Courbure discrète : théorie et applications

Rencontre organisée par :
Laurent Najman and Pascal Romon

18-22 novembre 2013

Pooran Memari

Geometric Aspects of the Space of Triangulations

Vol. 3, n^{o} 1 (2013), p. 141-150.

<http://acirm.cedram.org/item?id=ACIRM_2013__3_1_141_0>

Centre international de rencontres mathématiques
U.M.S. 822 C.N.R.S./S.M.F.
Luminy (Marseille) France

Actes des rencontres du C.I.R.M.
Vol. **3** n° 1 (2013) 141-150

Geometric Aspects of the Space of Triangulations

Pooran Memari

Abstract

These are the notes of my talk presented in the colloquium on discrete curvature at the CIRM, in Luminy (France) on November 21st, 2013, in which we study the space of triangulations from a purely geometric point of view and revisit the results presented in [21] and [20] **(joint works with Patrick Mullen, Fernando De Goes and Mathieu Desbrun)**. Motivated by practical numerical issues in a number of modeling and simulation problems, we first introduce the notion of a *compatible dual complex* (made out of convex cells) to a primal triangulation, such that a simplicial mesh and its compatible dual complex form what we call a *primal-dual triangulation.* Using algebraic and computational geometry results, we show that for simply connected domains, compatible dual complexes exist only for a particular type of triangulation known as weakly regular. We also demonstrate that the entire space of primal-dual triangulations, which extends the well known (weighted) Delaunay/Voronoi duality, has a convenient, geometric parameterization. We finally discuss how this parameterization may play an important role in discrete optimization problems such as optimal mesh generation, as it allows us to easily explore the space of primal-dual structures along with some important subspaces.

1. Introduction

Mesh generation traditionally aims at tiling a bounded spatial domain with simplices (triangles in 2D, tetrahedra in 3D) so that any two of these simplices are either disjoint or sharing a lower dimensional face. The resulting triangulation provides a discretization of space through both its primal (simplicial) elements *and* its dual (cell) elements. Both types of element are crucial to a variety of numerical techniques, finite element (FE) and finite volume (FV) methods being arguably the most widely used in computational science. To ensure numerical accuracy and efficiency, specific requirements on the size and shape of the primal (typically for FE) or the dual elements (typically for FV) in the mesh are often sought after.

Towards Generalized Primal/Dual Meshes. A growing trend in numerical simulation is the simultaneous use of primal *and* dual meshes: Petrov-Galerkin finite-element/finite-volume methods (FE/FVM, [3, 19, 24]) and exterior calculus based methods [4, 6, 13] use the ability to store quantities on both primal and dual elements to enforce (co)homological relationships in, e.g., Hodge theory. The choice of the dual, defined by the location of the dual vertices, is however not specified a priori. A very common dual to a triangulation in $\mathbb{R}^d$ is the cell complex which uses the circumcenters of each d-simplex as dual vertices. The barycentric dual, for which barycenters are used instead of circumcenters, is used for certain finite-volume computations, but it fails to satisfy both the orthogonality and the convexity conditions on general triangulations.

While the circumcentric Delaunay-Voronoi duality [25, 8] is one of the cornerstones of meshing methods and, as such, has been extensively used in diverse fields, more general dualities are often desired. Building on a number of results in algebraic and computational geometry, in [20] we present a more general primal-dual pairs of complexes, *primal-dual triangulations*, that we briefly describe in the first sections of this note.

Text presented during the meeting "Discrete curvature: Theory and applications" organized by Laurent Najman and Pascal Romon. 18-22 novembre 2013, C.I.R.M. (Luminy).

In the last part of this note, we show how the Weighted-Delaunay/Laguerre duality could help to provide lower error bounds on numerical computations [21]. While most previous meshing methods focused on designing well-shaped primal triangulations *or* dual complexes, we provide a unifying approach to mesh quality based on the placement of primal *and* orthogonal dual elements with respect to each other. In an effort to provide meshes most appropriate for fast, yet reliable computations, we propose functionals on primal-dual mesh pairs that offer formal bounds on the numerical error induced by the use of diagonal Hodge stars. We then demonstrate that meshes that minimize our functionals have desirable geometric and numerical properties. These resulting Hodge-optimized meshes offer a much-needed alternative to the traditional use of barycentric or circumcentric duals in discrete computations. Finally, the resulting set of meshing tools we introduce has wide applications: even when a specific connectivity is needed, some of our contributions can be applied to increase numerical robustness and accuracy of basic operators.

2. Preliminaries

We start by reviewing important notions that we build upon and extend in subsequent sections.

2.1. Complex, Subdivision, and Triangulation.

A **cell complex** in $\mathbb{R}^d$ is a set K of convex polyhedra (called cells) satisfying two conditions:

(1) Every face of a cell in K is also a cell in K, and
(2) If C and C' are cells in K, their intersection is a common face of both.

A **simplicial complex** is a cell complex whose cells are all simplices. The **body** $|K|$ of a complex K is the union of all its cells. When a subset P of $\mathbb{R}^d$ is the body of a complex K, then K is said to be a **subdivision** of P; if, in addition, K is a simplicial complex, then K is said to be a **triangulation** of P. For a set X of points in $\mathbb{R}^d$, a triangulation of X is a simplicial complex K for which each vertex of K is in X. In that case the body of K is the convex hull of X. Let us note that the triangulations we consider in this work, usually coming from a point set, they partition a simply connected domain in $\mathbb{R}^d$ (corresponding to the convex hull of the point set).

Also, in the definition of a triangulation of X, we do not require all the points of X to be used as vertices; a point $\mathbf{x_i} \in X$ is called *hidden* if it is not used in the triangulation. A triangulation with no hidden points is called a *full* triangulation.

2.2. Triangulations in $\mathbb{R}^d$ through Lifting in $\mathbb{R}^{d+1}$.

Let $\mathbf{X} = \{\mathbf{x_1}, \ldots, \mathbf{x_n}\}$ be a set of points in $\mathbb{R}^d$. A simple way of constructing a triangulation of $\mathbf{X}$ is through the following *lifting procedure*: take an arbitrary function $L : \mathbf{X} \longrightarrow \mathbb{R}$ called the *lifting function*; consider the points $(\mathbf{x_i}, L(\mathbf{x_i})) \in \mathbb{R}^{d+1}$, i.e., the points of $\mathbf{X}$ *lifted* onto the graph of L; in the space $\mathbb{R}^{d+1}$, consider $\mathrm{Conv}(L)$ the convex hull of vertical rays $\{(x_i, l)|\ l \geq L(\mathbf{x_i}), l \in \mathbb{R}, x_i \in X\}$; the bounded faces of $\mathrm{Conv}(L)$, i.e. faces which do not contain vertical half lines, form the **lower envelope** of the lifting L. If the function L is generic (see [11] Chap. 7), the orthogonal projection (onto the first d coordinates) of the lower envelope of L partitions the convex hull of $\mathbf{X}$ and produces a triangulation of $\mathbf{X}$.

It is clear that the above lifting procedure may produce triangulations for which not all points of $\mathbf{X}$ are vertices. A triangulation of a set $\mathbf{X}$ of points obtained through lifting is full (i.e., has no hidden points) if and only if all the points $(\mathbf{x_i}, L(\mathbf{x_i}))$ lie on the **lower envelope of** L (or, in other words, if function L can be extended, through linear interpolation in the triangles, to a **convex** piecewise-linear function).

Regular Triangulations : A triangulation obtained by orthogonally projecting the lower envelope of a lifting of $\mathbf{X}$ in $\mathbb{R}^{d+1}$ (onto the first d coordinates) is called a *regular triangulation* ([29], Definition 5.3).

2.3. Weighted Delaunay Triangulations.

A special choice for the lifting function produces the well-known and widely-used Delaunay triangulation (see [26, 22] for properties of Delaunay triangulations, and [25] for numerous applications). Indeed, let X be a set of points in $\mathbb{R}^d$. Consider the *lifting* of the points in X onto the surface of the paraboloid $h(\mathbf{x}) = \|\mathbf{x}\|^2$ in $\mathbb{R}^{d+1}$; i.e., each $\mathbf{x_i} = (a_1, \ldots, a_d) \in X$ gets mapped to $(\mathbf{x_i}, h_i) \in \mathbb{R}^{d+1}$ with $h_i = \|\mathbf{x_i}\|^2 = a_1^2 + \cdots + a_d^2$. Then the orthogonal projection of the lower envelope of this lifting partitions the convex hull of $\mathbf{X}$ and produces a (full) triangulation coinciding with the **Delaunay triangulation** of $\mathbf{X}$.

A regular triangulation (of the convex hull) of a point set $\mathbf{X}$ can now be seen as a generalization of the Delaunay triangulation of $\mathbf{X}$ as follows. We first define a **weighted point set** as a set $(\mathbf{X}, W) = (\mathbf{x_1}, w_1), \ldots, (\mathbf{x_n}, w_n)$, where $\mathbf{X}$ is a set of points in $\mathbb{R}^d$, and $\{w_i\}_{i \in [1,\ldots,n]}$ are real numbers called weights. The **weighted Delaunay triangulation** of $(\mathbf{X}, W)$ is then the triangulation of $\mathbf{X}$ obtained by projecting the lower envelope of the points $(\mathbf{x_i}, \|\mathbf{x_i}\|^2 - w_i) \in \mathbb{R}^{d+1}$. Note that a weighted Delaunay triangulation can now have hidden points.

Notice also that given a lifting function L and its values $l_i = L(\mathbf{x_i})$ at the points of X, one can always define weights to be the difference between the paraboloid and the function L, $w_i = \|\mathbf{x_i}\|^2 - l_i$. We conclude that for simply connected domains (i.e. convex hull of $\mathbf{X}$), the notions of regular triangulations and weighted Delaunay triangulations are *equivalent.* Let us note that as it is shown in [5], this equivalence is no longer true for domains with non trivial topology.

2.4. **Generalized Voronoi Diagrams vs. Weighted Delaunay Triangulation.** Delaunay triangulations (resp., weighted Delaunay triangulations) can also be obtained (or defined) from their dual *Voronoi diagrams* (resp, *power diagrams*). Let $(\mathbf{X}, W) = \{(\mathbf{x_i}, w_i)\}_{i \in I}$ be a weighted point set in $\mathbb{R}^d$. The power of a point $\mathbf{x} \in \mathbb{R}^d$ with respect to a weighted point $(\mathbf{x_i}, w_i)$ (sometimes referred to as the Laguerre distance) is defined as $d^2(\mathbf{x}, \mathbf{x_i}) - w_i$, where $d(.,.)$ stands for the Euclidean distance. Using this power definition, to each x_i we associate its weighted Voronoi region $V(\mathbf{x_i}, w_i) = \{\mathbf{x} \in \mathbb{R}^d |\ d^2(\mathbf{x}, \mathbf{x_i}) - w_i \leq\ d^2(\mathbf{x}, \mathbf{x_j}) - w_j, \forall j\}$. The power diagram of $(\mathbf{X}, W)$ is the cell complex whose cells are the weighted Voronoi regions.

Note that when the weights are all equal, the power diagram coincides with the Euclidean Voronoi diagram of $\mathbf{X}$. Power diagrams are well known to be dual to weighted Delaunay triangulations, as we review next.

The **dual** of the power diagram of $(\mathbf{X}, W)$ is the weighted Delaunay triangulation of $(\mathbf{X}, W)$. This triangulation contains a k-simplex with vertices $\mathbf{x_{a_0}}, \mathbf{x_{a_1}}, \ldots, \mathbf{x_{a_k}}$ in $\mathbf{X}$ if and only if $V(x_{a_0}, w_{a_0}) \cap V(x_{a_1}, w_{a_1}) \cap \cdots \cap V(x_{a_k}, w_{a_k}) \neq \emptyset$, $\forall k \geq 0$. While many other generalization of Voronoi diagrams exist, they do not form straight-edge and convex polytopes, and are thus not relevant here.

3. Compatible Dual Complexes of Triangulations

We now show that the notion of mesh duality can be extended so that the dual complex is defined geometrically, and independently from the triangulation—while the combinatorial compatibility between the triangulation and its dual is maintained.

Definition 1 (Simple Cell Complex)**.** A cell complex K in $\mathbb{R}^d$ is called **simple** if every vertex of K is incident to $d+1$ edges. K is called labeled if every d-dimensional cell of K is assigned a unique label; in this case, we write $K = \{C_1, \ldots, C_n\}$, where n is the number of d-dimensional cells of K, and C_i is the i-th d-dimensional cell.

Definition 2 (Compatible Dual Complex)**.** Let T be a triangulation of a set $\mathbf{X} = \{\mathbf{x_1}, \ldots, \mathbf{x_n}\}$ of points in $\mathbb{R}^d$; and $K = \{C_{i_1}, \ldots, C_{i_n}\}$ be a labeled simple cell complex, i.e. there is a one-to-one correspondence between x_p and C_{i_p}. K is called a **compatible dual complex** of T if, for every pair of points $\mathbf{x_p}$ and $\mathbf{x_q}$ that are connected in T, C_{i_p} and C_{i_q} share a face.

This compatibility between K and T is purely combinatorial, i.e., it simply states that the connectivity between points induced by K coincides with the one induced by T. Notice that the cell C_{i_p} associated to the point $\mathbf{x_p}$, does not necessarily contain $\mathbf{x_p}$ in its interior. Moreover, the edge $[\mathbf{x_p}, \mathbf{x_q}]$ and its *dual* $C_{i_p} \cap C_{i_q}$ are not necessarily *orthogonal* to each other, unlike most conventional geometric dual structures. Consequently, we can generalize the notion of mesh duality through the following definition:

Definition 3 (Primal-Dual Triangulation (PDT))**.** A pair (T, K) is said to form a d-dimensional **primal-dual triangulation** if T is a triangulation in $\mathbb{R}^d$ and K is a compatible dual complex of T. If every edge $[\mathbf{x_p}, \mathbf{x_q}]$ and its *dual* $C_{i_p} \cap C_{i_q}$ are orthogonal to each other, the pair (T, K) is said to form an **orthogonal primal-dual triangulation**.

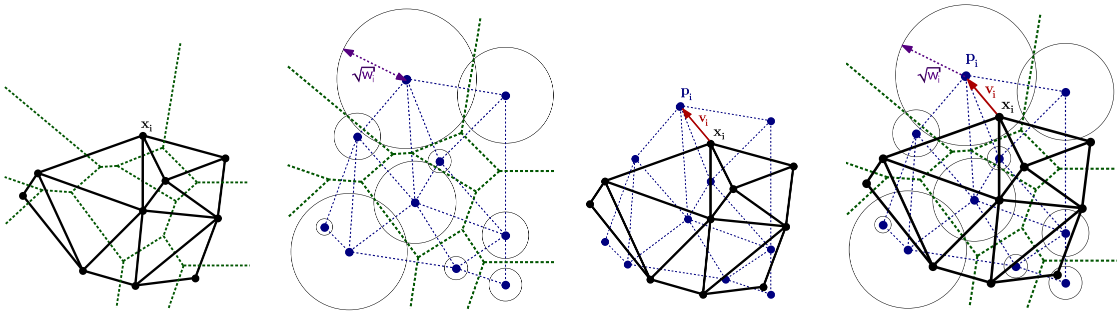

Figure 3.1: *Primal-Dual Triangulation: primal triangulation, dual complex, and combinatorially equivalent regular triangulation separately displayed for clarity.*

3.1. **Characterization of Primal-Dual Triangulations.** An immediate question is whether any triangulation can be part of a PDT. We first characterize the triangulations that admit a compatible dual complex through the following two definitions:

Definition 4 (Combinatorial Equivalence)**.** Two triangulations T and T' are combinatorially equivalent if there exists a labeling which associates to each point $\mathbf{x_i}$ in T a point $\mathbf{x_i'}$ in T' so that the connectivity between $\mathbf{x_i}$'s induced by T matches the connectivity between the $\mathbf{x_i'}$'s induced by T'.

Definition 5 (Combinatorially Regular Triangulations (CRT))**.** A triangulation T of a d-dimensional point set X is called a **combinatorially regular triangulation** if there exists a d-dimensional point set X' admitting a regular triangulation T' such that T and T' are combinatorially equivalent.

Remark: these CRT triangulations have been introduced in [16] under the name of *weakly* regular triangulations, since a displacement of their vertices suffices to make them regular. Figure 3.2 (after [16]) shows an example of a combinatorially regular triangulation which is not, itself, regular.

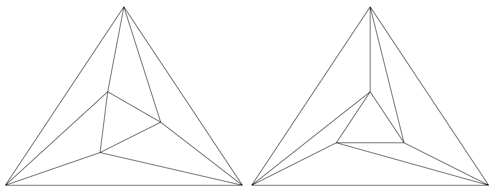

Figure 3.2: A regular triangulation (left), once deformed (right), becomes a combinatorially regular triangulation which is *not*, itself, regular.

Existence of PDTs in 2D. The 2D case is rather simple, due to this result mainly based on a classical theorem of Steinitz [27](see also [16] and [20]):

Proposition 6. *Any 2-dimensional triangulation is combinatorially regular.*

Therefore, every 2D triangulation T can be part of a PDT pair (T, K). However in higher dimensions (three and above), as we showed in [20], the situation is rather different:

Proposition 7. *For $d \geq 3$, there exist d-dimensional triangulations that do not admit any compatible dual complex.*

This is equivalent to the fact that there exist triangulations which are not combinatorially regular. The simplest non-combinatorially regular examples are the Brucker sphere and the Barnette sphere [9].

3.2. **PDT=CRT.** We claim that combinatorially regular triangulations are *the only ones that admit compatible dual complexes.* The proof revolves around a theorem due to Aurenhammer:

Every simple cell complex in $\mathbb{R}^d$, $d \geq 3$, is dual to a regular triangulation.

This theorem was proved in [2] through an iterative construction which is valid in any dimension $d \geq 3$. In [20], we used this theorem to prove the following theorem which surprisingly implies that in higher dimensions there are triangulations that do *not* admit a dual complex:

Theorem 8 (PDT Characterization)**.** *A d-dimensional triangulation T admits a compatible (not necessarily orthogonal) dual complex* if and only if *T is combinatorially regular.*

4. Parameterizing Primal-Dual Triangulations

We have established that primal-dual triangulations cover all dual complexes in $d \geq 3$; they also cover all $2D$ triangulations, but only triangulations which admit a dual in $d \geq 3$. We now focus on parameterizing the whole space of primal-dual triangulations with n points in $\mathbb{R}^d$ by simply adding parameters at the points. We then explore a geometric interpretation of this intrinsic parameterization as well as its properties.

The proof of Theorem 8 leads us very naturally to a parameterization of all the triangulations that admit a compatible dual complex:

Definition 9. A **parameterized primal-dual triangulation** is a primal-dual triangulation parameterized by a set of triplets $(\mathbf{x_i}, w_i, \mathbf{v_i})$, where $\mathbf{x_i}$ is the *position* in $\mathbb{R}^d$ of the ith node, w_i is a real number called the *weight* of x_i, and $\mathbf{v_i}$ is d-dimensional vector called the *displacement vector* of $\mathbf{x_i}$. The triangulation associated with the triplets $(\mathbf{x_i}, w_i, \mathbf{v_i})$ is defined such that its dual complex K is the *power diagram of weighted points* $(\mathbf{p_i}, w_i)$, where $\mathbf{p_i} = \mathbf{x_i} + \mathbf{v_i}$, see Fig 3.1.

The dual complex K can be seen as the generalized Voronoi diagram of the x_i's for the distance $d(\mathbf{x}, \mathbf{x_i}) = \|\mathbf{x} - \mathbf{x_i} - \mathbf{v_i}\|^2 - w_i$. When the vectors $\mathbf{v_i}$ are all null, the parameterized primal-dual triangulation T is regular, thus perpendicular to its dual K, and the pair (T, K) forms an orthogonal primal-dual triangulation. This proves that weighted Delaunay triangulations are sufficient to parameterize the set of all orthogonal primal-dual triangulations of a simply connected domain (see also [12]). The displacement vectors extend the type of triangulations and duals we can parameterize.

Characterizing the classes of *equivalent* triplets parameterizing the same PDT, and completed with constraints for the parameters in order to avoid redundancy between equivalent triplets, we find an efficient parameterization for the space of primal-dual triangulations (see [20] for details):

Theorem 10 (PDT Parametrization)**.** *There is a bijection between all primal-dual triangulations in $\mathbb{R}^d$ and sets of triplets $(\mathbf{x_i}, w_i, \mathbf{v_i})$, $1 \leq i \leq n$, where $\mathbf{x_i}, \mathbf{v_i} \in \mathbb{R}^d$, $w_i \in \mathbb{R}$ with $\sum_i w_i = 0$, $\sum_i \mathbf{v_i} = 0$, and $\sum_i \|\mathbf{x_i} + \mathbf{v_i}\|^2 = \sum_i \|\mathbf{x_i}\|^2$.*

Remark: using this parametrization, the particular case of Delaunay / Voronoi PDT of a set of points $\{\mathbf{x_i}\}_{i=1..n}$ is naturally parameterized by triplets $(\mathbf{x_i}, 0, 0)$. Note also that the condition $\sum_i w_i = 0$ may be replaced by $\min_i w_i = 0$, by simply subtracting the constant $\min_i w_i$ from all the weights of triplets (the PDT depending on weight differences will remain unchanged). This new condition implies that all the weights are positive which may be useful in some applications.

5. Applications in Mesh Optimization

In the previous section, we derived a natural parameterization of all non-orthogonal primal-dual structures of simply connected domains in $\mathbb{R}^d$. Besides the theoretical interest of these new primal-dual structures, we anticipate numerous applications. We believe that our results can benefit mesh optimization algorithms as we provide a particularly convenient way to explore a large space of primal-dual structures. We have already provided a first step in this direction by designing pairs of primal-dual structures that optimize accuracy bounds on differential operators using our parameterization [21], thus extending variational approaches designed to improve *either* primal (Optimal Delaunay Triangulations [28]) *or* dual (Centroidal Voronoi Tesselations) structures. In that work, we introduce Hodge-optimized triangulations (HOT), a family of well-shaped primal-dual pairs of complexes designed for fast and accurate computations in computer graphics. Other existing work most commonly employs barycentric or circumcentric duals: while barycentric duals guarantee that the dual of each simplex lies within the simplex, circumcentric duals are often preferred due to the induced orthogonality between primal and dual complexes. We instead promote the use of weighted duals ("power diagrams"). They allow much greater flexibility in the location of dual vertices while keeping primal-dual orthogonality, thus providing an invaluable extension to the usual choices of dual by only adding one additional scalar per primal vertex. Furthermore, we introduce a family of functionals on pairs of complexes that we derive from bounds on the errors induced by diagonal Hodge stars, commonly used in discrete computations. The minimizers of these functionals, called HOT meshes, are shown to be generalizations of Centroidal Voronoi Tesselations and Optimal Delaunay Triangulations, and to provide increased accuracy and flexibility for a variety of computational purposes. This approach is detailed in the following sections.

6. Hodge Optimized Triangulations

To demonstrate the advantages of using regular/power triangulations, we focus on a particularly relevant type of functionals measuring primal and dual properties. Recall that for an arbitrary primal element σ, the diagonal approximation of the Hodge star $\star$ [4] of a continuous differential form α assumes

$$\int_{*\sigma} \star\,\alpha \equiv \frac{|*\sigma|}{|\sigma|}\int_{\sigma}\alpha, \tag{6.1}$$

where $|.|$ denotes the Lebesgue measure (length, area, volume) of a simplex or cell. In other words, the discrete k^{th} Hodge star is encoded as a diagonal matrix $\star^k$ with $\forall i, \quad (\star^k)_{ii} := \frac{|*\sigma_i^k|}{|\sigma_i^k|}$, where σ_i^k (resp., $*\sigma_i^k$) is the i^{th} k-simplex (resp., k-cell) of the primal-dual triangulation $\mathcal{M} = (\mathcal{T}, \mathcal{D})$; the discrete Hodge star of a discrete primal k-form ω^k is then computed as $\star^k\,\omega^k$, and the extension to dual discrete forms (with, this time, $(\star^k)^{-1}$) is trivial (for further details see, e.g., [6]). We will link approximation error of diagonal Hodge stars to optimal transport.

6.1. **Basics of Optimal Transport.** The optimal transport problem seeks to determine the optimal way to move a pile of dirt M to a hole N of the same volume, where "optimal" means that the integral of the distances by which the dirt is moved is minimal. The notion of "distance" (i.e., cost of transport) may vary based on context. A common distance function defined between probability distributions in $\mathbb{R}^d$ with bounded support is the q-Wasserstein metric, defined as

$$W_q(\mu,\nu) = \left(\inf_{\pi\in\mathcal{P}(\mu,\nu)} \int_{\mathbb{R}^d\times\mathbb{R}^d} \|\,x-y\,\|^q \;\mathrm{d}\pi(x,y)\right)^{1/q}.$$

Let us recall here the Kantorovich-Rubinstein theorem, stating that for two distributions μ and ν with bounded support, the 1-Wasserstein distance between μ and ν can be rewritten as:

$$W_1(\mu,\nu) = \sup_{\substack{\text{continuous } \varphi:S\to\mathbb{R}\\ \text{Lips}(\varphi)\le\lambda}} \frac{1}{\lambda}\int_S \varphi(x)\,\mathrm{d}(\mu-\nu), \tag{6.2}$$

where $\text{Lips}(\varphi)$ represents the Lipschitz constant of function φ. This expression will be useful shortly to link optimal transport and approximation error of diagonal Hodge stars.

6.2. **Deriving Tight Bounds through Optimal Transport.** While computationally convenient, diagonal Hodge stars are obviously not very accurate: they are generally only exact for constant forms. We can quantify the induced inaccuracy of $\star^k$ by defining the *error density* e_i on the dual of the simplex σ_i as the average difference between the discrete approximation and the real Hodge star value:

$$\begin{aligned} e_i :=& \frac{1}{|*\sigma_i|}\left|\frac{|*\sigma_i|}{|\sigma_i|}\int_{\sigma_i}\omega - \int_{*\sigma_i}\star\omega\right| = \left|\frac{1}{|\sigma_i|}\int_{\sigma_i}\omega - \frac{1}{|*\sigma_i|}\int_{*\sigma_i}\star\omega\right| \\ =& \left|\int_{\sigma_i} f(x)\frac{d\sigma_i}{|\sigma_i|} - \int_{*\sigma_i} f(x)\frac{d*\sigma_i}{|*\sigma_i|}\right| = \left|\int_{\sigma_i\cup *\sigma_i} f(x)\left[\frac{d\sigma_i}{|\sigma_i|} - \frac{d\,*\sigma_i}{|*\sigma_i|}\right]\right|, \end{aligned}$$

where $f(x)$ is the component of ω on $d\sigma_i$. We deduce, using Eq.(6.2), that the *tightest bound* one can find on the Hodge star error density per simplex for an arbitrary λ-Lipschitz form is simply λ times the minimum cost over all transport plans between σ_i (seen as a uniform distribution over the mesh element that integrates to one) and $*\sigma_i$ (also seen as a uniform distribution integrating to one); i.e., with a slight abuse of notation,

$$e_i \le \lambda\, W_1(\sigma_i, *\sigma_i). \tag{6.3}$$

This formally establishes a link between Hodge star accuracy and optimal transport. Note that we only required ω to be Lipschitz continuous, a reasonable assumption in most graphics applications.

6.3. **Error Functionals on Meshes.** From these local error densities, we can assemble a total error by taking the $L_{p\geq 1}$ integral norm of the error over the mesh area by integrating the p^{th} power of the error density over each convex hull[1] of σ_i and its dual $*\sigma_i$. This directly yields:

$$E_p(\mathcal{M},\star^k) = \left(\sum_{\sigma_i\in\Sigma^k}\ \int_{\mathrm{CH}(\sigma_i\cup *\sigma_i)} {e_i}^p\right)^{\frac{1}{p}} = \left(\sum_{\sigma_i\in\Sigma^k} \frac{|\sigma_i||*\sigma_i|}{\binom{k}{d}}\ {e_i}^p\right)^{\frac{1}{p}},$$

since the convex hull $\mathrm{CH}(\sigma_i \cup *\sigma_i)$ is, up to a dimension factor, simply the product of the primal and dual volumes due to our primal/dual orthogonality assumption of mesh $\mathcal{M}$. Note that these convex hulls, coined "support volumes in" [15] and "diamonds" in [14, 6], tile the whole primal mesh, thus providing a proper volume integral.

From Eq. (6.3), we conclude that a tight bound for the p^{th} power of the total error is expressed as:

$$E_p(\mathcal{M},\star^k)^p \leq \frac{\lambda}{\binom{k}{d}} \sum_{\sigma_i\in\Sigma^k} |*\sigma_i||\sigma_i|\ W_1(\sigma_i,*\sigma_i)^p. \tag{6.4}$$

(Notice that $E_\infty(\mathcal{M},\star^k)$ is thus, up to the Lipschitz constant, bounded by the maximum of the minimum W_1 distance between primal and dual elements of the mesh as expected.) For notational convenience, we will denote by $\star^k\text{-}\mathrm{HOT}_{p,1}(\mathcal{M})$ the bound (with Lipschitz and dimension constants removed) obtained in Eq. (6.4); more generally, we will define

$$\boxed{\star^k\text{-}\mathrm{HOT}_{p,q}(\mathcal{M}) \equiv \sum_{\sigma_i\in\Sigma^k} |*\sigma_i||\sigma_i|\ W_q(\sigma_i,*\sigma_i)^p}$$

as relevant *functionals* (or energies) to construct meshes, as minimizing them will control the quality of the discrete Hodge stars. Let us note that most of our HOT energies are evaluated by splitting simplices/cells into canonical subsimplices for which closed-forms integral expressions $W(\mathbf{p},T)$ of simplex-T-to-point-$\mathbf{p}$ transport are easily found.

Discussion: Our HOT energies are archetypical, general-purpose examples of mesh quality measures imposed on both primal *and* dual meshes, but they are by no means unique: from the local error densities e_i, other energies can be formulated to target more specific errors occurring in mesh computations. Note also that the use of a 1-Wasserstein distance is notably less attractive numerically than a 2-Wasserstein distance. Fortunately, we can also provide a bound of the Hodge star error which, while less tight than the previously derived $\mathrm{HOT}_{p,1}$, will be particularly convenient to deal with computationally:

$$E_2(\mathcal{M},\star^k)^2 \leq \sum_{\sigma_i\in\Sigma^k} |*\sigma_i||\sigma_i|\ W_2(\sigma_i,*\sigma_i)^2 \equiv \star^k\text{-}\mathrm{HOT}_{2,2}(\mathcal{M}).$$

The reader may have noticed that the functional $\star^0\text{-}\mathrm{HOT}_{2,2}(\mathcal{M})$ is, in the case of equal weights, the well-known Centroidal Voronoi Tesselation (CVT) energy $(\sum_i\int_{V_i} \| \mathbf{x}-\mathbf{x}_i \|^2\, dV)$ for which several minimization techniques, from Lloyd iterations [7] to quasi-Newton methods [18], have been developed. L_p variants (i.e., $\star^0\text{-}\mathrm{HOT}_{2p,2}(\mathcal{M})$ for $p\geq 2$) were also explored recently [17]. However, these energies only correspond to $\star^0$, and are not as tight as $\mathrm{HOT}_{1,p}$. Our HOT energies can thus be seen as a direct generalization of the CVT-like functionals. Note finally that the Optimal Delaunay Triangulation (ODT) energy used in [1] can also be seen as a variant of $\star^d\text{-}\mathrm{HOT}_{2,2}(\mathcal{M})$ in $\mathbb{R}^d$ for which the dual mesh is restricted to be "barycentric"; alas, the resulting mesh will not necessarily lead to an orthogonal primal-dual triangulation—even if the resulting simplices were proven to be very close to isotropic.

[1]The reader may notice that when σ_i and $*\sigma_i$ do *not* intersect, the integration domain is no longer a convex set, but a signed union of subsimplices. For simplicity, we will still use the term "convex hull".

6.4. **General Minimization Procedure.** Given that both (continuous) vertex positions and (discrete) mesh connectivity need to be optimized, the task of finding HOT meshes is seemingly intractable. Thankfully, regular triangulations provide a good *parameterization* of the type of primal-dual meshes we wish to explore: one can simply *continuously optimize* both positions and associated weights to find a HOT mesh. However, HOT energies are not convex in general, and a common downfall of non-convex optimization is its propensity to settle into local minima. In our case, this can be nicely alleviated by starting the optimization from a mesh relatively near the optimal solution, i.e., for which the vertices are well spread out. We thus run, after initializing the domain with uniformly sampled vertices over the domain, a few iterations of CVT [7] or ODT [1] to quickly disperse the vertices and get mesh elements roughly similar in size. Once a good initial mesh has been found, we perform a gradient descent, or alternatively, an L-BFGS algorithm [23]. A linear search is performed to adapt the step size along the gradient or the quasi-Newton direction. This common minimization procedure works quite well without requiring anything else but an evaluation of our HOT energies and their gradients. Note that since the positions $\mathbf{x}_i$ and the weights w_i have very different scales, we proceed by alternatively minimizing our HOT energies with respect to vertex positions and weights.

7. Results

HOT meshes can be beneficial in a number of contexts in modeling of surfaces and volumes, as well as in simulation. A particularly common linear operator in mesh processing is the Laplacian Δ, be it in the plane or on a discrete surface. As its Discrete Exterior Calculus (DEC) expression for 0-forms is $\Delta = d_0^t \star^1 d_0$, our HOT energies for $\star^1$ should be particularly adapted to its accurate computation: the d_0 operator being exact, the only loss of accuracy rises from this particular Hodge star. A $\star^1$-HOT$_{2,2}$ mesh indeed results, on a discretization of a simple test domain with 200 vertices, in a 5% reduction of the condition number of the Laplacian matrix with Dirichlet boundary conditions compared to a CVT mesh. The result is much more dramatic for the Laplacian of *dual* 0-forms, where the condition number drops from 254 to 90 on the same example.

To some extent, our approach can even help to deal with situations where the primal triangulation is given and cannot safely be altered: for instance, moving vertices and/or changing the connectivity of a triangle mesh in $\mathbb{R}^3$ is potentially harmful, as it affects the surface shape. Still, the ability to optimize weights to drive the selection of the dual mesh is very useful. We can easily optimize primal-dual triangulations (meshes) by minimizing a functional (energy) with respect to weights. The connectivity is kept intact, regardless of the weights—only the position and shape of the compatible dual is optimized. Our 2D and 3D experiments [21] show that only optimizing the weights is particularly simple and beneficial on a number of meshes. Fig. 7.1 depicts a triangle mesh of a hand and its intrinsic dual before and after weight optimization, showing a drastic reduction in the number of negative dual edges—thus providing a practical alternative to the use of intrinsic Delaunay meshes advocated in [10].

As another illustrative example, Fig. 8.1 shows that even an optimized Delaunay triangulation (ODT mesh, 195K tets, 36K vertices)) with exceptionally high-quality tetrahedra [28] can be made significantly better centered (i.e., with dual vertices closer to the inside of their associated primal simplex) using a simple weight optimization. Note also that in this example the number of tetrahedra with a dual vertex outside of the primal tet dropped from 17041 on the ODT mesh to 5489 on the optimized mesh—a two third reduction of *outcentered* tetrahedra.

8. Conclusion

We introduced the notion of compatible dual complex for a given triangulation in $\mathbb{R}^d$, and discussed the conditions under which an arbitrary triangulation of a simply connected domain admits a compatible, possibly non-orthogonal dual complex. Note that our only requirement for the dual is that it is made out of *convex polytopes*, thus reducing the space of possible primal-dual pairs to a computationally-convenient subset for which basis functions and positive barycentric coordinates are easily defined. We also pointed out a link to a previously-introduced notion of weakly regular triangulation by Lee in the nineties, and that there are triangulations that do *not* admit a dual complex. In addition, the parameterization we derived for all non-orthogonal primal-dual structures, provides a particularly convenient way to explore a large space of triangulations

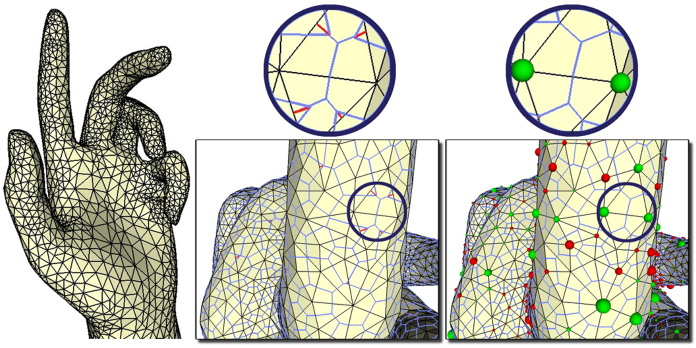

Figure 7.1: **Improving Dual Structure of a Surface Mesh**: For a given triangular mesh (left) there are several triangles whose circumcenter is far outside the triangle (center, lines drawn in red). By optimizing only the weights the new dual vertices are better placed inside the unchanged triangles (right) while keeping primal/dual orthogonality.

for which we anticipate numerous applications. As we have shown in the last sections of this note, our results can in particular benefit mesh optimization algorithms. In addition to applications in mesh optimization, modeling (as in computational biology) that uses *convex space tilings* could directly use our parameterization of PDTs. Clustering techniques based on k-means may also benefit from parameterizing clusters by more than just centers, as weights and vectors add more flexibility to the segmentation of input data.

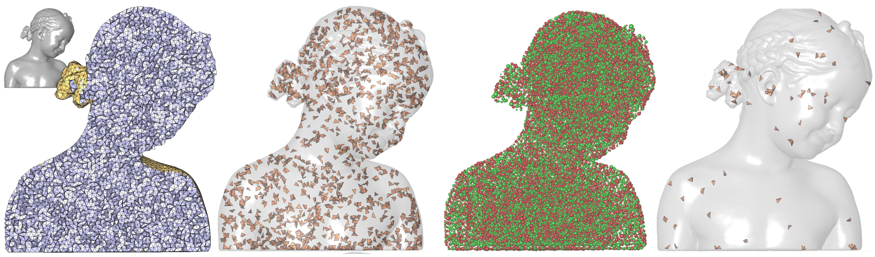

Figure 8.1: **Improving Dual Structure of 3D Meshes**: A the dual of a high-quality ODT mesh of the Bimba con Nastrino (a) can be optimized in terms of Hodge star operator's accuracy; by improving minimal dual edge length and self-centeredness. (c) Weights are displayed according to sign (red/green) and magnitude. When we single out the tetrahedra with a distance between weighted circumcenter and barycenter greater than 0.5% of the bounding box, one can see the optimized mesh (d) is significantly better than the original ODT (b), even if the primal triangulations are exactly matching.

References

[1] Pierre Alliez, David Cohen-Steiner, Mariette Yvinec, and Mathieu Desbrun. Variational tetrahedral meshing. *ACM Trans. on Graphics (SIGGRAPH)*, 24(3):617–625, July 2005.

[2] F. Aurenhammer. A criterion for the affine equivalence of cell complexes in $\mathbb{R}^d$ and convex polyhedra in $\mathbb{R}^{d+1}$. *Discrete and Computational Geometry*, 2(1):49–64, 1987.

[3] B. R. Baligaa and S. V. Patankarb. A Control Volume Finite-Element Method For Two-Dimensional Fluid Flow And Heat Transfer. *Numerical Heat Transfer*, 6:245–261, 1983.

[4] Alain Bossavit. *Computational Electromagnetism.* Academic Press, Boston, 1998.

[5] Fernando De Goes, Pierre Alliez, Houman Owhadi, Mathieu Desbrun, et al. On the equilibrium of simplicial masonry structures. *ACM Transactions on Graphics*, 32(4), 2013.

[6] Mathieu Desbrun, Eva Kanso, and Yiying Tong. Discrete differential forms for computational modeling. In A. Bobenko and P. Schröder, editors, *Discrete Differential Geometry.* Springer, 2007.

[7] Qiang Du, Vance Faber, and Max Gunzburger. Centroidal voronoi tessellations: Applications and algorithms. *SIAM Rev.*, 41:637–676, December 1999.

[8] H. Edelsbrunner. *Algorithms in Combinatorial Geometry.* Springer-Verlag, 1987.

[9] G. Ewald. *Combinatorial convexity and algebraic geometry.* Springer Verlag, 1996.

[10] Matthew Fisher, Boris Springborn, Alexander I. Bobenko, and Peter Schröder. An algorithm for the construction of intrinsic delaunay triangulations with applications to digital geometry processing. In *ACM SIGGRAPH Courses*, pages 69–74, 2006.

[11] I.M. Gelfand, M.M. Kapranov, and A.V. Zelevinsky. *Discriminants, resultants, and multidimensional determinants.* Springer, 1994.

[12] D. Glickenstein. Geometric triangulations and discrete Laplacians on manifolds. *Arxiv preprint math/0508188*, 2005.

[13] Leo J. Grady and Jonathan R. Polimeni. *Discrete Calculus: Applied Analysis on Graphs for Computational Science.* Springer, 2010.

[14] P. Hauret, E. Kuhl, and M. Ortiz. Diamond Elements: a finite element/discrete-mechanics approximation scheme with guaranteed optimal convergence in incompressible elasticity. *Int. J. Numer. Meth. Engng.*, 72(3):253–294, 2007.

[15] A. N. Hirani. *Discrete Exterior Calculus.* PhD thesis, Caltech, May 2003.

[16] C.W. Lee. Regular triangulations of convex polytopes. *Applied Geometry and Discrete Mathematics–The Victor Klee Festschrift (P. Gritzmann, B. Sturmfels, eds.), DIMACS Series in Discrete Mathematics and Theoretical Computer Science, Amer. Math. Soc*, 4:443–456, 1991.

[17] Bruno Lévy and Yang Liu. L_p Centroidal Voronoi Tesselation and its Applications. *ACM Trans. on Graph.*, 29(4), 2010.

[18] Yang Liu, Wenping Wang, Bruno Lévy, Feng Sun, DongMing Yan, Lin Lu, and Chenglei Yang. On Centroidal Voronoi Tessellation - energy smoothness and fast computation. *ACM Trans. on Graph.*, 28(4), 2009.

[19] S. F. McCormick. *Multilevel Adaptive Methods for Partial Differential Equations.* SIAM, 1989.

[20] Pooran Memari, Patrick Mullen, and Mathieu Desbrun. Parametrization of generalized primal-dual triangulations. In *Proceedings of the 20th International Meshing Roundtable*, pages 237–253. Springer, 2012.

[21] Patrick Mullen, Pooran Memari, Fernando de Goes, and Mathieu Desbrun. Hot: Hodge-optimized triangulations. *ACM Transactions on Graphics (TOG)*, 30(4):103, 2011.

[22] O.R. Musin. Properties of the Delaunay triangulation. In *Symposium on Computational Geometry*, page 426, 1997.

[23] J. Nocedal and S. J. Wright. *Numerical optimization.* Springer Verlag, 1999.

[24] A. Paluszny, S. Matthäi, and M. Hohmeyer. Hybrid finite element finite volume discretization of complex geologic structures and a new simulation workflow demonstrated on fractured rocks. *Geofluids*, 7:186–208, 2007.

[25] F. P. Preparata and M. I. Shamos. *Computational Geometry: An Introduction.* Springer-Verlag, 1985.

[26] VT Rajan. Optimality of the Delaunay triangulation in $\mathbb{R}^d$. *Discrete and Computational Geometry*, 12(1):189–202, 1994.

[27] E. Steinitz. Polyeder und raumeinteilungen. *Encyclopädie der mathematischen Wissenschaften*, 3(9):1–139, 1922.

[28] Jane Tournois, Camille Wormser, Pierre Alliez, and Mathieu Desbrun. Interleaving delaunay refinement and optimization for practical isotropic tetrahedron mesh generation. *ACM Trans. Graph.*, 28:75:1–75:9, July 2009.

[29] G.M. Ziegler. *Lectures on polytopes.* Springer, 1995.

CNRS-LTCI, Télécom-ParisTech • memari@telecom-paristech.fr