# HOT: Hodge-Optimized Triangulations

Patrick Mullen     Pooran Memari     Fernando de Goes     Mathieu Desbrun
Caltech

## Abstract

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. Previous 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 greater flexibility in the location of dual vertices while keeping primal-dual orthogonality, thus providing a valuable 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.

**Keywords:** Optimal triangulations, Discrete Exterior Calculus, Discrete Hodge Star, Optimal Transport.

**Links:** DL PDF WEB

## 1 Introduction

A vast array of modeling and simulation techniques assume that a mesh is given, providing a discretization of a 2D or 3D domain in simple triangular or tetrahedral elements. As the accuracy and stability of most computational endeavors heavily depend on the shape and size of the worst element [Shewchuk 2002], mesh element quality is often a priority when conceiving a mesh generation algorithm. Be it for finite-volume, finite-element, finite-difference, or less mainstream computational schemes, the need for good triangle or tetrahedron meshes is ubiquitous not only in computer graphics, but in computational sciences as well—and as computational power increases, so does the demand for effective meshing.

While generically "good" dual or primal elements can be obtained via Centroidal Voronoi Tesselations [Du et al. 1999] or Optimal Delaunay Triangulation [Alliez et al. 2005] respectively, an increasing number of numerical methods need strict control over *both* primal and dual meshes: from discrete differential operators in modeling (e.g., [Meyer et al. 2003]) to pressure solves in fluid simulation (as recently mentioned in [Batty et al. 2010]), the placement of primal elements with respect to their orthogonal dual elements is increasingly recognized as crucial to reliable computations. However, very little is available to quickly and effectively design such orthogonal primal-dual structures over complex domains. To address this lack of adequate meshing tools, we introduce a theoretical analysis of what makes a mesh and its dual numerically optimal in some common graphics contexts, along with practical algorithms to produce optimized primal-dual triangulations.

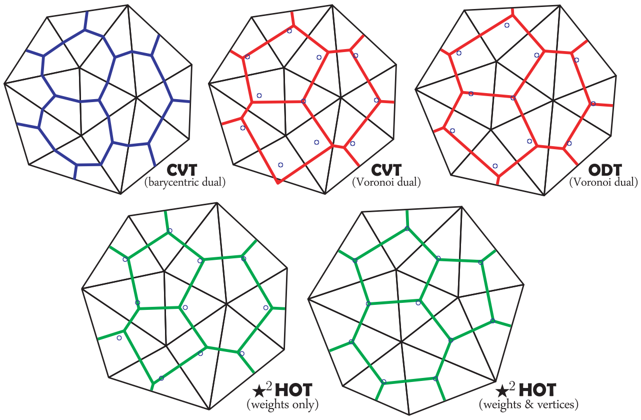

**Figure 1:** ***Primal/Dual Triangulations**: Using the barycentric dual (top-left) does not generally give dual meshes orthogonal to the primal mesh. Circumcentric duals, both in Centroidal Voronoi Tesselations (CVT, top-middle) and Optimal Delaunay Triangulations (ODT, top-right), can lead to dual points far from the barycenters of the triangles (blue points). Leveraging the freedom provided by weighted circumcenters, our Hodge-optimized triangulations (HOT) can optimize the dual mesh alone (bottom-left) or both the primal and dual meshes (bottom-right), e.g., to make them more self-centered while maintaining primal/dual orthogonality.*

### 1.1 Previous Work

Meshing complex 2D or 3D domains with high-quality elements has generated a tremendous number of research efforts. Bounds on numerical errors have resulted in the use of Delaunay triangulations [Edelsbrunner 1987] for finite-element computations, and Voronoi diagrams [Okabe et al. 2000] for finite-volume methods. However, the combined use of a primal mesh and its dual structure has increased over the last decade in both modeling and simulation, with quantities of both geometric (normals, mean and Gaussian curvatures, tangents) and physical (velocities, fluxes, circulations, vorticities) nature inherently located either on the primal mesh or its dual [Desbrun et al. 2007]. Calculations involving these primal and dual values in graphics were formalized in Discrete Exterior Calculus (DEC—see, e.g., [Hirani 2003]), now used in vision and image processing as well [Grady and Polimeni 2010].

**Delaunay/Voronoi pairs.** In the context of discrete differential geometric operators, Meyer et al. [2003] recommended a Voronoi (circumcentric) dual for tighter error bounds—but locally reverted to the barycentric dual when a dual vertex was not contained in its primal simplex. For fluid simulation, Perot and Subramanian [2007] and Elcott *et al.* [2007] advocated circumcentric duals as well, this time to ensure that pressure gradients between adjacent cells were parallel to the velocity samples stored on the common face. In DEC terminology, this simply means that the flux through a face and the circulation along its associated dual edge measure the same component of a vector field. Moreover, another advantage of the

Delaunay/Voronoi duality for fluid simulation exploited in [Elcott et al. 2007] is that the convexity and non-self-intersection of dual Voronoi cells make them ideal for the use of generalized barycentric coordinates [Warren et al. 2007]. Still, the seemingly natural choice of Delaunay/Voronoi triangulation is far from being without drawbacks. First and foremost, it it extremely difficult in practice to get "self-centered" Delaunay triangulations [Rajan 1994] for which each circumcenter lies inside its associated tetrahedron: failure to satisfy this property locally can lead to numerical degeneracies. Recent methods attempting to optimize meshes to avoid this issue remain impractical for complex domains [VanderZee et al. 2010]. A second drawback of a Delaunay/Voronoi pair is the inability to choose the positions of dual nodes locally without significantly degrading the primal mesh: having more flexibility in the placement of pressure samples would significantly improve the treatment of free surfaces in embedded boundary methods [Batty et al. 2010]. Consequently, and while abundantly vetted by theoretical guarantees, Delaunay/Voronoi triangulations are too restrictive in many practical situations. We will, instead, promote the use of arbitrary convex orthogonal primal/dual pairs to offer significantly more flexibility (see Fig. 1).

**Accuracy *vs* Efficiency.** Sparsity is crucial when dealing with large linear algebraic problems frequently encountered in geometry processing. Graphics literature is replete with low order methods using as-sparse-as-possible formulations for efficiency. While nonlinear and/or high-order methods have their own advantages and proponents, it is often highly desirable to find the simplest, fastest approximation valid for most applications, even if only to initialize a subsequent non-linear solver. In the context of DEC, this quest for efficiency often translates to the use of the so-called diagonal Hodge stars (that include the famous cotangent weights [Pinkall and Polthier 1993] widely used in geometry processing) over other discretizations of Hodge stars [Bossavit 1998; Auchmann and Kurz 2006; Wilson 2007] to approximate primal values based on dual values (and vice-versa); as inverse Hodge stars appear even in basic operators [Fisher et al. 2006; Elcott et al. 2007], having diagonal approximations can greatly increase computational efficiency. However, once a primal-dual triangulation is chosen, one has no control over the error incurred by diagonal approximations: efficiency may thus only be achieved on particularly good meshes. We will, instead, design meshes to minimize formal error bounds of diagonal Hodge stars, generally increasing accuracy without the additional costs associated with refinement (increasing system sizes) or higher-order Galerkin Hodge stars (decreasing the sparsity and making inversion more difficult).

### 1.2 Contributions

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—more precisely, regular triangulations along with their associated power diagrams—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 geometrical 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. Moreover, our work unveils an important connection between Hodge star accuracy and optimal transport. 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 improve condition numbers of basic operators as well as increase numerical robustness and accuracy.

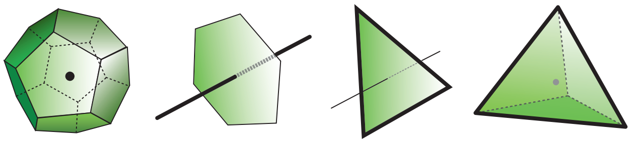

**Figure 2:** ***Duality***: *The dual of a triangulation in $\mathbb{R}^d$ associates to each $k$-simplex $\sigma^k$ a $(d-k)$-cell $*\sigma^k$ (here, $k=0,1,2,3$). Having $\sigma^k \cap *\sigma^k$ "well centered" within the primal simplex* and *its orthogonal dual cell is crucial to numerics in modeling and simulation.*

## 2 Preliminaries and Definitions

Before introducing our Hodge-based meshes, we first provide some background on orthogonal primal/dual triangulations and optimal transport as we will make heavy use of these notions throughout the paper.

### 2.1 Regular-Power Triangulations

The notion of dual for a triangulation $\mathcal{T}$ in $\mathbb{R}^d$ is well known and routinely used in graphics: each $d$-simplex is associated with a dual vertex (dual 0-cell), each $(d-1)$-simplex is associated with a dual edge (dual 1-cell) between the two dual vertices associated with the two adjacent $d$-simplices, etc. Primal vertices $\mathbf{x}_i$ are then associated with dual $d$-cells $V_i$, and the dual of $\mathcal{T}$ forms a cell complex $\mathcal{D}$. However, this concept of dual is abstract, as the location of the dual vertices are not specified a priori. A very common dual to a triangulation is the cell complex which uses the circumcenters of each $d$-simplex as dual vertices. If the initial triangulation is Delaunay (i.e., satisfying the empty circumsphere property [Edelsbrunner 1987]), this dual is simply the Voronoi diagram of the primal vertices, and its nice properties of non-self-intersection, convexity, and orthogonality of the primal and dual elements have led to its use in countless papers in graphics and computational sciences. The barycentric dual, for which barycenters are used instead of circumcenters (see Fig. 1), is also quite common in particular for finite-volume computations; however, it fails to satisfy both the orthogonality and the convexity conditions on general triangulations.

**Desirable primal-dual pairs.** Remaining agnostic with respect to the choice of a dual, we will call a **primal-dual triangulation** $\mathcal{M}$ in $\mathbb{R}^n$ any pair $(\mathcal{T}, \mathcal{D})$ with $\mathcal{T}$ being a triangulation in $\mathbb{R}^d$ and $\mathcal{D}$ a compatible dual complex of $\mathcal{T}$ (i.e., their respective adjacency matrices are transpose of each other). Moreover, if every edge $[\mathbf{x}_i, \mathbf{x}_j]$ of $\mathcal{T}$ and its *dual* $V_i \cap V_j$ in $\mathcal{D}$ are orthogonal to each other, the pair $(\mathcal{T}, \mathcal{D})$ is said to form an *orthogonal primal-dual triangulation*. Finally, we will denote as $*$ the operation of **duality** (see Fig. 2); that is, a primal simplex $\sigma$ will have its dual referred to as $*\sigma$ with the orientation induced by the primal orientation and the manifold orientation. For a more formal definition, see [Munkres 1984; Hirani 2003].

**Regular/power duality.** Delaunay/Voronoi primal-dual triangulations are restrictive in that they allow no choice on the dual once the primal mesh is fixed. A natural question to ask is: are there other primal-dual triangulations that satisfy the orthogonality, non-intersection, and convexity conditions we require? The answer is affirmative: the known duality between **regular triangulations** (also called weighted Delaunay triangulations) and **power diagrams** (also called Laguerre or weighted Voronoi diagrams) provides *all* the satisfactory orthogonal primal-dual triangulations [Glickenstein 2005]. This exact characterization of the primal/dual triangulations we seek will be particularly convenient as it will lead to a number of new geometric functionals measuring mesh quality; it will also yield straightforward generalizations of standard DEC operators without some of the most limiting factors that the Delaunay/Voronoi duality possesses.

Formally, a weighted point set is defined as a pair $(\mathbf{X}, W) =$

$\{(\mathbf{x}_1, w_1), \ldots, (\mathbf{x}_n, w_n)\}$, 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 **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 $\|\,\mathbf{x}-\mathbf{x}_i\,\|^2 - w_i$, where $\|\,.\,\|$ stands for the Euclidean distance. Using this power definition, to each $\mathbf{x}_i$ we can associate its weighted Voronoi region $V_i^w = \{\mathbf{x}\in\mathbb{R}^d |\ \|\,\mathbf{x}-\mathbf{x}_i\,\|^2 - w_i \leq \|\,\mathbf{x}-\mathbf{x}_j\,\|^2 - w_j, \forall j\}$. The power diagram of $(\mathbf{X}, W)$ is the cell complex whose cells are the weighted Voronoi regions. The dual of the power diagram of $(\mathbf{X}, W)$ is the regular triangulation of $(\mathbf{X}, W)$: this triangulation contains a $k$-simplex with vertices $\mathbf{x}_0, \mathbf{x}_1, \ldots, \mathbf{x}_k$ in $\mathbf{X}$ iff $\bigcap_{j=0}^{j=k} V_j^w \neq \emptyset$.

Note that in a regular triangulation, a point $\mathbf{x}_i \in \mathbf{X}$ can be hidden, i.e., it may not be used in the triangulation because its weighted Voronoi region is empty. Note also that when the weights are all equal, the power diagram coincides with the Euclidean Voronoi diagram of $\mathbf{X}$. Geometrically, one can think of the weight $w_i$ as the square of the radius of a unique circle centered at vertex $\mathbf{x}_i$; then there exists in each triangle a circle, centered at what is known as the **weighted circumcenter**, which is orthogonal to each of the circles centered at the vertices. All of these properties can be extended to the case where the weights are negative [Pedoe 1988], and thus regular triangulations and their associated power diagrams generalize the usual Delaunay/Voronoi duality nicely. This simple addition of a weight to each vertex allows us to conveniently parameterize the *entire* space of good (i.e. orthogonal, convex, and non-self-intersecting) primal-dual triangulations $\mathcal{M}$.

### 2.2 Basics of Optimal Transport

The optimal transport problem dates back to Gaspard Monge. (For a description of the vast literature on this topic, we refer the reader to [Villani 2009].) In essence, it 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 (one infinitesimal unit of volume at a time) is minimal. While Monge's variational formulation of the problem assumed that all the dirt at a point $x \in M$ must be moved through a point-to-point mapping $s$ to one location $s(x) \in N$, this restriction was relaxed by Kantorovich who replaced the mapping $s$ with a probability measure $\pi \in \mathcal{P}(M \times N)$ that specifies the joint measure of dirt-hole correspondences; i.e., $\pi$ is a *transport plan* between a probability measure $\mu$ on $M$ and a probability measure $\nu$ on $N$ with $\pi(\cdot \times N) = \mu$ and $\pi(M \times \cdot) = \nu$. This extension to the transport of measures marked a renewed interest in optimal transport as it proved general enough to apply to many scientific fields (for recent graphics-related applications, see [Mémoli 2011; Lipman and Daubechies 2010]).

**Wasserstein metric.** For measures the notion of "distance" (i.e., cost of transport) may vary based on context. A common distance function defined between probability measures in $\mathbb{R}^d$ with bounded support is the $q$-Wasserstein metric, defined as

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

To reuse the analogy mentioned above, if each measure is viewed as a unit amount of piled-up dirt, the metric is the minimum "cost" of turning one pile into the other, which is assumed to be the amount of dirt that needs to be moved times the $L_p$ distance it has to be moved. Because of this analogy, the metric is sometimes called the earth mover's distance. Note, as it will be crucial in Section 3, that by a direct application of the Hölder inequality for two measures of unit mass,

$$W_1(\mu,\nu) \leq W_2(\mu,\nu). \tag{1}$$

Finally, we will also need the Kantorovich-Rubinstein theorem, stating that for two measures $\mu$ and $\nu$ with bounded support, the 1-Wassertein distance between $\mu$ and $\nu$ can be rewritten as:

$$W_1(\mu,\nu) = \sup_{\substack{\varphi:\mathbb{R}^d\to\mathbb{R}\\ \mathrm{Lip}(\varphi)\leq\lambda}} \frac{1}{\lambda}\int_{\mathbb{R}^d} \varphi(x)\,\mathrm{d}(\mu-\nu), \tag{2}$$

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

## 3 Error Functionals for Diagonal Hodge Stars

To demonstrate the advantages of using regular/power triangulations, we focus on a particularly relevant type of functional measuring primal and dual properties. Recall that for an arbitrary primal element $\sigma$, the diagonal approximation of the Hodge star $\star$ [Bossavit 1998] of a continuous differential form $\alpha$ assumes

$$\int_{*\sigma} \star\,\alpha \approx \frac{|*\sigma|}{|\sigma|}\int_{\sigma}\alpha, \tag{3}$$

where $|.|$ denotes the Lebesgue measure (length, area, volume) of a simplex or cell. In other words, the discrete $k^{\text{th}}$ 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 $\sigma_i^k$ (resp., $*\sigma_i^k$) is the $i^{\text{th}}$ $k$-simplex (resp., $(d-k)$-cell) of the primal-dual triangulation $\mathcal{M} = (\mathcal{T}, \mathcal{D})$; the discrete Hodge star of a discrete primal $k$-form $\omega^k$ is then computed as $\star^k\,\omega^k$, and the extension to dual discrete forms (now with $(\star^k)^{-1}$) is trivial (for further details see, e.g., [Desbrun et al. 2007]).

### 3.1 Deriving Tight Bounds through Optimal Transport

While computationally convenient, diagonal Hodge stars are 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 a $k$-simplex $\sigma_i$ as the average difference between the discrete approximation and the exact Hodge star value:

$$e_i := \frac{1}{|*\sigma_i|}\left|\frac{|*\sigma_i|}{|\sigma_i|}\int_{\sigma_i}\omega - \int_{*\sigma_i}\star\omega\right| = \left|\frac{1}{|\sigma_i|}\int_{\sigma_i}\omega - \frac{1}{|*\sigma_i|}\int_{*\sigma_i}\star\omega\right|.$$

We now notice that due to the orthogonality of $\sigma$ and $*\sigma$, the component of $\omega$ along $\sigma$ is the same as the component of $\star\omega$ along $*\sigma$ (this is the same property that allows orthogonal primal-dual triangulations to admit a diagonal Hodge star in the first place). Writing this component as a scalar function $f(x)$, we can rewrite the integrals involved above as

$$\int_{\sigma_i}\omega \ = \int_{\sigma_i} f(x)\ d\mu_{\sigma_i} \quad \text{and} \quad \int_{*\sigma_i}\star\omega \ = \int_{*\sigma_i} f(x)\ d\mu_{*\sigma_i},$$

where $d\mu_{\sigma_i}$ and $d\mu_{*\sigma_i}$ are the volume forms of $\sigma_i$ and $*\sigma_i$ respectively. We can use these expressions to rewrite the error density as

$$e_i = \left|\int_{\sigma_i} f(x)\frac{d\mu_{\sigma_i}}{|\sigma_i|} - \int_{*\sigma_i} f(x)\frac{d\mu_{*\sigma_i}}{|*\sigma_i|}\right| = \left|\int_{\mathbb{R}^d} f(x)\ d(\bar{\mu}_{\sigma_i} - \bar{\mu}_{*\sigma_i})\right| \tag{4}$$

where now $d\mu_{\sigma_i}/|\sigma_i|$ and $d\mu_{*\sigma_i}/|*\sigma_i|$ are uniform probability distributions over $\sigma_i$ and $*\sigma_i$ respectively, and $d\bar{\mu}_{\sigma_i}$ and $d\bar{\mu}_{*\sigma_i}$ are their trivial extensions to $\mathbb{R}^d$—i.e., for any measurable set $S$ in $\mathbb{R}^d$,

$$\int_S d\bar{\mu}_{\sigma_i} = \int_{S\cap\sigma_i}\frac{d\mu_{\sigma_i}}{|\sigma_i|} \quad \text{and} \quad \int_S d\bar{\mu}_{*\sigma_i} = \int_{S\cap *\sigma}\frac{d\mu_{*\sigma_i}}{|*\sigma_i|}.$$

From Eq.(4) and Eq.(2), we deduce that the *tightest* bound one can find on the Hodge star error density per simplex for an arbitrary $\lambda$-Lipschitz form is simply $\lambda$ times the minimum cost over all transport plans between $\sigma_i$ (seen as a uniform probability measure over the mesh element) and $*\sigma_i$ (also seen as a uniform probability measure over the dual element); that is, with a slight abuse of notation,

$$e_i \leq \lambda \; W_1(\sigma_i, *\sigma_i). \tag{5}$$

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

### 3.2 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, i.e., by summing the integrals of the $p^{\text{th}}$ power of the error densities $e_i$ over local regions, specific to $\sigma_i$ and $*\sigma_i$, that tile the mesh. Such regions have been defined in previous work, coined "support volumes" in [Hirani 2003] and "diamonds" in [Hauret et al. 2007; Desbrun et al. 2007]: when $\sigma_i$ and $*\sigma_i$ intersect, these regions that we will refer to as $\diamond(\sigma_i \cup *\sigma_i)$ are just the convex hulls of $\sigma_i$ and $*\sigma_i$; in the general case, they become signed unions of convex hulls of the primal vertices of $\sigma_i$ and each boundary element of the dual cell $*\sigma_i$ . Using $\Sigma^k$ to denote the set of $k$-simplices of a triangulation, the total error is thus:

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

since the volume of the diamond $\diamond(\sigma_i \cup *\sigma_i)$ is, up to a dimension factor, simply the product of the primal and dual volumes due to our primal/dual orthogonality assumption of mesh $\mathcal{M}$.

From Eq. (5), we conclude that a tight bound for the $p^{\text{th}}$ power of the total error is expressed as:

$$E_p(\mathcal{M}, \star^k)^p \leq \frac{\lambda^p}{\binom{d}{k}} \sum_{\sigma_i \in \Sigma^k} |*\sigma_i||\sigma_i| \; W_1(\sigma_i, *\sigma_i)^p. \tag{6}$$

Notice that $E_\infty(\mathcal{M}, \star^k)$ is thus, up to the Lipschitz constant, bounded by the maximum of the $W_1$ distances between primal and dual elements of the mesh as expected. For notational convenience, we will denote by $\star^k$- $\text{HOT}_{p,1}(\mathcal{M})$ the bound (with Lipschitz and dimension constants removed) obtained in Eq. (6); more generally, we will define

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

as relevant functionals (or energies) to construct meshes, since minimizing them will control the quality of the discrete Hodge stars.

**Continuity of HOT functionals.** Because they are based on volume integrals, the HOT functionals are continuous over the space of regular/power triangulations. They are indeed continuous in the vertex positions of the primal and dual meshes, but also through primal mesh flips: an edge or face flip in a regular triangulation happens when a dual (power) edge vanishes. Hence the diamond weighting we use for our total error renders our HOT functionals continuous with respect to both vertices and weights. This will be particularly convenient when it comes time to optimize a mesh in order to minimize these functionals.

### 3.3 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 (see some examples in Section 5). In particular, linear combinations of HOT energies may be used if multiple Hodge stars are needed, for example when using Laplacians of $k$-forms with $k > 1$. Note also that the use of a 1-Wasserstein distance is notably less attractive numerically than a 2-Wasserstein distance as we will discuss in Section 4.4. Fortunately, we can also provide a bound of the Hodge star error which, while less tight than the previously derived $\text{HOT}_{p,1}$, will be particularly convenient to deal with computationally: the existence of optimal transport plans when the cost is the distance squared (i.e., $W_2$) being well studied, a useful bound on the Hodge star error can be derived using the inequality given in Eq. (1) to yield:

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

The reader may have noticed that the functional $\star^0$- $\text{HOT}_{2,2}(\mathcal{M})$ 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 [Du et al. 1999] to quasi-Newton methods [Liu et al. 2009], have been developed. $L_p$ variants (i.e., $\star^0$- $\text{HOT}_{2p,2}(\mathcal{M})$ for $p \geq 2$) were also explored recently [Lévy and Liu 2010]. However, these energies only correspond to $\star^0$, and are not as tight as $\text{HOT}_{1,p}$. 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 [Alliez et al. 2005] can also be seen as a variant of $\star^d$- $\text{HOT}_{2,2}(\mathcal{M})$ 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.

## 4 Hodge-Optimized Triangulations

In the remainder of this paper, we call a HOT triangulation any pair $\mathcal{M}$ consisting of a regular triangulation $\mathcal{T}$ and its associated power diagram $\mathcal{D}$ for which $\mathcal{T}$, $\mathcal{D}$, or both, have been optimized in order to reduce one (or a linear combination of) HOT functional(s). We now describe the basic computations involved in optimizing meshes for two particularly interesting (and unexplored) families of energies: $\text{HOT}_{2,2}$ and $\text{HOT}_{1,1}$.

### 4.1 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 optimize the continuous values of 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, finding a good non-optimal minimum is often enough to dramatically improve the mesh quality. We thus start our minimization process by initializing the domain with uniformly sampled vertices over the domain, and running a few iterations of CVT [Du et al. 1999] or ODT [Alliez et al. 2005] to quickly disperse the vertices and get mesh elements roughly similar in size: from such a decent initial mesh, an optimized mesh can be quickly obtained by performing a gradient descent, or alternatively (without much added implementation complexity), an L-BFGS algorithm [Nocedal and Wright 1999]—a particular quasi-Newton method where the (inverse) Hessian is approximated based on the $M$ previous steps (we use $M = 7$). A (binary or golden-ratio) linear search is performed to adapt the step size along the gradient or the quasi-Newton direction based on two simple tests (known as Wolfe conditions): the step size should be small enough to make sure the energy decreases, but large enough to induce a marked gradient change. This common minimization procedure works quite well without requiring anything else but an evaluation of our HOT energies and their gradients, which we will derive in closed-form from direct integration and/or application of the Reynolds theorem (see Appendices). Note finally that the positions $\mathbf{x}_i$ and the weights $w_i$ have very different scales (units of $m$ vs. $m^2$); we thus proceed by alternatively minimizing our HOT energies with respect to vertex positions and weights. After each step the connectivity is updated using the 2D or 3D regular triangulation package of CGAL [CGAL 2010]. Pseudocode of our general procedure is given in Fig. 3, but more specialized optimization techniques could most likely be devised; in particular, based on the HOT energy we wish to minimize, a few alternative minimization procedures may be simpler to implement or faster to converge. We will point out some such special cases shortly.

While both position *and* weight are optimized by default, HOT optimizations are relevant even if only one of these optimizations is performed. For instance, if one has a given (possibly non-flat) triangulation, vertices could be held fixed while weights are optimized to better one or more of the Hodge stars. Similarly, weights could be kept fixed, e.g. in contexts where they represent power or capacity of the nodes, and a best node placement is sought after—or simply in cases where a given connectivity needs to be maintained. We will discuss some useful variants in Section 5.

**Boundary Handling.** As in any variational method, boundary conditions can significantly affect the results. Except for the work of Alliez et al. [Alliez et al. 2005; Tournois et al. 2009; Sieger et al. 2010], we found very little about boundary handling in previous related work in graphics; for instance, recent papers focusing on the CVT energy like [Lévy and Liu 2010; Liu et al. 2009] only discuss how to partition a given domain into well-shaped Voronoi cells, providing no insight on dealing with the difficult issue of generating good simplices at the domain boundary. While boundary treatment may be context dependent (fixing vertices or even weights [Cheng et al. 2008] at the boundary being two of the most desirable options), we experimented with a very simple procedure to handle boundaries gracefully for all Hodge stars. We first make sure that each dual vertex $\mathbf{c}$ of a boundary $d$-simplex $T$ is associated with a "ghost" dual vertex $\hat{\mathbf{c}}$ used to enforce that dual edges at the boundary never have negative lengths: $\hat{\mathbf{c}}$ is put at the projection of $\mathbf{c}$ onto the boundary face of $T$ if $\mathbf{c}$ is within $T$, and put on top of $\mathbf{c}$ otherwise. We also alter the definition of the energy to become HOT $/|\mathcal{M}|$, i.e., we simply divide the energy by the total area: as volume-shrinkage is no longer rewarded, minimizing the HOT "volume density" makes the optimization steps behave well even at the boundary. We left the evaluations performed to check the Wolfe conditions unchanged: we did not alter the connectivity of the current triangulation (for efficiency reasons) or clamp vertices to the domain during the line search. Once the Newton step is done, however, a vertex is reprojected on the domain's boundary if (a) the vertex has an unbounded weighted Voronoi region, (b) it is outside the domain, or (c) it has one or more of its adjacent triangles' circumcenters *or* barycenters outside the domain. This approach is simple and it lets the vertices move freely from the inside to the boundary and vice-versa. We will not incorporate the term $|\mathcal{M}|$ in our derivations (to avoid cluttering the explanations) as it is a trivial alteration of our procedure.

### 4.2 Weighted Circumcenters

Most of the derivations involved in this section rely on the crucial relation linking primal and dual vertices in regular triangulations. We will use $\mathbf{c}(\sigma)$ to denote the weighted circumcenter of simplex $\sigma$, i.e., the unique intersection of the mutually-orthogonal affine spaces supporting the primal simplex $\sigma^k$ and its weighted dual $*\sigma^k$ (see Fig. 4). Of particular importance are the circumcenters of the $d$-simplices for a mesh $\mathcal{T}$ in $\mathbb{R}^d$: these form the vertices of its (weighted) dual complex $\mathcal{D}$. For a $k$-simplex $\sigma^k$, if $\mathbf{x}_i$ is any of the vertices of $\sigma^k$, the (weighted) circumcenter is expressed as:

$$\mathbf{c}(\sigma^k) = \mathbf{x}_i + \frac{1}{2k!|\sigma^k|} \sum_{\mathbf{x}_j \in \sigma^k} \left(|\mathbf{x}_i - \mathbf{x}_j|^2 + w_i - w_j\right) \boldsymbol{\sigma}_{\jmath}^{k} \qquad (7)$$

where $\boldsymbol{\sigma}_{\jmath}^{k}$ denotes the inward-pointing normal of the face of $\sigma^k$ opposite to $\mathbf{x}_j$ weighted by the volume of the face. With this general formula, weighted circumcenters are easy to differentiate, both with respect to vertices and weights. Notice that when the weights of $\sigma^k$ are all equal, one finds the expression for the (Voronoi) circumcenter used in [Alliez et al. 2005]. Armed with this useful identity, we can now formulate the various HOT energies.

```
// MESH OPTIMIZATION
// Input: vertices $\mathbf{x}^0 = \{\mathbf{x}_i\}$ and weights $w^0 = \{w_i\}$,
// and a HOT functional $E(\mathbf{x}, w)$.
$n \leftarrow 0$
repeat
    Compute $E(\mathbf{x}^n, w^n)$ // See Appendices A and B
    // Optimize $\mathbf{x}$
    Pick step direction $\mathbf{d}^x$ for $E(\mathbf{x}^n, w^n)$
    Find $\alpha$ satisfying Wolfe's condition(s)
    $\mathbf{x}^{n+1} \leftarrow \mathbf{x}^n + \alpha\, \mathbf{d}^x$     // Vertex updates
    Update regular triangulation
    // Optimize $w$
    Pick step direction $\mathbf{d}^w$ for $E(\mathbf{x}^{n+1}, w^n)$
    Find $\beta$ satisfying Wolfe's condition(s)
    $w^{n+1} \leftarrow w^n + \beta\, \mathbf{d}^w$     // Weight updates
    Update regular triangulation
    $n \leftarrow n + 1$
until (convergence criterion met)
```

**Figure 3:** *Basic pseudocode of our HOT optimization. Step directions are picked as gradient descent or quasi-Newton steps.*

### 4.3 HOT$_{2,2}$ Meshes

When a $W_2$-based transport cost is used, the HOT functionals are quite easy to compute in closed form. Indeed, a direct application of Pythagoras' theorem reveals that an optimal transport plan to move the normalized uniform measure for a simplex $\sigma$ to its orthogonal dual $*\sigma$ can be achieved by splitting the plan into two stages: first, optimally transport the measure from $\sigma$ to its (weighted) circumcenter $\mathbf{c}(\sigma)$, then from $\mathbf{c}(\sigma)$ to the dual cell $*\sigma$. The fact that the circumcenter $\mathbf{c}(\sigma)$ is at the intersection of the mutually-orthogonal affine spaces supporting $\sigma$ and $*\sigma$ makes the cost of the direct optimal transport and the sum of transport costs via $\mathbf{c}(\sigma)$ equal. The optimal transport cost is thus directly expressible as we now detail.

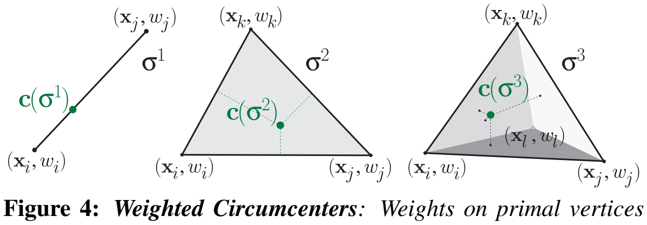

**Figure 4:** ***Weighted Circumcenters***: *Weights on primal vertices determine the placement of the weighted circumcenters for each edge (left), triangle (middle), and tetrahedron (right). Moreover, the orthogonal projection of the (weighted) circumcenter $\mathbf{c}(\sigma)$ onto one of $\sigma$'s faces falls on the (weighted) circumcenter of that face.*

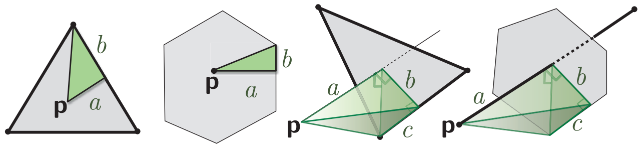

**Figure 5:** ***Splitting Mesh Elements****: Most of our HOT energies are evaluated by splitting simplices/cells into canonical subsimplices (in green) for which closed-form integral expressions $W(\mathbf{p},T)$ of simplex-$T$-to-point-$\mathbf{p}$ transport are easily found. Notations used for 2D (left) and 3D (right) in Sections 4.3 and 4.4 are indicated.*

**Energy computations.** For both $\star^0$ and $\star^d$ in dimension $d\!=\!2,3$, $\mathrm{HOT}_{2,2}$ energies can be easily computed by splitting weighted Voronoi $d$-cells or primal $d$-simplices into canonical simplices for which a closed form for the optimal transport cost to a point is easy to obtain—see this splitting in Fig. 5. For instance, for a right triangle $T$ with width $a$ and height $b$ in 2D, the integral over the triangle of the squared distance to the vertex $\mathbf{p}$ adjacent to $a$ is:

$$W(\mathbf{p},T)=\int_0^a\int_0^{\frac{bx}{a}}(x^2+y^2)\,dy\,dx=\frac{a^3b}{4}+\frac{ab^3}{12}$$

The equivalent formula in 3D for the bi-orthogonal tetrahedron $T$ split shown in Fig. 5(right) is now:

$$W(\mathbf{p},T)=\int_0^a\int_0^{\frac{bx}{a}}\int_0^{\frac{cy}{b}}(x^2+y^2+z^2)\,dzdydx=\frac{abc}{10}\left(a^2+\frac{b^2}{2}+\frac{c^2}{6}\right).$$

Costs for transport from edges are simpler to derive as they only involve 1D integrals, and the other remaining stars are just combinations of transport over edges, areas, and volumes. For completeness, all the transport costs needed in 2D and 3D can be found in Appendix A.

**$\star^0$- $\mathrm{HOT}_{2,2}$ in $\mathbb{R}^d$.** As mentioned in Section 3.3, this energy turns out to be the well-known CVT energy [Du et al. 1999] when all the weights are equal. For this particular case, various optimization strategies have been proposed, such as Lloyd iterations (to get near the optimal mesh) followed by a quasi-Newton method (to accelerate convergence) [Liu et al. 2009], and our approach mimics these strategies in this case. In the general case, however, the gradient with respect to a vertex position does not only contain the usual centroidal term from CVT, but also extra terms based on weight differences. As for the gradient with respect to a weight, its expression is surprisingly simple: it is simply the Laplacian (using the weighted version of the cotangent formula) of the weights, meaning that a $\star^0$- $\mathrm{HOT}_{2,2}$ mesh will have harmonic weights. We can then conclude that a CVT mesh is, in fact, a $\mathrm{HOT}_{2,2}$ mesh with Neumann boundary condition for the weights; other non-trivial boundary conditions will lead to $\mathrm{HOT}_{2,2}$ that are not CVT.

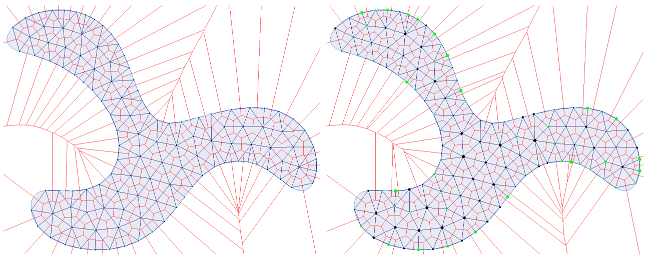

**Figure 6:** ***ODT vs. $\star^2$-HOT$_{2,2}$****: An ODT mesh (left) and a $\star^2$-HOT mesh (right) are computed for the same 2D shape; while the primal triangulations are rather similar, the dual vertices end up closer to the simplices' barycenters for the HOT mesh thanks to the additional freedom provided by the weights.*

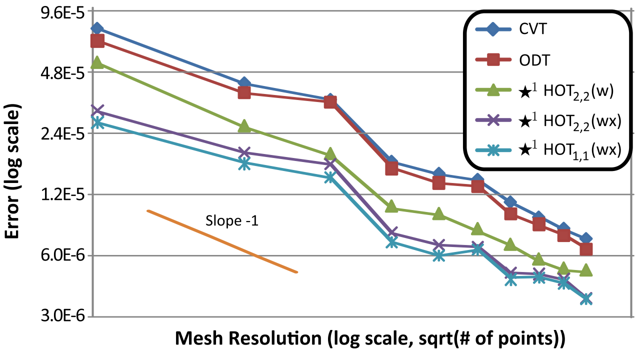

**Figure 7:** ***Laplace equation****: We plot the $L^2$ errors (with respect to the solution on a very fine mesh) of the solutions of a Laplace equation on a circle with Dirichlet boundary conditions $f(u,v)=(u^2\!+\!v^2)\sin(u)\cos(v)$ for CVT, ODT, and HOT meshes with weight optimization only (w) or vertex and weight optimization (wx). While the log-log plot of the errors as a function of mesh size shows, as expected, that the convergence rate is not affected, our HOT meshes still noticeably outperform both CVT and ODT meshes, with $W_1$ only marginally better than $W_2$.*

**$\star^d$- $\mathrm{HOT}_{2,2}$ in $\mathbb{R}^d$.** Although seemingly the "dual" version of the $\star^0$ case, this energy requires specific derivations that we now go over. Computing the energy gradient with respect to weights is made easier if one notices that $\partial\mathbf{c}(\sigma)/\partial w_i$ is independent of the weights: Eq. (7) is indeed linear in the weights of the mesh in any dimension. Consequently, the gradient of the energy (which is quadratic in $\mathbf{c}$) with respect to a weight is linear in the weights of the mesh, offering alternative optimization approaches. One can for instance solve for the optimal weights directly by finding the weights that make all gradients zero: this is simply achieved via a global, sparse linear system collecting all the gradient equations and equating them to zero (see Appendix A; notice that this system is simply a Poisson problem). A slightly slower approach—but easier to implement and parallelize—is to compute the optimal weight $w_i^*$ for each vertex assuming that the other weights are unchanged, and update $w_i \leftarrow (w_i^*+w_i)/2$ (moving straight to the optimal value may lead to overshooting, hence the half-way update). The gradient with respect to vertices is detailed in Appendix A, and we optimize vertex positions as sketched in Fig. 3. A $\star^d$- $\mathrm{HOT}_{2,2}$ result in 2D can be seen in Fig. 6.

**$\mathrm{HOT}_{2,2}$ for Other Stars.** Be it in 2D ($d=2$) or 3D ($d=3$), the $\mathrm{HOT}_{2,2}$ functionals for $\star^k$ (where $k=1\ldots d-1$) can be derived using the circumcenter formula of Eq. 7. However, they do not simplify in the same way as those for $\star^0$ and $\star^d$ did above. Consequently, a direct application of the general HOT algorithm (Fig. 3) is called for, and it performs as expected.

### 4.4 $\mathrm{HOT}_{1,1}$ Meshes

While the $\mathrm{HOT}_{1,1}$ functionals provide the tightest $L_1$ bounds on Hodge star errors, their use of the 1-Wasserstein distance makes the equations more difficult: optimal transport plans are often much less obvious to determine, and their costs more difficult to compute. In fact, the $\mathrm{HOT}_{1,1}$ energy for $\star^0$ is directly related to facility location problems [Okabe et al. 2000] that are commonplace in operations research, as it amounts to find the location (sometimes called the (continuous, or integral) geometric median) which minimizes the integrated Euclidean distance to the interior of a polygonal region. Thankfully, closed forms of many of the energies can still be found (albeit, with more difficulty) as detailed in Appendix B. For a few of the energies, in particular those for $\star^1$ and $\star^2$ in 3D, closed forms are not easily calculated, and numerical quadrature may be the only practical approach for their computation. However, in practice we found the $\mathrm{HOT}_{1,1}$ energies to not give significant

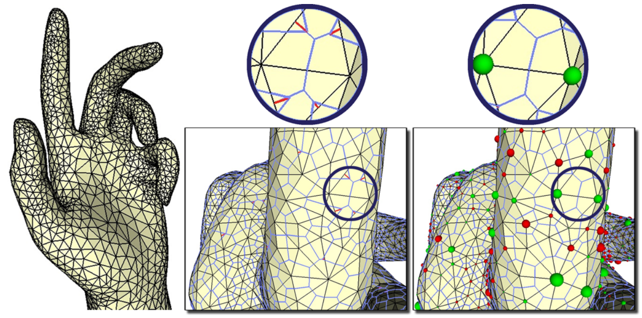

**Figure 8:** ***Surface Weight Optimization****: For a given triangular mesh (left) there are several triangles whose circumcenter is far outside the triangle (center, lines drawn in red). By optimizing only the weights the new dual vertices are better placed inside the unchanged triangles (right) while keeping primal/dual orthogonality.*

improvements over their HOT$_{2,2}$ counterparts, and thus may only prove useful when the tightest formal bounds are required (Fig. 7).

### 4.5 Discussion

In many ways, HOT meshes can be seen as a generalization of CVT meshes. However, one must be careful with the term "Centroidal Voronoi Tesselation," as being centroidal is a only necessary condition of a CVT energy minimum: for instance, a regular grid is centroidal, and yet the CVT energy is not at a local minimum. Similarly, having each weighted circumcenter at the barycenter of its associated triangle is not sufficient to minimize the $\star^d$-HOT$_{2,2}$ functional in $\mathbb{R}^d$: the functional also captures the error distribution throughout the domain. A HOT mesh for $\star^k$ tries instead to strike a balance between being "centroidal" or "medial" (i.e., with each $k$-simplex being "self-centered" for $W_p$), and having each $k$ cell being of the same volume. In 2D, most of these energies are globally minimized for a perfect hexagonal tiling of the plane; however, this is no longer true in 3D and above, as an equilateral simplex no longer tiles $\mathbb{R}^{d>2}$. Consequently, while geometric functionals could be easily derived to simply force a mesh to be centroidal or medial (in the generalized diamond-based sense), HOT functionals also favor uniform sizing of the optimal mesh.

## 5 Applications and Results

HOT meshes can be beneficial in a number of contexts in modeling of surfaces and volumes, as well as in simulation. We mention a few examples to demonstrate the generality of our approach and provide numerical experiments. We also discuss variants and extensions.

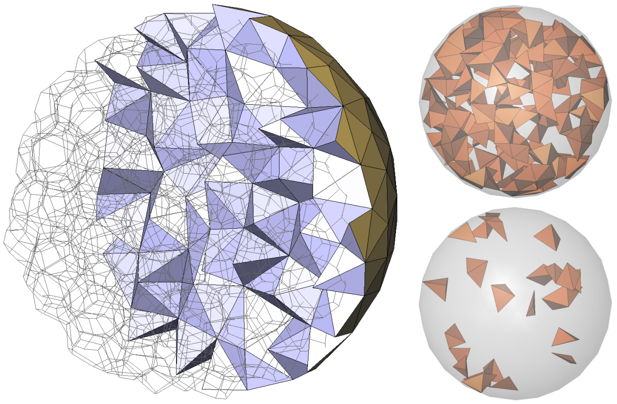

**Figure 9:** ***HOT$_{2,2}$ Sphere****: Optimizing an ODT mesh of a sphere for both weights and vertex positions results in a nice mesh (left) with 30 tetrahedra whose dual vertex is outside of the tet (bottom-right), compared to 206 in the original ODT mesh (top-right).*

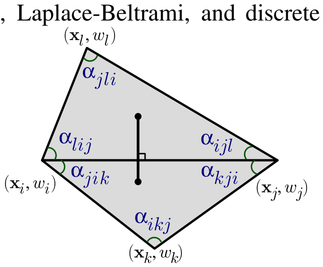

**HOT DEC.** Most of the DEC methods in graphics (including the huge literature on Laplacian, Laplace-Beltrami, and discrete conformal parameterization) can be directly adapted to work on HOT meshes as well. In fact, if one computes the diagonal Hodge star using the usual ratio of (signed) volume of a simplex and its dual, then no modification is required. If, instead, closed form formulae are used to express each diagonal Hodge star, they only need to be modified to include the contribution due to the weights. For instance, the traditional Hodge star $(\star^1)_{ij} = \big(\cot(\alpha_{ikj}) + \cot(\alpha_{jli})\big)/2$ for a one-form between vertex $i$ and vertex $j$ becomes:

$$
\begin{aligned}
(\star^1)_{ij} =& \frac{1}{2}\Big( \cot\alpha_{ikj} + \cot\alpha_{jli} \\
&+ (w_i - w_k)\frac{\cot\alpha_{kji}}{||\mathbf{x}_i - \mathbf{x}_j||^2} + (w_j - w_k)\frac{\cot\alpha_{jik}}{||\mathbf{x}_i - \mathbf{x}_j||^2} \\
&+ (w_i - w_l)\frac{\cot\alpha_{ijl}}{||\mathbf{x}_i - \mathbf{x}_j||^2} + (w_j - w_l)\frac{\cot\alpha_{lij}}{||\mathbf{x}_i - \mathbf{x}_j||^2}\Big).
\end{aligned}
$$

These changes can be accommodated seamlessly in existing codes, and allow for much greater flexibility: weights can be, for instance, optimized (with fixed connectivity or not) to locally "displace" dual vertices onto an immersed boundary [Batty et al. 2010] through a least-square fit. Vertices can be optimized as well, for instance in applications requiring local remeshing to maintain good numerics.

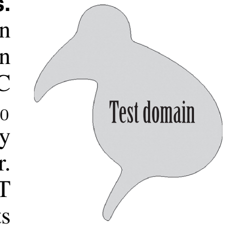

**Laplace & Laplace-Beltrami Operators.** A particularly common linear operator in mesh processing is the Laplacian $\Delta$, be it in the plane or on a discrete surface. Its DEC expression for 0-forms being $\Delta = d_0^t \star^1 d_0$ and the $d_0$ operator being exact, the only loss of accuracy rises from the Hodge star. Consequently, meshes minimizing the HOT energy for $\star^1$ should be appropriate for its accurate computation, as evidenced by Fig. 7 where up to 65% error reduction is achieved compared to CVT. In fact, [Glickenstein 2005] and [Wardetzky et al. 2007] were the first to recognize the importance of orthogonal primal/dual meshes to ensure good numerical qualities of the Laplacian. A $\star^1$-HOT$_{2,2}$ mesh indeed results, on a 200V discretization of the test domain depicted in the inset, in a 5% reduction of the condition number of the Laplacian matrix with Dirichlet boundary conditions compared to a CVT mesh (much greater improvements are witnessed when compared to arbitrary, non-optimized meshes). 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. This is partially due to an increase of the minimum dual edge length (going from 2.0e$^{-3}$ for CVT to 1.5e$^{-2}$ on the same mesh), providing an alternative approach to removing short dual edges presented recently in [Sieger et al. 2010]. Similar improvements were found for the Laplace-Beltrami operator of the surface mesh in Fig. 8.

**Improving Dual Structure.** We often have to deal with situations where the 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 find the weights to, e.g., minimize the $L_2$ distance squared between weighted circumcenters (defined in Eq. 7 through an equation that is linear in the weights) and triangle barycenters using a single linear solve. The connectivity is kept intact, regardless of the weights—only the position and shape of the compatible dual $\mathcal{D}$ is optimized. Although one cannot guarantee that the resulting dual will be flawless (self-centered and non-self-intersecting), it will be improved compared to the original circumcentric dual. Even for HOT energies, our 2D and 3D tests show that only optimizing the weights is particularly simple and beneficial on a number of meshes. Fig. 8 depicts a triangle mesh of a hand and its intrinsic dual before and after weight $\star^2$-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 [Fisher et al. 2006]. Similarly, Fig. 11 shows that even an ODT mesh with exceptionally high-quality tetrahedra [Tournois et al. 2009] can be made significantly better centered 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 HOT mesh—a two third reduction of "outcentered" tetrahedra. As a final illustrative example we show results on a 3D sphere (Fig. 9). Starting from an ODT mesh and optimizing only the weights drops the number of outcentered tetrahedra from 206 to 52, while allowing the optimization to also move the positions of the vertices further reduces this number to 30, resulting in the mesh shown in the figure. On the other hand, if a weighted Delaunay mesh is undesirable, optimizing only the positions still reduces the number of outcentered tetrahedra to 118, almost half of the original ODT mesh, while still using a circumcentric (Voronoi) dual.

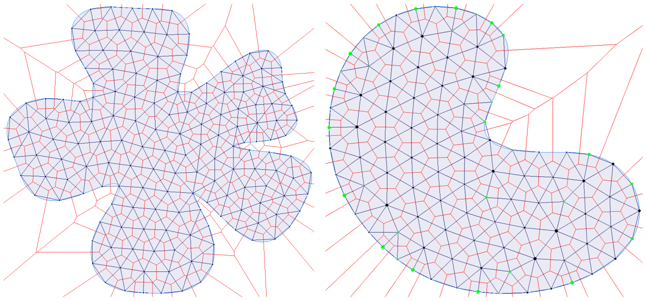

**Figure 10:** ***HOT***$_{1,1}$ ***Meshes****: A "Medial Voronoi Tesselation" (i.e., a $\star^0$-HOT$_{1,1}$ mesh) has vertices near the integral geometric median of each Voronoi cell (left); $\star^1$-HOT$_{1,1}$ mesh tends to have primal and dual edges intersecting near their midpoints (right, weights shown as balls with color/size indicating sign/magnitude).*

**Accuracy & Extensions.** While we described archetypical primal-dual HOT energies, one can use regular triangulations and power diagrams to derive other relevant energies. Even in the context of Hodge star accuracy, we point out that the "diamond weighting" proposed in Section 3 can be modified if one wishes to improve a *particular* Hodge star (and *not* its inverse): for instance, the discrete Hodge star between 0-forms and $d$-forms in $\mathbb{R}^d$ should use a weighting equal to 1, while the inverse Hodge star should use the volume of the local $d$-cells. Similarly, one may minimize a linear combination of HOT energies if multiple Hodge stars need to be optimized simultaneously. Designing new energies based on targeted numerical tasks should be straightforward—although continuity and convexity of these functionals will need to be studied on a case-by-case basis. Nevertheless, our $\star^k$-HOT energies lead consistently to a $5\%$ to $35\%$ $L^1$- and $L^\infty$- improvement on both $\star^k$ and $(\star^k)^{-1}$ for linear and non-linear functions alike on 2D non-convex domains like depicted in the inset earlier—even if the error minimization is not run to convergence. As for the 3D Bimba mesh in Fig. 11, our $\star^3$-optimization of only the weights already reduces both the $L^1$ and $L^2$ norm of $\star^3$-errors for linear functions by $16\%$.

**HOT$_{1,1}$ vs. HOT$_{2,2}$.** While slower to converge when the 1-Wasserstein distance is used, HOT$_{1,1}$ and HOT$_{2,2}$ meshes are visually quite similar. Numerical tests, similarly, do not demonstrate major differences: a simple Laplace's equation with Dirichlet boundary conditions on various mesh sizes clearly indicate that HOT$_{1,1}$ are slightly better than HOT$_{2,2}$, but both are significantly better than CVT or ODT (see Fig. 7). HOT$_{1,1}$ meshes can, in fact, be slightly worse than their $W_2$ equivalents when their accuracy is tested using polynomial test functions. It is therefore unclear that using the $W_1$ cost is worth the added computational burden for graphics applications, despite offering tighter theoretical bounds.

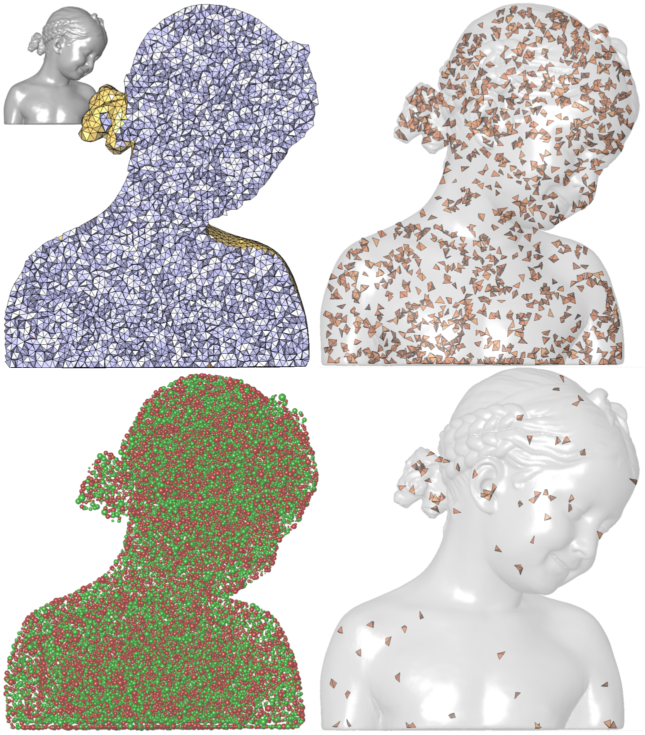

**Figure 11:** ***3D Weight Optimization****: A high-quality ODT mesh of the Bimba con Nastrino (top left cross-section; 195K tets, 36K vertices) can be $\star^3$-optimized by a few (30) iterations of our weight optimization, thus improving minimal dual edge length and self-centeredness (bottom left; weights are displayed according to sign (red/green) and magnitude (radius)). 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 HOT mesh (bottom right) is significantly better than the original ODT (top right), even if the primal triangulations are exactly matching. If we further increase the visualization threshold to the point when the HOT mesh has a single "bad" tetrahedron, the non-weighted original Bimba mesh then exhibits 192 such tetrahedra.*

## 6 Future Work

Several future directions are ripe for exploration. For instance, formulating other functionals based on particular numerical tasks (such as eigenvalue problems) or other families of functions (other than just Lipschitz) could be of interest. In fact, the induced symmetries of our HOT meshes may improve other operators as well. Deriving $L^p$-based functionals (using the quadratures pointed in [Lévy and Liu 2010]) or incorporating a sizing field in the functionals should be relatively straightforward. We also believe that a sustained effort to produce better optimizations for HOT-like energies is in order to ensure efficient, industrial-strength implementation. Finally, as always in meshing, providing a richer set of boundary conditions would also extend the number of potential applications, thus helping the adoption of HOT meshes. Combining HOT optimization with feature protection through boundary weights as proposed in [Cheng et al. 2008] could offer a practical extension of our approach in this direction.

**Acknowledgements** The authors thank David Cohen-Steiner and Omid Amini for early support and feedback, as well as Pierre Alliez for data. This research was partially funded through NSF grants (CCF-0811373, CMMI-0757106, and CCF-1011944), and by the generous support of Pixar Animation Studios.

## References

ALLIEZ, P., COHEN-STEINER, D., YVINEC, M., AND DESBRUN, M. 2005. Variational tetrahedral meshing. *ACM Trans. on Graphics (SIGGRAPH) 24*, 3 (July), 617–625.

AUCHMANN, B., AND KURZ, S. 2006. A geometrically defined discrete Hodge operator on simplicial cells. *IEEE Trans. Magn. 42*, 4, 643–646.

BATTY, C., XENOS, S., AND HOUSTON, B. 2010. Tetrahedral embedded boundary methods for accurate and flexible adaptive fluids. *Computer Graphics Forum (Eurographics) 29* (May), 695–704(10).

BOSSAVIT, A. 1998. *Computational Electromagnetism*. Academic Press, Boston.

CGAL, 2010. Computational Geometry Algorithms Library (release 3.8). http://www.cgal.org.

CHENG, S.-W., DEY, T. K., AND LEVINE, J. 2008. Theory of a practical Delaunay meshing algorithm for a large class of domains. In *Algorithms, Architecture and Information Systems Security*, B. Bhattacharya, S. Sur-Kolay, S. Nandy, and A. Bagchi, Eds., vol. 3 of *World Scientific Review*, 17–41.

DESBRUN, M., KANSO, E., AND TONG, Y. 2007. Discrete differential forms for computational modeling. In *Discrete Differential Geometry*, A. Bobenko and P. Schröder, Eds. Springer.

DU, Q., FABER, V., AND GUNZBURGER, M. 1999. Centroidal voronoi tessellations: Applications and algorithms. *SIAM Rev. 41* (December), 637–676.

EDELSBRUNNER, H. 1987. *Algorithms in Combinatorial Geometry*. Springer-Verlag.

ELCOTT, S., TONG, Y., KANSO, E., SCHRÖDER, P., AND DESBRUN, M. 2007. Stable, circulation-preserving, simplicial fluids. *ACM Trans. on Graphics 26*, 1 (jan), 4.

FISHER, M., SPRINGBORN, B., BOBENKO, A. I., AND SCHRÖDER, P. 2006. An algorithm for the construction of intrinsic delaunay triangulations with applications to digital geometry processing. In *ACM SIGGRAPH Courses*, 69–74.

GLICKENSTEIN, D. 2005. Geometric triangulations and discrete Laplacians on manifolds. *Arxiv preprint math/0508188.*

GRADY, L. J., AND POLIMENI, J. R. 2010. *Discrete Calculus: Applied Analysis on Graphs for Computational Science*. Springer.

HAURET, P., KUHL, E., AND ORTIZ, M. 2007. 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.

HIRANI, A. N. 2003. *Discrete Exterior Calculus*. PhD thesis, Caltech.

LÉVY, B., AND LIU, Y. 2010. $L_p$ Centroidal Voronoi Tesselation and its Applications. *ACM Trans. on Graph. 29*, 4.

LIPMAN, Y., AND DAUBECHIES, I. 2010. Surface Comparison With Mass Transportation. In *ArXiv preprint 0912.3488.*

LIU, Y., WANG, W., LÉVY, B., SUN, F., YAN, D., LU, L., AND YANG, C. 2009. On Centroidal Voronoi Tessellation - energy smoothness and fast computation. *ACM Trans. on Graph. 28*, 4.

MÉMOLI, F. 2011. A spectral notion of Gromov-Wasserstein distances and related methods. *Journal of Applied and Computational Harmonic Analysis 30*, 3, 363–401.

MEYER, M., DESBRUN, M., SCHRÖDER, P., AND BARR, A. H. 2003. Discrete differential-geometry operators for triangulated 2-manifolds. In *Visualization and Mathematics III*, Springer-Verlag, H.-C. Hege and K. Polthier, Eds., 35–57.

MUNKRES, J. R. 1984. *Elements of Algebraic Topology*. Addison-Wesley.

NOCEDAL, J., AND WRIGHT, S. J. 1999. *Numerical optimization*. Springer Verlag.

OKABE, A., BOOTS, B., SUGIHARA, K., AND CHIU, S. N. 2000. *Spatial tessellations: Concepts and applications of Voronoi diagrams*, 2nd ed. Probability and Statistics. Wiley.

PEDOE, D. 1988. *Geometry, a comprehensive course*, 2nd ed. Dover Publications.

PEROT, J. B., AND SUBRAMANIAN, V. 2007. Discrete calculus methods for diffusion. *Journal of Computational Physics 224*, 59–81.

PINKALL, U., AND POLTHIER, K. 1993. Computing discrete minimal surfaces and their conjugates. *Experimental Mathematics 2(1)*, 15–36.

RAJAN, V. 1994. Optimality of the Delaunay triangulation in $\mathbb{R}^d$. *Discrete and Computational Geometry 12*, 1, 189–202.

SHEWCHUK, J. 2002. What is a Good Linear Element? Interpolation, Conditioning, and Quality Measures. In *Proc. of the 11th Int. Meshing Roundtable*, 115–126.

SIEGER, D., ALLIEZ, P., AND BOTSCH, M. 2010. Optimizing voronoi diagrams for polygonal finite element computations. In *Proceedings of the 19th International Meshing Roundtable*. 335–350.

TOURNOIS, J., WORMSER, C., ALLIEZ, P., AND DESBRUN, M. 2009. Interleaving delaunay refinement and optimization for practical isotropic tetrahedron mesh generation. *ACM Trans. Graph. 28* (July), 75:1–75:9.

VANDERZEE, E., HIRANI, A. N., GUOY, D., AND RAMOS, E. 2010. Well-centered triangulation. *SIAM Journal on Scientific Computing 31*, 6, 4497–4523.

VILLANI, C. 2009. *Optimal Transport: Old and New*. Fundamental Principles of Mathematical Sciences, 338. Springer-Verlag.

WARDETZKY, M., MATHUR, S., KÄLBERER, F., AND GRINSPUN, E. 2007. Discrete laplace operators: no free lunch. In *Symposium on Geometry processing*, 33–37.

WARREN, J., SCHAEFER, S., HIRANI, A., AND DESBRUN, M. 2007. Barycentric coordinates for convex sets. *Advances in Computational Mathematics 27*, 3, 319–338.

WILSON, S. O. 2007. Cochain algebra on manifolds and convergence under refinement. *Topology and its Applications 154*, 9, 1898 – 1920.

## A HOT$_{2,2}$ Energies

In this appendix, we give explicit 2D and 3D formulations of HOT$_{2,2}$ energies and their gradients in a form most amenable for direct implementation. In any dimension $d$, HOT$_{2,2}$ energies can be expressed as a function of the *signed distances* between the (weighted) circumcenters of $n-$ and $(n+1)$-simplices, for $0 \leq n \leq d-1$ (considering the primal vertices as circumcenters of 0-simplices). This is based on the simple geometric observation that the (weighted) circumcenter of an $(n+1)$-simplex projects orthogonally to the (weighted) circumcenters of its $n$-subsimplices. We will make use of this property when deriving closed-form expressions of the HOT$_{2,2}$ energies.

**Signed Distances between Circumcenters** Let $\mathbf{c}_{0\ldots n}$ denote the (weighted) circumcenter of simplex $[\mathbf{x}_0, \ldots, \mathbf{x}_n]$. For any $0 \leq n \leq d-1$, the signed distance between the circumcenter of the $n$-simplex $[\mathbf{x}_0, \ldots, \mathbf{x}_n]$ to the circumcenter of the $(n+1)$-simplex $[\mathbf{x}_0, \ldots, \mathbf{x}_{n+1}]$ is defined as the Euclidean distance between $\mathbf{c}_{0\ldots n}$ and $\mathbf{c}_{0\ldots(n+1)}$ and its sign is positive if the simplex $[\mathbf{x}_0, \ldots, \mathbf{x}_n, \mathbf{c}_{0\ldots(n+1)}]$ has the same *orientation* as $[\mathbf{x}_0, \ldots, \mathbf{x}_n, \mathbf{x}_{n+1}]$, and negative otherwise.

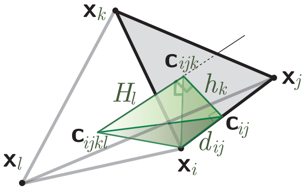

We further denote by $d_{ij}$ the signed distance between the circumcenter of the 0-simplex $[\mathbf{x}_i]$ to the circumcenter of the 1-simplex $[\mathbf{x}_i, \mathbf{x}_j]$, i.e., the distance between $\mathbf{c}_i (= \mathbf{x}_i)$ and $\mathbf{c}_{ij}$, with a positive sign if $(\mathbf{x}_i - \mathbf{c}_{ij})$ has the same orientation as $(\mathbf{x}_i - \mathbf{x}_j)$, and negative otherwise. Note that $d_{ji}$ thus corresponds to the signed distance between $\mathbf{x}_j$ and $\mathbf{c}_{ij}$. It is easy to see that: $d_{ij} = \frac{|e_{ij}|^2 + w_i - w_j}{2|e_{ij}|}$, $d_{ji} = \frac{|e_{ij}|^2 + w_j - w_i}{2|e_{ij}|}$, where $e_{ij} = \mathbf{x}_j - \mathbf{x}_i$. Going further up in dimension, we denote by $h_k$ the signed distance between $c_{ij}$ and $c_{ijk}$ in a triangle $t_{ijk} = [\mathbf{x}_i, \mathbf{x}_j, \mathbf{x}_k]$. We have

$$h_k = \frac{\cot \beta_k |e_{ij}|}{2} + \frac{\cot \beta_i w_j + \cot \beta_j w_i}{2|e_{ij}|} - \frac{w_k |e_{ij}|}{4|t_{ijk}|}$$

where $\beta_k$ is the angle at $\mathbf{x}_k$ in triangle $t_{ijk}$. Finally, we denote by $H_l$ the signed distance between $\mathbf{c}_{ijk}$ and $\mathbf{c}_{ijkl}$ in tetrahedron $T_{ijkl}$.

Through the cell and simplex splitting explained in Section 4.3, we can use the integral forms of $W(\mathbf{p}, T)$ given earlier, resulting in closed-form expressions of all the HOT$_{2,2}$ energies for every triangle $t_{ijk}$ and tetrahedron $T_{ijkl}$ as a function of the signed distances $d_{ij}$, $h_k$, and $H_l$ between circumcenters as follows:

**2D HOT$_{2,2}$ Formulas:**

$$\star^0\text{- HOT}_{2,2}(t_{ijk}) = \sum_{i,j,k} \Big( \frac{d_{ij}^3 h_k}{4} + \frac{d_{ij} h_k^3}{12} \Big).$$

$$\star^1\text{- HOT}_{2,2}(t_{ijk}) = \sum_{i,j,k} \Big( \frac{d_{ij}^3 h_k}{6} + \frac{d_{ij} h_k^3}{6} \Big).$$

$$\star^2\text{- HOT}_{2,2}(t_{ijk}) = \sum_{i,j,k} \Big( \frac{d_{ij}^3 h_k}{12} + \frac{d_{ij} h_k^3}{4} \Big).$$

**3D HOT$_{2,2}$ Formulas:**

$$\star^0\text{- HOT}_{2,2}(T_{ijkl}) = \sum_{i,j,k,l} \frac{1}{5} \Big( \frac{H_l^3 h_k d_{ij}}{12} + \frac{H_l h_k^3 d_{ij}}{4} + \frac{H_l h_k d_{ij}^3}{2} \Big).$$

$$\star^1\text{- HOT}_{2,2}(T_{ijkl}) = \sum_{i,j,k,l} \frac{1}{3} \Big( \frac{H_l^3 h_k d_{ij}}{12} + \frac{H_l h_k^3 d_{ij}}{4} + \frac{H_l h_k d_{ij}^3}{6} \Big).$$

$$\star^2\text{- HOT}_{2,2}(T_{ijkl}) = \sum_{i,j,k,l} \frac{1}{3} \Big( \frac{H_l^3 h_k d_{ij}}{6} + \frac{H_l h_k^3 d_{ij}}{4} + \frac{H_l h_k d_{ij}^3}{12} \Big).$$

$$\star^3\text{- HOT}_{2,2}(T_{ijkl}) = \sum_{i,j,k,l} \frac{1}{5} \Big( \frac{H_l^3 h_k d_{ij}}{2} + \frac{H_l h_k^3 d_{ij}}{4} + \frac{H_l h_k d_{ij}^3}{12} \Big).$$

**Vertex Position Optimization** In order to find the optimized position of the vertices of the mesh, we only need the derivative of the signed distances between (weighted) circumcenters with respect to vertices. Some of them may be easily derived from the formulas we provided before in this appendix. For instance:

$$\frac{\partial d_{ij}}{\partial \mathbf{x}_i} = -\frac{(\mathbf{x}_j - \mathbf{x}_i)}{2|e_{ij}|} + \frac{(w_i - w_j)(\mathbf{x}_j - \mathbf{x}_i)}{2|e_{ij}|^3}.$$

$$\frac{\partial d_{ji}}{\partial \mathbf{x}_i} = -\frac{(\mathbf{x}_j - \mathbf{x}_i)}{2|e_{ij}|} + \frac{(w_j - w_i)(\mathbf{x}_j - \mathbf{x}_i)}{2|e_{ij}|^3}.$$

More generally, we can derive all other formulas by using Eq. 7, which defines the equation for the vector from $\mathbf{x}_i$ to the (weighted) circumcenter of any simplex incident to $\mathbf{x}_i$. Through repeated uses of Pythagoras' theorem, one can then easily differentiate the signed distances between (weighted) circumcenters with respect to $\mathbf{x}_i$.

**Weight Optimization** The weight optimization of each HOT$_{2,2}$ energy can be easily done using the following simple formulas:

$$\frac{\partial d_{ij}}{\partial w_i} = \frac{1}{2|e_{ij}|}, \ \frac{\partial d_{ji}}{\partial w_i} = -\frac{1}{2|e_{ij}|},$$

$$\frac{\partial h_k}{\partial w_i} = \frac{\cot \theta_j}{2|e_{ij}|}, \ \frac{\partial h_j}{\partial w_i} = \frac{\cot \theta_k}{2|e_{ik}|}, \ \frac{\partial h_i}{\partial w_i} = -\frac{|e_{jk}|}{4|t_{ijk}|}.$$

While the derivative of the weighted circumcenter with respect to primal vertices does not have a short expression, the derivative with respect to the weights can be easily computed in any dimension: this derivative at a vertex is proportional to the inverse of the distance to the opposite facet in the outward normal direction of that facet. In 2D, this results in $\partial \mathbf{c}^w_{t_{ijk}} / \partial w_i = e^{\perp}_{jk} / (4|t_{ijk}|)$, while in 3D, $\partial \mathbf{c}^w_{T_{ijkl}} / \partial w_i = \mathbf{n}_i(T_{ijkl}) / (12\ |T_{ijkl}|)$, where $\mathbf{n}_i(T_{ijkl})$ is the normal (weighted by its area) of the facet opposite to $\mathbf{x}_i$ in tetrahedron $T_{ijkl}$.

**$\star^d$- HOT$_{2,2}$-Optimal Weights in $\mathbb{R}^d$:** In this particular case, there is a linear expression for the optimal $w_i^*$'s when all other weights are considered fixed. In 2D, we get:

$$w_i^* = \frac{2 \sum_{j \in \Omega(i)} \Big( \cot(\beta_k) + \cot(\beta_l) \Big) w_j + 4 \sum_{t_{ijk}} (\mathbf{c}_t - \mathbf{b}_t) \cdot e^{\perp}_{jk}}{\sum_{t_{ijk}} \frac{\|e_{jk}\|^2}{|t_{ijk}|}}$$

where $\beta_k$ is the angle at $\mathbf{x}_k$ in triangle $t_{ijk}$, and $\mathbf{b}_t$ is the barycenter of the triangle $t_{ijk}$. In 3D, we have instead:

$$w_i^* = \Bigg( \sum_{T_{ijkl} \in \Omega(i)} \Big[ w_j \cot(\alpha_{kl}) |e_{kl}| + w_k \cot(\alpha_{jl}) |e_{jl}| + w_l \cot(\alpha_{jk}) |e_{jk}| + 2(\mathbf{c}_T - \mathbf{b}_T) \cdot \mathbf{n}_i(T_{ijkl}) \Big] \Bigg) \Big/ \Bigg( \sum_{T_{ijkl} \in \Omega(i)} \frac{2\ |t_{jkl}|^2}{3\ |T_{ijkl}|} \Bigg)$$

where $\alpha_{kl}$ is the dihedral angle at edge $e_{kl}$.

## B HOT$_{1,1}$ Energies

In this appendix, we give explicit 2D and 3D formulations of HOT$_{1,1}$ energies for those which are practical to compute.

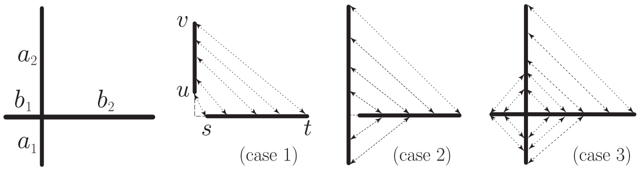

**Figure 12:** $W_1$ ***transport between orthogonal edges****: For two arbitrary edges parameterized by the length $a_1, a_2, b_1$, and $b_2$ (left), the transport plan (sampled in dotted lines) and cost of the normalized measure from one edge to the other can be computed in closed-form for the three possible configurations displayed above. Notice that the measure of an edge is sometimes evenly transported to two locations on the other edge (cases 2 and 3).*

**HOT$_{1,1}$ for $\star^0$ and $\star^d$ in $\mathbb{R}^2$ and $\mathbb{R}^3$.** Using the cell and simplex splitting used in Section 4.3, we can derive closed-form expressions for both $\star^0$- HOT$_{1,1}(\mathcal{M})$ and $\star^d$- HOT$_{1,1}(\mathcal{M})$ as these two cases only involve transport plans between convex regions and points. In 2D we can express the integral over a right triangle with width $a > 0$ and height $b > 0$ to the point $\mathbf{p}$ adjacent to $a$ (see Fig. 5(left)) as:

$$\begin{aligned} W(\mathbf{p}, T) &= \int_0^a \int_0^{\frac{bx}{a}} \sqrt{x^2+y^2}\, dy\, dx \\ &= \frac{1}{6}\left[ab\sqrt{a^2+b^2} + a^3 \sinh^{-1}\left(\frac{b}{a}\right)\right] \end{aligned}$$

The gradients of these energies are then expressed as:

$$\frac{\partial W(\mathbf{p},T)}{\partial a} = \frac{1}{6}\left[b\sqrt{a^2+b^2} + 3a^2 \sinh^{-1}\left(\frac{b}{a}\right)\right]$$

and

$$\frac{\partial W(\mathbf{p},T)}{\partial b} = \frac{1}{3} a\sqrt{a^2+b^2}.$$

The 3D counterpart for a biorthogonal tetrahedron and one of its vertices (see Fig. 5(right)) becomes:

$$\begin{aligned} W(\mathbf{p}, T) =& \int_0^a \int_0^{\frac{bx}{a}} \int_0^{\frac{cy}{b}} \sqrt{x^2+y^2+z^2}\, dz dy dx \\ =& \frac{1}{24}\Bigg( abc\sqrt{a^2+b^2+c^2} + ab\,(b^2+3a^2) \sinh^{-1}\Big(\frac{c}{\sqrt{a^2+b^2}}\Big) \\ & +3a^4\left[\tan^{-1}\Big(\frac{b}{c}\Big) - \tan^{-1}\Big(\frac{l\sqrt{a^2+b^2+c^2}}{ac}\Big)\right] \\ & +a^4\left[\tan^{-1}\Big(\frac{c}{b}\Big) - \tan^{-1}\Big(\frac{ac}{l\sqrt{a^2+b^2+c^2}}\Big)\right] \Bigg). \end{aligned}$$

The algorithm in Fig. 3 can then be applied directly using these expressions. As expected for $\star^0$, we obtain what could be called a "Medial Voronoi Tesselation": each vertex is on (or very near) the geometric median of its weighted Voronoi regions, see Fig. 10. (Note that when the weights are all constant, a Lloyd-like algorithm could also be used, for which each vertex is repeatedly moved to the median of its own Voronoi cell; we found, however, the quasi-Newton method with Wolfe conditions to be more efficient as in the HOT$_{2,2}$ case.) Similarly for $\star^d$, the resulting meshes have their dual complex centered at the median of each triangle.

**$\star^1$- HOT$_{1,1}$ in 2D.** The optimal transport cost from a primal edge to a dual edge for the 1-Wasserstein distance can be computed in closed form with some effort. We first compute the optimal transport cost between two separated, orthogonal edges (depicted in Fig. 12) parameterized as:

$$W(s,t,u,v) = \int_s^t \sqrt{x^2 + \Big(\frac{v-u}{t-s}(x-s)+u\Big)^2}\, dx \qquad (8)$$

where $t > s \geq 0$ and $v > u \geq 0$, for which a closed-form expression can be found in Mathematica. The general cost $E$ is then computed as one of three possible edge configurations depicted in Fig. 12 (with $a = a_1 + a_2, b = b_1 + b_2$):

- $b_1 < 0$ & $a_1 < 0$: (case 1)
  $E = \frac{a}{2} W(|b_1|, b_2, |a_1|, a_2)$
- $b_1 < 0$ & $a_1 > 0$: (case 2)
  $E = \frac{a}{2}(W(|b_1|, 2a_1 b/a + |b_1|, 0, a_1) + W(2a_1 b/a + |b_1|, b_2, a_1, a_2))$
- Otherwise: (case 3)
  $$\begin{aligned} E = \frac{a}{2}(&2W(0, b_1, 0, ab_1/b) + W(d_1, 2a_1 b/a - b_1, ab_1/b, a_1) \\ &+ W(2a_1 b/a - b_1, b_2, a_1, a_2)). \end{aligned}$$

The optimality of these transport plans can be seen by noting that, by the triangle inequality of the Euclidean metric, transport lines for the optimal transport plans cannot cross (see Chapter 8 of [Villani 2009]). This uniquely defines the optimal transport plan for case 1, while the density splitting in cases 2 and 3 follow from this and the symmetry of the regions in which the density is split.

## C Reynolds Transport Theorem

The Reynolds theorem provides an alternate, and often more intuitive way to differentiate integrated quantities like our HOT energies, which may also prove useful for deriving expressions of future energies as well. It states that the rate of change of the integral of a scalar function $F$ within a volume $V$ is equal to the volume integral of the instantaneous changes of $F$ occurring within the volume, plus the surface integral of the rate at which $F$ is being transported through the surface $\partial V$ (bounding $V$) to and from the surrounding region; i.e.,

$$\frac{d}{dt}\bigg|_{t=t_0} \int_{V_t} F(\mathbf{x}, t)\, dV = \int_{V_{t_0}} \frac{\partial F}{\partial t}(\mathbf{x}, t_0)\, dV + \int_{\partial V_{t_0}} F(\mathbf{x}, t_0)\, \dot{\mathbf{x}} \cdot \mathbf{n}\, dA$$

where $\mathbf{n}$ is the outward unit normal vector to the boundary $\partial V_{t_0}$. The term $\dot{\mathbf{x}}$ should be understood as the derivative of a point on the boundary with respect to $t$, so that the second term of the equation evaluates the transport of $F$ through the boundary as $t$ is varied.