Discrete Exterior Calculus for Meshes with Concyclic Polygons

Alexander Schier, Reinhard Klein

January 24, 2023

Discrete exterior calculus (DEC) is a numerical method for solving partial differential equations on meshes with applications in computer graphics, numerics, and physical simulations. It discretizes PDEs in a way such that important integral theorems hold exactly instead of being approximated. The drawback of the method is that it has strong requirements on the mesh. While the narrow range of admissible meshes mentioned in the original work could be widened, current methods still exclude an essential category of meshes, i.e., meshes with concyclic triangle pairs. Such meshes are common, as many synthetic meshes, e.g., triangulations of CAD models, handcrafted 3D models, and even results of surface meshing algorithms, contain such triangle pairs. Our paper describes an approach that allows us to use meshes with arbitrary triangulations that may contain concyclic triangle pairs by defining DEC operators for concyclic polygons.

Keywords: Discrete Exterior Calculus, Differential Forms, Discrete Differential Geometry, Mesh Processing

Published in: Computer Aided Geometric Design, Vol 101, 2023
DOI: `https://doi.org/10.1016/j.cagd.2023.102170`

1. Introduction

Discrete exterior calculus (DEC) introduced by Hirani in 2003 [Hir03] is a numerical scheme for solving PDEs on meshes that found a wide range of applications in the computer graphics community. The scheme *defines* all quantities on mesh elements like vertices, edges, and faces instead of approximating them on these elements using a more general function. Typical computer graphics algorithms work with given meshes instead of creating application-specific meshes, so the scheme is well-suited for such applications.

Important integral theorems like Stokes' theorem or the divergence theorem hold *exactly* on mesh elements in the DEC scheme, but they are also restricted to the mesh elements. While a more general scheme must allow integrating over arbitrary parts of a manifold, DEC restricts itself to evaluations on discrete mesh elements, but there it is numerically exact. For example, in more general schemes, one could approximate a line integral over any path on the manifold, and DEC restricts the evaluation of such integrals to a chain of mesh edges, but the integration along these edges is numerically exact. This feature is highly desirable in computer graphics, as the input often already is discretized as a mesh, which on the one hand, should not be changed anymore, and on the other hand, does not provide more information than the mesh geometry itself.

An applied introduction to discrete exterior calculus can be found in the SIGGRAPH Course by Crane et al. 2013 [CdGDS13], and the earlier course notes from SIGGRAPH Asia 2008 [GWDS08], which also present many uses of DEC in computer graphics. Further applications include fluid simulations, like Darcy flow [HNC15], liquid surrounded by air [ATW15] and the surface Navier-Stokes equation [NRV17]. DEC also was used in the heat method for geodesic calculations [CWW13] and the transport of tangent vector fields [SSC19]. In vector field processing, DEC was used for Hodge decomposition of vector fields [ZDWT19], and computing 3D cross fields [CC19].

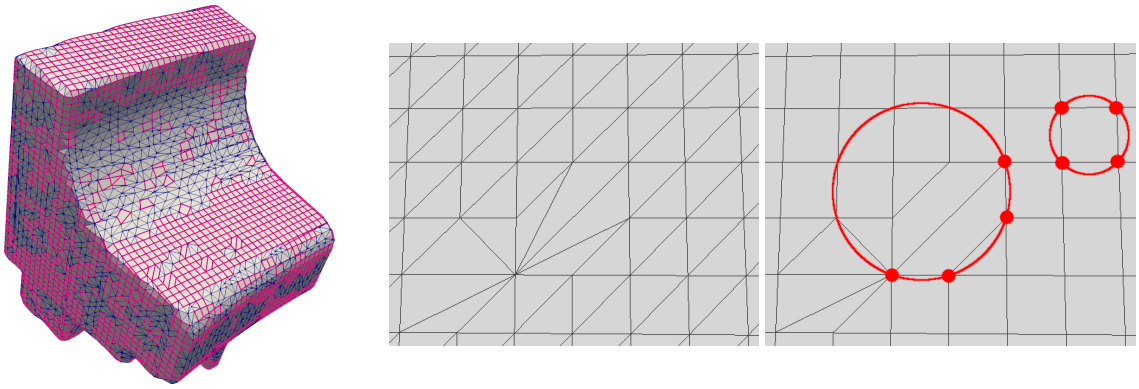

Figure 1: Left: The triangulation of the fandisk model contains many concyclic triangle pairs, both on its flat and at its curved parts that can be merged into concyclic polygons. Right: By removing the edges between concyclic triangles, we obtain a polygonal mesh that does not contain any concyclic polygon pairs without remeshing other parts of the model.

For a long time, DEC was thought to have strong requirements on the mesh, i.e., that the definition of the circumcentric Hodge star operators requires the mesh to be *well-centered* which means that no triangle is allowed to have an obtuse angle. This requirement later could be weakened to allow the use of Delaunay meshes, and finally, the scheme was shown to work on general triangulations which may induce negative dual edges, as long as the dual mesh is non-degenerate, i.e., all dual cells must have a non-zero area, and all dual edges must have a non-zero length. Unfortunately, there are many meshes for which the dual mesh contains degenerate edges. Some meshes also lead to degenerate dual faces.

A prime example of such a mesh is the triangulation of a rectangular grid. For a triangle with a right angle, the circumcenter lies on the midpoint of the edge opposite to the right angle, and, consequently, for triangulated rectangular grids, both circumcenters lie at the same point, and the length of the dual edge is zero. One could consider if DEC could be extended for rectangular quads to handle such meshes. However, the class of non-admissible meshes is much larger as not only right-angled triangle pairs cause problems, but each pair of *concyclic* triangles induces a degenerate dual edge. The problem is also not limited to triangle *pairs*, but an arbitrary number of triangles may be concyclic, i.e., have the same circumcenter. Notable sources of concyclic meshes are triangulations of rectangular grids, e.g., CAD and handcrafted 3D models created using quads. Even standard surface meshing algorithms like marching cubes [LC87] and dual contouring [JLSW02] naturally create triangulations with a degenerate dual mesh for some ways of how the surface intersects the grid used by the algorithm.

Fig. 1 shows the fandisk model ([BZK09]), remeshed by merging concyclic triangles into concyclic polygons. The flat portions of the model have mostly rectangular quads, which probably stem from the used triangulation algorithm. The curved parts of the model also contain concyclic polygons, of which many are far from being rectangular. The right side of the figure shows a part of the mesh that contains both rectangles and non-rectangular quads, which were created by merging concyclic triangles into polygons.

A usual way to deal with such restrictions is remeshing, or as suggested by previous literature, using perturbation of the vertex positions to break up concyclic triangle pairs. However, there are two reasons against this approach. First, one often does not want to modify a model because it must not lose accuracy or is used in a larger pipeline in which several algorithms use the same model, and the results must match later on. Second, one usually does not know the function values at the positions of the perturbed vertices. A 3D scan, for example, provides a point cloud with properties like vertex colors and normals. When moving a vertex to improve mesh quality, the color and normal at the new position are unknown.

We show in section 8 that changing the vertex coordinates to break concyclic triangle pairs would require perturbations that are significantly smaller than the original mesh resolution for the calculations to converge to the correct solution. On the other hand, we observed that iterative solvers do not converge when the perturbations are too small. Our novel scheme avoids the problem by merging the concyclic triangle pairs into concyclic polygons until the polygon forms no concyclic pair with any adjacent face. We show the necessary changes in the primal and dual meshes and how the definitions of the DEC operators need to be

adapted when using primal meshes that contain concyclic polygons, so we can keep using circumcentric duality with all its advantages. In addition, we show that it is easy to adapt existing software to support our new approach.

The main contributions of the paper are:

- We describe the effect of concyclic triangle pairs on the circumcentric dual meshes, usually used for discrete exterior calculus.
- We provide a thorough analysis of affected DEC operators and why they are ill-defined on such meshes.
- We propose a scheme that allows us to use DEC on concyclic meshes and define non-singular operators by using an intrinsic polygonal mesh.
- We show the convergence of our scheme by solving a Poisson problem with a known solution on meshes with many concyclic triangle pairs and an eigenvalue problem on spheres with concyclic triangulations.
- We define a new operator that maps between concyclic triangle and polygon meshes to allow using the novel polygonal operators with unchanged triangle meshes.
- Finally, we show how existing DEC software can easily be adapted to support our proposed method.

2. Previous work

Discrete exterior calculus was first described by Hirani in 2003 [Hir03] for well-centered meshes, i.e., meshes in which each triangle contains its circumcenter. This property is equivalent to the requirement that no triangle is allowed to contain non-acute angles. In general, most triangulations are *not* well-centered, and for most shapes, it is hard to create a well-centered triangulation. An approach for remeshing was described by Mullen et al. in 2011 [MMdGD11] who propose a method to create Hodge-optimized triangulations for the use with circumcentric dual meshes.

In 2013 Hirani et al. described a construction using negative volumes [HKV13] that weakened the requirements using signed volumes similar to [Gli05], allowing the use of discrete exterior calculus with Delaunay meshes. As there are well-known remeshing methods for creating Delaunay meshes and many engineering meshes are already designed to have the Delaunay property, this method allows the use of discrete exterior calculus in many more applications than before.

Another DEC modification for dealing with general meshes is using a barycentric dual mesh, and either a Hodge star operator based on Whitney-forms or a Galerkin Hodge star [MHS16]. The drawback of this approach is that while the discrete Hodge star operator is a sparse matrix, it does not have a sparse inverse, so DEC operations involving the inverse operator are computationally costly.

In 2018 Hirani et al. published a Corrigendum [HKV18] to the Delaunay Hodge star paper stating that the DEC construction in the paper can be used for an even broader category of meshes. In a follow-up paper, Mohamed et al. described a method for using DEC with circumcentric duality for general meshes [MHS18]. DEC can be used for almost every mesh with their construction, but one issue is still not solved by any previous approach. The last mesh requirement for using DEC with a circumcentric dual mesh is that the dual mesh is not allowed to have edges of length zero. At first, this limitation looks like a theoretical problem, but meshes affected by this problem are more common than one may think. The problem is acknowledged in the literature but not solved without remeshing or perturbation as suggested, e.g., in [MHS18].

This paper will show how to extend the discrete exterior calculus scheme for meshes containing concyclic triangle pairs that induce dual edges of length zero to lift this last restriction.

3. The Prevalence of Concyclic Triangle Pairs

We examined the prevalence of concyclic triangle pairs in several data sets by testing if the angles at the corners opposite to an edge sum up to a value in the range $[\pi-\varepsilon, \pi+\varepsilon]$ with an epsilon of $0.0001 \cdot 2\pi$. We chose to test for an ε larger than the machine precision because the condition number of the DEC matrices depends on ε, so almost concyclic triangle pairs are also numerically relevant. For example, in our experiments,

iterative solvers like BiCGStab did not converge with a too small ε. Therefore near-concyclic triangle pairs should be treated the same as concyclic triangle pairs to increase numeric stability.

An extreme example for a concyclic mesh would be a plane with a triangulated rectangular grid, which occurs, for example, in CAD models. In such a plane, all triangles are concyclic with the circumcenters lying on the shared edge. We found the Stanford models to contain only a few concyclic triangles, but in the COSEG [WAVK+12] data sets, we found a large percentage of concyclic triangles. While finding 35% concyclic triangles in the set of chairs is no surprise, the goblet models also contained many concyclic triangle pairs with 25% of all triangles being part of a concyclic triangle pair, even in the curved areas of the models. The handcrafted models in the 3D model repository of Keenan Crane[1] unsurprisingly also contain a relatively large number of concyclic triangles. The Nefertiti model[2] contains only a small percentage of concyclic triangles due to the very fine resolution, which captures a lot of surface noise in the reconstruction. Still, it contains 40 concyclic triangle pairs.

A table with the number of concyclic triangle pairs in different data sets is provided in B.

3.1. Sources of Concyclic Triangle Pairs

Many synthetic meshes are created from quads. One design criterion of such meshes is to optimize the direction of the mesh edges to align with the curvature and the edges of the object, which are often orthogonal in man-made objects. For example, the results of many CAD operations are rectangular quads, and many deformations of these quads retain the sum of opposite angles because of symmetry reasons. Hand-crafted models often use quads because they have advantages for modeling, and it is easier to work with them. Good quadrangulations allow for simple refinement, mesh cutting, and other operations along features lines, so they are easier to handle than triangle meshes. Such quad meshes are often highly structured, e.g., designed to align with straight edges of the modeled object, which often leads to concyclic or even rectangular faces.

Another source for concyclic triangle pairs are standard surface triangulation algorithms, like marching cubes [LC87] and dual contouring [JLSW02], which intersect the surface with an orthogonal grid. This choice leads to regular structures in the resulting meshes and, for axis-aligned surfaces, even to right triangles or, in the case of dual contouring, rectangular quads.

Even when creating randomized meshes, like the example meshes with vertices in general position in section 9, one gets a number of nearly concyclic triangles pairs when they are not explicitly broken up by adding further perturbation until no concyclic triangle pairs remain.

4. Discrete Exterior Calculus for Surface Meshes

For simplicity, we use the definition of a 2D simplicial complex as a discretization of a 2-manifold mesh, in which two triangles only intersect in a shared edge, and there are no hanging edges or isolated vertices. For the general definition in arbitrary dimensions and more details, see [Hir03]. In the following, we assume that the *primal* mesh is non-degenerate, i.e., that it does not contain edges of length zero and no triangles with three colinear vertices, but may contain concyclic triangle pairs.

We adopt the usual notation of using $\sigma^k \in K^k$ for k-simplices, but use $\star\sigma^k$ to indicate the corresponding dual (n-k)-cell to better distinguish the notation for dual cells from the notation $\star\omega^k$ for discrete dual differential forms. We denote differential k-forms on the manifold by $\hat{\omega}^k$ and write $\omega^k \in \Omega^k$ for the corresponding *discrete* differential forms, i.e., the vector of the differential k-form $\hat{\omega}^k$ integrated over all k-simplices in the mesh is $\omega^k = \left(\left\langle \hat{\omega}^k, \sigma_1^k \right\rangle, \ldots, \left\langle \hat{\omega}^k, \sigma_n^k \right\rangle\right)^T$.

4.1. Dual Meshes

The usual choice in discrete exterior calculus is to use a *circumcentric* dual mesh, which has a simple construction and yields sparse matrices for the DEC operators. The disadvantage of using a circumcentric dual mesh is that the dual vertex corresponding to a triangle is not guaranteed to lie inside the triangle. When

[1] https://www.cs.cmu.edu/~kmcrane/Projects/ModelRepository/

[2] http://nefertitihack.alloversky.com/

a triangle contains an obtuse angle, the circumcenter lies in the other half-space than the obtuse angle with respect to the edge opposite to the angle and when the triangle contains a right angle, the circumcenter lies exactly on the edge, cf. Fig. 2. Other choices like barycentric duality [MHS16] avoid the problem, but they have drawbacks like that the inverse for the sparse matrix of the barycentric Hodge star is not sparse itself, and thus the computational complexity is worse for calculations that involve the inverse Hodge star operator.

4.2. Types of Triangle Pairs

When constructing dual edges, we have to consider four types of triangle pairs:

- A *well-centered* triangle pair consists of two triangles that contain their own circumcenter. All triangle angles are acute, and the dual edge intersects the primal edge shared by the two triangles at its midpoint. The original DEC definition by Hirani [Hir03] required a mesh solely consisting of such triangles.
- A *Delaunay* triangle pair consists of two triangles with the Delaunay property [Del34], i.e., the circumcircles of the triangles do not contain the apex of the other triangle. One of the triangles has an obtuse angle at its apex in such a pair, and both dual vertices lie inside the non-obtuse triangle. In 2013, Hirani et al. [HKV13] described how to handle such meshes.
- A *non-delaunay* triangle pair has at least one triangle for which the interior of the circumcircle contains the apex of the other triangle. Because the apex of the other triangle is inside the circumcircle, the circumcenter is nearer to the apex of the other triangle than to the apex of its own triangle, and thus the dual edge between the circumcenters is flipped. Mohamed et al. [MHS18] showed how to handle the non-Delaunay case.
- A *concyclic* triangle pair consists of two triangles that share the same circumcenter, which happens when the angles opposite to the shared edge sum to π. When both triangles have a right angle at their apex, the circumcenter lies on the shared edge. Otherwise, one triangle has an acute, and the other one has an obtuse angle at its apex, and the circumcenter lies inside the triangle with the acute angle. As the two dual vertices lie at the center of the same circle, the length of the dual edge is 0. This last case is addressed in this paper.

To see that the list is exhaustive, one can consider the possible locations of the circumcenters of a triangle pair and the resulting dual edge.

- The dual edge either crosses the primal edge or lies entirely in one half-space with respect to the primal edge. We consider the ambiguous case of a dual vertex lying on the primal edge as an intersection.
- The dual edge either has a positive length, a negative length (flipped edge), or length zero (degenerate edge).

From these two properties we get the 3×2 combinations shown in Fig. 2. Note that except for the well-centered triangle pairs, the dual edge is a geodesic that can intersect more than one triangle as the circumcenters can lie far from their triangle. As we will see in the next section, this does not pose a problem, as we project the dual edge into the tangent plane of the directly adjacent triangles using only its length and do not need to unfold the intersected triangles explicitly.

How the different combinations affect the dual mesh, e.g., by causing overlaps, is depicted in Fig. 11. As described in [MHS18], such overlaps can be handled gracefully by using negative volumes and we will see in the following that using negative volumes also works for concyclic triangle pairs and that positive and negative areas cancel out exactly for such triangle pairs. Therefore, we can avoid numerically unstable computations by skipping both the positive and the negative contributions to the area of the dual cell.

4.3. Circumcentric Duality on curved 2-Manifolds

By definition, the circumcenter of each triangle lies in the tangential plane of the triangle. For a well-centered triangle, it lies inside the triangle itself and thus also on the 2-manifold mesh. The circumcenters of triangles that are not well-centered lie outside the triangles and, in general, not on the mesh surface. For an example,

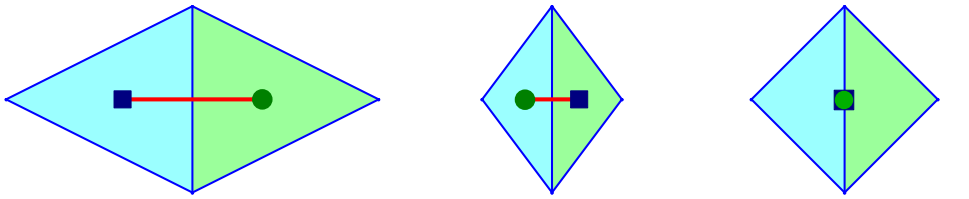

(a) well-centered triangles (positive dual edge)

(c) non-delaunay triangles (negative dual edge)

(e) concyclic triangles (degenerate dual edge)

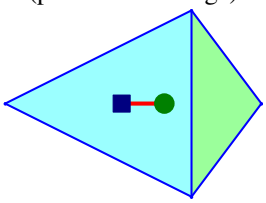

(b) Delaunay triangles (positive dual edge)

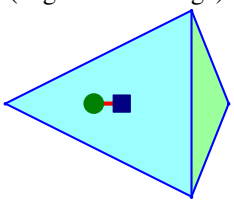

(d) non-delaunay triangles (negative dual edge)

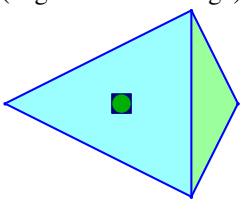

(f) concyclic triangles (degenerate dual edge)

Figure 2: Discrete exterior calculus was initially defined using well-centered triangles that contain their own circumcenter as shown in (a). The definition was later extended for Delaunay triangulations (b), which guarantees positive dual edges. Recently, this mesh requirement was weakened to allow for more general triangulations by accounting for flipped edges when computing dual cells so the configurations (c) and (d) are admissible. Collapsing dual edges as depicted in (e) and (f) are still a problem in the circumcentric DEC scheme.

see Fig. 3, which shows a non-delaunay triangle pair with a flipped dual edge and the projection of the circumcenters onto the plane of their neighbor triangle. Previous works like [HKV13] and [MHS18] did not address the problem explicitly but showed planar examples assuming that the method would extend to general 2-manifold meshes. Indeed, these approaches and our extension are well-defined for curved 2-manifold meshes, as unfolding the geodesic of the dual edge into the plane yields a straight edge [MMP87]. Therefore we do not need to compute the geodesic path explicitly when only the length of the dual edge is required. We do not even need to unfold the two triangles adjacent to the primal edge when we split the dual edge into two segments, each from the primal edge midpoint to one of the triangle circumcenters. The length of the dual edge is then the sum of the signed lengths of the two segments.

5. Meshes with Degenerate Dual Edges

We call a dual edge induced by a concyclic triangle pair *degenerate*, and we call the corresponding primal edge *removable*. The crucial property why a dual edge is degenerate is its length being 0, which causes many operations to be ill-defined for such edges. We will show that the corresponding primal edges can be removed such that the DEC operators on the remaining primal mesh stay consistent and the operations in the dual mesh become well-defined.

In the following sections, we will show why some operators are ill-defined on meshes containing degenerate dual edges. See A for the proofs that the other operators are still well-defined. We will see later in section 6 how to deal with the operators affected by degenerate dual edges.

5.1. Evaluation of (dual) forms

We assumed that the primal mesh is non-degenerate, and thus all operations that do not involve the dual mesh are *well-defined*. In particular, the evaluation of differential 1-forms on primal edges is unchanged even when they are removable. The evaluation of a dual 0-form on dual vertices is not changed, as all dual vertices are well-defined, even when the edge between them is degenerate. Evaluating a *dual* 1-form $\hat{\omega}^1_d$ on a degenerate

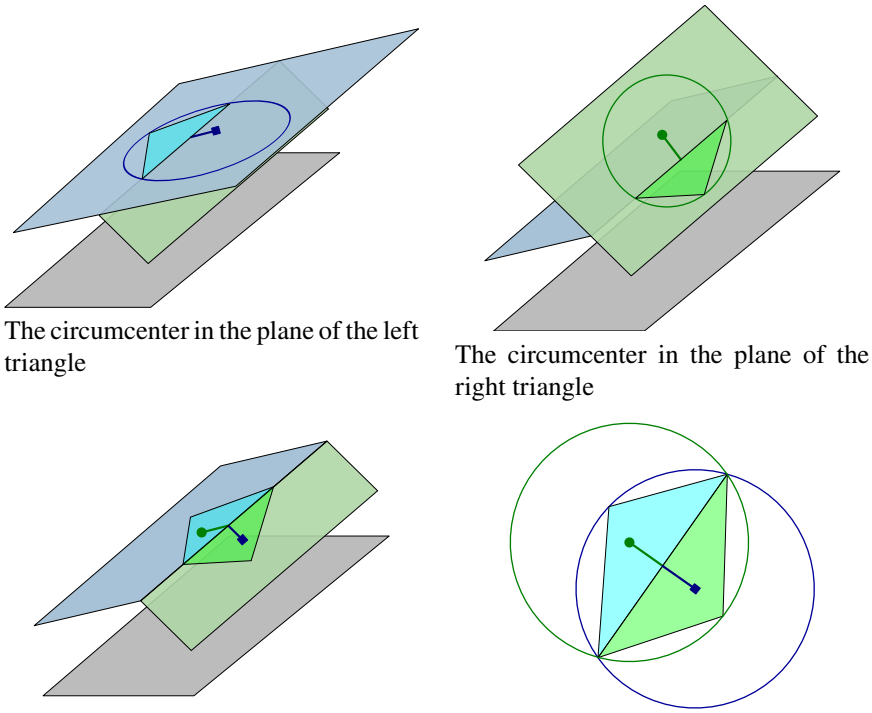

The circumcenter in the plane of the left triangle

The circumcenter in the plane of the right triangle

The circumcenters projected back on the the mesh

The kinked edge unfolds to a straight geodesic

Figure 3: The circumcenter lies in the tangential plane of a triangle. When it is outside of the triangle itself, in general, it does not lie on the 2-manifold mesh, and we project it onto the plane of the neighbor triangle using the known length of the edge segment and the fact that the segment is orthogonal to the edge between the triangles.

dual edge $\star\sigma^1$ yields 0 as the integration interval has length 0

$$\langle \hat{\omega}_d^1, \star\sigma^1 \rangle = \int_{\star\sigma^1} \hat{\omega}_d^1 = 0 \quad \forall \hat{\omega}_d^1 \in \Omega_d^1 \tag{1}$$

and the integral is *well-defined* and corresponds to the limit $|\star\sigma^1| \to 0$ naturally. Integration over dual cells also remains unchanged, although we will see later that some of the elementary triangles that make up the cell can have an area of zero.

5.2. Differential Operators

As the differential operators are purely combinatoric, they are not affected by metric properties like the length of dual edges. The primal differential operators do not require a dual mesh at all, and for the dual differential operators, we prove in A that they are not affected by degenerate dual edges.

5.3. Hodge Stars

The DEC Hodge star operators $\star_k$ map between discrete primal k-forms ω^k and discrete dual $(n-k)$-forms ω_d^{n-k}, such that for primal simplices σ_i^k and the corresponding dual cells $\star\sigma_i^k$ the following holds

$$(\star_k \omega^k)_i = \langle \star_k \hat{\omega}^k, \star\sigma_i^k \rangle = \frac{|\star\sigma_i^k|}{|\sigma_i^k|} \langle \hat{\omega}^k, \sigma_i^k \rangle = (\omega_d^{n-k})_i \tag{2}$$

and the corresponding equation for the inverse Hodge star operators $\star_k^{-1}$ is

$$(\star_k^{-1}\omega_d^{n-k})_i = \langle \star_k^{-1}\hat{\omega}_d^{n-k}, \sigma_i^k\rangle = \frac{|\sigma_i^k|}{|\star\sigma_i^k|}\langle \hat{\omega}_d^{n-k}, \star\sigma_i^k\rangle = (\omega^k)_i \tag{3}$$

where $|\omega^k|$ is the signed volume of a k-simplex and $|\star\omega^{n-k}|$ is the signed volume of the corresponding dual cell. The primal Hodge star is always well-defined as all σ_i^k have positive lengths and areas, but it may yield inconsistent dual 0-forms, assigning different values to two vertices that lie on the same point of the mesh. For example, consider two concyclic triangles with the same area and a discrete primal 2-form ω^2, which assigns different values to these triangles. Then the dual vertices coincide at the same point on the mesh, but the dual 0-form $\star\omega^2$ assigns different values to the dual vertices. On the other hand, we will see that some of the Hodge star operators are only well-defined when the primal mesh does not contain concyclic triangles as they induce degenerate edges in the dual mesh. As before, we limit our analysis to 2D simplicial complexes and find that two of the Hodge stars are not well-defined on meshes with concyclic triangle pairs:

- The primal-to-dual Hodge star for 2-forms $\star_2$ is only well-defined for specific 2-forms.
- The dual-to-primal Hodge star for 1-forms $\star_1^{-1}$ is ill-defined for all dual 1-forms.

Primal-to-dual Hodge star for 2-forms

For a degenerate dual edge $[\star\sigma_i^2, \star\sigma_j^2]$, the dual vertices lie at the same point on the manifold and thus

$$\frac{1}{|\sigma_i^2|}\langle\hat{\omega}^2, \sigma_i^2\rangle = \frac{1}{|\star\sigma_i^2|}\langle\star_2\hat{\omega}^2, \star\sigma_i^2\rangle = \frac{1}{|\star\sigma_j^2|}\langle\star_2\hat{\omega}^2, \star\sigma_j^2\rangle = \frac{1}{|\sigma_j^2|}\langle\hat{\omega}^2, \sigma_j^2\rangle$$

As $\star\sigma_i^2$ and $\star\sigma_j^2$ are vertices we have by definition $\frac{1}{|\star\sigma_i^2|} = \frac{1}{|\star\sigma_j^2|} = 1$ and $\langle\star_2\hat{\omega}^2, \star\sigma_i^2\rangle = \langle\star_2\hat{\omega}^2, \star\sigma_j^2\rangle$ holds because the two dual vertices have the same position on the manifold. Therefore it follows that $\star_2$ is only well-defined for discrete 2-forms ω^2 with $\left(\star_2\omega^2\right)_i = \frac{(\omega^2)_i}{|\sigma_i^2|} = \frac{(\omega^2)_j}{|\sigma_j^2|} = \left(\star\omega^2\right)_j$ and for such 2-forms we have

$$\langle\hat{\omega}^2, \sigma_i^2\rangle + \langle\hat{\omega}^2, \sigma_j^2\rangle = \rho|\sigma_i^2 \cup \sigma_j^2| \quad \text{with} \quad \rho := \frac{\langle\hat{\omega}^2, \sigma_i^2\rangle}{|\sigma_i^2|} = \frac{\langle\hat{\omega}^2, \sigma_j^2\rangle}{|\sigma_j^2|} \tag{4}$$

The physical interpretation of Eq. 4 is that the discrete 2-form describes the integral over a density ρ that must be the same in the triangles σ_i^2 and σ_j^2. For general 2-forms, that do not satisfy Eq. 4, the Hodge star $\star_2$ is *ill-defined*, because the discrete form $\star\omega^2$ contains different values for two 0-simplices, which correspond to the same point on the manifold.

Dual-to-primal Hodge star for 1-forms

For the evaluation of primal 1-forms $\star_1^{-1}\hat{\omega}_d^1$ on removable primal edges σ^1, we have

$$\langle\star_1^{-1}\hat{\omega}_d^1, \sigma^1\rangle = \frac{|\sigma^1|}{|\star\sigma^1|}\langle\hat{\omega}_d^1, \sigma_d^1\rangle \tag{5}$$

and as $|\star\sigma^1| = 0$, the Hodge star $\star_1^{-1}$ is *ill-defined* because the denominator is 0.

5.4. Discrete Flat and Sharp operators

Hirani defined in [Hir03] several discrete sharp and flat operators for mapping vector fields to discrete 1-forms and vice versa, of which most are not affected by degenerate dual edges. Only the *dual-primal-primal* flat operator $\flat_{dpp}$ is ill-defined. A vector field can be evaluated on a 1-simplex σ^1 using $\flat_{dpp}$ by

$$\left\langle X^{\flat_{dpp}}, \sigma^1\right\rangle = \sum_{\sigma^n \succ \sigma^1} \frac{|\star\sigma^1 \cap \sigma^n|}{|\star\sigma^1|} X(\sigma^n)\cdot\vec{\sigma}^1 \tag{6}$$

which is clearly *ill-defined* when σ^1 is a removable edge, because $|\star\sigma^1| = 0$. The sharp operators $\sharp_{pp}$ and $\sharp_{pd}$ are still *well-defined* for exact forms, but for our method we would need a new interpolation function inside concyclic polygons.

5.5. Wedge Product

The two definitions of the primal-primal wedge product in [Hir03, section 7.1] are still *well-defined* for triangles with removable edges on their boundary. The weighting factor in the first definition is, for a 2D simplicial complex, the length of the dual edge intersected with the primal triangle. As the length of a degenerate dual edge is 0, the weighting factor for triangle edges $v_{\tau(k)}$ that are removable is 0 as well, such that they do not contribute to the weighted sum. The second definition of the wedge product is also not affected because it does not involve any metric terms.

5.6. Laplace Operators

The DEC Laplacians are composed of dual and primal derivative operators, and the Hodge stars needed to map between forms defined on the primal and dual mesh. We now show how the results for the operators described in the previous sections affect the Laplace operators.

The Laplace operator for primal 0-forms is defined as $L_0 := \star_0^{-1}\mathbf{d}_1 \star_1 \mathbf{d}_0$ and, as its definition does not contain operators which are ill-defined on meshes with degenerate dual edges, it is *well-defined*. This Laplacian corresponds to the area-weighted finite differences Laplace operator [GRS17], for which the combinatoric differential operators $\mathbf{d}_0$ and $\mathbf{d}_1$ define the used finite differences stencil, the Hodge star $\star_1$ contains the ratio between the primal and dual edge lengths and the Hodge star $\star_0^{-1}$ defines the area weighting.

On a triangulated rectangular grid, for example, the differential operators yield a 7-point stencil from the 6 adjacent edges at each inner vertex, of which the 2 diagonal edges are removable. We will see in section 7.2 that the cotangent weights of the diagonal edges are 0 and thus, both our scheme, which removes these edges because the adjacent triangles are concyclic, and the definition of the cotangent Laplacian, yield the well-known 5-point finite differences stencil for a rectangular grid.

The Laplacian for primal 2-forms $L_2 := \mathbf{d}_1 \star_1^{-1} \mathbf{d}_1\star_2$ and the Laplacian for dual 0-forms $L_0^{\text{dual}} := \star_2\mathbf{d}_1\star_1^{-1}\mathbf{d}_0$ are *ill-defined*, because their definitions involve the ill-defined inverse Hodge star $\star_1^{-1}$. The Laplace operator for dual 2-forms $L_2^{\text{dual}} := \mathbf{d}_1 \star_1 \mathbf{d}_1\star_0^{-1}$ is *well-defined* as all required operators are well-defined. The Laplace operator for primal 1-forms $L_1 = \star_1^{-1}\mathbf{d}_0 \star_2 \mathbf{d}_1 + \mathbf{d}_0 \star_0^{-1} \mathbf{d}_1\star_1$ and the dual Laplace operator $L_1^{\text{dual}} = \star_1\mathbf{d}_1 \star_0^{-1} \mathbf{d}_1 + \mathbf{d}_0 \star_2 \mathbf{d}_1\star_1^{-1}$ are both *ill-defined* as their definitions involve $\star_1^{-1}$.

6. Collapsing dual edges

We have shown that several important DEC operators are ill-defined on meshes that contain concyclic triangle pairs because of degenerate edges in the dual mesh and will now adapt the DEC scheme to deal with such meshes. Our solution to the problem is to collapse the edge in the dual mesh by removing the corresponding edge in the primal mesh, such that two concyclic primal triangles are merged into a concyclic polygon.

In the following, we describe the necessary changes and how the discrete differential forms and operators need to be adapted, and we will show that the previously ill-defined operators are well-defined on the changed mesh.

6.1. Changes in the Mesh

The deletion of a removable primal edge and its corresponding degenerate dual edge affects the following mesh elements:

- The primal and the corresponding dual edge are removed.
- The adjacent primal triangles are merged into a concyclic quadrilateral.
- The two dual vertices corresponding to the primal triangles are merged.

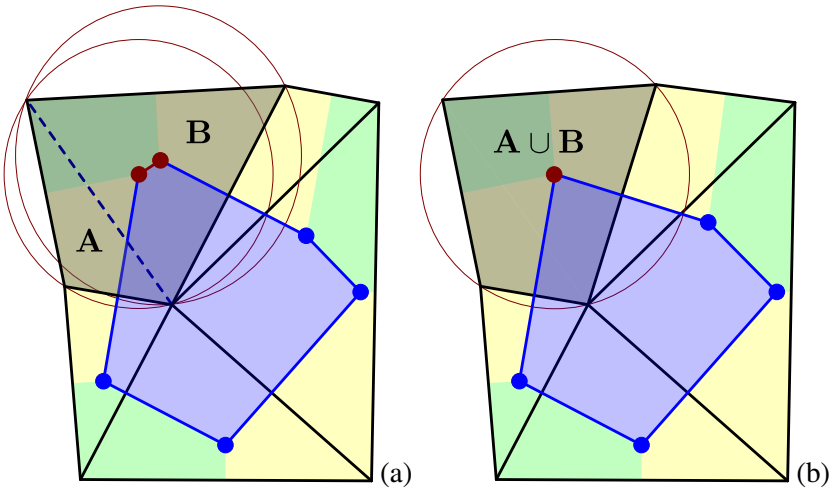

Figure 4: When the length of the dual edge (dark red) becomes zero, the primal triangles A and B that are dual to the endpoints of the edge become concyclic. We remove the primal edge (dashed dark blue) between triangle A and triangle B, such that the dual edge collapses into a single dual vertex corresponding to the quad $A \cup B$. As the vertices of the quad are concyclic, the dual mesh is still circumcentric.

There is no change in the primal vertices, and even though degenerate edges are removed from the boundary of adjacent dual cells, the removal does not change the shape of the dual cell as the lengths of the removed edges were zero. Note that more than two triangles can be concyclic, and then all these triangles are merged into the same concyclic polygon. In many real-world applications, most of these concyclic polygons are quadrilaterals that only consist of two merged triangles, but in general, an arbitrary number of triangles may be merged.

The virtue of our method is that the dual edges that intersect the boundary edges of the polygon are still orthogonal to the boundary, as the circumcenter of the polygon is the same as the circumcenters of the merged triangles. Thus we can keep using a circumcentric dual mesh and do not need to change much in the algorithms. In the following, we describe the changes for a concyclic triangle pair. The extension to concyclic polygons that consist of more than two triangles is straightforward.

6.2. Concyclic Polygons on 2-Manifolds

When removing edges, the 3D shape of the resulting polygons is, in general, no longer uniquely defined as the 2-manifold polygon shape is interpolated from more than three points. For concyclic triangle pairs on curved manifolds, there are two points of view.

In the exterior space, i.e., the 3D space in which the 2-manifold is embedded, only co-planar triangle pairs can be concyclic, as the circumcircles are defined in the tangential planes of the triangles. The more relevant perspective, as DEC operations are defined on the surface of the 2-manifold, is looking at *intrinsically* concyclic polygons, which may not be concyclic in 3D, but are concyclic on the surface itself. An example is the triangle pair in Fig. 3.

We can determine if a triangle pair is intrinsically concyclic without projecting the circumcenters by using the angle sum at the opposite triangle corners, which is an intrinsic surface metric and does not depend on the dihedral angle at the shared edge. A numerically more robust way is to use the sum of cotangents at these corners as a cotangent sum of 0 is equivalent to the dual edge being degenerate.

To avoid the problem that the 3D shape of the interior of a polygon is no longer uniquely defined, we do not remove the edges from the mesh itself but only skip them in computations. So all lengths, areas, and angles of the mesh elements stay the same, and we only change the adjacency information used in the DEC operators. The information lost due to this change describes the collapsed dual edges, which had, by definition, length 0 before and are thus not relevant for computations. In the primal mesh, the evaluation of a 1-form on removed edges is undefined, and the evaluation of 2-forms on concyclic polygons is the sum of

the evaluations of the 2-form on the original triangles. We will discuss in section 10 in detail how to use such a virtual polygonal mesh for computations on triangle meshes. We will now show that removing edges between concyclic triangle pairs cannot result in isolated vertices.

Lemma 6.1. In a non-degenerate mesh, an inner vertex cannot be adjacent only to removable edges.

Proof. We consider the possibly non-flat 1-ring of n triangles around an inner vertex, and assume that all edges adjacent to the vertex are removable. As the two angles opposite to a removable edge sum up to π by definition, the sum of all angles opposite to the inner vertex is $n\pi$. As the 1-ring consists of n triangles, the total angle sum also is $n\pi$, so the angle sum around the inner vertex must be 0. This is a contradiction to the mesh being non-degenerate. □

We can prove the same result for boundary vertices, but we have to consider that dual edges corresponding to boundary edges are degenerate when the angle opposite to the boundary edge is $\frac{\pi}{2}$.

Lemma 6.2. In a non-degenerate mesh, a boundary vertex cannot be adjacent only to removable edges.

Proof. We consider the possibly non-flat 1-ring of n triangles around a boundary vertex and assume that all edges adjacent to the boundary vertex have a degenerate dual edge.

Then the angles opposite to each inner edge sum up to π, and the angles opposite to the boundary edges adjacent to the vertex, are $\frac{\pi}{2}$. A triangle strip with n triangles has $n-1$ inner edges and two boundary edges adjacent to the vertex, so the total angle sum at the angles opposite to the boundary vertex is $(n-1)\pi + 2\frac{\pi}{2} = n\pi$, and thus the angles at the boundary vertex must be 0. This is a contradiction to the mesh being non-degenerate. □

6.3. Consistent Discrete Dual Forms

When allowing degenerate dual meshes, the space of discrete dual k-forms is no longer $\mathbb{R}^{|K_d^k|}$ as some vectors cannot be the result of integrating a differential form over the k-simplices of the mesh. In particular, the result of integrating a 1-form over a degenerate dual edge can never be non-zero, and the function values at the vertices of such an edge must be the same for all 0-forms, as the vertices lie at the same point on the manifold.

Definition 6.1 (Consistent 0-form)**.** We call a discrete dual differential 0-form $\omega_d^0 \in \mathbb{R}^{|K_d^0|}$ *consistent*, if for each degenerate dual edge $\star\sigma_i^1 = [\star\sigma_j^2, \star\sigma_k^2]$ in the mesh $\left(\omega_d^0\right)_j = \langle\hat{\omega}_d^0, \star\sigma_j^2\rangle = \langle\hat{\omega}_d^0, \star\sigma_k^2\rangle = \left(\omega_d^0\right)_k$.

Definition 6.2 (Consistent 1-form)**.** We call a discrete dual differential 1-form $\omega_d^1 \in \mathbb{R}^{|K_d^1|}$ *consistent*, if for each degenerate dual edge $\star\sigma_i^1 = [\star\sigma_j^2, \star\sigma_k^2]$ of the mesh $\left(\omega_d^1\right)_i = \langle\hat{\omega}_d^1, \star\sigma_i^1\rangle = 0$.

The second equation follows from Eq. 1 and the first must hold as the dual vertices $\star\sigma_j^2$ and $\star\sigma_k^2$ lie at the same point on the manifold. We do not need such a definition for dual 2-forms, because we assumed the primal mesh to be non-degenerate and thus all dual cells are non-degenerate and no two primal vertices lie at the same point. It is easy to see that the space of consistent discrete dual k-forms for a given mesh is a vector space, which is a subspace of $\mathbb{R}^{|K_d^k|}$. In the following, we assume that all forms we work with are consistent.

6.4. Mapping discrete triangle mesh forms to discrete polygonal mesh forms

In the following, we describe how the discrete differential forms need to be adapted for the changed mesh.

- **Discrete primal 0-forms** are not affected by the edge removal as the primal vertices are unchanged.
- **Discrete primal 1-forms** are changed by removing the value of the form integrated over the removed edges from the vector.

- **Discrete primal** 2**-forms** are adapted by using the sum of the evaluation of the form on the concyclic triangles σ_i^2 and σ_j^2 as the value for the polygon $\sigma_i^2 \cup \sigma_j^2$:

$$\omega^2_{i\cup j} = \left\langle \hat{\omega}^2, \sigma_i^2 \cup \sigma_j^2 \right\rangle := \left\langle \hat{\omega}^2, \sigma_i^2 \right\rangle + \left\langle \hat{\omega}^2, \sigma_j^2 \right\rangle = \omega_i^2 + \omega_j^2 \tag{7}$$

  By merging triangles, we lose the information about the different densities in the triangles, but the integral of $\hat{\omega}^2$ over the polygon is the same as the integral of $\hat{\omega}^2$ over the two triangles.
- **Discrete dual** 0**-forms** are adapted by removing the entry for one of the dual vertices belonging to the collapsed edge from the vector, as we remove one of the vertices of the edge. If the form was consistent before, both entries had the same value.
- **Discrete dual** 1**-forms** are changed by removing the entry for the collapsed edge from the vector, as it is removed from the mesh. In consistent forms, the removed entry were 0.
- **Discrete dual** 2**-forms** are unchanged, as the shape of the dual cells is unchanged.

6.5. Adapting the Operators

After changing the mesh and adapting existing discrete differential forms to the new mesh, we also need to change the operator matrices. For the definition of the operators we follow [Hir03] and only explain the required changes.

The Hodge star operators

The Hodge star $\star_0$ is unchanged. In the Hodge star matrix $\star_1$, the rows and columns corresponding to the removed edges are deleted. As no primal or dual vertex position was changed, the entries for other edges are unchanged. In the Hodge star $\star_2$, we replace the rows and columns of the merged triangles with a single row and column for the concyclic polygon, and the value of the diagonal entry is the sum of the two original values. The dual Hodge star operators $\star_0^{-1}$, $\star_1^{-1}$, and $\star_2^{-1}$ are the inverted matrices of the corresponding primal Hodge star operators, like before.

The Differential Operators

In the matrix of the discrete differential operator for primal 0-forms $\mathbf{d}_0$, the rows for the removed edges are deleted, so the operator describes the coboundary of the vertices in the changed mesh. In the matrix of the discrete differential operator $\mathbf{d}_1$ for primal 1-forms, the rows of the merged triangles are removed, and a row for the concyclic polygon containing the sum of the two removed rows is added, for which the new index is chosen to be consistent with the index in the adapted primal 2-forms. As the merged triangles orient the removed inner edge differently, the row entries for the edge cancel out, such that new row correctly describes the boundary of the polygon. The dual discrete differential operators are defined as before by the transpose of the primal operators.

Sharp and Flat Operators

The flat operators are already well-defined for non-removable primal edges and non-degenerate dual edges, so they do not need to be changed and the sharp operators require a new interpolation scheme for the gradient inside a concyclic polygon, which we leave for future work.

The Wedge Product

We define the evaluation of the primal-primal wedge product between two 1-forms α^1 and β^1 on a concyclic polygon $\sigma_i^2 \cup \sigma_j^2$ by the sum the evaluations on the original triangles

$$\left\langle \alpha^1 \wedge \beta^1, \sigma_i^2 \cup \sigma_j^2 \right\rangle := \left\langle \alpha^1 \wedge \beta^1, \sigma_i^2 \right\rangle + \left\langle \alpha^1 \wedge \beta^1, \sigma_j^2 \right\rangle \tag{8}$$

This approach works both with the original definition in [Hir03] section 7.1 and the alternative definition in section 7.2 and is consistent with the summation in Eq. 7. The inner edges do not contribute to the result of either of the two definition, because the evaluation of the form on inner edges is added twice but with opposite signs.

7. The DEC Laplacian for Concyclic Polygons

It is known [DKT08] that the DEC Laplacian is equivalent to the cotangent Laplacian [PP93], which uses $\cot(\phi) + \cot(\theta)$ for edge weights, where ϕ and θ are the angles at the triangle corners opposite to the edge. We will show in the following sections that this identity still holds for our new DEC scheme for concyclic polygons.

7.1. Discrete Laplacians

There are many different discrete Laplacians, which each have different advantages and disadvantages. Wardetzky et al. provide an overview in [WMKG07], which especially compares discrete Laplacians by the properties of being symmetric, having local support, having linear precision, fulfilling the maximum principle, being positive semi-definite, and their convergence properties. The paper also shows that the cotangent Laplacian does not guarantee positivity, i.e., it has negative weights for non-delaunay triangle pairs and for concyclic triangle pairs, the edge weight is 0 and the Laplace matrix becomes singular, which is exactly the problem we are solving in this paper.

For general polygons, there are a few notable polygonal Laplacians. Alexa and Wardetzky extended in 2011 the cotangent Laplacian to polygons [AW11], and a simple polygonal Laplacian for possibly non-convex and non-planar faces was recently defined by Bunge et al. [BHKB20]. The Diamond Laplace by Bunge et al. [BBA21] constructs quadrilateral diamonds at polygon edges and can easily be generalized for arbitrary polyhedral meshes in 3D. De Goes defines a whole framework for discrete differential operators on polygonal meshes [DGBD20] based on an extension of the virtual element method. While these Laplacians are useful for calculations on polygonal surfaces, mixing them with the discrete exterior calculus complicates the scheme that previously could be defined purely in terms of DEC operators. Our scheme does not aim to replace these polygonal Laplacians, but uses polygons to resolve DEC singularities in meshes with concyclic triangle pairs.

7.2. Equivalence to the Cotangent Laplacian

We will now show that the DEC Laplacian defined using the adapted operators is equivalent to the cotangent Laplacian. We prove that the result of both Laplacians does not depend on removable edges by showing that inner edges of triangulated concyclic polygons do not contribute to the result of either of the Laplacians. For the cotangent Laplacian, it is sufficient to show that the weight of the edge between a concyclic triangle pair is 0, as the weight of each edge solely depends on its two adjacent triangles.

Lemma 7.1. The cotangent weight of the diagonal edge in a triangulated concyclic quadrilateral is 0, and thus the edge does not contribute to the result of the cotangent Laplacian.

Proof. The weight of the diagonal edge is defined as the sum of the cotangents at the corners opposite to the edge and we have $\cot(\phi) + \cot(\pi - \phi) = 0$, because the angle sum of these corners is π when the two triangles form a concyclic quadrilateral (cf. Euclid's elements book 3, proposition 22 [Fit07]). □

Lemma 7.2. An edge inside a concyclic polygon does not contribute to the result of the DEC Laplacian.

Proof. The DEC Laplacian is defined as $L = \star_2^{-1}\mathbf{d}_1\star_1\mathbf{d}_0$ and the operator $\star_1$ maps the value of the evaluation of a primal 1-form on a removable edge σ^1 to 0 (cf. Eq.19). Thus the edge does not influence the result of the DEC Laplacian. □

Theorem 7.3. The DEC Laplacian for meshes with concyclic polygons and the cotangent Laplacian are equivalent.

Proof. We know from [DKT08] that the edge weights of both Laplacians are the same in non-degenerate meshes. By 7.1 and 7.2, it follows that the weights for removable edges are also the same, i.e., zero, in both Laplacians. Thus the Laplacians are also equivalent for meshes with concyclic triangle pairs. □

7.3. A Note on Implementation

For implementation, it is inefficient and unnecessary to build the DEC operators and afterward remove edges and apply the changes mentioned in the previous sections. Instead, we recommend building the data structures directly by first finding concyclic triangle pair and then creating a copy of the mesh with degenerate dual edges and the corresponding primal edges removed. The polygonal mesh is then used for calculating the combinatoric operators, and the original mesh is used for the geometric properties. These modifications are mainly an additional preprocessing step that allows the other algorithms used for building DEC structures to skip removable edges and require only minor changes in existing software.

8. Evaluation

8.1. Numerical Error

Most of the mesh changes do not make a numerical difference for the mesh elements that already were non-degenerate, so many operators do not lose accuracy. No primal vertices were added or removed, and the collapsed dual edges did not contribute to calculations before. There are still two points at which the scheme causes a loss of accuracy.

By *removing primal edges*, we lose accuracy in operations for which the integral over a removable edge was non-zero. An example is the gradient operator. As $\mathbf{d}_0\omega^0$ is evaluated on fewer edges, the sharp operator interpolates the gradient $\left(\mathbf{d}_0\omega^0\right)^\sharp$ using fewer edges, which can reduce the accuracy of the result.

When *merging primal triangles*, we sum the discrete 2-forms for these triangles using equation 7, which results in a coarser approximation of the 2-form when equation 4 does not hold for the original discrete 2-form, as a discrete polygonal 2-form cannot represent the different densities of the merged triangles.

The *condition number* of the Hodge star matrix $\star_0$ is unchanged as the shape of the dual cells do not change and thus the areas $|\star\sigma^0|$ do not change. The Hodge star $\star_1^{-1}$ matrix becomes non-singular and the entries for the non-degenerate edges are unchanged. Near-singular entries can be removed by collapsing dual edges shorter than an $\varepsilon > 0$. The condition of the Hodge star matrix $\star_2$ can change because merging triangles results in larger primal cells and thus smaller matrix entries $(\star_2)_{ii} = 1/|\sigma_i^2|$.

9. Convergence

We test the convergence first by solving a Poisson problem on different types of meshes and then by solving an eigenvalue problem on spheres with concyclic triangulations.

9.1. Poisson Problem

For testing convergence, we solve a Poisson problem for Franke's function on the unit square $[0,1]\times[0,1]$:

$$\begin{aligned} u(x,y) = &\frac{3}{4}\exp\left(-\frac{(9x-2)^2}{4} - \frac{(9y-2)^2}{4}\right) + \frac{3}{4}\exp\left(-\frac{(9x+1)^2}{49} - \frac{(9y+1)}{10}\right) \\ &+ \frac{1}{2}\exp\left(-\frac{(9x-7)^2}{4} - \frac{(9y-3)^2}{4}\right) - \frac{1}{5}\exp\left(-(9x-4)^2 - (9y-7)^2\right) \end{aligned} \tag{9}$$

We compare our results with the analytic solution on four different types of synthetic meshes, which are shown in Fig. 6.

- Well-centered meshes that work well with the original DEC definition in [Hir03].
- Triangulations of rectangular grids, in which the diagonals are removable edges that induce degenerate dual edges.
- Meshes that are fan triangulations of hexagonal tilings that consist of 4 concyclic triangles and 6 concyclic vertices in each hexagon.
- Meshes with vertices in general position that contain no concyclic triangle pairs and can be processed using negative volumes as described in [MHS18]). For our example meshes, we used a rectangular grid and perturbed the inner vertices by a random distance between $-0.25h$ and $+0.25h$, where h is the width of the grid cells (cf. Fig. 6).

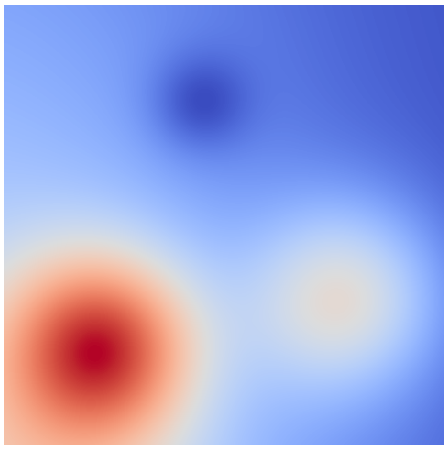

Figure 5: The solution of the Poisson problem in Eq. 9.

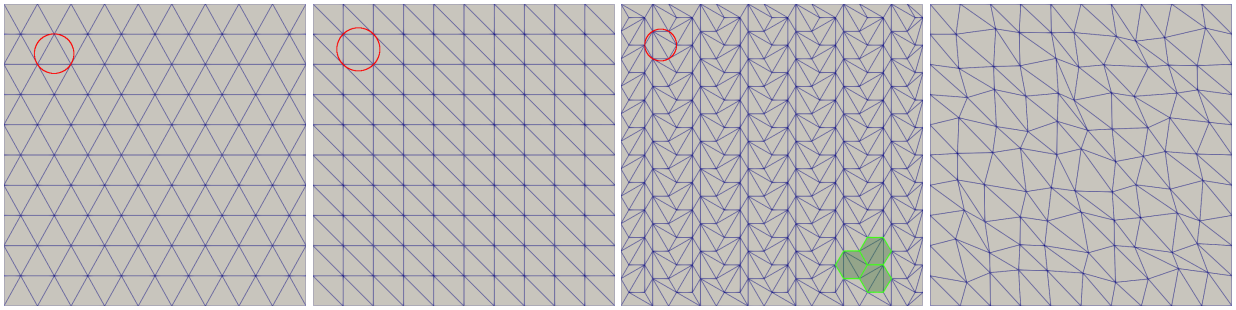

Figure 6: We tested the convergence of a Poisson problem on four types of meshes. The first mesh is a well-centered triangulation that is suited for every DEC variant. The second mesh is a triangulation of a rectangular grid, in which the diagonals induce degenerate dual edges, and the third type is a triangulation of a hexagonal grid that contains four concyclic triangles per hexagon. The fourth mesh is a distorted triangulated rectangular mesh with vertices in general position.

For solving, we used an LU-solver. We tried solving with an iterative solver, but it did not converge for near-concyclic meshes like we created using the perturbation approach. We did not examine precisely how large the perturbation must be for the solver to converge because large enough perturbations yield inaccurate results (cf. Fig. 7) and thus are no viable solution.

The convergence plots for solving the problem on different triangulations are depicted in Fig. 7. We chose to plot the error against the median inner edge length, as it has proven to be the most robust measure for a fair comparison between the different meshes. The convergence behavior is the same when using other metrics, like the minimum or average edge length, but we observe different absolute errors for the different mesh types. Computations using non-degenerate meshes together with classic DEC and using concyclic meshes with our novel polygonal DEC converge under mesh refinement with the same rate. Triangle DEC calculations on the perturbed meshes also converge as long as the perturbation is chosen significantly smaller than the edge lengths but diverge for finer meshes.

Perturbations of 10^{-7} and less could, in practice, not guarantee that dual edges have a length the is numerically greater than 0 when using double precision for the vertex coordinates of the primal mesh. Furthermore, many mesh libraries and data formats assume single precision accuracy for vertex coordinates, which may require writing own import and export functions and carefully checking which floating-point precision is used in all functions and data formats.

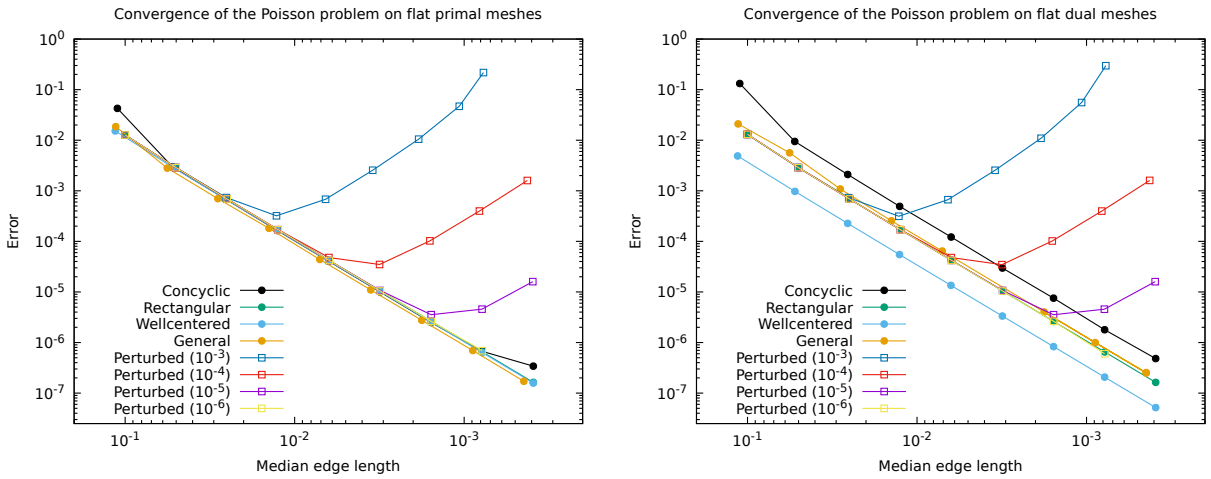

Figure 7: The convergence of the Poisson problem in Eq. 9 on different types of meshes.

9.2. Spherical Harmonics

For evaluating our operator on a mesh that is not intrinsically flat, we opted to use the same eigenvalue problem as De Goes [DGBD20] and Bunge [BHKB20] and solve $Lx = \lambda_4 Y_2^4$ where $Y_2^4 = (x^2 - y^2)(7y^2 - 1)$ is a spherical harmonic function, which is an eigenfunction of the Laplacian on the unit sphere with the eigenvalue $\lambda_4 = -20$.

We solved the problem on spheres with quadrilateral and hexagonal tessellations, for which triangulations contain concyclic triangle pairs. The problem can be formulated using DEC operators as

$$\text{primal:} \qquad (\star_0^{-1}\mathbf{d}_1 \star_1 \mathbf{d}_0)\omega^0 = \lambda\mathbf{Y}^0 \qquad \text{with} \qquad \mathbf{Y}^0 = \left(Y_2^4\left(\sigma_1^0\right), \ldots, Y_2^4\left(\sigma_n^0\right)\right)^T \tag{10}$$

$$\text{dual:} \qquad (\star_2\mathbf{d}_0^T \star_1^{-1} \mathbf{d}_1^T)\omega_d^0 = \lambda\mathbf{Y}_d^0 \qquad \text{with} \qquad \mathbf{Y}_d^0 = \left(Y_2^4\left(\star\sigma_1^2\right), \ldots, Y_2^4\left(\star\sigma_n^2\right)\right)^T \tag{11}$$

The vectors were normalized by dividing by the mass matrix weighted norm $\|x\|_M = \sqrt{x^T M x}$, with $M = \star_0$ for primal 0-forms and $M = \star_2^{-1}$ for dual 0-forms.

We measure the errors weighted with the corresponding mass matrix as

$$E = \sqrt{\left(\mathbf{Y}^0 - \omega^0\right)^T \star_0 \left(\mathbf{Y}^0 - \omega^0\right)} \qquad \text{and} \qquad E_{\text{dual}} = \sqrt{\left(\mathbf{Y}_d^0 - \omega_d^0\right)^T \star_2^{-1} \left(\mathbf{Y}_d^0 - \omega_d^0\right)} \tag{12}$$

Fig. 8 shows the convergence of our method. For comparison, we include results using the methods of Alexa ([AW11]), De Goes ([DGBD20]), and Bunge ([BHKB20] , [BBA21]). Note that the mesh resolution (which is the inverse of the mean edge length) of the data points for our method is different, as we triangulated the polygonal meshes and thus have a different mean edge length in our input meshes.

For [DGBD20], [BHKB20], and [AW11] we used the comparison data from [DGBD20], which was kindly provided by Fernando De Goes. The results in [DGBD20] deviate from the results obtained using Bunge's code as the original code does not use the mass matrix when normalizing the spherical harmonics vector. However, after changing the normalization of the spherical harmonics vector to use the mass matrix weighted norm $\|\cdot\|_M$ instead of using the 2-norm like Bunge's code, we were able to reproduce De Goes' results for [BHKB20].

We computed the results for the diamond Laplace ourselves using the code from [BBA21] with the benchmark code from [BHKB20] (both provided by the authors as open-source software) together with our modification for using the mass matrix weighted norm. We could not use the results in [BBA21] directly, as they solved for Y_3^{-1} instead of Y_2^4 and also normalized the spherical harmonics vector using the 2-norm, which cannot be compared directly with the results of the other methods.

9.3. Comparison with using a Minimal Edge Length

Another simple method for fixing the degenerate operators that comes into mind is enforcing a minimal dual edge length in the operators (without changing the mesh itself) when the length of the dual edge is zero or

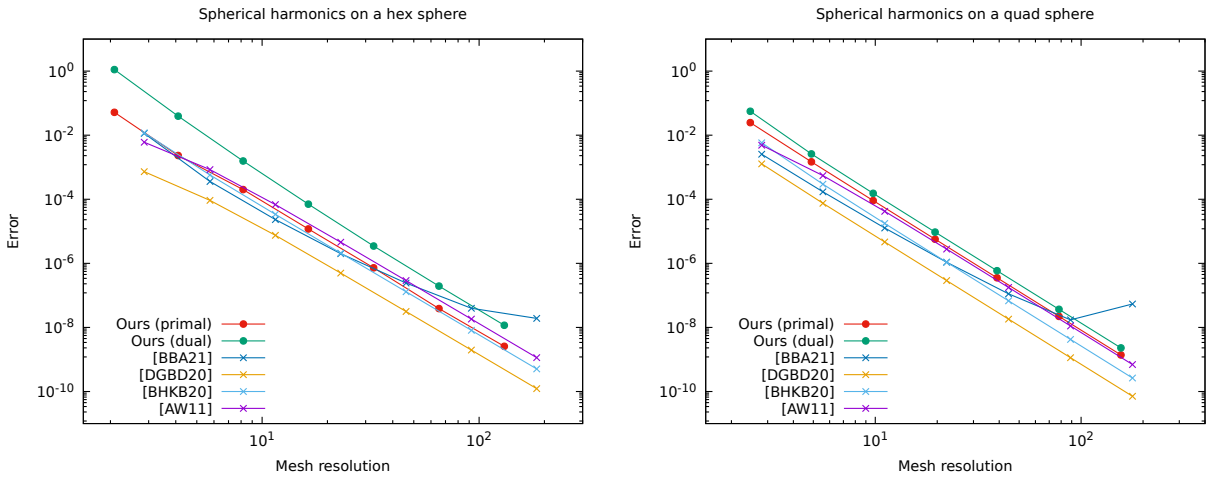

Figure 8: We solve the Poisson problem on a sphere using our Laplacian with the spherical harmonic function $Y_2^4 = (x^2 - y^2)(7y^2 - 1)$ scaled with the corresponding eigenvalue as right hand side and compare the resulting eigenvector with the analytic solution. We compare our results with other polygonal Laplacians, by De Goes ([DGBD20] with $\lambda = 1.0$), Bunge ([BHKB20] and [BBA21]), and Alexa ([AW11]). The x-axis is the mesh resolution given by the inverse of the average edge length.

in a numerically unstable range by replacing $|\star\sigma^1|$ in all calculations with ε (and $-\varepsilon$ respective for negative edges), if abs($|\star\sigma^1|$) $< \varepsilon$. This change directly affects the Hodge star operators $\star_1$, and $\star_1^{-1}$ and makes the latter well-defined and numerically stable for a large enough ε. It also indirectly affects $\star_0$ and $\star_0^{-1}$ as the area $|\star\sigma^0|$ of dual cells adjacent to degenerate dual edges changes.

It is easy to see that a large enough ε will result in numerically stable operations, but on the other hand, the corresponding (virtual) geometry is not well-defined. Assume, for example, a flat dual cell with a degenerate dual edge on its boundary. The elementary triangle angle opposite the degenerate dual edge on the boundary of the dual cell is zero. When we arbitrarily assign the dual edge a nonzero length ε, this angle becomes nonzero, and thus the angle sum at the center vertex gets larger than 2π, so the corresponding virtual geometry is no longer flat.

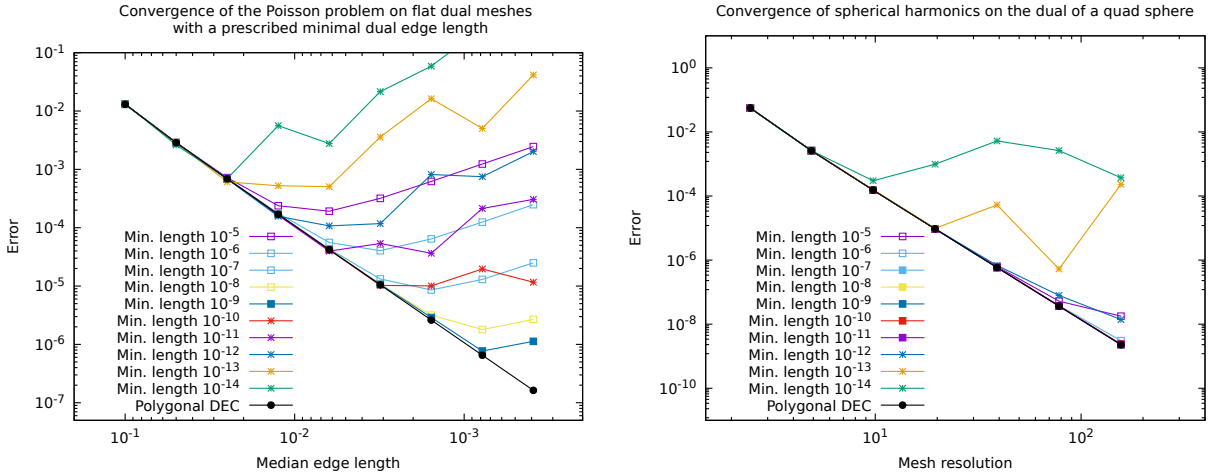

Figure 9: The results of solving the problems from previous sections with operators that use a minimal dual edge length. Left: Plot of the Poisson problem in section 9.1. Right: Plot of the spherical harmonics problem from section 9.2.

To compare using a minimal edge length with our approach, we solved the Poisson problem from section 9.1 and the eigenvalue problem from section 9.2 using a prescribed minimal dual edge length together with the method of [MHS18] for using DEC with general triangular meshes.

Fig. 9 shows the results of solving these problems on the triangulated rectangular grid and the quad sphere, respectively. It can be seen that using a minimal dual edge length yields similar results as perturbing the mesh directly, i.e., the error diverges under mesh refinement when the minimal dual edge length ε stays the same. We also observe that a too small ε results in numerically unstable results. In our experiments, the best choice for a minimal dual edge length was 10^{-9} for the Poisson problem and between 10^{-7} and 10^{-11} for the Eigenvalue problem.

9.4. Mean Curvature

As we remove primal edges, another relevant benchmark is calculating the mean curvature on the primal mesh because the DEC Laplacian, which uses values from the 1-ring around a vertex, is approximated using a smaller neighborhood. Mean curvature can be calculated as $H = \frac{1}{2}(\kappa_1 + \kappa_2) = \frac{1}{2}\|\Delta x + \Delta y + \Delta z\|$ and we can approximate it using the discrete 0-form Laplacian $\Delta_0 = \star_2^{-1}\mathbf{d}_1 \star_1 \mathbf{d}_0$ as described in section 5.6. For the benchmark, we use a triangulated UV sphere (see Fig. 10 right) as it allows for simple refinement and the original quad faces are concyclic at any mesh resolution. We refine U and V simultaneously, so the aspect ratio of the quads stays approximately the same, and each sphere consists of $(U-2)\cdot U$ quads and $2\cdot U$ triangles at the poles. As the faces of a UV sphere do not have uniform size, we weigh the curvature error at a primal vertex with the area of its dual cell and calculate the root mean square error as

$$E_{\text{RMSE}} = \sqrt{\sum_{i=0}^{|K^0|} \frac{e_i^2}{|K^0|}} \quad \text{with} \quad e_i = \star_{ii}^0 \left(\frac{1}{R} - H\left(\sigma_i^0\right)\right)$$

The result is plotted in Fig. 10, and we observe that the curvature calculation using our novel method converges with a rate of 2.

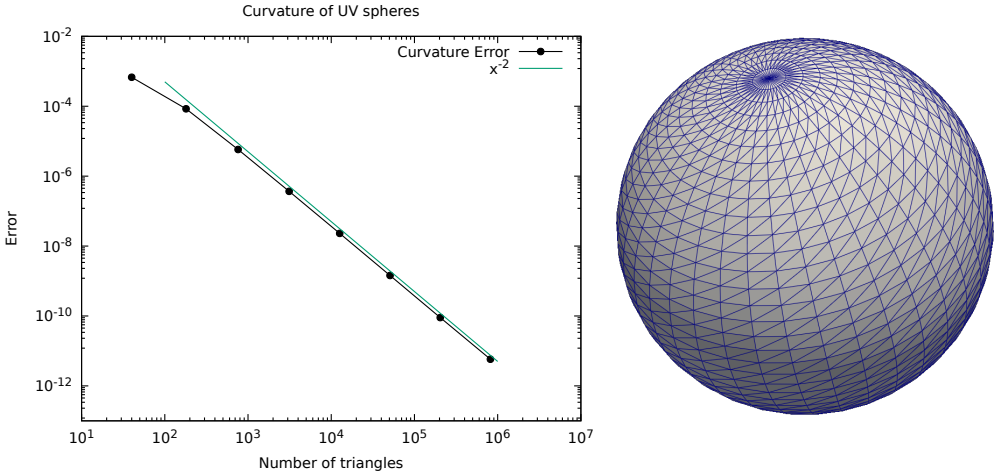

Figure 10: Left: The weighted root mean square error of the curvature calculated using our DEC Laplacian on a UV sphere. Right: A triangulation of a UV sphere with 40×40 subdivisions.

10. Polygonal DEC for Triangle Meshes

In many use-cases, one would like to use the (concyclic) triangle mesh directly instead of the corresponding polygonal mesh. We show how the DEC scheme for concyclic polygons can be used for calculations on triangle meshes.

- **0-Forms:** Discrete *primal* 0-forms defined on the triangle mesh and discrete primal 0-forms defined on the polygonal mesh are identical, as the set of primal vertices is unchanged. Using the triangle mesh provides, in some cases, more accuracy because the mesh contains more information about the original

shape. To use a discrete *dual* 0-form defined for the polygonal mesh with the corresponding triangle mesh, we assign the function values of merged vertices to both original vertices in the triangle mesh. To use a discrete dual 0-form defined on the triangle mesh with the polygonal mesh, we can ignore the values of removed vertices, because they are the same as the value of the remaining vertex of the corresponding dual edge.

- **1-Forms:** For using discrete *primal* 1-forms defined for the polygonal mesh with the triangle mesh, the values for the primal edges which are only present in the triangle mesh are set to 0, which is consistent with the polygonal scheme in which the edge does not exist. The evaluation of *dual* 1-forms on degenerate dual edges is 0 due to Eq. 1. In the other direction, the values for edges only present in the triangle mesh are removed from the vector to convert discrete triangle mesh forms into discrete polygonal mesh forms.
- **2-Forms:** For discrete *primal* 2-forms, we define a new operator that maps between discrete forms on polygons and discrete forms on the triangles that were merged into these polygons. Discrete *dual* 2-forms are unchanged, as the dual cells are unchanged.

10.1. Mapping between Triangle Mesh 2-Forms and Polygon Mesh 2-Forms

For a given triangle mesh $K = (K^0, K^1, K^2)$, we construct a corresponding polygonal mesh $P = (P^0, P^1, P^2)$ in which concyclic triangles are merged into polygons, where K^k and P^k are the vertex, edge and triangle/polygon sets of the two meshes. We have $P^0 = K^0$ as the vertices are unchanged, $P^1 \subseteq K^1$ because the polygonal mesh P is constructed by removing edges from K, and for each polygon $p^2 \in P^2$ we store the corresponding triangles $\{\sigma^2 \in K^2 | \sigma^2 \subseteq p^2\}$.

Similar to the Hodge star, which maps between primal and dual forms, we can define an operator $\hexagon$ that maps between forms acting on triangle meshes and forms acting on polygon meshes. Except for 0-forms, dual 2-forms where it is the identity, the operator has no full inverse, but we can define a right inverse $\hexagon^{-1}$.

The operator $\hexagon_1 \in \mathbb{R}^{|P^1| \times |K^1|}$ is a rectangular matrix, which removes the entries for edges that are not present in the polygonal mesh from a triangle mesh 1-form and the operator $\hexagon_1^{-1} = \hexagon_1^T$ maps polygon mesh 1-forms to triangle mesh 1-forms by setting the entries of the new edges to 0 and is a right inverse, but in general no left inverse for $\hexagon_1$.

For 2-forms defined on triangle meshes, $\hexagon_2 \in \mathbb{R}^{|P^2| \times |K^2|}$ sums the values of the form evaluated on the triangles inside a polygon (cf. equation 7) by

$$\left(\hexagon_2\right)_{ij} := \begin{cases} 1 & \text{if } \sigma_j^2 \subseteq p_i^2 \\ 0 & \text{otherwise} \end{cases} \tag{13}$$

In the other direction, 2-forms defined on the polygonal mesh are mapped to 2-forms acting on triangles by the $\hexagon_2^{-1} \in \mathbb{R}^{|K^2| \times |P^2|}$ operator, which distributes the value of the form evaluated on a polygon over the corresponding triangles, such that Eq. 4 holds:

$$\left(\hexagon_2\right)_{ij}^{-1} := \begin{cases} \frac{|\sigma_i^2|}{|p_j^2|} & \text{if } \sigma_i^2 \subseteq p_j^2 \\ 0 & \text{otherwise} \end{cases} \tag{14}$$

The operator $\hexagon_2^{-1}$ is the right inverse of $\hexagon_2$, and in the subspace of triangle mesh 2-forms that satisfy Eq. 4 also a left inverse.

Given a DEC operator $O^{(P)}$ on the polygonal mesh, we can use these operators to define triangle DEC operators, which use an intrinsic polygonal mesh:

$$\mathbf{O}^{(T)} = \hexagon^{-1} \mathbf{O}^{(P)} \hexagon. \tag{15}$$

We can, for example, define a 2-form Laplacian for meshes that contain concyclic triangle pairs using the differential and Hodge star operators for the corresponding polygonal mesh as $L_{2,\text{dual}}^{(T)} = \hexagon_2^{-1} \mathbf{d}_1 \star_1^{-1} \mathbf{d}_0 \star_2 \hexagon_2$.

11. Implementation in existing DEC code

A typical DEC implementation has two important steps for building the necessary operator matrices (cf. [ES05]):

- Calculate the differential operators as adjacency matrices.
- Calculate the Hodge star operators using the ratios of the primal and dual simplex lengths, areas, and volumes.

Our approach adds a preprocessing step:

- Detect concyclic triangle pairs and mark the edge between them so it can either be explicitly deleted beforehand or skipped in the further calculations.

After our preprocessing, we adapt the other two steps in the following way. When calculating the differential operators $\mathbf{d}_0$ and $\mathbf{d_1}$, we omit the entries for removable edges. One option is to create the polygonal mesh explicitly and use the standard algorithms unchanged. The other option is to only mark the edges as removable and skip them when iterating over adjacent edges in other algorithms. The Hodge star for polygonal 2-forms is defined as $\star_2^{(P)} = \hexagon_2 \star_2^{(T)} \hexagon_2^T$, which is equivalent to defining the volume form used in the Hodge star definition (see section 5.3) as $(|p_1|, \dots, |p_{|P^2|}|)^T = \hexagon_2(|\sigma_1^2|, \dots, |\sigma^2_{|K^2|}|)^T$. Similarly, we calculate $\star_1^{(P)} = \hexagon_1 \star_1^{(T)} \hexagon_1^T$, which is equivalent to removing the rows and columns corresponding to edges that do not exist in the polygonal mesh. The Hodge star operator $\star_0$ is the same for both meshes, as the vertices are the same, and the area of the dual cells is not changed by removing degenerate dual edges.

12. Limitations

Discrete exterior calculus is defined for simplicial complexes of arbitrary dimension, but we only described our method for 2D surface meshes. An extension to 3D would be desirable, especially for physics simulations on tetrahedral meshes, e.g., for 3D fluid simulations. In a tetrahedral mesh, degenerate dual edges are induced by pairs of tetrahedra with cospherical vertices that cause similar problems as the concyclic triangles in 2D meshes. It should be possible to extend the scheme for simplicial complexes of arbitrary dimension by merging n-simplices, for which the enclosing hyperspheres have the same center, into cospherical n-polyhedra and collapsing the corresponding dual edges of length 0.

Another topic for future work is consistently defining the operators that need interpolation of the gradient inside concyclic polygons.

13. Conclusion

We have shown that meshes with concyclic triangles pose a problem for discrete exterior calculus, which is not handled by previous approaches for using DEC with general meshes. After we proved for each operator if it is well-defined on meshes with concyclic triangles, we analyzed how the mesh can be changed, such that the new operators are well-defined. By removing the inner edges of concyclic polygons in the triangulation, we exploited that the cause of the problem, which is concyclic triangles having the same circumcenter, also ensures that the resulting polygon allows for a circumcentric dual mesh. Thus, the definition of the dual mesh still enforces orthogonality between primal and dual edges.

We evaluated meshes from several data sets and determined the prevalence of concyclic triangles to show that our extension to the DEC is relevant for different categories of meshes, like triangulated CAD models, hand-crafted 3D meshes, and even triangulated surfaces in 3D scans. We tested the numerical convergence of our novel scheme by solving a Poisson problem and an eigenvalue problem. We compared the result to the results on well-centered meshes optimal for DEC calculations, general meshes treated by previous approaches, and slightly perturbed meshes proposed to work around the limitations of other approaches. Our tests include rectangular meshes common in real-world applications and meshes with a vast number of concyclic triangles. We found that our method performed on these meshes as well as when using meshes with better triangulations and outperformed the perturbation approach, which does not converge under mesh refinement.

14. Acknowledgements

We like to thank Fernando De Goes for providing us with the example meshes and benchmark data used in [DGBD20] for comparison and helpful hints for the evaluation and Bunge et al. for providing the source code from [BHKB20] and [BBA21]. We also like to thank the anonymous reviewers for their valuable feedback.

References

[ATW15] Ryoichi Ando, Nils Thuerey, and Chris Wojtan, *A stream function solver for liquid simulations*, ACM Transactions on Graphics (TOG) **34** (2015), no. 4, 1–9.

[AW11] Marc Alexa and Max Wardetzky, *Discrete laplacians on general polygonal meshes*, ACM Transactions on Graphics (TOG) **30** (2011), no. 4, 1–10.

[BBA21] Astrid Bunge, Mario Botsch, and Marc Alexa, *The diamond laplace for polygonal and polyhedral meshes*, Computer Graphics Forum, vol. 40, Wiley Online Library, 2021, pp. 217–230.

[BHKB20] Astrid Bunge, Philipp Herholz, Misha Kazhdan, and Mario Botsch, *Polygon laplacian made simple*, Computer Graphics Forum, vol. 39, 2020, pp. 303–313.

[BZK09] David Bommes, Henrik Zimmer, and Leif Kobbelt, *Mixed-integer quadrangulation*, ACM Transactions On Graphics (TOG) **28** (2009), no. 3, 1–10.

[CC19] Etienne Corman and Keenan Crane, *Symmetric moving frames*, ACM Transactions on Graphics (TOG) **38** (2019), no. 4, 1–16.

[CdGDS13] Keenan Crane, Fernando de Goes, Mathieu Desbrun, and Peter Schröder, *Digital geometry processing with discrete exterior calculus*, ACM SIGGRAPH 2013 Courses (New York, NY, USA), SIGGRAPH '13, Association for Computing Machinery, 2013.

[CWW13] Keenan Crane, Clarisse Weischedel, and Max Wardetzky, *Geodesics in heat: A new approach to computing distance based on heat flow*, ACM Transactions on Graphics (TOG) **32** (2013), no. 5, 1–11.

[Del34] Boris Delaunay, *Sur la sphere vide*, Izv. Akad. Nauk SSSR, Otdelenie Matematicheskii i Estestvennyka Nauk **7** (1934), no. 793-800, 1–2.

[DGBD20] Fernando De Goes, Andrew Butts, and Mathieu Desbrun, *Discrete differential operators on polygonal meshes*, ACM Transactions on Graphics (TOG) **39** (2020), no. 4, 110–1.

[DKT08] Mathieu Desbrun, Eva Kanso, and Yiying Tong, *Discrete differential forms for computational modeling*, Discrete differential geometry, Springer, 2008, pp. 287–324.

[ES05] Sharif Elcott and Peter Schröder, *Building your own dec at home*, ACM SIGGRAPH 2005 Courses, 2005, pp. 8–es.

[Fit07] Richard Fitzpatrick, *Euclid's elements of geometry*, Euclidis Elementa, 2007.

[Gli05] David Glickenstein, *Geometric triangulations and discrete laplacians on manifolds*, arXiv preprint math/0508188v1 (2005).

[GRS17] Michael Griebel, Christian Rieger, and Alexander Schier, *Upwind schemes for scalar advection-dominated problems in the discrete exterior calculus*, Transport Processes at Fluidic Interfaces, Birkhäuser, 2017, pp. 145–175.

[GWDS08] Eitan Grinspun, Max. Wardetzky, Mathieu Desbrun, and Peter Schröder, *Discrete differential geometry: an applied introduction*, ACM SIGGRAPH Courses Notes. ACM, New York (2008).

[Hir03] Anil Nirmal Hirani, *Discrete exterior calculus*, Ph.D. thesis, California Institute of Technology, 2003.

[HKV13] Anil N Hirani, Kaushik Kalyanaraman, and Evan B VanderZee, *Delaunay hodge star*, Computer-Aided Design **45** (2013), no. 2, 540–544.

[HKV18] _____, *Corrigendum to "delaunay hodge star"[comput. aided des. 45 (2013) 540–544]*, Computer-Aided Design **96** (2018), 59–60.

[HNC15] Anil N Hirani, Kalyana B Nakshatrala, and Jehanzeb H Chaudhry, *Numerical method for darcy flow derived using discrete exterior calculus*, International Journal for Computational Methods in Engineering Science and Mechanics **16** (2015), no. 3, 151–169.

[JLSW02] Tao Ju, Frank Losasso, Scott Schaefer, and Joe Warren, *Dual contouring of hermite data*, ACM Transactions on Graphics (TOG) **21** (2002), no. 3, 339–346.

[LC87] William E Lorensen and Harvey E Cline, *Marching cubes: A high resolution 3d surface construction algorithm*, Proceedings of the 14th annual conference on Computer graphics and interactive techniques, 1987, pp. 163–169.

[MHS16] Mamdouh S Mohamed, Anil N Hirani, and Ravi Samtaney, *Comparison of discrete hodge star operators for surfaces*, Computer-Aided Design **78** (2016), 118–125.

[MHS18] _____, *Numerical convergence of discrete exterior calculus on arbitrary surface meshes*, International Journal for Computational Methods in Engineering Science and Mechanics **19** (2018), no. 3, 194–206.

[MMdGD11] Patrick Mullen, Pooran Memari, Fernando de Goes, and Mathieu Desbrun, *Hot: Hodge-optimized triangulations*, ACM Transactions on Graphics **30** (2011), no. 4, Art–No.

[MMP87] Joseph SB Mitchell, David M Mount, and Christos H Papadimitriou, *The discrete geodesic problem*, SIAM Journal on Computing **16** (1987), no. 4, 647–668.

[NRV17] Ingo Nitschke, Sebastian Reuther, and Axel Voigt, *Discrete exterior calculus (dec) for the surface navier-stokes equation*, Transport processes at fluidic interfaces, Birkhäuser, 2017, pp. 177–197.

[PP93] Ulrich Pinkall and Konrad Polthier, *Computing discrete minimal surfaces and their conjugates*, Experimental mathematics **2** (1993), no. 1, 15–36.

[SSC19] Nicholas Sharp, Yousuf Soliman, and Keenan Crane, *The vector heat method*, ACM Transactions on Graphics (TOG) **38** (2019), no. 3, 1–19.

[WAVK+12] Yunhai Wang, Shmulik Asafi, Oliver Van Kaick, Hao Zhang, Daniel Cohen-Or, and Baoquan Chen, *Active co-analysis of a set of shapes*, ACM Transactions on Graphics (TOG) **31** (2012), no. 6, 1–10.

[WMKG07] Max Wardetzky, Saurabh Mathur, Felix Kälberer, and Eitan Grinspun, *Discrete laplace operators: no free lunch*, Symposium on Geometry processing, Aire-la-Ville, Switzerland, 2007, pp. 33–37.

[ZDWT19] Rundong Zhao, Mathieu Desbrun, Guo-Wei Wei, and Yiying Tong, *3d hodge decompositions of edge-and face-based vector fields*, ACM Transactions on Graphics (TOG) **38** (2019), no. 6, 1–13.

Appendix

A. Well-defined Operators

For completeness, we show which operators are well-defined on all meshes with concyclic triangle pairs.

A.1. Differential Operators

The differential operator for dual 0-forms $\mathbf{d}_0$ is defined by the evaluation of the resulting dual 1-form $\mathbf{d}_0\hat{\omega}^0_d$ along a dual edge $\star\sigma^1 = [\star\sigma^2_i, \star\sigma^2_j]$ using Stokes' theorem

$$\langle \mathbf{d}_0\hat{\omega}^0_d, \star\sigma^1\rangle = \langle \hat{\omega}^0_d, \partial(\star\sigma^1)\rangle = \langle \hat{\omega}^0_d, \star\sigma^2_j\rangle - \langle \hat{\omega}^0_d, \star\sigma^2_i\rangle \tag{16}$$

where $\partial\left(\star\sigma^1_d\right)$ is the dual 0-chain that assigns the values $+1$ and -1 to the vertices of the dual edge.

The left hand side is 0 as integrating any dual 1-form over a degenerate dual edge $\star\sigma^1$ yields 0 (see Eq. 1) and the right hand is 0, because $\star\sigma^2_i$ and $\star\sigma^2_j$ lie at the same point on the manifold. So the differential operator for dual 0-forms is *well-defined*.

The differential operator for dual 1-forms is defined by the evaluation of the resulting dual 2-form on the dual cells $\star\sigma^0$ using Stokes' theorem

$$\langle \mathbf{d}_1\hat{\omega}^1_d, \star\sigma^0\rangle = \langle \hat{\omega}^1_d, \partial(\star\sigma^0)\rangle = \sum_{\star\sigma^1_i \prec \star\sigma^0} \langle \hat{\omega}^1_d, \star\sigma^1_i\rangle \tag{17}$$

where $\star\sigma^1_i \prec \star\sigma^0$ means that $\star\sigma^1_i$ is a dual edge on the boundary of the dual cell $\star\sigma^0$. When one of the dual edges $\star\sigma^1_i$ is degenerate, we have to consider the effect on the integration.

From Eq. 1 we know that the contribution of the degenerate dual edge to the right-hand side is 0. The left-hand side is not affected by the edge, because the two elementary triangles spanned by the edge $\star\sigma^1$ and the center of the dual cell (cf. Fig. 11f) have an area of 0 and thus do not contribute to the surface integral. So the differential operator for dual 1-forms is *well-defined*.

A.2. Hodge Star Operators

For the 0-form Hodge star $\star_0$, Eq. 2 is equivalent to approximating the integral over the dual cell by multiplying the area of the dual cell $\star\sigma^0$ with the function value at the primal vertex σ^0.

$$\langle \star_0\hat{\omega}^0, \star\sigma^0\rangle = |\star\sigma^0|\langle \hat{\omega}^0, \sigma^0\rangle \tag{18}$$

It is easy to see, that the definition is *well-defined* for arbitrary dual cells $\star\sigma^0$. For the evaluation of dual 1-forms $\star_1\hat{\omega}^1$ on degenerate edges $\star\sigma^1$ we have

$$\langle \star_1\hat{\omega}^1, \star\sigma^1\rangle = \frac{|\star\sigma^1|}{|\sigma^1|}\langle \hat{\omega}^1, \sigma^1\rangle = 0 \tag{19}$$

where the left hand side is 0 for all 1-forms $\hat{\omega}^1$ because of Eq. 1 and the right hand side is 0 because the length the dual edge $\star\sigma^1$ is 0. Furthermore we have $|\sigma^1| > 0$ because we assumed the primal mesh to be non-degenerate. Thus the Hodge star $\star_1$ is *well-defined*. The Hodge star $\star_2^{-1}$ is *well-defined*, as the primal triangles are not affected by degenerate dual edges. For the evaluation of $\star_2^{-1}\hat{\omega}^0_d$ on the vertices of degenerate dual edges $\star\sigma^1 = [\star\sigma^2_i, \star\sigma^2_j]$, we have

$$\frac{1}{|\sigma^2_i|}\langle \star_2^{-1}\hat{\omega}^0_d, \sigma^2_i\rangle = \langle \hat{\omega}^0_d, \star\sigma^2_i\rangle = \langle \hat{\omega}^0_d, \star\sigma^2_j\rangle = \frac{1}{|\sigma^2_j|}\langle \star_2^{-1}\hat{\omega}^0_d, \sigma^2_j\rangle$$

By this result, the dual Hodge star $\star_2^{-1}$ yields primal 2-forms $(\star_2^{-1}\omega^0_d)$ for which the primal Hodge star $\star_2$ is well-defined, as Eq. 4 holds. For the Hodge star $\star_0^{-1}$ we have

$$\langle \star_0^{-1}\hat{\omega}^2_d, \sigma^0\rangle = \frac{1}{|\star\sigma^0|}\langle \hat{\omega}^2_d, \star\sigma^0\rangle \tag{20}$$

To see that the operator is well-defined, we must show that all dual cells have a non-zero area.

Lemma A.1. The area of a dual cell corresponding to a vertex of a non-degenerate primal mesh is non-zero.

Proof. For a dual cell to have an area of zero, all of its vertices must be colinear. As the primal edges are orthogonal to the dual edges, all primal edges adjacent to the corresponding primal vertex must be parallel, which is only possible when the edge lengths are 0. This would be a contradiction to the assumption that the primal mesh is non-degenerate. □

By Eq. 20 and Lemma A.1, the Hodge star $\star_0^{-1}$ is *well-defined*.

A.3. Flat Operators

- **PPP-flat:** The discrete *primal-primal-primal* flat does not involve the dual mesh and thus is *well-defined* for all primal edges.
- **PPD-flat:** The *primal-primal-dual* flat is *well-defined* and the evaluation of the 1-form $X^{\flat_{ppd}}$ on removable edges σ^1 yields $\langle X^{\flat_{ppd}}, \sigma^1 \rangle = \int_{\sigma^1} \bar{X} \cdot (\star\vec{\sigma}^1) = 0$ because of the dot product with $\star\vec{\sigma}^1 = \vec{0}$.
- **PDP-flat:** The *primal-dual-primal* flat averages the vector field on the two adjacent triangles and does not involve the dual edge, so it is *well-defined*.
- **PDD-flat:** The *dual-primal-dual* flat is *well-defined* and the evaluation on degenerate dual edges $\star\sigma^1$ yields
  $\langle X^{\flat_{pdd}}, \star\sigma^1 \rangle = \sum_{\sigma^0 \prec \sigma^1} X(\sigma^0) \cdot (\star\vec{\sigma}^1) = 0$ because of the dot product with $\star\vec{\sigma}^1 = \vec{0}$.
- **DPD-flat:** The *dual-primal-dual* flat is *well-defined* and the evaluation on degenerate dual edges $\star\sigma^1$ yields
  $\langle X^{\flat_{dpd}}, \star\sigma^1 \rangle = \sum_{\sigma^2 \succ \sigma^1} X(\sigma^2) \cdot (\star\sigma^1 \cap \sigma^2) = 0$ because the edge segment $(\star\sigma^1 \cap \sigma^2)$ has length 0 for degenerate dual edges $\star\sigma^1$.

Table 1: Overview which operators are well-defined on meshes that contain degenerate dual edges

	primal	dual
$\mathbf{d}_0$	✓	✓
$\mathbf{d}_1$	✓	✓
$\star_0$	✓	✓
$\star_1$	✓	✗
$\star_2$	(✗)[1]	✓
$\sharp_{pp}$	(✓)[2]	N/A
$\sharp_{pd}$	(✓)[2]	N/A

	primal	dual
$\flat_{ppp}$	✓	N/A
$\flat_{dpp}$	N/A	✗
$\flat_{ppd}$	✓	N/A
$\flat_{pdp}$	✓	N/A
$\flat_{pdd}$	✓	N/A
$\flat_{dpd}$	N/A	✓

1) Except for primal 2-forms that satisfy Eq. 4.
2) Requires a new interpolation scheme for concyclic polygons.

B. Prevalence of Concyclic Triangles Pairs

We examined the prevalence of concyclic triangle pairs for some common 3D model sets.

Table 2: The number of concyclic triangles in different data sets

Name	Number of models	Average number of faces	Average number of concyclic faces	Percent of concyclic faces
Regular Plane (5000 triangles)	1	5,000	5,000	100
Stanford Bunny	1	3,851	0	0
Stanford Bunny (fine)	1	70,580	2	0.003
Stanford Dragon	1	40,348	0	0
Stanford Armadillo	1	345,944	56	0.016
Stanford Happy Buddah	1	123,056	4	0.003
Fandisk	1	18,886	5,088	26.941
Nefertiti	1	2,018,232	40	0.002
COSEG Four-Legged	20	15,227	4	0.023
COSEG Tele-Aliens	200	8,321	357	4.295
COSEG Goblets	12	5,245	1,338	25.518
COSEG Large-Vases	300	2,533	137	5.408
COSEG Chair	20	19,048	6,730	35.329
Snakeboard (Crane)	6	12,091	1,562	12.916
ISS interior (NASA)	1	305,870	77,887	25.464
Gale Crater (NASA)	1	526,342	4,412	0.838
San Diego Conference center (Crane)	1	39,498	13,854	35.075

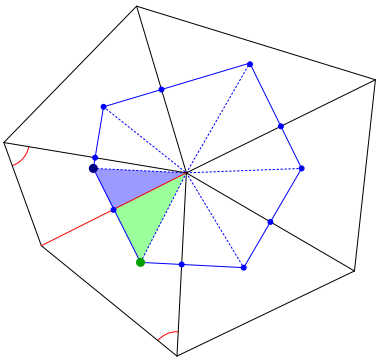

(a) Well-Centered Triangles

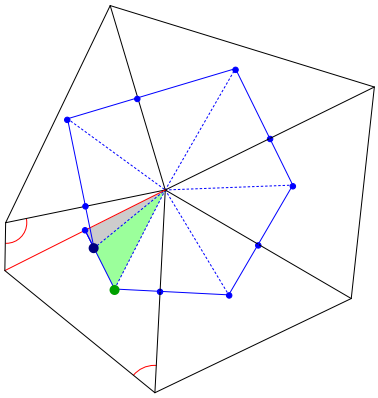

(b) Delaunay Triangle Pair

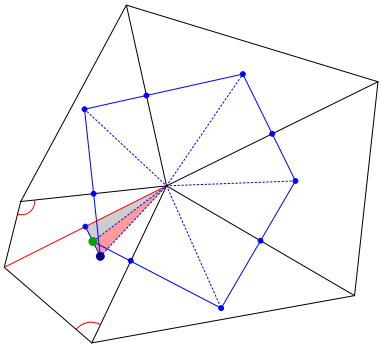

(c) Non-Delaunay Triangle Pair

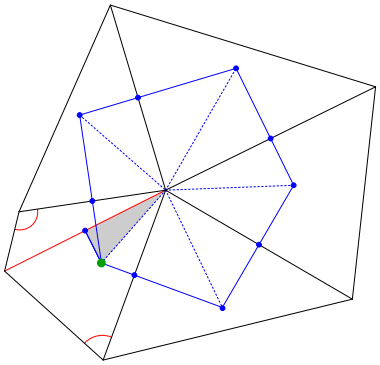

(d) Concyclic Triangle Pair

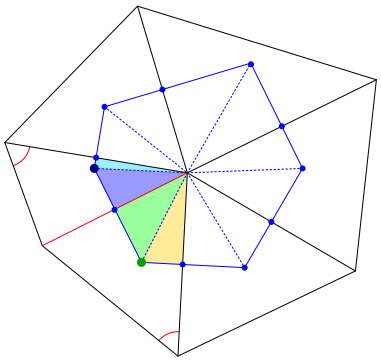

(e) A triangle pair contributes four elementary triangles to the dual cell

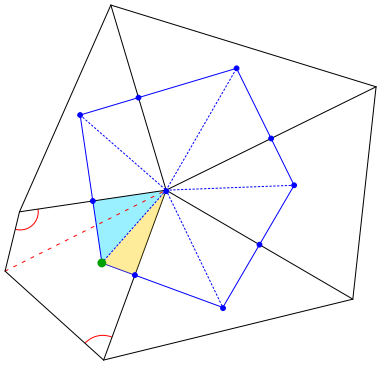

(f) Removing a primal edge collapses the adjacent elementary dual triangles

Figure 11: A dual cell can be partitioned into elementary triangles, defined by the dual cell center, the midpoint of a primal edge, and the circumcenter of an adjacent primal triangle, such that each primal edge contributes two elementary triangles to the cell. (a) When all primal triangles are well-centered, the elementary triangles are a partition of the dual cell. (b) A Delaunay triangle pair may have one elementary triangle with a negative area, which cancels out with a part of the other elementary triangle (gray). (c) A triangle pair that is not Delaunay contributes a negative area to the dual cell area (red). (d) A concyclic triangle pair does not contribute to the area of the dual cell as the two elementary triangles cancel out exactly (gray). (e) The two primal triangles adjacent to a primal edge contribute four elementary triangles to the dual cell. (f) When a removable primal edge is deleted, the adjacent elementary triangles collapse. As they were canceling out each other before, the shape of the dual cell does not change.