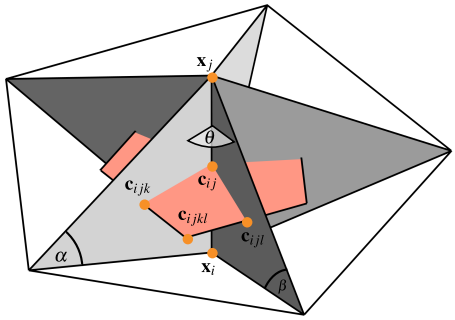

Figure 2: *Computing the measure of the dual $\star(i,j)$ (red) of edge (i,j) can be accomplished on a per-tetrahedron basis. The area can either be expressed with respect to triangle and tetrahedra circumcenters or the dihedral angle θ and the interior triangle angles α and β.*

Let $\mathbf{c}_{ijk}$ be the circumcenter of triangle (i,j,k) (see Figure 1). Then we have

$$\mathrm{Vol}(\star(i,j)) = (\mathbf{c}_{ijk'} - \mathbf{c}_{ijk})^\mathsf{T} \mathbf{n}_{ij}, \tag{12}$$

where k,k' are the vertices opposite of edge (i,j) in the triangles incident on this edge and $\mathbf{n}_{ij}$ the normal of the oriented edge (i,j) (see inset below). We can split the segment between the two circumcenters at the edge midpoint $\mathbf{c}_{ij}$

$$\mathbf{c}_{ijk'} - \mathbf{c}_{ijk} = \mathbf{c}_{ijk'} - \mathbf{c}_{ij} + \mathbf{c}_{ij} - \mathbf{c}_{ijk} = (\mathbf{c}_{ijk'} - \mathbf{c}_{ij}) - (\mathbf{c}_{ijk} - \mathbf{c}_{ij}) \tag{13}$$

even if $\mathbf{c}_{ij}$ is not part of the segment. This leads to

$$\begin{aligned} \mathrm{Vol}(\star(i,j)) &= (\mathbf{c}_{ij} - \mathbf{c}_{ijk'})^\mathsf{T} \mathbf{n}_{ji} + (\mathbf{c}_{ij} - \mathbf{c}_{ijk})^\mathsf{T} \mathbf{n}_{ij} \\ &= r_{ijk'} \cos \phi^{k'}_{ij} + r_{ijk} \cos \phi^{k}_{ij} \\ &= r_{ijk'} \cos \theta_{k'} + r_{ijk} \cos \theta_k, \end{aligned} \tag{14}$$

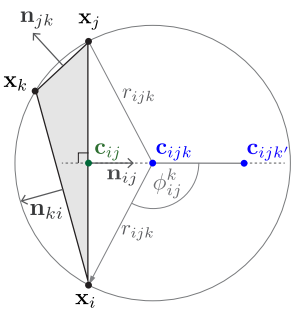

where $r_{ijk}, r_{ijk'}$ are the circumradii and the angles $\phi^k_{ij}, \phi^{k'}_{ij}$ are define by the edges $\mathbf{x}_i - \mathbf{c}_{ijk}, \mathbf{x}_i - \mathbf{c}_{ijk'}$ and the normals $\mathbf{n}_{ij}$. Note that $\mathbf{n}_{ji} = -\mathbf{n}_{ij}$. The angle ϕ^k_{ij} is less than $\pi/2$ if $\mathbf{c}_{ijk}$ is inside the triangle (see Figure 1) and larger than $\pi/2$ if the circumcenter is outside (see the inset). Moreover, ϕ^k_{ij} is identical to θ_k by the inscribed angle theorem: θ_k is half the angle formed by the edge (i,j) with the circumcenter $\mathbf{c}_{ijk}$. The primal edge can be expressed as $2r_{ijk}\sin\theta_k$ as well as $2r_{ijk'}\sin\theta_{k'}$, yielding

$$w_{ij} = \frac{r\cos\phi^k_{ij}}{2r_{ijk}\sin\phi^k_{ij}} + \frac{r'\cos\phi^{k'}_{ij}}{2r_{ijk'}\sin\phi^{k'}_{ij}} = \frac{1}{2}\cot\phi^k_{ij} + \frac{1}{2}\cot\phi^{k'}_{ij}. \tag{15}$$

5. Cell-based assembly for the dual construction

For the dual construction we have so far assumed that the mesh is Delaunay. It turns out that the dual construction leads to an operator that is identical to the one obtained with the primal construction. This suggests that we can use the dual operator for an arbitrary mesh. In the case where an edge in the mesh is not Delaunay, the corresponding dual edge has to be considered inverted, i.e., its measure is prescribed a negative sign. In fact, Eq. 12 already provides the way to compute this signed length for triangulations. Moreover, it suggests that the measure of the dual edge and thus the edge weight can be computed on a per-triangle basis, with the two summands being the appropriately signed contributions of the two triangles to the measure of the dual edge. We denote the contribution of triangle (i,j,k) to the dual edge measure $\mathrm{Vol}(\star(i,j))$ as $\mathrm{Vol}(\star(i,j))_k$. A way to compute this quantity that generalizes to higher dimensions is to consider $\mathrm{Vol}(\star(i,j))_k$ the signed base of the oriented triangle $(\mathbf{x}_i, \mathbf{c}_{ij}, \mathbf{c}_{ijk})$. Noting that $\frac{1}{2}\mathrm{Vol}(i,j)$ is the height of this triangle we find

$$\mathrm{Vol}(\star(i,j))_k = \frac{2\,|\mathbf{c}_{ij}-\mathbf{x}_i,\ \mathbf{c}_{ijk}-\mathbf{x}_i|}{\mathrm{Vol}(i,j)}, \tag{16}$$

where we have used $|\cdot,\cdot|$ to denote a 2×2 determinant.

Now our goal is to derive an analogous computation for the dual construction in the tetrahedral case. Here, $\star(i,j)$ is a planar polygon perpendicular to the edge (i,j), and we want to compute its area from the appropriately signed areas of the quadrilaterals contributed by the tetrahedra incident on (i,j); see Figure 2 for an illustration and notation. It is important to be aware that the quadrilaterals may be outside the tetrahedron, in which case their appropriate sign may be negative or positive, depending on how a quadrilateral intersects the half-spaces defined by the triangles.

We decompose the quadrilateral $(\mathbf{c}_{ij}, \mathbf{c}_{ijk}, \mathbf{c}_{ijkl}, \mathbf{c}_{ijl})$ into two triangles $(\mathbf{c}_{ij}, \mathbf{c}_{ijk}, \mathbf{c}_{ijkl})$ and $(\mathbf{c}_{ij}, \mathbf{c}_{ijkl}, \mathbf{c}_{ijl})$. These triangles provide contributions a_{ijkl} and a_{ijlk} to $\mathrm{Vol}(\star(i,j))$, and we have

$$\mathrm{Vol}(\star(i,j)) = \sum_{(i,j,k,l)} a_{ijkl} + a_{ijlk}, \tag{17}$$

where the sum runs over all tetrahedra incident on (i,j), see Figures 2 and 3. Figure 4 illustrates how the signed areas work in a 2D example: The contribution of two triangles to a dual cell area can be split into the sum of four signed areas, two for each triangle.

The plane of the dual cell is perpendicular to the edge (i,j) and intersects the edge at its midpoint. Similar to the planar case discussed above, we can compute the signed area by considering the triangle forming the base of a tetrahedron and dividing the volume of this tetrahedron by the height over this base. For that, we form two tetrahedra from our two triangles $(\mathbf{c}_{ij}, \mathbf{c}_{ijk}, \mathbf{c}_{ijkl})$ and $(\mathbf{c}_{ij}, \mathbf{c}_{ijkl}, \mathbf{c}_{ijl})$ by adding the vertex $\mathbf{x}_j$ (see Figure 3) and denote their signed volumes as v_{ijkl} and v_{ijlk}, respectively, which can be computed as determinants. We provide source code with this submission, making sure the order of the edge vectors is correct, so that the determinant provides the right sign. The height over the base is, like in the planar case, $\frac{1}{2}\mathrm{Vol}(i,j)$. Therefore, we have

$$a_{ijkl} = 6\frac{v_{ijkl}}{\mathrm{Vol}(i,j)}. \tag{18}$$

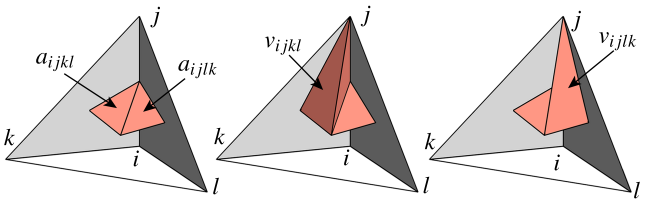

Figure 3: *Contributions of a tetrahedron to $\star(i,j)$ (left). The quadrilateral is made up of two triangles connecting circumcenters of the tetrahedron, adjacent faces and the edge (i,j). The (signed) areas can be computed using the volume of v_{ijkl}, v_{ijlk} the tetrahedra (center, right)*

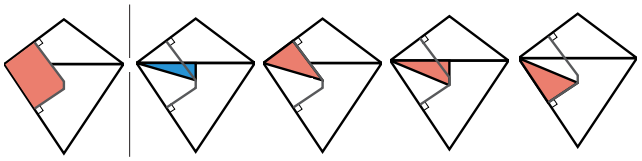

Figure 4: *The parts of a dual cell associated with two adjacent triangles can be computed by summing up signed areas associated to each individual triangle. This also works if a dual edge does not intersect its primal edge, like in this example.*

By plugging into $w_{ij} = \mathrm{Vol}(\star(i,j))/\mathrm{Vol}(i,j)$, we get the per-tetrahedron accumulation formula:

$$w_{ij} = \sum_{(i,j,k,l)} w_{ijkl} + w_{ijlk} \tag{19}$$

$$\text{with}\quad w_{ijkl} = 6\frac{v_{ijkl}}{\mathrm{Vol}(i,j)^2}. \tag{20}$$

This expression reveals the fact that the coefficients of the dual Laplacian are, like for the primal case, rational expressions in vertex coordinates.

The formulation using volumes has an additional advantage: the volumes of the cells $\star i$ dual to vertices can be computed in the same process as

$$\mathrm{Vol}(\star i) = \sum_{(i,j,k,l)} v_{ijkl}. \tag{21}$$

These volumes are necessary for the circumcentric definition of a lumped mass matrix, which turned out to work well in our numerical experiments (see Section 8).

It is also possible to derive an expression for w_{ijkl} that only depends on edge lengths and angles of the tetrahedron, similarly to the *cotan* expression, showing the intrinsic nature of the weights of the dual construction. The derivation is elementary but tedious; we provide it in Appendix B. One expression involving the cotangent of the dihedral angle θ at the edge (i,j), as well as the angles α and β inside the triangles incident on (i,j) opposite this edge is

$$w_{ijkl} = \frac{\mathrm{Vol}(i,j)}{8}\cot\theta\left(2\frac{\cot\alpha\cot\beta}{\cos\theta} - \cot^2\alpha - \cot^2\beta\right). \tag{22}$$

The angles θ, α, β are depicted in Figure 2.

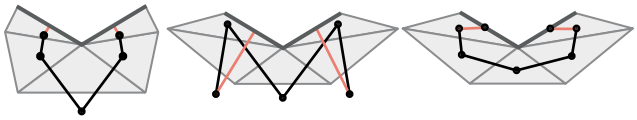

Figure 5: *Constructing the dual edge at the boundary can lead to negative weights whenever circumcenters do not reside inside their primal elements (center). Using barycenters instead can remedy this situation (right). In some situations, however, the dual cells can be extended towards the boundary in a sensible manner (left).*

6. Boundary conditions

For many geometric problems involving the Laplace operator, the discretization of boundary conditions is crucial. This is especially the case for tetrahedral meshes embedded in 3-space, because they always exhibit boundary faces (in contrast to the case of two dimensional meshes immersed in 3-space).

Most commonly we have to handle Dirichlet and Neumann boundary conditions. As an example application we consider minimization of Dirichlet energy (not to be confused with Dirichlet boundary condition, which is a different mathematical construct) over a domain V. To obtain a solution to this minimization problem subject to Dirichlet boundary conditions we have to solve

$$\begin{aligned} &\Delta f = 0 \\ &\text{s.t. } f(x) = g_1(x) \text{ for } x \in \partial V. \end{aligned} \tag{23}$$

Neumann boundary conditions do not explicitly prescribe values at the boundary but rather the derivative of the sought function in normal direction:

$$\begin{aligned} &\Delta f = 0 \\ &\text{s.t. } \mathbf{n}_x^\top \nabla f(x) = g_2(x) \text{ for } x \in \partial V. \end{aligned} \tag{24}$$

Here, $\mathbf{n}_x$ represents the normal at the boundary ∂V.

Implementing Dirichlet boundary conditions Dirichlet boundary conditions can be implemented in a straightforward manner for both, dual and primal discrete Laplacians. Since values at the boundary are prescribed and dual cells are well defined for interior vertices, they can be implemented in the discrete setting by replacing rows of the operator matrix corresponding to boundary vertices by rows of the identity matrix and solving the resulting linear system with an appropriate right hand side containing the constraints.

Neumann boundary conditions: dual construction. Discretization using the dual approach offers a particularly easy treatment of this type of boundary conditions. The central idea of the dual discretization is to employ Stokes' theorem to relate the integral of the Laplacian over a dual cell to the normal derivative integrated over the cell boundary (see Eq. 8). At the mesh boundary dual cells are not closed, and it is not immediately clear how to treat incomplete cells. Neumann boundary conditions, however, prescribe exactly this information. Consequently we can assemble the dual Laplacian from its cells, as described in Section 5, without any special treatment of the mesh boundary. Neumann boundary conditions now appear as the sparse right hand side of the linear system

$$\mathbf{L}\mathbf{f} = -\mathbf{g}.$$

Neumann boundary conditions: primal construction. To implement Neumann boundary conditions for the primal construction we can explicitly integrate the gradient of the piecewise linear function in normal direction at the boundary. Consider a tetrahedron t for which the first face is part of the boundary and face normals and areas are denoted by $\mathbf{n}_i$ and $\mathbf{a}_i$, respectively. The gradient on t of any piecewise linear function is given by a linear mapping [CPS15]:

$$\mathbf{G}_0 = -\frac{a_0}{3\mathrm{Vol}(t)} \mathbf{n}_0^\top \begin{pmatrix} a_0\mathbf{n}_0 & a_1\mathbf{n}_1 & a_2\mathbf{n}_2 & a_3\mathbf{n}_3 \end{pmatrix}. \tag{25}$$

To discretize Eq. 24 using the primal operator, we can formulate the linear equality constraints using Eq. 25, introduce Lagrange multipliers and solve the corresponding KKT system. Due to the coincidence of dual and primal construction in the 2D setting, this explicit handling is not necessary for triangle meshes. In this case boundary conditions are just implemented following the line of argument from the previous paragraph on the dual construction.

Boundary dual cells. From a geometric point of view, the construction of a dual mesh is usually not obvious at the boundary. However, constructing the dual per cell and adding the individual contributions gives a practicable solution to this problem. Dual cells at boundary vertices are completed by projecting circumcenters onto boundary facets. Figure 5 (left) illustrates the situation in 2D. This is one pragmatic solution among many and might introduce problems in several ways. If the mesh is Delaunay, the interior of all tetrahedral circumspheres contains no points, and for convex boundaries this allows for very large circumspheres at the boundary, whose centers can be arbitrarily far away. This results in dual faces of large area that consequently lead to large coefficients in the Laplacian and the mass matrix. Moreover, negative volumes can appear in the mass matrix (Figure 5 (center)) and negative coefficients w_{ij} can lead to a violation of the maximum principle and an indefinite Laplacian (see Section 7). This might not only produce visible artifacts in practice and prevent the use of certain linear solvers but can also hinder the applicability of theoretical results that require certain properties of the Laplacian or a positive definite mass matrix. If this is not desired, we propose to alter the construction for boundary elements. To this end, we replace the circumcenters of the tetrahedra and the faces with centroids if the circumcenters are not located in the interior of their primal objects (see Figure 5). This strategy is reminiscent of the mixed finite element/finite volume strategy described by Meyer et al. [MDSB03] and guarantees a positive mass matrix and positive coefficients w_{ij}.

7. Properties for tetrahedral meshes

It is useful to analyze discrete Laplacians based on a set of properties that the continuous counterpart satisfies [WMKG07]. These properties are: locality, symmetry, linear precision, positive semidefiniteness, and the existence of a maximum principle, induced in the discrete case by all edge weights being non-negative. Since both constructions generate coefficients $w_{ij} = w_{ji}$ that are non-zero only if (i, j) is a mesh edge, they lead to local and symmetric operators. In fact, both constructions satisfy an extended no-

	Primal	Dual	
		Delaunay	non-Delaunay
(Strong) Locality	✓		✓
Linear precision	✓		✓
NSD	✓	(✓)	✗
Max. principle	✗	(✓)	✗

Table 1: *Properties of discrete Laplacians for tetrahedral meshes. The properties of the primal Laplacian are independent of Delaunay-ness. The dual Laplacian is NSD and has a maximum principle for Delaunay meshes without boundary; for meshes with boundary the properties depend on boundary treatment.*

tion of locality, which we discuss below. Interestingly, both constructions have linear precision, while the other properties can be attributed to either construction. A summary of the results is provided in Table 1. In retrospect, this suggests that the good behavior of the *cotan* Laplacian for triangulations is really due to the coincidence of the primal and dual construction.

7.1. Locality

The common notion of locality for discrete Laplacians is related to local support in terms of the mesh graph: for each row i of the matrix $\mathbf{L}$, the indices of the non-zero coefficients correspond to mesh vertices at graph (edge) distance r from vertex i, where r is a fixed constant, independent of the mesh size. In most cases $r = 1$ meaning only coefficients w_{ij} for which (i, j) is an edge of the mesh are non-zero. The *cotan* and the uniform Laplacians are examples of local operators by this definition, while the Laplacian introduce by Hildebrand et al. [HP11] is not.

Interestingly, this formulation fails to imply another natural locality property: a local change to the geometry affects the operator only locally. Concretely, the smooth Laplacian can be defined in terms of the derivatives of the metric tensor. This means the Laplacian is a linear function of the local geometry. An appropriate discrete analog is that the coefficients associated with mesh edges should be linear combinations (with fixed coefficients) of values that can be computed from individual simplices incident on the edge.

This definition of locality is very convenient for the actual computation of the weights, because it allows looping over the tetrahedra and distributing the contributions to the edge weights. In this way, there is no need for data structures that enable constant time access to adjacency information for the mesh elements. Moreover, it also implies that local geometric or combinatorial changes to the mesh have only local effect on the operator matrix: if a vertex is moved, only the adjacent tetrahedra are affected, and so, only the edge weights in the vertex star (edges of the adjacent tetrahedra) can possibly change. Changing the combinatorics, likewise, can only affect the modified region. We make use of this property below when we discuss the effect of flips on the Dirichlet energy.

As we have shown, the primal as well as the dual construction are local in this strong sense. It is worth pointing out that this is not necessarily the case for all constructions of discrete Laplace operators, in particular if they depend on the dual mesh. For example, Herholz et al. [HKA15] perform a global optimization with the goal of finding positive weights via finding an immersed orthogonal dual. This global optimization leads to coefficients that are non-vanishing only on mesh edges, hence local in the usual sense, but non-local in the strong definition above.

7.2. Linear precision

Linear precision means that $\mathbf{Lf}$ is exact on interior vertices if $\mathbf{f}$ is sampled from a linear function. Constant functions are part of the null space by construction. Together with the coordinate functions they form a basis of all linear functions on the mesh. Linear precision can therefore be formulated as

$$\sum_{(i,j)\in\mathcal{M}} w_{ij}(\mathbf{x}_j - \mathbf{x}_i) = \mathbf{0} \tag{26}$$

for all interior vertices i. In particular, this means that interior vertices in a planar triangle mesh, or tetrahedral mesh in $\mathbb{R}^3$, are reproduced by the operator $\mathbf{L}$.

In turns out that both constructions lead to Laplacians with linear precision:

Primal. One way to derive the operator is as the gradient of the volumes enclosed by the tetrahedra [Cra19], similar to the *cotan* Laplacian for triangles derived from the gradient of the area of the triangulation. Clearly, the gradient of the volume for a tetrahedral mesh in a 3-dimensional subspace vanishes for interior vertices.

Dual. Notice that

$$w_{ij}(\mathbf{x}_j - \mathbf{x}_i) = \frac{\mathrm{Vol}(\star(i,j))}{\mathrm{Vol}(i,j)}(\mathbf{x}_j - \mathbf{x}_i) = \mathrm{Vol}(\star(i,j))\mathbf{n}_{\star(i,j)} \tag{27}$$

is the normal vector of the face $\star(i, j)$ weighted with the area of this face, i.e. the area vector of the dual face. The set of all dual faces for the edges emanating from vertex i form the boundary of the cell $\star i$. By Stokes' theorem the area vectors of closed surfaces sum to zero.

7.3. Non-negativity of edge weights

Many differential equations involving the Laplacian satisfy a maximum principle. For example, the heat equilibrium equation attains extremal heat values on the domain boundary. If we discretize this problem, a sufficient but not necessary condition for the discrete solution to have the same property is that all edge weights are positive [WMKG07].

If all edge weights w_{ij} are non-negative, then $-\mathbf{L}$ is a (singular) *M-matrix*, implying a discrete maximum principle for a variety of discrete approximations of differential equations [BP94]. We argue that for equations such as the above-mentioned heat equilibrium problem, the continuous version satisfies a strong maximum principle, meaning that it applies to any subdomain of the solution. This, in turn, really makes the M-matrix property and, thus, non-negativity of edge weights necessary if we want to capture the same behaviour in the discrete setting.

The *cotan* Laplacian for triangles has the following property: an

interior edge has positive edge weight if (and only if) the edge has the (intrinsic) Delaunay property [BS07]. As it turns out, the implication "interior Delaunay-edge ⇒ non-negative weight" is connected to the dual construction of the Laplacian and generalizes in this way to higher dimensions. Indeed, if an edge has the Delaunay property, its dual face is an element of the Voronoi diagram, which is immersed, so the area is non-negative. The converse implication "interior edge with non-negative weight ⇒ Delaunay-edge" is not true anymore for tetrahedral meshes, which can be easily verified by an example.

The geometry of a dual mesh at the boundary is not immediately obvious, and several strategies can be employed. Consequently the sign of the weight for boundary edges depends on the specific technique. If we compute the edge weights by summing up over the contributions from each triangle or tetrahedron, then the weight of boundary edges may be negative, whether they have the Delaunay property or not. This is well-known for triangle meshes and, unsurprisingly, the treatment of the boundary is often not quite clear [CWW13] and the topic of current research [SGWJ18].

How the Delaunay property for interior edges relates to the weights resulting from the primal construction on tetrahedral meshes is unclear. Experimentally, there is no simple connection between the sign of the edge weight and the geometry. The only sufficient condition for positive edge weights we are aware of is that all dihedral angles are acute. This is immediately obvious from Eq. 7. It is well known, on the other hand, that the condition that all dihedral angles are acute is too rigid to be practical for tetrahedral meshes [KPP12]. Therefore the Delaunay property is, in contrast to the two dimensional case, not sufficient for interior edges to have positive weights for the primal construction.

7.4. Positive semidefiniteness

Noting that the Dirichlet energy $\int \|\nabla f\|^2$ is non-negative and approximated by -$\mathbf{f}^T\mathbf{L}\mathbf{f}$ suggests that -$\mathbf{L}$ should be a positive semidefinite (PSD) matrix and consequently the Laplacian $\mathbf{L}$ should be negative semidefinite (NSD).

The primal construction always leads to a matrix $\mathbf{L}$ that is NSD, regardless of the shape of the simplices or boundaries. This simply follows from the observation that the construction yields the exact Dirichlet energy for piecewise linear functions.

There is no obvious connection of the dual construction to NSD. Indeed, we find that for arbitrary tetrahedral meshes the resulting matrix $\mathbf{L}$ may be indefinite. A sufficient condition for NSD in the dual construction is that all edge weights are positive. In this case $-\mathbf{L}$ is diagonally dominant, because its diagonal elements are $\mathbf{L}_{ii} = -\sum_{(i,j)} w_{ij}$. For tetrahedral meshes in $\mathbb{R}^3$ this is difficult even for Delaunay meshes, because the boundary edges may still have negative coefficients unless we employ special construction rules at the boundary, as detailed in Section 6.

8. Experimental results

We conducted a series of numerical experiments to empirically assess the properties of both operators. The first two experiments are based on known properties of smooth solutions for spherical shells or solid balls in 3D. Then we test whether the operators can be robustly computed on real world tetrahedral meshes using standard floating point arithmetic. Lastly, as an analogy of triangle meshes in 3D we consider the tetrahedralization of a 3-sphere in $\mathbb{R}^4$.

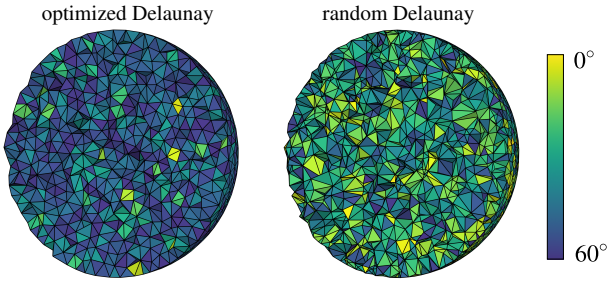

Figure 6: *We run numerical experiments on optimized (left) and random (right) Delaunay meshes. The color shows the minimal dihedral angle per tetrahedron.*

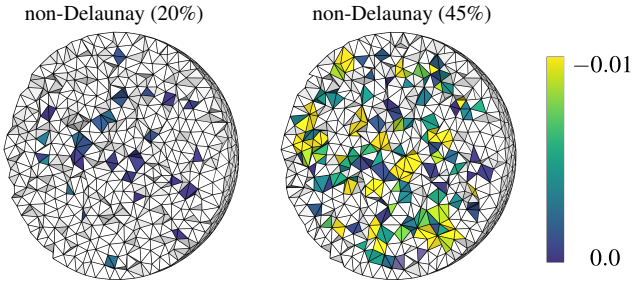

Figure 7: *The dual operator exhibits negative coefficients for non-Delaunay meshes. In this image we color-code the minimal edge coefficient over all edges of the tetrahedra. Positive weights are depicted in white. The mesh on the left contains about* 20% *non-Delaunay tetrahedra and the right one about* 45%.

8.1. Dirichlet energy minimization

In this experiment we consider the solution $\mathbf{f}$ of the Dirichlet problem on a shell of the unit 3-ball. Constraining the inner and outer

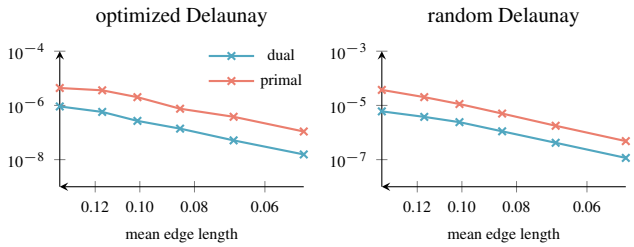

Figure 8: *Variance of the solution of a Dirichlet problem with radial symmetry (Eq. 28) relative to mean edge length for optimized Delaunay meshes (left) and random Delaunay meshes containing slivers (right).*

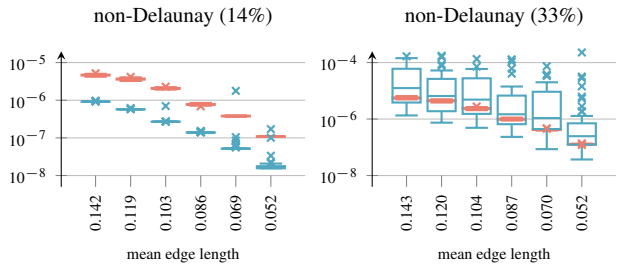

Figure 9: *Variance of the solution for the radially symmetric Dirichlet problem on non-Delaunay meshes relative to mean edge length. Box-plots visualize the behavior for for 50 different random samples of the weights. The effect on mean edge lengths due to changing combinatorics is negligible. A larger fraction of non-Delaunay tetrahedra cause larger variation in the dual Laplacian and also generally larger variance.*

shell to different values, we expect that any interior concentric shell has a constant value. We then take the variance of values on an interior shell as a measure of the quality of the solution.

To generate tetrahedral meshes for this problem we use CGAL [AJR*20] and constrain three spheres at radii $0.5, 0.75$, and 1. This results in tetrahedralizations with sets of vertices on the shells, which we index by $\mathcal{B}_{0.5}, \mathcal{B}_{0.75}$ and $\mathcal{B}_1$. The discrete Dirichlet problem subject to the boundary conditions (c.f. Section 6) is then

$$\begin{aligned} \underset{\mathbf{f}\in\mathbb{R}^n}{\arg\min} \quad & -\mathbf{f}^\mathsf{T}\mathbf{L}\mathbf{f} \\ \text{s.t.} \quad & f_i = 0 \text{ for } i \in \mathcal{B}_1 \\ & f_i = 1 \text{ for } i \in \mathcal{B}_{0.5}. \end{aligned} \tag{28}$$

In the solution, we consider the variance of values at vertices in $\mathcal{B}_{0.75}$:

$$\mathrm{Var}(\{f_i | i \in \mathcal{B}_{0.75}\}).$$

In order to understand the dependence of the variance relative to the characteristics of the mesh we generate different types of tetrahedral meshes. The "nicest" version of the mesh is CGAL's output, to which we refer as *optimized Delaunay*. A considerably worse type of Delaunay mesh results from replacing all vertices that are not on one of the three shells by uniform random samples and then recomputing the Delaunay tetrahedralization – we refer to these meshes as *random Delaunay*. Figure 6 illustrates the mesh quality of a mesh from this set as compared to an optimized Delaunay mesh.

In Figure 8 we show results of the Dirichlet experiment for optimal vs. non-optimal Delaunay meshes (see inset for a representative visualization of the solution). In both cases the variance for the dual operator is about an order of magnitude lower compared to the primal version. Both operators perform better on high quality Delaunay meshes, but we observe convergence for random Delaunay meshes.

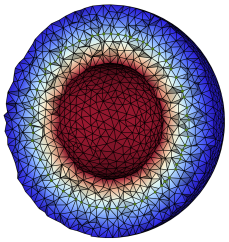

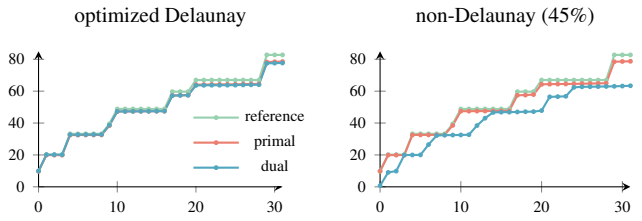

Figure 10: *The smallest 32 eigenvalues (solutions for λ in Eq. (29)) from the reference solution along with discrete approximations computed using dual and primal Laplace operators. For Delaunay meshes, the spectrum of both operators matches the analytical solution quite closely. For the non-Delaunay case, the spectrum diverts from the reference for the dual operator.*

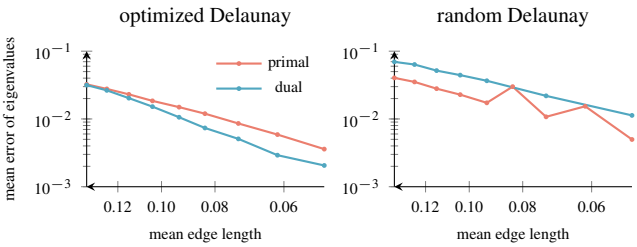

Figure 11: *We measure the mean deviation of the first 32 eigenvalues of the Laplacian from the analytic reference solution. Both the primal and the dual operators empirically converge. For optimized Delaunay meshes the dual operator converges faster while for random Delaunay meshes the primal has an advantage.*

To understand the behaviour of both operators on non-Delaunay meshes, we compute weighted (regular) Delaunay tetrahedralizations using randomly assigned weights. By varying the distribution of weights we obtain meshes with different percentages of non-Delaunay tetrahedra. Figure 7 depicts two meshes from this set. Tetrahedra are color-coded with respect to the smallest coefficients w_{ij} across all six edges; tetrahedra containing only edges with positive coefficients are depicted in white. If only a few tetrahedra are non-Delaunay, the amount of negative coefficients is still relatively low, since there are enough positive contributions to $\mathrm{Vol}(\star(i,j))$. As the number of non-Delaunay tetrahedra increases, we see more negative coefficients emerge.

We found that the numerical solution of the Dirichlet problem on random weighted Delaunay meshes of the spherical shell depends not only on the percentage of non-Delaunay tetrahedra but also on the individual instances for the weight vectors. This has two reasons: (1) The combinatorics of the mesh depend on the weights; (2) even for constant combinatorics the dual mesh change with the weights, and so the coefficients of the dual operator vary. In Figure 9, we capture this effect using box-plots. In general, the dual Laplacian is affected more by non-Delaunay tetrahedra, exhibits larger variation with varying weights and becomes worse than the primal one for a larger fraction of non-Delaunay tetrahedra.

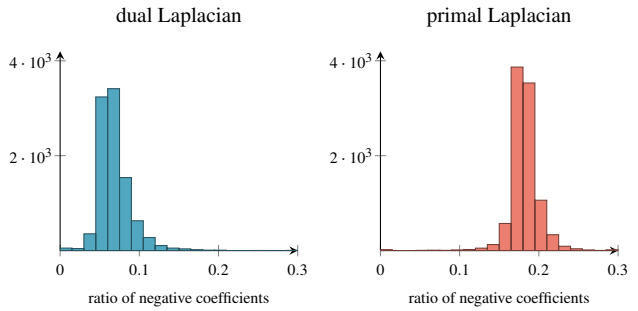

Figure 12: *For each of the 10000 tetrahedral meshes in the dataset provided by Hu et al. [HZG*18], we evaluate the fraction of negative coefficients w_{ij} for the dual and the primal Laplacians and show a histogram of the results. The amount of negative coefficients is systematically smaller for the dual Laplacian.*

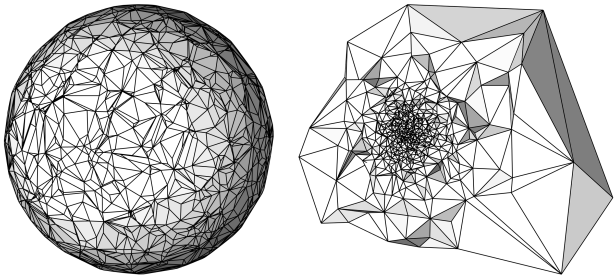

Figure 13: *Visualizations of the tetrahedral convex hull of the random uniform point samples on the unit 3-sphere with 4000 samples. Left image shows a cut along the xyz plane in $\mathbb{R}^4$ resulting in a polyhedral 2-sphere in $\mathbb{R}^3$ consisting of triangles and planar quadrilaterals, depending on how the plane intersects the tetrahedra. The image on the right results from first stereographically projecting the 3-sphere onto the xyz space resulting in a tetrehedral mesh in $\mathbb{R}^3$ and then selecting only the tetrahedra that are fully contained in the negative x halfspace of $\mathbb{R}^3$.*

8.2. 3D eigenmodes

Similar to spherical harmonics, the eigenmodes of a 3-ball $\mathcal{B}^3$ with constrained boundary are known analytically as the solutions to the equation

$$-\Delta f = \lambda f \text{ on } \mathcal{B}^3 \text{ with } f|_{\partial\mathcal{B}^3} = 0. \tag{29}$$

The discrete spectrum can be analytically expressed in terms of the Bessel functions. Given a discrete Laplacian and mass matrix $\mathbf{L}, \mathbf{M} \in \mathbb{R}^{n\times n}$, the discrete analogue to Eq. 29 can be formulated as the generalized eigenvalue problem

$$-\mathbf{L}_{II}\mathbf{f}_I = \lambda \mathbf{M}_{II}\mathbf{f}_I, \tag{30}$$

where the set I references interior vertices and $\mathbf{L}_{II}$ and $\mathbf{M}_{II}$ are constructed by extracting the appropriate rows and columns of $\mathbf{L}$ and $\mathbf{M}$, respectively. The meshes for the discrete problem are constructed similarly to the previous problem, only there is no need for constraining vertices to interior shells, and the whole ball is meshed.

The numerical results capture the analytic values remarkably well, even for relatively small meshes with 3500 vertices, as demonstrated in Figure 10. For non-Delaunay meshes the dual operator is not guaranteed to be negative semidefinite. Consequently, the spectrum might contain negative eigenvalues (Figure 10, right). These additional values do not merely constitute an offset to the spectrum as eigenvalues might be missing or severely distorted. For Delaunay meshes we empirically observe convergence of the lower end of the spectrum. In Figure 11 we illustrate the mean deviation of the 32 smallest eigenvalues with respect to the analytic solution. For optimized Delaunay meshes the dual operator converges faster than the primal one with increasing mesh resolution, while for more general Delaunay meshes we see that the primal operator captures the spectrum more faithfully.

8.3. Numerical robustness

Implementing both operators requires special care in handling degenerate meshes containing elements with volumes that are close to zero. In this context we investigate the question whether one of the operators is more susceptible to almost degenerate situations. To this end, we construct both operators for all 10,000 meshes from the dataset provided by Hu et al. [HZG*18] without explicitly handling degenerate situations. In all cases our code does not generate infinite (NaN or ±INF) double values. This merely hints at the fact that the computation of both operators might be numerically equally stable. The test set is far from representative and does not explicitly contain defective meshes.

On the same test set we also evaluate the fraction of negative coefficients w_{ij}, see Figure 12. Here we only consider the coefficients where the vertices i and j are not both part of the boundary. We can see that, even though the meshes are generally not Delaunay, the dual operator still produces significantly fewer negative coefficients compared to the primal operator.

8.4. Unit 3-sphere in $\mathbb{R}^4$

For a triangle mesh, the result of applying the Laplacian to the coordinate vector yields *mean curvature normals*, i.e.

$$(\mathbf{LX})_i = c\, m_i H_i \mathbf{n}_i, \tag{31}$$

where $\mathbf{n}_i \in \mathbb{R}^d$ is the unit (mean curvature) normal vector at vertex i, H_i is the approximate discrete mean curvature in the region associated with vertex i, and m_i is the area or, more generally, measure of this region. The constant c depends on the intrinsic dimension of the mesh (but not the co-dimension), i.e. $c = 2$ for triangle meshes and $c = 3$ for tetrahedral meshes. In fact, Crane [Cra19] uses this relation to derive the *cotan* Laplacian in arbitrary intrinsic dimension.

We use this relation as a test case by measuring the mean curvatures and normal directions resulting from different unit tetrahedral 3-spheres in $\mathbb{R}^4$, comparing the primal to the dual construction of the operator matrix $\mathbf{L}$.

Generating the tetrahedral 3-spheres requires choosing a set of

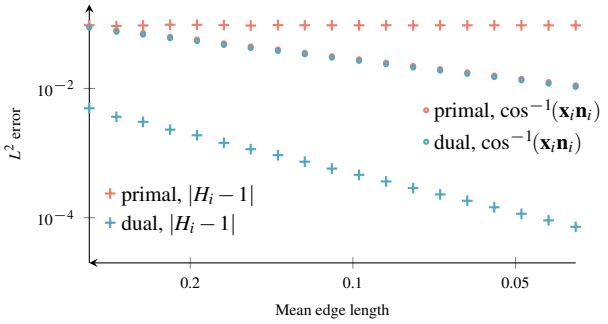

Figure 14: *Convergence behavior of the mean curvatures and angular normal deviation for the primal and dual construction of the discrete Laplacian applied to tetrahedral 3-spheres. Vertices are sampled uniformly from $\mathcal{S}^3$. While normal directions are converging for both operators, point-wise mean curvature estimations converge only for the dual construction of the Laplacian.*

unit vectors in $\mathbb{R}^4$ and a tetrahedral mesh with this vertex set. We focus here on vertex sets that are uniform random samples of the 3-sphere. This can be done by drawing 4 random variates from a normal distribution and then normalizing the resulting 4-vector [Mul59]. Taking the convex hull of the point set obtained in this way provides a tetrahedral mesh (see Figure 13). This mesh is *extrinsically* Delaunay, i.e., all tetrahedra are elements of the Delaunay triangulation of the point set in $\mathbb{R}^4$. The tetrahedral mesh is not necessarily *intrinsically* Delaunay. This situation is similar to considering the convex hull of a 2-sphere in $\mathbb{R}^3$. Nonetheless, we find that the dual construction generates no negative edge weights.

We compute the unit mean curvature normals as

$$\mathbf{n}_i = \frac{(\mathbf{LX})_i}{\|(\mathbf{LX})_i\|} \tag{32}$$

and compare them to the exact normal at $\mathbf{x}_i$ (which is $\mathbf{x}_i$ itself on the unit sphere) based on the angle $\cos^{-1}(\mathbf{x}_i\mathbf{n}_i)$. Figure 14 shows that normal directions converge for both constructions.

For the point-wise curvatures

$$H_i = \frac{\|(\mathbf{LX})_i\|}{3m_i} \tag{33}$$

we have tried different definitions of the volume m_i associated with vertex i. It turns out that the volume of the circumcentric dual cell leads to significantly better results than other options, such as the barycentric cell, or using the non-diagonal mass matrix resulting from the FEM formulation. This is the case for both the primal and the dual construction. The results for $m_i = \mathrm{Vol}(\star i)$ is shown in Figure 14. Only the discrete tetrahedral Laplacian based on the dual construction shows convergence of mean curvature.

We have tried various other ways to define the vertex locations of the mesh, such as using smoothing of the vertex positions or refinement procedures. While different distributions lead to different constants in the errors for point-wise mean curvature, the general behavior that the error is constant for the primal constructions and converges for the dual construction is consistent. Note that it has been established for triangulations that the convergence of point-wise mean curvatures cannot be expected for the *cotan* Laplacian [HPW06], even if the normals or the operator itself do converge.

Similarly to the experiments in 3D, we have generated tetrahedralizations in 4D that are not Delaunay. However, we find that for the constructions we tried, not even the normals of the tetrahedra are converging, so the convergence of the Laplacian or any derived quantities cannot be expected.

9. Discussion

The preceding discussion reveals how the properties discussed by Wardetzky et al. [WMKG07] really emerge from the primal or dual construction of the operator. The coincidence of these constructions for triangles explains, in retrospect, the favorable features of this operator. The distinction explains why we cannot expect a similarly good and simple discretization for tetrahedral meshes.

The conditions for a "perfect" operator, i.e. one that has all desirable properties of its continuous counterpart, however, are similar to the triangle case: the dual construction leads to such a perfect operator on Delaunay meshes, because it has positive edge weights and $-\mathbf{L}$ is therefore PSD, in addition to being local, symmetric, and linearly precise by construction. The edge weights on the boundary are not necessarily positive. This characterization is identical to the situation for triangles.

Accepting that edges at the boundary may have negative weights implies that it is enough that the mesh is *constrained* Delaunay (see the discussion of boundary edges in the context of intrinsic Delaunay triangulations [BS07]). This is convenient because is it the default output of mesh generation tools such as TetGen [Si15].

If PSD for general meshes is required, the primal construction provides the desired operator. This construction, however, only leads to a perfect operator if all dihedral angles are acute – a requirement that is too rigid for practical applications.

The fact that both operators have linear precision for any tetrahedral mesh has interesting consequences: Any affine combination of the primal and dual operator matrix is a local, symmetric discrete Laplacian with linear precision. Varying the affine weights locally gives rise to a large class of operators with these properties.

It seems fruitful to generalize the class of immersed orthogonal meshes by using weighted Delaunay triangulations. In fact, we believe the following statements are true in this context: the primal discrete Laplacian can have positive coefficients only for weighted Delaunay tetrahedralizations, and if it has positive coefficients then they can always be interpreted as those of the dual construction for an appropriately weighted Delaunay tetrahedralization.

Discrete Laplacians are commonly used to compute the Dirichlet energy $-\mathbf{f}^\mathsf{T}\mathbf{Lf}$ of a function $\mathbf{f} \in \mathbb{R}^n$, defined on the vertices of the mesh. It is known that for a given fixed set of points in the plane, the Delaunay triangulation minimizes the Dirichlet energy for any $\mathbf{f}$ among all possible triangulations of the point set [Rip90]. This result no longer holds for tetrahedral meshes and the primal Laplacian [Ale19]. Considering $-\mathbf{f}^\mathsf{T}\mathbf{Lf}$ for the *dual* construction of a tetrahedral Laplacian, simply ignoring that it may become negative, leads to interesting observations: bistellar flips either increase or

decrease the energy, independent of **f**, similar to the case for the primal construction. In fact, this is a consequence of the strong locality property, symmetry, and linear precision – properties that both operators have. The geometric conditions for the flips bear similarity with the Delaunay condition in that for co-spherical points the Dirichlet energy is independent of the combinatorics. Still, Delaunay flips do not necessarily decrease the Dirichlet energy. Rather, we find experimentally that a 3-2 flip always decreases the energy. We leave proving this observation and possibly generalizing Rippa's result to higher dimensions using the dual construction for future work.

The *cotan* Laplacian has been used for several decades on a variety of meshes and generated many useful results. Only recently there is a trend to use the intrinsic Delaunay triangulation [BS07, FSSB07], because of the beneficial properties in some numerical computations [SSC19]. One could say that the *cotan* Laplacian profits from using Delaunay triangulations, but the effect only shows up in specific scenarios. Our experiments with a variety of meshes show that the situation is aggravated for tetrahedral meshes. Here, in all computations we tried, using Delaunay meshes is significantly better than allowing for non-Delaunay tetrahedra. Fortunately, this is not a strong restriction, because the constrained Delaunay triangulations commonly generated by meshing tools are sufficient for the good properties resulting from the dual construction. Constrained Delaunay triangulations only give up the Delaunay property on the boundary and this is where the edge weights are difficult to control in any case.

Acknowledgements

We thank Keenan Crane for insightful comments on the properties of the *cotan* Laplacian on tetrahedral meshes. This research was partially supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy – The Berlin Mathematics Research Center MATH+ (EXC-2046/1, project ID: 390685689) and the Personalized Health and Related Technologies (PHRT) SwissHeart grant.

References

[AJR*20] ALLIEZ P., JAMIN C., RINEAU L., TAYEB S., TOURNOIS J., YVINEC M.: 3D mesh generation. In *CGAL User and Reference Manual*, 5.0.2 ed. CGAL Editorial Board, 2020. 9

[Ale19] ALEXA M.: Harmonic triangulations. *ACM Trans. Graph. 38*, 4 (July 2019). URL: https://doi.org/10.1145/3306346.3322986, doi:10.1145/3306346.3322986. 3, 11, 13

[All65] ALLENDOERFER C. B.: Generalizations of theorems about triangles. *Mathematics Magazine 38*, 5 (1965), 253–259. doi:10.1080/0025570X.1965.11975649. 13

[Aur87] AURENHAMMER F.: Power diagrams: properties, algorithms and applications. *SIAM Journal on Computing 16*, 1 (1987), 78–96. 2

[AW11] ALEXA M., WARDETZKY M.: Discrete laplacians on general polygonal meshes. *ACM Trans. Graph. 30*, 4 (July 2011). URL: https://doi.org/10.1145/2010324.1964997, doi:10.1145/2010324.1964997. 2

[BC14] BARGTEIL A. W., COHEN E.: Animation of deformable bodies with quadratic bézier finite elements. *ACM Trans. Graph. 33*, 3 (June 2014). doi:10.1145/2567943. 2

[BHKB20] BUNGE A., HERHOLZ P., KAZHDAN M., BOTSCH M.: Polygon laplacian made simple. *Computer Graphics Forum 39*, 2 (2020), 303–313. 2

[BP94] BERMAN A., PLEMMONS R. J.: *Nonnegative Matrices in the Mathematical Sciences*. Society for Industrial and Applied Mathematics, 1994. doi:10.1137/1.9781611971262. 7

[BS07] BOBENKO A. I., SPRINGBORN B. A.: A discrete Laplace-Beltrami operator for simplicial surfaces. *Discrete & Computational Geometry 38*, 4 (2007), 740–756. doi:10.1007/s00454-007-9006-1. 8, 11, 12

[CPS15] CHERN A., PINKALL U., SCHRÖDER P.: Close-to-conformal deformations of volumes. *ACM Trans. Graph. 34*, 4 (July 2015), 56:1–56:13. URL: http://doi.acm.org/10.1145/2766916, doi:10.1145/2766916. 3, 6

[Cra19] CRANE K.: The n-dimensional cotangent formula, 2019. URL: https://www.cs.cmu.edu/~kmcrane/Projects/Other/nDCotanFormula.pdf. 3, 7, 10

[CWW13] CRANE K., WEISCHEDEL C., WARDETZKY M.: Geodesics in heat: A new approach to computing distance based on heat flow. *ACM Trans. Graph. 32*, 5 (Oct. 2013). URL: https://doi.org/10.1145/2516971.2516977, doi:10.1145/2516971.2516977. 8

[dGMMD14] DE GOES F., MEMARI P., MULLEN P., DESBRUN M.: Weighted triangulations for geometry processing. *ACM Trans. Graph. 33*, 3 (June 2014). URL: https://doi.org/10.1145/2602143, doi:10.1145/2602143. 2

[DMSB99] DESBRUN M., MEYER M., SCHRÖDER P., BARR A. H.: Implicit fairing of irregular meshes using diffusion and curvature flow. In *Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques* (USA, 1999), SIGGRAPH '99, ACM Press/Addison-Wesley Publishing Co., p. 317–324. URL: https://doi.org/10.1145/311535.311576, doi:10.1145/311535.311576. 3

[Dzi88] DZIUK G.: Finite elements for the beltrami operator on arbitrary surfaces. In *Partial Differential Equations and Calculus of Variations*, Hildebrandt S., Leis R., (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 1988, pp. 142–155. doi:10.1007/BFb0082865. 3

[EGH00] EYMARD R., GALLOUËT T., HERBIN R.: Finite volume methods. In *Techniques of Scientific Computing (Part 3), Solution of Equations in $\mathbb{R}^n$ (Part 3)*, Lions J. L., Ciarlet P., (Eds.), vol. 7 of *Handbook of Numerical Analysis*. Elsevier, 2000, pp. 713–1018. doi:10.1016/S1570-8659(00)07005-8. 4

[Fie11] FIEDLER M.: *Matrices and Graphs in Geometry*. Encyclopedia of Mathematics and its Applications. Cambridge University Press, 2011. doi:10.1017/CBO9780511973611. 13

[FSSB07] FISHER M., SPRINGBORN B., SCHRÖDER P., BOBENKO A. I.: An algorithm for the construction of intrinsic delaunay triangulations with applications to digital geometry processing. *Computing 81*, 2-3 (2007), 199–213. doi:10.1007/s00607-007-0249-8. 12

[Gli07] GLICKENSTEIN D.: A monotonicity property for weighted delaunay triangulations. *Discrete & Computational Geometry 38*, 4 (2007), 651–664. 2

[Hir03] HIRANI A. N.: *Discrete exterior calculus*. PhD thesis, California Institute of Technology, 2003. 2, 4

[HKA15] HERHOLZ P., KYPRIANIDIS J. E., ALEXA M.: Perfect laplacians for polygon meshes. In *Proceedings of the Eurographics Symposium on Geometry Processing* (2015), SGP '15, Eurographics Association, p. 211–218. doi:10.1111/cgf.12709. 2, 7

[HP11] HILDEBRANDT K., POLTHIER K.: On approximation of the Laplace–Beltrami operator and the Willmore energy of surfaces. *Computer Graphics Forum 30*, 5 (2011), 1513–1520. 7

[HPW06] HILDEBRANDT K., POLTHIER K., WARDETZKY M.: On the convergence of metric and geometric properties of polyhedral surfaces. *Geometriae Dedicata 123*, 1 (2006), 89–112. doi:10.1007/s10711-006-9109-5. 11

[HZG*18] HU Y., ZHOU Q., GAO X., JACOBSON A., ZORIN D., PANOZZO D.: Tetrahedral meshing in the wild. *ACM Trans. Graph. 37*, 4 (July 2018), 60:1–60:14. doi:10.1145/3197517.3201353. 10

[KBT17] KOSCHIER D., BENDER J., THUEREY N.: Robust extended finite elements for complex cutting of deformables. *ACM Trans. Graph. 36*, 4 (July 2017). URL: https://doi.org/10.1145/3072959.3073666, doi:10.1145/3072959.3073666. 2

[KPP12] KOPCZYŃSKI E., PAK I., PRZYTYCKI P.: Acute triangulations of polyhedra and $\mathbb{R}^n$. *Combinatorica 32*, 1 (2012), 85–110. doi:10.1007/s00493-012-2691-2. 8

[MDSB03] MEYER M., DESBRUN M., SCHRÖDER P., BARR A. H.: Discrete differential-geometry operators for triangulated 2-manifolds. In *Visualization and Mathematics III* (Berlin, Heidelberg, 2003), Hege H.-C., Polthier K., (Eds.), Springer Berlin Heidelberg, pp. 35–57. 3, 6

[MKB*08] MARTIN S., KAUFMANN P., BOTSCH M., WICKE M., GROSS M.: Polyhedral finite elements using harmonic basis functions. *Computer Graphics Forum 27*, 5 (2008), 1521–1529. doi:10.1111/j.1467-8659.2008.01293.x. 2

[MMdGD11] MULLEN P., MEMARI P., DE GOES F., DESBRUN M.: Hot: Hodge-optimized triangulations. *ACM Trans. Graph. 30*, 4 (July 2011). URL: https://doi.org/10.1145/2010324.1964998, doi:10.1145/2010324.1964998. 2

[MTPS08] MEZGER J., THOMASZEWSKI B., PABST S., STRASSER W.: Interactive physically-based shape editing. In *Proceedings of the 2008 ACM Symposium on Solid and Physical Modeling* (New York, NY, USA, 2008), SPM '08, Association for Computing Machinery, p. 79–89. doi:10.1145/1364901.1364915. 2

[Mul59] MULLER M. E.: A note on a method for generating points uniformly on n-dimensional spheres. *Commun. ACM 2*, 4 (Apr. 1959), 19–20. URL: https://doi.org/10.1145/377939.377946, doi:10.1145/377939.377946. 11

[PP93] PINKALL U., POLTHIER K.: Computing discrete minimal surfaces and their conjugates. *Experim. Math. 2* (1993), 15–36. 3

[RGTC98] ROTH S., GROSS M. H., TURELLO S., CARLS F. R.: A bernstein-bézier based approach to soft tissue simulation. *Computer Graphics Forum 17*, 3 (1998), 285–294. doi:10.1111/1467-8659.00275. 2

[Rip90] RIPPA S.: Minimal roughness property of the delaunay triangulation. *Comput. Aided Geom. Des. 7*, 6 (Oct. 1990), 489–497. doi:10.1016/0167-8396(90)90011-F. 11

[RWP06] REUTER M., WOLTER F.-E., PEINECKE N.: Laplace-beltrami spectra as a shape-dna of surfaces and solids. *Computer-Aided Design 38*, 4 (2006), 342–366. Symposium on Solid and Physical Modeling 2005. doi:https://doi.org/10.1016/j.cad.2005.10.011. 2

[SCV14] SOLOMON J., CRANE K., VOUGA E.: Laplace-beltrami: The swiss army knife of geometry processing. Symposium on Geometry Processing Graduate School (Cardiff, UK, 2014), 2014. 2

[SDG*19] SCHNEIDER T., DUMAS J., GAO X., BOTSCH M., PANOZZO D., ZORIN D.: Poly-spline finite-element method. *ACM Trans. Graph. 38*, 3 (Mar. 2019). URL: https://doi.org/10.1145/3313797, doi:10.1145/3313797. 2

[SGWJ18] STEIN O., GRINSPUN E., WARDETZKY M., JACOBSON A.: Natural boundary conditions for smoothing in geometry processing. *ACM Trans. Graph. 37*, 2 (May 2018). URL: https://doi.org/10.1145/3186564, doi:10.1145/3186564. 8

[SHD*18] SCHNEIDER T., HU Y., DUMAS J., GAO X., PANOZZO D., ZORIN D.: Decoupling simulation accuracy from mesh quality. *ACM Trans. Graph. 37*, 6 (Dec. 2018). URL: https://doi.org/10.1145/3272127.3275067, doi:10.1145/3272127.3275067. 2

[Si15] SI H.: Tetgen, a delaunay-based quality tetrahedral mesh generator. *ACM Trans. Math. Softw. 41*, 2 (Feb. 2015). URL: https://doi.org/10.1145/2629697, doi:10.1145/2629697. 11

[Sor06] SORKINE O.: Differential representations for mesh processing. *Computer Graphics Forum 25*, 4 (2006), 789–807. 2

[SSC19] SHARP N., SOLIMAN Y., CRANE K.: Navigating intrinsic triangulations. *ACM Trans. Graph. 38*, 4 (2019). 12

[WBH*07] WARDETZKY M., BERGOU M., HARMON D., ZORIN D., GRINSPUN E.: Discrete quadratic curvature energies. *Computer Aided Geometric Design 24*, 8 (2007), 499 – 518. Discrete Differential Geometry. doi:https://doi.org/10.1016/j.cagd.2007.07.006. 4

[WKS*11] WEBER D., KALBE T., STORK A., FELLNER D. W., GOESELE M.: Interactive deformable models with quadratic bases in bernstein-bezier-form. *The visual computer 27*, 6-8 (2011), 473–483. 2

[WMKG07] WARDETZKY M., MATHUR S., KÄLBERER F., GRINSPUN E.: Discrete laplace operators: No free lunch. In *Proceedings of the Fifth Eurographics Symposium on Geometry Processing* (Goslar, DEU, 2007), SGP '07, Eurographics Association, p. 33–37. 1, 2, 3, 6, 7, 11

[XZ99] XU J., ZIKATANOV L.: A monotone finite element scheme for convection-diffusion equations. *Math. Comput. 68*, 228 (Oct. 1999), 1429–1446. doi:10.1090/S0025-5718-99-01148-5. 3

Appendix A: Computation of normals and measures for simplices in higher dimension

For tetrahedra immersed in a higher-dimensional space (or generalizations to higher dimensions) we believe the following details on computations are useful, as they are not widely known or published.

Consider the matrix of edge vectors

$$\mathbf{E} = \begin{pmatrix} \mathbf{x}_1 - \mathbf{x}_0 \\ \mathbf{x}_2 - \mathbf{x}_0 \\ \mathbf{x}_3 - \mathbf{x}_0 \end{pmatrix}. \tag{34}$$

When the coordinates $\mathbf{x}_i$ are elements of $\mathbb{R}^3$, we have that $\mathrm{Vol}(i,j,k,l) = \frac{1}{6}\det\mathbf{E}$, and the *area vectors* $\mathrm{Vol}(i,j,k)\mathbf{n}_{ijk}$ are proportional to the columns of $\mathbf{E}^{-1}$ [All65, Fie11, Ale19]. A computationally useful way to handle sub-simplices, i.e. tetrahedra in $\mathbb{R}^4$, is to consider the singular value decomposition

$$\mathbf{E} = \mathbf{U}\begin{pmatrix} \sigma_0 & & & \\ & \sigma_1 & & \\ & & \sigma_2 & \\ & & & 0 \end{pmatrix}(\mathbf{v}_0, \mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3). \tag{35}$$

Then we have $\mathrm{Vol}(i,j,k,l) = \frac{1}{6}\sigma_0\sigma_1\sigma_2$. Moreover, $\mathbf{v}_3$ is the normal of the tetrahedron (i,j,k,l) in $\mathbb{R}^4$. This immediately suggests that the area vectors are the columns of the *pseudoinverse* $\mathbf{E}^+$. This approach based on the SVD of edge vectors generalizes to arbitrary sub-simplices and co-dimensions. Note that it provides no orientation of the co-space, like the cross product for computing the normal from the edge vectors.

Appendix B: Derivation of the intrinsic dual edge weights

For the tetrahedron shown in Figure 2, the contribution of one tetrahedron to the dual $\mathrm{Vol}(\star(i,j))$ is given by the area of the quadrilateral spanned by $\mathbf{c}_{ij}$, $\mathbf{c}_{ijk}$, $\mathbf{c}_{ijkl}$ and $\mathbf{c}_{ijl}$. With α, β and θ we can

calculate its area using edge lengths:

$$d_k = \|\mathbf{c}_{ijk} - \mathbf{c}_{ij}\|, \tag{36}$$
$$d_l = \|\mathbf{c}_{ijl} - \mathbf{c}_{ij}\|, \tag{37}$$
$$h_k = \|\mathbf{c}_{ijkl} - \mathbf{c}_{ijk}\|, \tag{38}$$
$$h_l = \|\mathbf{c}_{ijkl} - \mathbf{c}_{ijl}\|. \tag{39}$$

As in the two dimensional case (see Section 4.3), we get $\angle(\mathbf{c}_{ij}, \mathbf{c}_{ijk}, \mathbf{x}_j) = \alpha$ and $\angle(\mathbf{c}_{ij}, \mathbf{c}_{ijl}, \mathbf{x}_j) = \beta$, and therefore

$$d_k = \frac{\mathrm{Vol}(i,j)}{2}\cot\alpha, \tag{40}$$
$$d_l = \frac{\mathrm{Vol}(i,j)}{2}\cot\beta. \tag{41}$$

By construction $(\mathbf{c}_{ijk}\mathbf{c}_{ij})$ is orthogonal to $(\mathbf{c}_{ijkl}\mathbf{c}_{ijk})$ and $(\mathbf{c}_{ijl}\mathbf{c}_{ij})$ is orthogonal to $(\mathbf{c}_{ijkl}\mathbf{c}_{ijl})$:

$$(\mathbf{c}_{ijk} - \mathbf{c}_{ij}) \cdot (\mathbf{c}_{ijk} - \mathbf{c}_{ijkl}) = 0, \tag{42}$$
$$(\mathbf{c}_{ijl} - \mathbf{c}_{ij}) \cdot (\mathbf{c}_{ijkl} - \mathbf{c}_{ijl}) = 0. \tag{43}$$

Using the fact that the four points $\mathbf{c}_{ij}, \mathbf{c}_{ijk}, \mathbf{c}_{ijkl}, \mathbf{c}_{ijl}$ lie in a plane, we define a coordinate system with $\mathbf{c}_{ij}$ as the origin, one axis going through $\mathbf{c}_{ijk}$ and a second axis perpendicular to this direction in the plane of the four points. In this coordinate system, we express the points of the quadrilateral:

$$\mathbf{c}_{ij} = (0,\ 0), \tag{44}$$
$$\mathbf{c}_{ijk} = (d_k,\ 0), \tag{45}$$
$$\mathbf{c}_{ijl} = (\cos\theta d_l,\ \sin\theta d_l), \tag{46}$$
$$\mathbf{c}_{ijkl} = (d_k,\ h_k). \tag{47}$$

Explicitly calculating the dot products in (42) and (43) using the coordinates established above yields a linear system that allows us to solve for h_k. Analogously, we also get an expression for h_l:

$$h_k = \frac{d_l - d_k\cos\theta}{\sin\theta}, \tag{48}$$
$$h_l = \frac{d_k - d_l\cos\theta}{\sin\theta}. \tag{49}$$

We obtain the following formula for $\mathrm{Vol}(\star(i,j))$:

$$\mathrm{Vol}(\star(i,j)) = \frac{1}{2}\sum_{(i,j,k,l)} d_k h_k + d_l h_l \tag{50}$$
$$= \frac{1}{2}\sum_{(i,j,k,l)} d_k\frac{d_l - d_k\cos\theta}{\sin\theta} + d_l\frac{d_k - d_l\cos\theta}{\sin\theta} \tag{51}$$
$$= \sum_{(i,j,k,l)} \frac{2d_k d_l - (d_k + d_l)^2\cos(\theta)}{2\sin\theta} \tag{52}$$
$$= \sum_{(i,j,k,l)} \frac{\mathrm{Vol}(i,j)^2(2\cot\alpha\cot\beta - (\cot^2\alpha + \cot^2\beta)\cos\theta)}{8\sin\theta} \tag{53}$$
$$= \sum_{(i,j,k,l)} \frac{\mathrm{Vol}(i,j)^2}{8}\cot\theta\Big(2\frac{\cot\alpha\cot\beta}{\cos\theta} - (\cot^2\alpha + \cot^2\beta)\Big). \tag{54}$$