

MESHES OPTIMIZED FOR DISCRETE EXTERIOR CALCULUS (DEC)

SARAH C. MOUSLEY*, MICHAEL DEAKIN †, PATRICK KNUPP ‡, AND SCOTT MITCHELL §

Abstract. We study the optimization of an energy function used by the meshing community to measure and improve mesh quality. This energy is non-traditional because it is dependent on both the primal triangulation and its dual Voronoi (power) diagram. The energy is a measure of the mesh's quality for usage in Discrete Exterior Calculus (DEC), a method for numerically solving PDEs. In DEC, the PDE domain is triangulated and this mesh is used to obtain discrete approximations of the continuous operators in the PDE. The energy of a mesh gives an upper bound on the error of the discrete diagonal approximation of the Hodge star operator. In practice, one begins with an initial mesh and then makes adjustments to produce a mesh of lower energy. However, we have discovered several shortcomings in directly optimizing this energy, e.g. its non-convexity, and we show that the search for an optimized mesh may lead to mesh inversion (malformed triangles). We propose a new energy function to address some of these issues.

1. Introduction. Like finite element or finite volume methods, Discrete Exterior Calculus (DEC) is a method for numerically solving partial differential equations (PDEs). To numerically solve a PDE using DEC, one first divides the function domain up into non-overlapping triangles, creating a mesh as in Figure 1.1. Using the mesh, the continuous operators in the PDE are transformed into matrices. These matrices are used to create a linear system of equations. The numbers in the solution vector of the linear system give approximations of the function's values at the vertices of the mesh. Our work is *not* about DEC implementation, but rather focuses on understanding how one should divide up the domain to create a mesh that is well-suited for DEC. We refer the reader to [2] and [3] for an introduction to DEC theory.

In 2011, Mullen et al. [5] introduced a mesh quality function called the HOT (Hodge-Optimized Triangulation) energy. The HOT energy of a mesh gives an upper bound on the error of the diagonal discrete approximation of the Hodge star operator, an operator that appears in some PDEs. In addition to this theoretical bound, in practice there is preliminary evidence that meshes that have low HOT energy yield better DEC numerical solutions to PDEs involving Hodge star operators. In [5] the Laplace equation is solved using DEC on a HOT optimized mesh, yielding more accurate solutions than those produced by using meshes optimized for other well-known energy functions, such as the Centriodal Voronoi Tessellation (CVT) energy and the Optimal Delaunay triangulation (ODT) energy.

*University of Illinois at Urbana-Champaign for Mathematics, mousley2@illinois.edu

†Sandia National Laboratories, mdeakin@sandia.gov

‡Sandia National Laboratories, knupp.patrick@gmail.com

§Sandia National Laboratories, samitch@sandia.gov



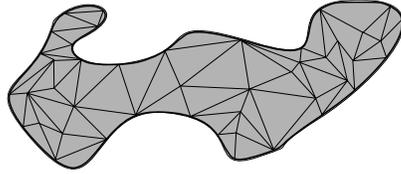


FIG. 1.1. A meshed domain

Our contributions:

- We analyze the optimization algorithm presented in [5] for obtaining low HOT energy meshes. The algorithm begins with a regular triangulation, a type of triangulation where to each vertex there is an associated weight and the weights are used to determine the edges (and triangles). Then adjustments to vertices positions and weights are made. When these adjustments are made, the triangulation may no longer be regular. In fact, it may not even be a triangulation. The HOT energy is only defined for regular meshes, thus one must extrapolate the HOT energy function to labeled abstract triangulations. In Section 3, we propose a different method for extrapolating HOT than is done in [5]. Additionally, in Section 4 we discuss how the extrapolated HOT energy of a non-regular labeled abstract triangulation compares to the HOT energy of the corresponding regular triangulation.
- We study the landscape of the HOT energy function. We discovered features of the HOT energy that may make it poorly suited for optimization (Sections 5 and 6).
- We developed software to study the HOT energy. Our software can compute the HOT energies of a mesh. When the vertex labels (weights) are all equal, we can also compute finite difference and analytic derivatives of the HOT energies with respect to the positions of the vertices.

2. Terminology and notation. In this section, we explain terminology and notation needed for understanding the discussions in subsequent sections.

Triangulations: By a *mesh* or *triangulation*, we will mean a decomposition of a 2D-domain into non-overlapping, non-degenerate triangles, though the ideas presented here can be extended to 3D-domains. An *abstract triangulation* is a cell-complex obtained from a triangulation by adjusting the vertices positions, keeping all the face relationships of the original triangulation. An abstract triangulation is allowed to have overlapping triangles. A *labeled (abstract) triangulation* is an (abstract) triangulation in which every vertex is labeled by a non-negative number. Let \mathcal{T} be a labeled abstract triangulation. Let $\{(\mathbf{x}_i, w_i)\}$ denote the labeled vertices of \mathcal{T} . We call a triangulation \mathcal{T}' that is dual to the power diagram (weighted Voronoi diagram) corresponding to $\{(\mathbf{x}_i, w_i)\}$ a *regular triangulation with respect to* $\{(\mathbf{x}, w_i)\}$. If the edge relationships of \mathcal{T} and \mathcal{T}' differ, we will call \mathcal{T} *non-regular*.

A *Delaunay triangulation* corresponding to a set of points $\{\mathbf{x}_i\}$ is a regular triangulation corresponding to $\{(\mathbf{x}_i, 0)\}$. Equivalently, a triangulation is Delaunay if the circumcircle of each triangle contains no vertices of the triangulation in its interior.

Parameters associated to labeled triangles: Let \triangle be a triangle embedded in \mathbb{R}^2 with vertices x_1, x_2, x_3 , labeled by w_1, w_2 , and w_3 respectively. We now define

some measurements associated to Δ that will appear in formulas for computing the HOT energy of a mesh.

For $i = 1, 2, 3$, let C_i be the circle of radius $\sqrt{w_i}$ centered at x_i . There exists a unique circle C such that C intersects all of the C_i orthogonally. The center of C is the *weighted circumcenter* of Δ and we denote it by $w(\Delta)$. If all the vertex labels are equal, the weighted circumcenter is the center of the circumcircle of Δ .

For $i \neq j$, let e_{ij} denote the oriented edge in Δ from x_i to x_j . Define the *weighted circumcenter* of e_{ij} to be the orthogonal projection of $w(\Delta)$ to e_{ij} and denote it by $w(e_{ij})$. Define d_{ij} to be the signed distance from x_i to $w(e_{ij})$, where the sign is taken to be positive if the vector from x_i to $w(e_{ij})$ points in the same direction as e_{ij} and otherwise is negative. The d_{ij} can be computed by the following formulas:

$$d_{ij} = \frac{|e_{ij}|^2 + w_i - w_j}{2|e_{ij}|} \quad \text{and} \quad d_{ji} = \frac{|e_{ij}|^2 + w_j - w_i}{2|e_{ij}|}. \quad (2.1)$$

Additionally, for pairwise distinct i, j, k , we define h_{ikj} to be the signed distance from e_{ij} to $w(\Delta)$, where the sign is positive if x_k and $w(\Delta)$ lie on the same side of e_{ij} , and otherwise is negative. When Δ is understood, we will often simplify notation and use h_k instead of h_{ikj} . The following formula computes h_k :

$$h_k = \frac{|e_{ij}| \cot \beta_k}{2} + \frac{w_j \cot \beta_i + w_i \cot \beta_j}{2|e_{ij}|} - \frac{w_k |e_{ij}|}{4a(\Delta)}, \quad (2.2)$$

where $a(\Delta)$ denotes the area of Δ . See also [4] for rearranged expressions.

Dual cells: Consider an edge σ of a labeled abstract triangulation \mathcal{M} . For each triangle Δ in which σ is an edge, consider the straight line segment from $w(\sigma)$ to $w(\Delta)$. We call the concatenation of all such line segments the edge *dual* to σ and denote it by $*\sigma$. In the case where \mathcal{M} is a regular triangulation, $*\sigma$ is an edge in the associated power diagram. If \mathcal{M} is regular, then we can also define duals to vertices and triangles. See [2] for these definitions.

HOT energy defined: Mullen et al. [5] define the HOT energy function of a regular mesh \mathcal{M} as follows. For $i = 0, 1, 2$,

$$\star^i\text{-HOT}_{p,p}(\mathcal{M}) = \sum_{\sigma \in \Sigma^i} |*\sigma| |\sigma| W_p(\mu_\sigma, \mu_{*\sigma})^p,$$

where Σ^i is the collection of all i -dimensional cells in \mathcal{M} , $|\cdot|$ is the volume measure, μ_σ and $\mu_{*\sigma}$ are the probability measures associated to the primal cell σ and the dual cell $*\sigma$ respectively, and W_p is the p -Wasserstein metric.

The energy $\star^i\text{-HOT}_{p,p}(\mathcal{M})$ gives an upper bound on the error of the discrete diagonal approximation of the Hodge star operator \star^i , which is a particular function of differential forms. The operator \star^i intakes an i -form and outputs an $(n-i)$ form, where n is the dimension of the mesh (in our case $n = 2$). More precisely, given an i -form ω , for each i -cell in \mathcal{M} , DEC uses the approximation

$$\int_{*\sigma} \star^i(\omega) \approx \frac{|*\sigma|}{|\sigma|} \int_{\sigma} \omega.$$

A total error is then computed by taking a weighted sum of the error in each of these integral approximations (cells with larger volume are weighted heavier). The energy $\star^i\text{-HOT}_{p,p}(\mathcal{M})$ gives an upper bound on the total error.

The following formulas from Appendix A in [5] can be used to compute $\star^1\text{-HOT}_{2,2}(\mathcal{M})$. Let σ be an edge contained in \mathcal{M} , say connecting x_i and x_j , and opposite x_k in a triangle Δ in \mathcal{M} . Define

$$\star^1\text{-HOT}_{2,2}(\sigma, \Delta) = \frac{1}{3} (d_{ij}^3 h_k + d_{ij} h_k^3 + d_{ji}^3 h_k + d_{ji} h_k^3). \quad (2.3)$$

Then define

$$\star^1\text{-HOT}_{2,2}(\sigma) = \sum \star^1\text{-HOT}_{2,2}(\sigma, \Delta), \quad (2.4)$$

where the sum is taken over all triangles Δ in \mathcal{M} containing σ . Finally, define

$$\star^1\text{-HOT}_{2,2}(\mathcal{M}) = \sum_{\sigma \in \Sigma^1(\mathcal{M})} \star^1\text{-HOT}_{2,2}(\sigma). \quad (2.5)$$

We will also use the notation

$$\star^1\text{-HOT}_{2,2}(\Delta) = \sum \star^1\text{-HOT}_{2,2}(\sigma, \Delta),$$

where the sum is taken over all edges σ in Δ . Similar explicit formulas for $\star^i\text{-HOT}_{2,2}$, $i = 0, 2$, can be found in [5].

Optimization algorithm:

To be self-contained, below we restate the mesh optimization algorithm from [5].

```

//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 ← 0
repeat
  Compute  $E(\mathbf{x}^n, w^n)$ 
  // 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 ← n + 1
until (convergence criterion met)

```

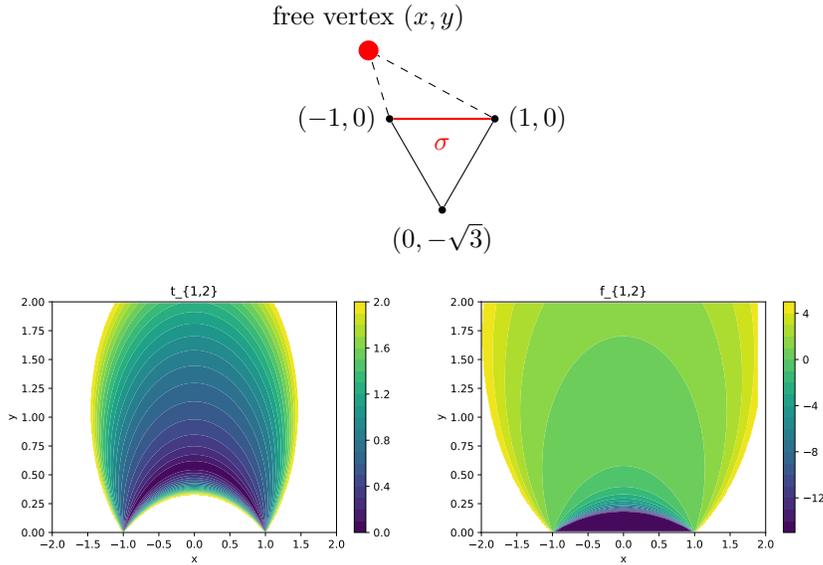


FIG. 3.1. Contour plots of the extrapolated \star^1 -HOT $_{2,2}$ energy of σ as (x, y) is moved. The formula-based extrapolation $f_{1,2}(\sigma) \rightarrow -\infty$ as (x, y) moves towards σ , while our theory-based extrapolation $t_{1,2}(\sigma) \rightarrow \infty$.

3. Extending HOT energy to non-regular cell complexes. The energy \star^i -HOT $_{p,p}$ is a function whose domain is the set of regular triangulations. Because making adjustments to vertex positions and weights (without updating to a regular triangulation) may result in a cell complex that is *non-regular*, it is perhaps unclear what it means for α to be valid step distance in the optimization algorithm present in Section 2. In this section, we discuss a method in the literature for extending the notion of \star^i -HOT $_{2,2}$ energy to non-regular labeled abstract triangulations. We then propose a different technique to extend the \star^i -HOT $_{2,2}$ function.

Consider the labeled abstract triangulation \mathcal{M}' obtained by moving vertex \mathbf{x}_i to a new point \mathbf{x}'_i , preserving all the original face relationships in \mathcal{M} . Let \mathcal{M}'' denote the regular mesh corresponding to $\{(\mathbf{x}_j, w_j) : j \neq i\} \cup \{(\mathbf{x}'_i, w_i)\}$. We emphasize that there is no guarantee that \mathcal{M}' and \mathcal{M}'' will be the same cell complex. Moreover, it may not make sense to compute \star^i -HOT $_{p,p}(\mathcal{M}')$, since \mathcal{M}' may not be regular. However, we can extrapolate Eqs. (2.3) and (2.4) to obtain a function $f_{1,2}$ whose domain is the class of labeled abstract triangulations. By extrapolate, we mean that for regular triangulations, $f_{1,2}$ and \star^1 -HOT $_{2,2}$ give the same answer. Now we can compute $f_{1,2}(\mathcal{M}')$. The function $f_{1,2}$ is the extension of \star^1 -HOT $_{2,2}$ used in the mesh optimization algorithm in [5] (see Section 2).

We found examples of non-regular labeled abstract triangulations for which $f_{1,2}$ yields *negative* values (see Figure 3.1). Negative values are problematic from both the conceptual and practical standpoints. Conceptually, the energy is defined as a product of positive lengths and Wasserstien distances, and its true minimum is zero. From a practical standpoint, we worry that an optimization step may progress towards a position with a negative extrapolated value, and restoring the Delaunay property would increase the calculated energy. We propose studying an alternative extrapolation of \star^1 -HOT $_{2,2}$, which always yields non-negative output.

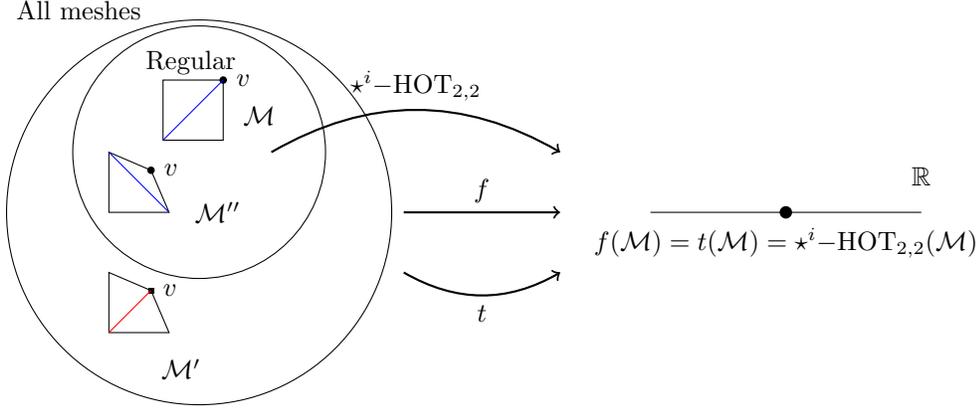


FIG. 3.2. The domain of $\star^1\text{-HOT}_{2,2}$ is the set of regular meshes. The functions $f_{1,2}$ and $t_{1,2}$ are extensions of $\star^1\text{-HOT}_{2,2}$.

Let \mathcal{M} be labeled abstract triangulation and σ an edge in \mathcal{M} . Let μ_σ and $\mu_{*\sigma}$ denote the probability measures corresponding to σ and $*\sigma$ respectively. Define

$$t_{1,2}(\mathcal{M}) = \sum_{\sigma \in \Sigma^1(\mathcal{M})} |\sigma| * \sigma |W_2(\mu_\sigma, \mu_{*\sigma})|^2. \quad (3.1)$$

We think of $f_{1,2}$ as extending the $\star^1\text{-HOT}_{2,2}$ formulas and $t_{1,2}$ as extending the Optimal Transport theory behind $\star^1\text{-HOT}_{2,2}$. See Figure 3.2 for an illustration of the extrapolation possibilities.

We derived the following formulas that can be used to compute $t_{1,2}(\mathcal{M})$. For an edge σ which is contained in two triangles,

$$|\sigma| * \sigma |W_2(\mu_\sigma, \mu_{*\sigma})|^2 \leq \frac{\text{sgn}(h_k + h_\ell)}{3} (d_{ij}^3 h_k + d_{ij} h_k^3 + d_{ji}^3 h_k + d_{ji} h_k^3 + d_{ij}^3 h_\ell + d_{ij} h_\ell^3 + d_{ji}^3 h_\ell + d_{ji} h_\ell^3), \quad (3.2)$$

where x_i and x_j are the endpoints of σ and x_k and x_ℓ are the vertices opposite σ . A similar statement holds if σ is contained in just one triangle.

Idea behind derivation of (3.2): Let x_0 be the intersection point in \mathbb{R}^2 of the lines containing σ and $*\sigma$. Then

$$W_2(\mu_\sigma, \mu_{*\sigma})^2 \leq \int_{\mathbb{R}^2} d(x_0, x)^2 d|\mu_\sigma - \mu_{*\sigma}|. \quad (3.3)$$

Inequality (3.3) can be obtained by modifying the proof of Theorem 6.15 in [6], which uses the fact that for any points $x_0, x, y \in \mathbb{R}^2$

$$d(x, y)^2 \leq (d(x, x_0) + d(x_0, y))^2 \leq 2(d(x, x_0)^2 + d(x_0, y)^2).$$

Because of how we chose x_0 , we have that $d(x, y)^2 = d(x, x_0)^2 + d(y, x_0)^2$. Thus, we are able to drop the factor of 2 that appears in the bound for $W_2(\sigma, *\sigma)^2$ obtained by

applying Theorem 6.15 verbatim. Now to obtain Ineq. (3.2), compute the integral in Ineq. (3.3). We suspect the inequality in (3.2) can be replaced with an equality.

When $f_{1,2}$ and $t_{1,2}$ differ: $f_{1,2}(\sigma)$ and $t_{1,2}(\sigma)$ agree if $|\ast\sigma| = h_k + h_\ell$ and disagree if $|\ast\sigma| = -(h_k + h_\ell)$. Those are the only possibilities since $|\ast\sigma| = |h_k + h_\ell|$. Figure 3.3 illustrates all the configurations of weighted circumcenters where $|\ast\sigma| = -(h_k + h_\ell)$.

Because $f_{1,2}$ and $t_{1,2}$ are both extensions of $\star^1\text{-HOT}_{2,2}$, they agree when the labeled abstract triangulation is a regular triangulation with respect to its weights. Observe that if “with respect to its weights” is removed from the previous statement, the statement is false as Figure 3 shows.

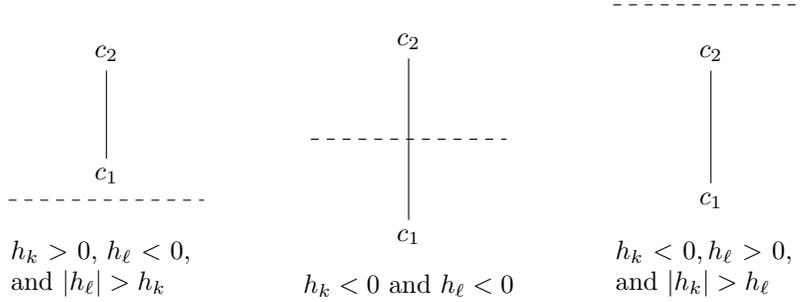


FIG. 3.3. The dotted line represents the line containing edge $\sigma = e_{ij}$. We assume x_k lies above the dotted line and x_ℓ below the dotted line. Let c_1 and c_2 denote the weighted circumcenters of triangles $x_i x_j x_k$ and $x_i x_j x_\ell$ respectively. The solid line is $\ast\sigma$.

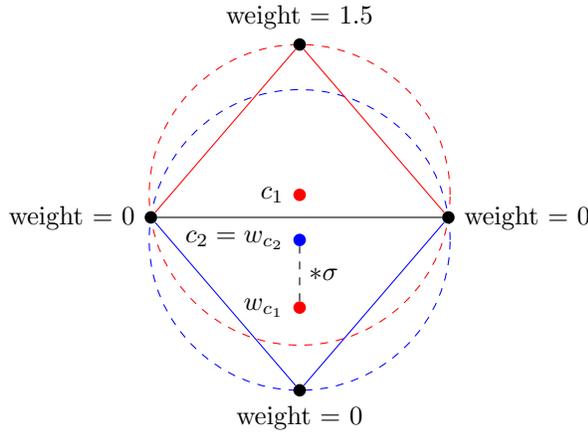


FIG. 3.4. This labeled triangulation has a weighted circumcenter configuration as in Figure 3.3. It is not the regular triangulation with respect to the weights indicated. However, the triangulation is the regular triangulation with respect to any set of all equal weights (the triangulation is Delaunay).

4. Updating triangulation: How does the energy change?. Let $\{(x_i, w_i)\}$ be a set of weighed points and \mathcal{M} the corresponding regular mesh. Let \mathcal{M}' be the labeled abstract triangulation with labeled points $\{(x'_i, w'_i)\}$ and all edge and face

relationships inherited from \mathcal{M} , i.e.

$$x'_i x'_j \in \Sigma^1(\mathcal{M}') \Leftrightarrow x_i x_j \in \Sigma^1(\mathcal{M}).$$

Let \mathcal{M}'' be the regular mesh corresponding to $\{(\mathbf{x}'_i, w'_i)\}$. The set-up we have just described arises in the mesh optimization algorithm in Section 2, where the (\mathbf{x}'_i, w'_i) are chosen so that Wolfe's conditions are satisfied. In particular,

$$f_{1,2}(\mathcal{M}') \leq \star^i\text{-HOT}_{2,2}(\mathcal{M}).$$

That is, the algorithm always moves towards a lower energy. However, \mathcal{M}' is replaced with \mathcal{M}'' before the next loop begins. In [5] there is no discussion of how $\star^i\text{-HOT}_{2,2}(\mathcal{M}'')$ compares to $f_{1,2}(\mathcal{M}')$ and $\star^i\text{-HOT}_{2,2}(\mathcal{M})$. Indeed, if $\star^i\text{-HOT}_{2,2}(\mathcal{M}'')$ is greater than $f_{1,2}(\mathcal{M}')$, or worse $\star^i\text{-HOT}_{2,2}(\mathcal{M})$, this would raise serious flags about how successful we can expect the mesh optimization algorithm to be.

As an initial step towards addressing this concern, we present two examples in Figures 4.1 and 4.2. In both examples for $i = 1, 2$ and Example 2 for $i = 0$, we find that

$$\star^i\text{-HOT}_{2,2}(\mathcal{M}'') < f_{i,2}(\mathcal{M}').$$

For $i = 0$ in Example 1,

$$f_{0,2}(\mathcal{M}') < \star^0\text{-HOT}_{2,2}(\mathcal{M}'') < \star^0\text{-HOT}_{2,2}(\mathcal{M}).$$

Thus, these examples are in support of the optimization algorithm.

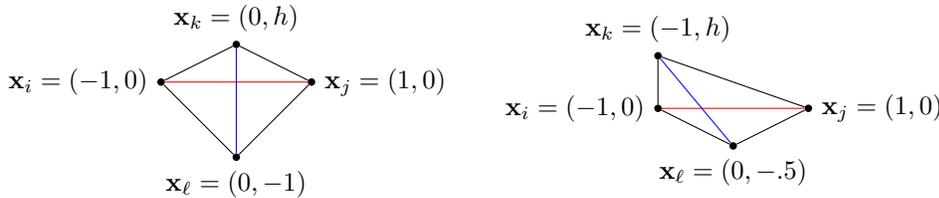


FIG. 4.1. *Non-Delaunay, Delaunay. Example 1 (left), Example 2 (right). In both examples, vertices x_i, x_j and x_l remain fixed, and x_k is a free vertex whose y -coordinate h is contained in $(0, 1)$. All vertex weights are 0.*

5. Non-convexity of $\star^i\text{-HOT}_{2,2}$ energy function. In optimization problems, working with a convex function is desirable because a local minimum of a convex function is also a global minimum. In [5], it is asserted that the HOT energy function is not in general convex, but no evidence is provided for the claim. In this section, we provide the evidence, confirming the claim.

Given a function $f : X \rightarrow \mathbb{R}_{\geq 0}$, we say f has *convex contours* if for all $c \in \mathbb{R}$, the set

$$\{x \in X : f(x) \leq c\}$$

is convex. If f is convex, then f has convex contours. Thus, if we can show that the contours of $\star^i\text{-HOT}_{2,2}$ are non-convex, we will have shown that $\star^i\text{-HOT}_{2,2}$ is non-convex function.

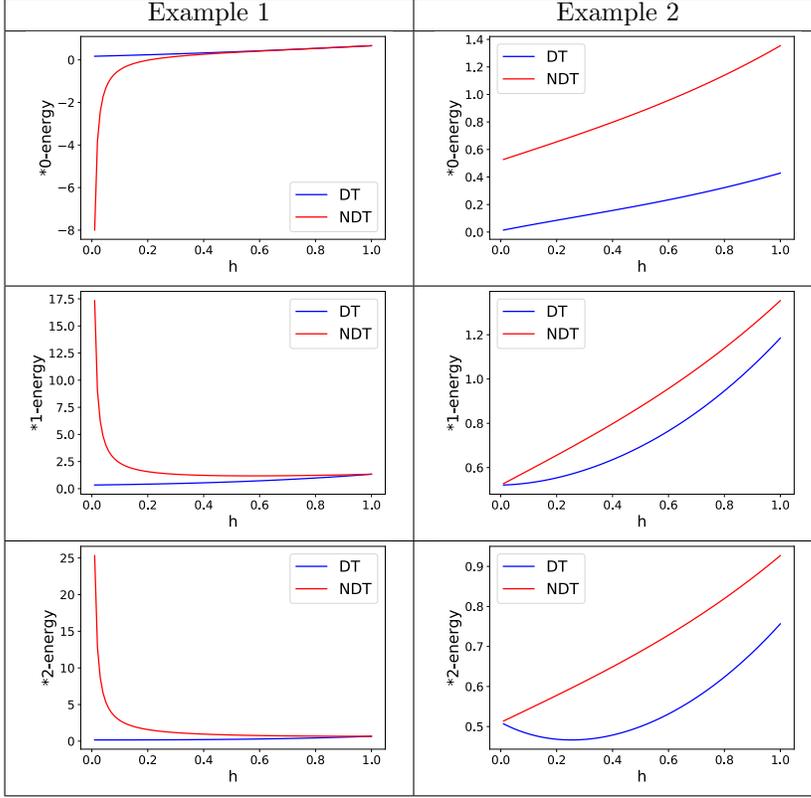


FIG. 4.2. Energies for example meshes in Fig. 4.1. By \star^p -energy for the Delaunay mesh, we mean $\star^p\text{-HOT}_{2,2}(\Delta_{i\ell k}) + \star^p\text{-HOT}_{2,2}(\Delta_{k j \ell})$ and for the non-Delaunay mesh, we mean $f_{p,2}(\Delta_{ijk}) + f_{p,2}(\Delta_{ij\ell})$.

Consider the Delaunay mesh \mathcal{M} corresponding to the set of points $\{(0, -1), (0, 1), (4, 0), (8, 0), (6, 0)\}$ as shown in Figure 5.1. For $(x, y) \in \mathbb{R}^2$, let $\mathcal{M}(x, y)$ denote the abstract triangulation obtained from \mathcal{M} by moving the vertex at $(6, 0)$ to (x, y) , keeping the face relationships of \mathcal{M} . Let $\mathcal{M}_D(x, y)$ be the Delaunay triangulation corresponding to $\mathcal{M}(x, y)$. The bottom right subfigure shows a contour plot for the function

$$(x, y) \mapsto \star^1\text{-HOT}_{2,2}\mathcal{M}_D(x, y),$$

and the top right subfigure shows a contour plot for the function

$$(x, y) \mapsto f_{1,2}\mathcal{M}(x, y).$$

Observe that the contours of both functions are non-convex, thus $\star^i\text{-HOT}_{2,2}$ and $f_{1,2}$ in general are not convex. The colors in the contour plot may be hard to see. So as further proof, Figure 5.1 (bottom left) plots h versus $\star^1\text{-HOT}_{2,2}(\mathcal{M}_D(h, 0))$. Observe that this curve is non-convex, establishing once again that that, in general, $\star^1\text{-HOT}_{2,2}$ is not convex.

THEOREM 5.1. $\star^i\text{-HOT}_{2,2}$ and $f_{1,2}$ are, in general, not convex.

5.1. Minimizing max energy triangle. There are other reasons convex contours are desirable, which we now discuss. Let v be a vertex in a mesh \mathcal{M} . Let

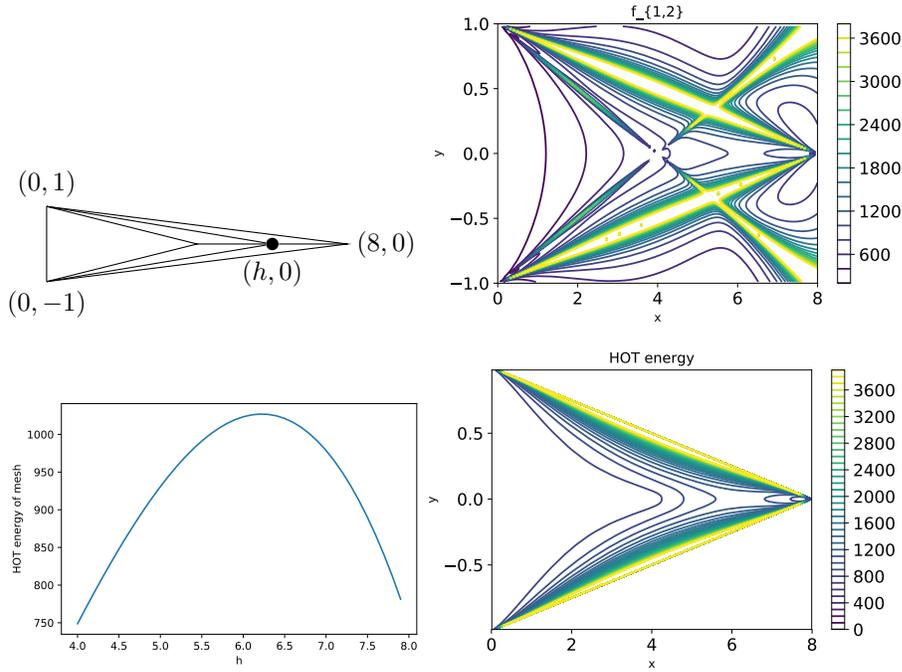


FIG. 5.1. This example shows that in general \star^1 -HOTA_{2,2} and $f_{1,2}$ are not convex

$\{v_1, v_2, \dots, v_k\}$ be the vertices adjacent to v , listed in the order they appear in the boundary of $\mathcal{P}(v)$, the patch of triangles in which v is a vertex. For $1 \leq j \leq k$ and $(x, y) \in \mathbb{R}^2$, let $\Delta_j(x, y)$ denote the triangle with vertices v_j, v_{j+1} and (x, y) . Fix $i = 0, 1, 2$. Define

$$f_j(x, y) = \star^i\text{-HOTA}_{2,2} \Delta_j(x, y).$$

If the contours of f_j are convex, then f_j is called a quasi-convex function. As discussed in [1], if all the f_j are quasi-convex, then generalized linear program (GLP) algorithms can be used to solve

$$\min_{(x,y) \in \mathbb{R}^2} \max_j \star^i\text{-HOTA}_{2,2} \Delta_j(x, y).$$

We studied the contours of the f_j and demonstrated that they are not convex. See Figure 5.2.

6. Barriers to mesh inversion. Consider a vertex v in a mesh \mathcal{M} . Let $\mathcal{P}(v)$ denote the *patch* for v , the union of the triangles in \mathcal{M} containing v . Note $\mathcal{P}(v)$ is star-shaped, containing a *kernel* sub-polygon from which all the edges of $\mathcal{P}(v)$ are visible without obstruction. Consider the optimization step where it is v 's turn to move. Let $\mathcal{M}(x, y)$ denote the labeled abstract triangulation obtained by moving v to (x, y) . If the new position for v lies outside the kernel of $\mathcal{P}(v)$, then we say \mathcal{M} has been *inverted*. An inverted triangulation is not truly a triangulation. Indeed, to obtain a triangulation we must change edge relationships between some of the vertices; otherwise we will have overlapping triangles.

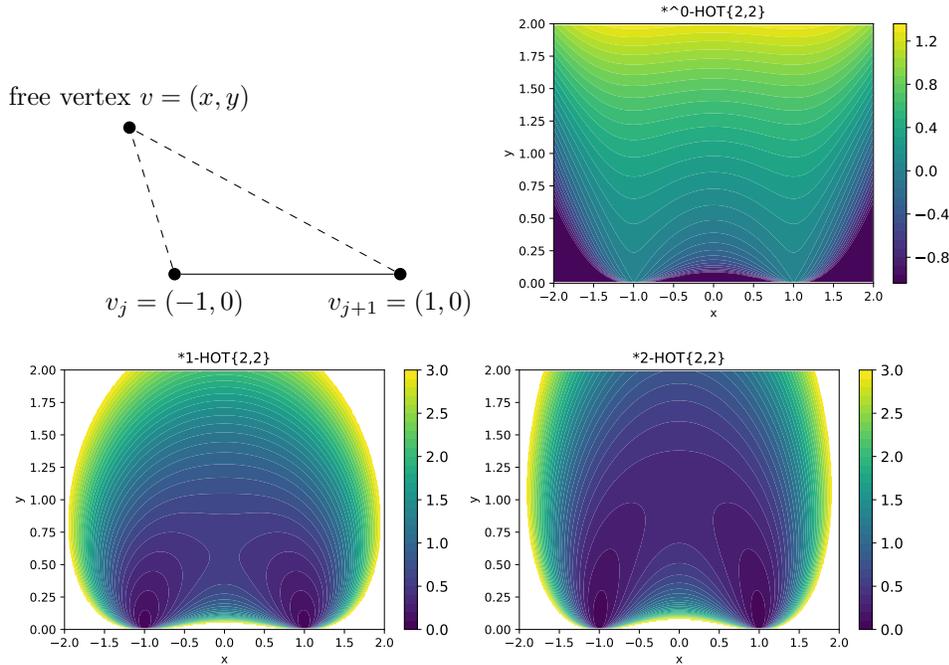


FIG. 5.2. Contour plots of the function f_j

In some applications, it is desirable to maintain the connectivity of the original mesh. In those situations, it is desirable that the energy function being optimized goes to infinity as v approaches the boundary of the kernel of $\mathcal{P}(v)$. This way, when we optimize the mesh energy, the mesh is protected against inversion. We will say that an energy function \mathcal{E} has a *barrier* if for every regular mesh \mathcal{M} , we have $\mathcal{E}(\mathcal{M}(x, y)) \rightarrow \infty$ as we move (x, y) to the boundary of the kernel of $\mathcal{P}(v)$.

We have shown the extrapolation of the \star^i -HOT $_{2,2}$ energy discussed in Section 3 do not have barriers. However, in some sense they come very close to having a barrier. We will make this precise in Theorem 6.2, but first consider the following examples.

Let $\mathcal{M}(x, y)$ be the abstract triangulation from Figure 5.1. Consider the kernel K of $(6, 0)$ in $\mathcal{M}(6, 0)$; that is, the polygon with vertices $(4, 0)$, $(8, 0)$, $(\frac{16}{3}, -\frac{1}{3})$ and $(\frac{16}{3}, \frac{1}{3})$. Observe that as (x, y) moves from inside K towards any point on the boundary of K except $(4, 0)$ and $(8, 0)$, the energy $f_{1,2}(\mathcal{M}(x, y)) \rightarrow \infty$. However, as (x, y) moves towards $(4, 0)$ or $(8, 0)$, the energy $f_{1,2}(\mathcal{M}(x, y))$ approaches a finite value (in fact a local minimum), demonstrating that $f_{1,2}$ does not have a barrier. Notice that both these local minima are bad because they result in degenerate triangles. Moreover, numerical optimization may be unstable here due to the close proximity of these minima to barriers.

Figure 5.2 gives another example showing that $f_{1,2}$ does not have a barrier. There the energy goes to infinity as (x, y) approaches any point on the x -axis except $(-1, 0)$ or $(1, 0)$. In these two cases, the energy approaches a finite value.

At first glance, it may seem like discussing the existence of barriers for the extrapolations $f_{i,2}$ and $t_{1,2}$ instead of barriers for \star^i -HOT $_{2,2}$ is the wrong discussion to be having since, after all, we seek HOT optimized meshes. However, given that

each loop of the mesh optimization algorithm is moving vertices using the \star^i -HOT_{2,2} extension $f_{i,2}$, studying barriers of the extended functions are, in fact, exactly the thing to be discussing.

THEOREM 6.1. *The extrapolated functions $f_{i,2}$ and $t_{1,2}$ do not have barriers, $i = 0, 1, 2$.*

For $t_{1,2}$ we now make a more technical, but more informative, statement that shows exactly why $t_{1,2}$ does not have a barrier, telling us what approach directions yield finite energy.

THEOREM 6.2. *let \mathcal{M} be a regular mesh and x_k an interior vertex of \mathcal{M} . Let $\mathcal{M}(x, y)$ denote the labeled abstract triangulation resulting from moving x_k to (x, y) , keeping all other vertices and all weights fixed, maintaining the face relationships of \mathcal{M} . Fix a coordinate system in \mathbb{R}^2 so that some edge opposite x_k is contained in the x -axis. For each such edge e_{ij} , say connecting the points $(s_i, 0)$ and $(s_j, 0)$ where $s_i < s_j$, define*

$$p_{ij}(x) = (s_j - s_i)(x - s_i)(x - s_j) - w_i(x - s_j) + w_j(x - s_i) - w_k(s_j - s_i).$$

Fix $c \in \mathbb{R}$. If $p_{ij}(c) \neq 0$ for some edge e_{ij} , then $t(\mathcal{M}(c, y)) \rightarrow \infty$ as $y \rightarrow 0$. Otherwise, $t(\mathcal{M}(c, y))$ limits to a finite value.

7. A modified energy function $\overline{\text{HOT}}$. We propose adjusting the HOT energy function to produce a new energy function $\overline{\text{HOT}}$ to address some of the issues we have identified with the HOT energy.

New energy defined: Let \mathcal{M} be a regular triangulation and σ an edge in \mathcal{M} connecting vertices x_i and x_j . Let Δ be a triangle in \mathcal{M} in which σ is an edge and let x_k denote the vertex opposite σ in Δ . First define the energy of σ relative to Δ as follows:

$$\star^1\text{-}\overline{\text{HOT}}_{2,2}(\sigma, \Delta) = \left(\frac{1}{a(\Delta)^2} \right) \frac{1}{3} (d_{ij}^3 h_k + d_{ij} h_k^3 + d_{ji}^3 h_k + d_{ji} h_k^3).$$

Now define the $\overline{\text{HOT}}$ energy of σ to be

$$\star^1\text{-}\overline{\text{HOT}}_{2,2}(\sigma) = \sum \star^1\text{-}\overline{\text{HOT}}_{2,2}(\sigma, \Delta),$$

where the sum is taken over all triangles Δ in \mathcal{M} in which σ is an edge. Finally define

$$\star^1\text{-}\overline{\text{HOT}}_{2,2}(\mathcal{M}) = \sum_{\sigma \in \Sigma^1(\mathcal{M})} \star^1\text{-}\overline{\text{HOT}}_{2,2}(\sigma).$$

Similarly, we modify the definitions of \star^0 -HOT_{2,2} and \star^2 -HOT_{2,2} by dividing them by the squared areas of triangles to produce new energy functions \star^0 - $\overline{\text{HOT}}_{2,2}$ and \star^2 - $\overline{\text{HOT}}_{2,2}$.

Features of new energy: $\overline{\text{HOT}}$ has desirable properties that HOT does not have.

1. For $i = 1, 2$, the formula-based extension of \star^i - $\overline{\text{HOT}}_{2,2}$ to the class of labeled abstract triangulations in which all vertices have equal weights has a barrier. In this class, the energy \star^0 - $\overline{\text{HOT}}_{2,2}$ goes to $-\infty$ as a vertex is moved to the boundary of its patch, which is bad since we seek energy minima. We suspect modifying the theory based extrapolation HOT instead of the formula based extrapolation will yield barriers for all stars.

We conjecture that the formula-based extension of $\star^i\text{-}\overline{\text{HOT}}_{2,2}$ has a barrier for $i = 1, 2$ even if the vertex weights are unequal, but have not completed the verification of this claim.

2. $\overline{\text{HOT}}$ is scale invariant. This means that triangles of the same shape but different size contribute equally to $\overline{\text{HOT}}$. In the HOT energy, the larger triangle makes a larger contribution.
3. $\overline{\text{HOT}}$ is dimensionless. Observe that the units of the HOT energy are length to the fourth power. By dividing terms by squares of triangle areas to obtain $\overline{\text{HOT}}$, we create an unitless energy function.

8. Conclusion. Our work presented in this paper raises many questions about the landscape of the HOT energy function and about how best to optimize it. In future work, we plan to address the following.

1. Is using the theory-based or the formula-based extrapolation of the HOT energy function better for obtaining a HOT optimized mesh?
2. Does updating the (labeled abstract) triangulation to a regular triangulation after all the vertices and weights have been moved always reduce the system energy? Or is it possible that updating increases the system energy?
3. Just because $\star^i\text{-}\text{HOT}_{2,2}$ is non-convex, it does not necessarily mean there are multiple local minima. If there are multiple minima, it does not necessarily mean that converging to one of them is dramatically worse than converging to another. We will study the existence of multiple local minima, and the energy differences between them.
4. Do $\overline{\text{HOT}}$ optimized meshes yield good solutions to PDEs?

Acknowledgements. This research was supported in part by an appointment with the NSF Mathematical Sciences Summer Internship Program sponsored by the National Science Foundation, Division of Mathematical Sciences (DMS). This program is administered by the Oak Ridge Institute for Science and Education (ORISE) through an interagency agreement between the U.S. Department of Energy (DOE) and NSF. ORISE is managed by ORAU under DOE contract number DE-SC0014664. This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research (ASCR), Applied Mathematics Program. Sandia National Laboratories is a multi-mission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC., a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525. Sandia National Laboratories is a multi-mission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC., a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

REFERENCES

- [1] N. AMENTA, M. BERN, AND D. EPPSTEIN, *Optimal point placement for mesh smoothing*, J. Algorithms, 30 (1999), pp. 302–322.
- [2] M. DESBRUN, A. N. HIRANI, M. LEOK, AND J. E. MARS DEN, *Discrete exterior calculus*, arXiv preprint math/0508341, (2005).
- [3] A. GILLETTE, *Notes on discrete exterior calculus*, tech. rep., Tech. Rep., University of Texas at Austin, Austin, Texas, USA, 2009.
- [4] F. D. GOES, P. MEMARI, P. MULLEN, AND M. DESBRUN, *Weighted triangulations for geometry processing*, ACM Transactions on Graphics (TOG), 33 (2014), p. 28.

- [5] P. MULLEN, P. MEMARI, F. DE GOES, AND M. DESBRUN, *Hot: Hodge-optimized triangulations*, ACM Transactions on Graphics (TOG), 30 (2011), p. 103.
- [6] C. VILLANI, *Optimal transport*, vol. 338 of Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences], Springer-Verlag, Berlin, 2009. Old and new.