Subdivision Directional Fields

BRAM CUSTERS, Utrecht University/TU Eindhoven
AMIR VAXMAN, Utrecht University

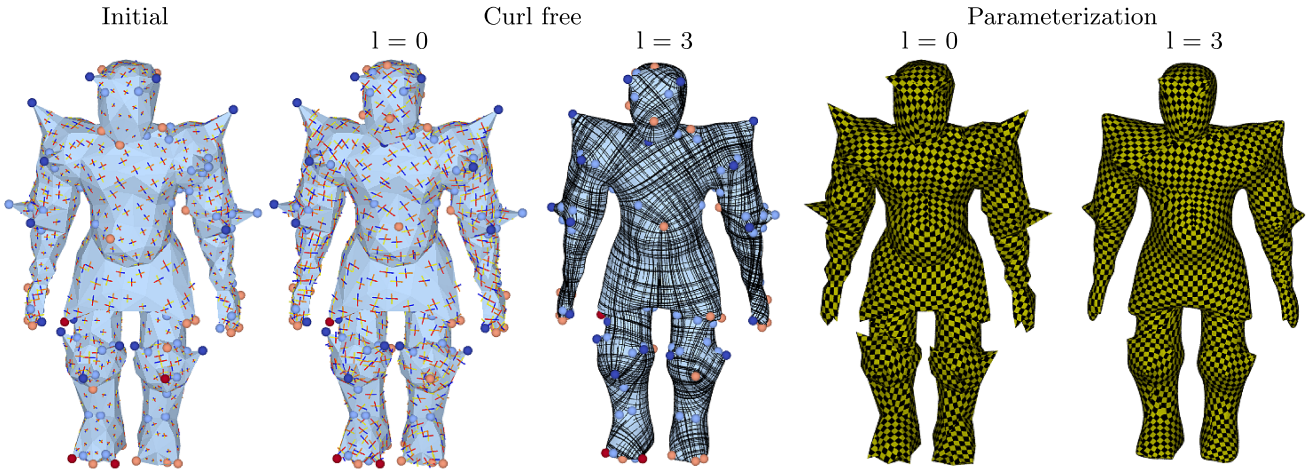

Fig. 1. Rotationally seamless parameterization with a subdivision directional field. An initial field (left) is optimized for low curl at the coarsest level $l = 0$. We subdivide the field to fine level $l = 3$ (center) and then solve for a seamless parameterization in both levels (right). Our subdivision preserves curl and thus results in a low integration error in both levels. The coarse-level optimization takes 7.5 seconds, the subdivision 7.6 seconds, and the parameterization 7.0 seconds, for a total of 22.1 seconds. This is a speedup of about two orders of magnitude compared to running the curl optimization directly on the fine level, taking 1,438.7 seconds.

We present a novel linear subdivision scheme for face-based tangent directional fields on triangle meshes. Our subdivision scheme is based on a novel coordinate-free representation of directional fields as halfedge-based scalar quantities, bridging the mixed finite-element representation with discrete exterior calculus. By commuting with differential operators, our subdivision is structure preserving: it reproduces curl-free fields precisely and reproduces divergence-free fields in the weak sense. Moreover, our subdivision scheme directly extends to directional fields with several vectors per face by working on the branched covering space. Finally, we demonstrate how our scheme can be applied to directional-field design, advection, and robust earth mover's distance computation, for efficient and robust computation.

CCS Concepts: • **Computing methodologies** → **Mesh geometry models**; **Mesh models**; *Shape analysis*;

Additional Key Words and Phrases: Directional fields, vector fields, subdivision surfaces, differential operators

Authors' addresses: B. Custers, Applied Geometric Algorithms group, TU Eindhoven, Groene Loper 5, Office MF 4.103, 5612 AZ Eindhoven, the Netherlands; email: b.a.custers@tue.nl; A. Vaxman, Geometric Computing Group, Utrecht University, Princetonplein 5, De Uithof, 3584 CC, Utrecht, The Netherlands; email: a.vaxman@uu.nl.

0730-0301/2020/01-ART11 $15.00
https://doi.org/10.1145/3375659

ACM Reference format:
Bram Custers and Amir Vaxman. 2020. Subdivision Directional Fields. *ACM Trans. Graph.* 39, 2, Article 11 (January 2020), 20 pages.
https://doi.org/10.1145/3375659

1 INTRODUCTION

Directional fields are central objects in geometry processing. They represent flows, alignments, and symmetry on discrete meshes. They are used for diverse applications such as meshing, fluid simulation, texture synthesis, and architectural design. There is then great value in devising robust and reliable algorithms that design and analyze such fields. In this article, we work with piecewise-constant tangent directional fields, defined on the faces of a triangle mesh. A directional field is the assignment of several vectors per face, where the most commonly used fields comprise single vectors. The piecewise-constant face-based representation of directional fields is a mainstream representation within the (mixed) *finite-element method* (FEM), where the vectors are often gradients of piecewise-linear functions spanned by values on the vertices.

Working with a fine-resolution smooth (and good-quality) mesh is often essential to get good results with methods that produce piecewise-constant directional fields (PCDFs). However, working on a fine mesh is also computationally expensive and often wasteful—the desired directional fields are smooth and mostly defined by a sparse set of features such as sinks, sources, and vortices.

A classical way to bridge this gap is to work with a multiresolution representation, based on a nested hierarchy of meshes. A

popular way to generate this representation is to use *subdivision surfaces*. Subdivision surfaces are generated by operators that comprise a set of stencils, often linear and stationary (with a fixed stencil), that are used to recursively refine functions defined on meshes (and consequently the vertex positions). These operators can be used to prolong and restrict functions between coarse and fine levels, allowing for *multigrid* field computation. We consider the *limit surface* as the target domain on which we compute the fields and represent the degrees of freedom of the computation by the coarse control mesh through subdivision.

To be able to work with hierarchical directional fields on subdivision surfaces, one needs to define specialized subdivision operators. A necessary requirement for obtaining consistent results is that the subdivision operators are *structure preserving*. In other words, the differential and topological properties of the directional fields are preserved under subdivision. This can be achieved by designing subdivision operators that commute with differential operators. Unfortunately, differential operators on piecewise-constant face-based fields are defined with the metric and the embedding of the mesh (e.g., face areas and normals) built in. As a result, these quantities have complicated and nonlinear expressions in the linearly subdivided vertex coordinates. Creating linear stationary subdivision operators directly on face-based directional fields is then a challenging task. Recently, de Goes et al. [2016b] devised a method for subdivision vector-field processing for differential forms in the *discrete exterior calculus* (DEC) setting. The differential quantities in DEC are inherently separated into combinatorial and metric operators; due to this, it is possible to define a stationary subdivision scheme for differential forms that commutes with the combinatorial part alone, as introduced in Wang et al. [2006].

Inspired by this insight, we introduce a coordinate-free representation for face-based fields, allowing us to decompose the face-based differential operators into independent combinatorial and metric components. With this decomposition, we define linear stationary subdivision operators for such fields. Our scheme naturally extends to branched covering spaces, where we then apply it to directional fields with an arbitrary number of vectors per face.

2 RELATED WORK

2.1 Directional Fields

Tangent directional fields on discrete meshes have been researched extensively in recent years. The important aspects of their design and analysis are summarized in two relevant surveys: de Goes et al. [2016a] focus on differential properties of mostly single vector fields, with an emphasis on different discretizations on meshes, whereas Vaxman et al. [2016] focus on discretization and representation of directional fields (with N vectors at every given tangent plane) and their applications.

The fundamental challenge of working with directional fields is how to discretize and represent them. The most common discretization considers one directional object per face, or alternatively piecewise-constant elements (e.g., Bommes et al. [2009]; Crane et al. [2010]; Tong et al. [2003]; Wardetzky [2006]). This representation conforms with the classic piecewise-linear paradigm of the FEM and admits a dimensionality-correct cohomological structure, when mixing conforming and nonconforming elements [Wardetzky 2006]. Moreover, the natural tangent planes, as a supporting plane to the triangles in the mesh, allow for simple representations of N-directional fields [Crane et al. 2010; Diamanti et al. 2014; Ray et al. 2008]. However, the representation is only C^0 smooth, and makes it difficult to define discrete operators of higher order including derivatives of directional fields, such as the Lie bracket [Azencot et al. 2013; Mullen et al. 2011; Sageman-Furnas et al. 2019], or Killing fields [Ben-Chen et al. 2010].

An alternative approach to single-vector field processing is DEC [Crane et al. 2013; Hirani 2003], which represents vector fields as 1-forms, discretized as scalars on oriented edges. DEC enjoys the benefit of representing fields in a coordinate-free manner, which allows for a decomposition of the differential operators into combinatorial and metric components. This is beneficial for the subdivision scheme we work with in this article. However, DEC is not as of yet defined to work with general N-directional fields, and, when using linear Whitney forms, it still suffers from discontinuities at edges and vertices. We note that alternative approaches exist that use vertex-based definitions [Knöppel et al. 2013; Liu et al. 2016; Sharp et al. 2019; Zhang et al. 2006], representing directional fields on tangent spaces defined at vertices. While enjoying better continuity, a full suite of differential operators has not yet been studied for them—in particular, differential operators that define discrete sequences, necessary for a correct Helmholtz-Hodge decomposition [Poelke and Polthier 2016; Wardetzky 2006].

2.2 Multiresolution Vector Calculus

Directional fields are important for applications such as meshing [Bommes et al. 2009; Kälberer et al. 2007; Zadravec et al. 2010], simulations on surfaces [Azencot et al. 2015], parameterization [Campen et al. 2015; Diamanti et al. 2015; Myles and Zorin 2012], and nonphotorealistic rendering [Hertzmann and Zorin 2000]. An underlying objective in all of these applications is to obtain fields that are as smooth as possible. Nevertheless, as demonstrated in Vaxman et al. [2016], directional fields are subject to aliasing and noise artifacts quite easily for coarse meshes. Using fine meshes alleviates this problem to some extent but incurs a price of increased computational overhead, especially for nonlinear methods. For this, a smooth and low-dimensional representation for smooth directional fields on fine meshes, such as the one we introduce, is much needed.

The most prevalent approach to low-dimensional smooth processing on fine meshes is to use some refinable multiresolution hierarchy. This paradigm is extensively employed in the FEM literature when using either refined elements (*h-refinement*) or higher-order basis functions (*p-refinement*) [Babuška and Suri 1994]. This has also been applied to vector fields in planes and in volumes [Schober and Kasper 2007]. A major difference in which our subdivision method departs from both of these approaches is that the geometry of the target limit surface is different from that of the control coarse mesh. As such, using p- or h-refinement directly on the coarse cage is susceptible to committing the so-called variational crime [Strang and Fix 2008], where the function space and the computation domain are mismatched.

A more closely related prominent approach to refinable spaces is *Isogeometric Analysis* [Hughes et al. 2005]. The premise is computation over refinable B-spline basis functions, replacing the piecewise-linear FEM functions. The setting promotes integration

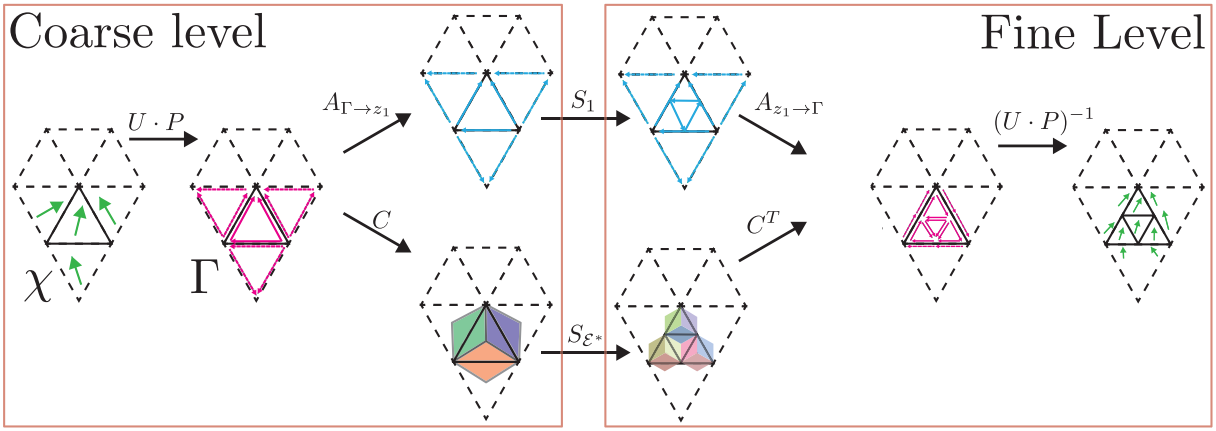

Fig. 2. SHM pipeline. A face-based tangent field in the space $\mathcal{X}$ is converted to the equivalent halfedge representation in space Γ (Section 5). The halfedge form is further separated into a DEC 1-form z_1 and a nonconforming function ϵ that is the half-curl of the field (Section 5.3). They are individually subdivided (Section 6) and assembled back to a field on a finer mesh.

over the target (smooth) domain and therefore is theoretically correct and structure preserving. However, they rely on quadrature rules to perform the complicated integrals that involve the basis functions. Methods such as those of Jüttler et al. [2016] and Nguyen et al. [2014] employ subdivision rules for evaluation on the limit surface but then design approximative quadrature rules for the exact integrals, tailored to fit specific differential operators.

A recent work by de Goes et al. [2016b] utilizes subdivision for 1-forms (first introduced in Wang et al. [2006]) as a means to represent vector fields in recursively refinable spaces. By doing so, they efficiently emulate the IGA premise in a linear setting and directly on the discrete meshes. This technique substitutes coarse inner-product matrices with inner-product matrices *restricted* from the fine domains, encoding fine-mesh geometry on the coarse mesh. Using subdivision matrices as prolongation operators is akin to collapsing a single V-cycle in a multigrid setting [Brandt 1977]. The essence of the technique is to design stationary 1-form subdivision operators that commute with the discrete differential operators. This is made possible as DEC operators are purely combinatorial.

Unfortunately, their approach does not readily extend to face-based piecewise-constant fields. The effect of stationary subdivision methods on triangle areas and normals is not linear, which makes it difficult to establish the required commutation rules. Our article introduces a novel representation of face-based fields using halfedge-based forms, which can be readily subdivided using stationary operators. As such, we introduce a metric-free subdivision method for face-based directional fields that guarantees structure preservation.

Directional fields. Much less has been explored in the literature about differential operators on directional fields. In Bommes et al. [2009] and Kälberer et al. [2007], directional fields are used as candidate gradients for functions on branched covering spaces. Diamanti et al. [2015] further define PolyCurl, which encodes the curl of N-directional fields. They then optimize for curl-free fields. However, we are not aware of any study of general directional calculus and its applications to geometry processing. We provide a branched subdivision scheme, and subsequently a multiresolution representation and a calculus suite for directional fields.

2.3 Subdivision Surfaces in Geometry Processing

Subdivision surfaces are popular objects in geometry processing, and are methods of choice for shape design for animation [Liu et al. 2014] and architectural geometry [Liu et al. 2006]. Their most popular utility is that of multiresolution (or just coarse-to-fine) mesh editing. In the context of simulation, they have been applied to fluid simulation [Stam 2003], thin-shell design [Cirak et al. 2002], and surface deformation [Grinspun et al. 2002; Thomaszewski et al. 2006]. The latter work also uses the folded V-cycle approach to work on the coarse mesh with the limit surface metric; nevertheless, they work with quadrature as well to approximate the exact solution.

3 CONTRIBUTIONS

The main contributions of our article are summarized as follows.

Halfedge forms (Section 5). We define a novel coordinate-free representation for piecewise-constant vector fields (PCVF) on faces. The essence of this representation is to consider their projection on the halfedges defining each triangle. We prove the equivalence of this representation to that of face-based fields, and show that these *halfedge forms* can be represented as the combination of a DEC 1-form and edge-based curl, which is consistent with the case where the 1-form is exact (the gradient of some scalar function). Halfedge forms are then a new type of 1-form that bridges mixed-FEM representation with that of DEC.

Subdivision vector fields (Section 6). Given the coordinate-free halfedge-form representation, we introduce a subdivision scheme to face-based vector fields with the following properties:

- Coarse gradient fields are subdivided into fine gradient fields, where the underlying scalar function is refined using a vertex-based scalar subdivision method.
- The curl of a subdivided vector field, as a scalar function, is a refinement of of the curl of the coarse vector field.

We depict the subdivision pipeline in Figure 2.

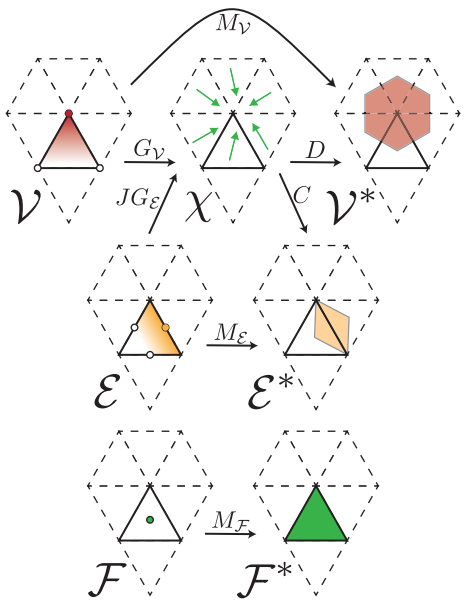

Fig. 3. The FEM function spaces and associated differential operators.

Subdivision directional fields (Section 7). Since we work with face-based fields, we show how our subdivision readily extends to N-directional fields, where there are N vectors per face, by reducing this case into working with single-vector fields in *branched* spaces.

We apply this *structure-preserving* subdivision to several applications in Section 8: earth mover's distance computation, seamless parameterization, vector field design, and operator-based advection. The common advantage that our method provides is the ability to process vector fields on subdivided meshes (with many triangles), considering only the degrees of freedom spanned by the coarse control mesh. By doing this, we save both time and memory.

We denote our face-based directional-field subdivision framework as the *subdivision halfedge-form method* (SHM).

4 BACKGROUND

We introduce a new discrete representation for vector fields that bridges mixed FEM and DEC. For this, we require an extensive amount of background on these spaces. Nevertheless, for the sake of compactness, we mostly introduce these well-known notions in the notation and formulation we use and little else; see Table 1 for our notations, Table 2 for the definitions of the discrete differential operators, and Figure 3 for the FEM space that we work in. We refer the reader to de Goes et al. [2016a] and Wardetzky [2006] for a more comprehensive account of the operators in FEM, and to Desbrun et al. [2005] for the operators in DEC. For compactness, we reduce the polysemous "FEM" to only mean the conforming/nonconforming piecewise-linear finite-element representation, to distinguish it from DEC, which is in essence another type of finite-element representation.

4.1 Function Spaces

We work with a triangle mesh $\mathcal{M} = (V, E, F)$ of arbitrary genus, and with or without boundaries. As we combine FEM and DEC formulations, we need to streamline notation at the expense of conventionality. We define $\mathcal{V}$ as the space of piecewise-linear (conforming) vertex-based functions, corresponding to 0-forms with linear Whitney forms in DEC and $\mathcal{S}_h$ in FEM. We further define $\mathcal{E}$ as the space of piecewise-linear mid-edge (nonconforming) functions, also known as the *Crouzeix-Raviart* elements [Crouzeix and Raviart 1973], corresponding to $\mathcal{S}_h^*$ in FEM. We define $\mathcal{F}$ as the space of piecewise-constant functions on faces, corresponding with dual 2-forms in DEC. We define the corresponding integrated (weak) function spaces on vertices as $\mathcal{V}^*$ (corresponding to dual 0-forms, integrated over Voronoi areas), on edges as $\mathcal{E}^*$ (integrated over edge diamond areas), and on faces as $\mathcal{F}^*$ (corresponding to primal 2-forms in DEC). Finally, we denote the space of face-based PCDFs of degree N, defined on the tangent spaces spanned by the supporting planes to the faces, as $\mathcal{X}^N$. The latter is in accordance with the conventional notation. We introduce our operators to the classic case of $N = 1$ and then generalize our constructions to N-directional fields in Section 7. For case $N = 1$, we omit the power and just use $\mathcal{X}$, the space of PCVF.

Orientation. We choose an arbitrary (but fixed) orientation for every edge in the mesh. This orientation consistently defines both source and target vertices (*primal orientation*), and left and right faces for each edge (*dual orientation*; corresponding with the CCW orientation of every face). For instance, in our notation, we use e_{ik} and get $left(e_{ik}) = ikl = t_2$ and $right(e_{ik}) = ijk = t_1$ (Figure 4). For edge e and adjacent face f, we define $s_{e,f} = \pm 1$ as the sign encoding the orientation (positive if $f = left(e)$, i.e. e is oriented CCW with respect to the face normal of f). DEC 1-forms depend on the direction and sign of the edge, so they are denoted as *oriented quantities*. Quantities in $\mathcal{E}^*$ depend on the direction of the edge on which they are defined, but not on the specific sign (whether e_{ik} or e_{ki}), and thus we denote them as *unsigned quantities*. For a face f, we define $J_{|f} = [\hat{n}_f \times]$ as the operator that performs the rotation around its normal $\hat{n}_f$.

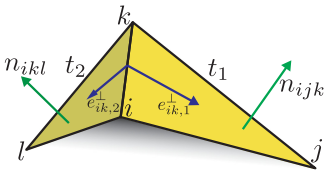

Fig. 4. Our notation for a single flap.

Mixing spaces. It is well known [Polthier and Preuß 2003] that the discrete differential FEM operators preserve the structure of differential operators in the discrete setting. In other words, we have a sequence $\text{Image}(G_{\mathcal{V}}) \subset \ker(C_{\mathcal{E}^*})$ (gradient fields are curl free) and a (dual) sequence $\text{Image}(JG_{\mathcal{E}}) \subset \ker(D_{\mathcal{V}})$ (rotated *cogradient* fields are divergence free). This *structure-preserving* property is essential to the correct and stable behavior of differential equations discretized with such operators. Note that the entire formulation can be done in a dual manner by switching conforming and nonconforming spaces and operators. However, we restrict ourselves to conforming gradients and nonconforming rotated cogradients. As such, we omit the space-indicating

Table 1. List of Notations and Symbols

Notation	Dimensions	Explanation
$\mathcal{V}, \mathcal{V}^*$	$\lvert V\rvert$	Primal and dual vertex-based PL conforming functions.
$\mathcal{E}, \mathcal{E}^*$	$\lvert E\rvert$	Primal and dual midedge-based PL nonconforming functions.
$\mathcal{F}, \mathcal{F}^*$	$\lvert F\rvert$	Primal and dual face-based PC functions.
$\mathcal{X}$	$3\lvert F\rvert$	Piecewise-constant face-based vector fields.
$\mathcal{Z}_1$	$\lvert E\rvert$	Edge-based DEC 1-forms.
Γ	$2\lvert F\rvert$	Halfedge forms.
P	$\lvert\Gamma\rvert \times \lvert\mathcal{X}\rvert$	Projection operator $\mathcal{X} \to \Gamma$ (Equation (10)).
U	$2\lvert F\rvert \times 3\lvert F\rvert$	Unpacking operator for halfedge form in each face, respecting null-sum (Equation (9)).
$M_{\mathcal{V}}$	$\lvert V\rvert \times \lvert V\rvert$	Mass matrix for vertices (Voronoi areas).
$M_{\mathcal{E}^*}$	$\lvert E\rvert \times \lvert E\rvert$	Mass matrix for integrated edge quantities (*inverse* diamond areas).
$M_{\mathcal{F}}$	$\lvert F\rvert \times \lvert F\rvert$	Mass matrix for face (triangle areas).
M_Γ	$2\lvert F\rvert \times 2\lvert F\rvert$	Mass matrix for halfedge forms (packed cotangent weights; Equation (14)).
$M_{\mathcal{X}}$	$3\lvert F\rvert \times 3\lvert F\rvert$	Mass matrix for piecewise-constant face-based vector fields. This amounts to repeating $M_{\mathcal{F}}$ three times per triangle.
$S^l_{\mathcal{P}}$	$\lvert\mathcal{P}^{l+1}\rvert \times \lvert\mathcal{P}^l\rvert$	Subdivision matrix for space $\mathcal{P}$ from level l to $l+1$. We use $\mathcal{P} \in \{\mathcal{V}, \mathcal{E}^*, \mathcal{F}^*, \mathcal{Z}_1, \Gamma\}$.
$\mathbb{S}^l_{\mathcal{P}}$	$\lvert\mathcal{P}^l\rvert \times \lvert\mathcal{P}^0\rvert$	Aggregated subdivision matrix from levels 0 to l.
$\mathbb{M}^0_{\mathcal{P}}$	$\lvert\mathcal{P}^0\rvert \times \lvert\mathcal{P}^0\rvert$	Restricted mass matrix from some level l to level 0 (Equation (5)).
$A_{\mathcal{Z}_1 \to \Gamma}$	$\lvert\Gamma\rvert \times \lvert E\rvert$	Assigning a 1-form value from an edge to its halfedges. $A_{\Gamma \to \mathcal{Z}_1} = (A_{\mathcal{Z}_1 \to \Gamma})^T$ sums the halfedge values to a single 1-form per edge.
$A_{\mathcal{E}^* \to \mathcal{F}^*}$	$\lvert F\rvert \times \lvert E\rvert$	Summing integrated edge quantities to the adjacent faces.
W	$\lvert\Gamma\rvert \times 2\lvert E\rvert$	Computing edge mean 1-form z_1 and half-curl ϵ (Equation (16)).
$\mathcal{P}^N$	$\lvert\mathcal{P}\rvert^N$	N-branched space of functions in $\mathcal{P}$ (e.g., N-directional fields are in $\mathcal{X}^N$).

For differential operators see Table 2. Superscript capital N is for number of vectors, and small l is for subdivision level.

subscripts, and just use D for (conforming) divergence and C for (nonconforming) curl.

Helmholtz-Hodge decomposition. Mixing conforming and nonconforming operators is essential to have a dimensionality-consistent Hodge decomposition [Wardetzky 2006]. For a closed surface without boundary, there is a well-defined Helmholtz-Hodge decomposition of $\mathcal{X}$ as follows:

$$\mathcal{X} = \text{Image}\,(G_{\mathcal{V}}) \oplus \text{Image}\,(JG_{\mathcal{E}}) \oplus \mathcal{H}_{\mathcal{X}}. \tag{1}$$

$\text{Image}(G_{\mathcal{V}})$ is the space of vectors fields that are gradients of functions in $\mathcal{V}$, $\text{Image}(JG_{\mathcal{E}})$ is the space of rotated cogradients of functions in $\mathcal{E}$, and $\mathcal{H}_{\mathcal{X}} = \ker(C) \bigcap \ker(D)$ is the space of PCVF *harmonic fields*. The space of harmonic fields has the correct dimension $2g$, where g is the genus of the mesh.

Inner products. Inner products on the function spaces are represented as mass matrices M, where two elements u, v in column vector form have the inner product $\langle u, v\rangle_{\mathcal{P}} = u^T M_{\mathcal{P}} v$ in some function space $\mathcal{P}$. $M_{\mathcal{X}}$ is the mass matrix of space $\mathcal{X}$, comprising diagonal values of triangle areas for each component of the vector field, and we further define $M_{\mathcal{V}}$ to be the diagonal matrix of *Voronoi areas* of every vertex. We define $M_{\mathcal{E}}$ to be the diagonal matrix of *diamond areas* supported on each edge (see Figure 3). Mass matrices for dual spaces are inverted mass matrices of the corresponding primal spaces. We note that $M_{\mathcal{V}}$ and $M_{\mathcal{E}}$ are in fact *lumped* versions of the FEM mass matrices. This lumping is done to make them diagonal and thus have simple inverses. We denote the L_2 norm of space $\mathcal{P}$ by $||u||_{\mathcal{P}} = \sqrt{\langle u, u\rangle_{\mathcal{P}}}$.

Hodge Laplacian. The integrated discrete Hodge Laplacian is obtained from minimizing the Dirichlet energy of vector fields and has the following form [Brandt et al. 2017]:

$$L_{\mathcal{X}} = C^T M_{\mathcal{E}^*} C + D^T M_{\mathcal{V}^*} D.$$

Its null space contains the harmonic vector fields. The pointwise version is $M_{\mathcal{X}}^{-1} L_{\mathcal{X}}$.

4.2 Discrete Exterior Calculus

DEC function spaces. The setup of DEC [Desbrun et al. 2005] on surface meshes is an alternative to the piecewise-constant representation. Instead of representing vectors explicitly, DEC works with primal and dual k-forms, where primal 0-forms are (pointwise) vertex-based functions, primal 1-forms are (integrated) edge-based functions (representing vectors), and primal 2-forms are (integrated) face-based functions. The space of primal 0 forms $\mathcal{Z}_0$, with the interpolation of linear *Whitney forms*, identifies with $\mathcal{V}$. The space of 1-forms $\mathcal{Z}_1$ comprises scalars on edges, representing oriented quantities. Such quantities are oriented in the sense that when a scalar z is attached to edge e_{ik}, then the corresponding scalar for the edge e_{ki} is $-z$. Note that the FEM space $\mathcal{E}^*$ does not have this property or edge sign dependence, and therefore it does not identify with $\mathcal{Z}_1$. The space of 2-forms $\mathcal{Z}_2$ identifies with $\mathcal{F}^*$ (note the dual space, as elements in $\mathcal{Z}_2$ are integrated).

The space of dual 0-forms $\mathcal{Z}_0^*$ are integrated vertex-based quantities and identifies with $\mathcal{V}^*$. Similarly, $\mathcal{Z}_2^*$ identifies with $\mathcal{F}$. Dual 1-forms in the space $\mathcal{Z}_1^*$ are defined on the union of the orthogonal duals to the edges. For edge ik in triangles ijk and ikl, the dual e^*_{ik} is the two perpendicular bisectors to e_{ik} from the center of the

Table 2. Operators Per Representations

Operator	FEM		DEC		Γ	
	Spaces	Formulation	Spaces	Formulation	Spaces	Formulation
Primal gradient	$\mathcal{V} \to \mathcal{X}$	$G_{\mathcal{V}}$	$\mathcal{V} \to \mathcal{Z}_1$	d_0	$\mathcal{V} \to \Gamma$	$d_{0,\Gamma} = U^{-1} A_{\mathcal{Z}_1 \to \Gamma} d_0$
Dual rotated gradient	$\mathcal{E} \to \mathcal{X}$	$JG_{\mathcal{E}}$	$\mathcal{F}^* \to \mathcal{Z}_1$	$(M_1)^{-1} d_1^T$	$\mathcal{E} \to \Gamma$	$-(M_\Gamma)^{-1} C_\Gamma^T$
Divergence	$\mathcal{X} \to \mathcal{V}^*$	$D = G_{\mathcal{V}}^T M_{\mathcal{X}}$	$\mathcal{Z}_1 \to \mathcal{V}^*$	$d_0^T M_1$	$\Gamma \to \mathcal{V}^*$	$D_\Gamma = (d_{0,\Gamma})^T M_\Gamma$
Curl	$\mathcal{X} \to \mathcal{E}^*$	C	$\mathcal{Z}_1 \to \mathcal{F}^*$	d_1	$\Gamma \to \mathcal{V}^*$	$C_\Gamma = C \cdot U$
Primal Laplacian	$\mathcal{V} \to \mathcal{V}^*$	$L_{\mathcal{V}} = G_{\mathcal{V}}^T M_{\mathcal{X}} G_{\mathcal{V}}$	$\mathcal{V} \to \mathcal{V}^*$	$d_0^T M_1 d_0$	$\mathcal{V} \to \mathcal{V}^*$	$(d_{0,\Gamma})^T M_\Gamma d_{0,\Gamma}$
Dual Laplacian	$\mathcal{E} \to \mathcal{E}^*$	$L_{\mathcal{E}} = (JG_{\mathcal{E}})^T M_{\mathcal{X}} JG_{\mathcal{E}}$	$\mathcal{F}^* \to \mathcal{F}$	$d_1 M_1^{-1} d_1^T$	$\mathcal{E} \to \mathcal{E}^*$	$C_\Gamma M_\Gamma^{-1} (C_\Gamma)^T$
Hodge Laplacian	$\mathcal{X} \to \mathcal{X}$	$L_{\mathcal{X}} = G_{\mathcal{V}} M_{\mathcal{V}^*} G_{\mathcal{V}}^T M_{\mathcal{X}} +$ $JG_{\mathcal{E}} M_{\mathcal{E}^*} (JG_{\mathcal{E}})^T M_{\mathcal{X}}$	$\mathcal{Z}_1 \to \mathcal{Z}_1$	$L_1 = d_0 M_{\mathcal{V}^*} d_0^T M_1 +$ $M_1^{-1} d_1^T M_{\mathcal{F}}^* d_1$	$\Gamma \to \Gamma$	$L_\Gamma = d_{0,\Gamma} M_{\mathcal{V}^*} (d_{0,\Gamma})^T \cdot M_\Gamma +$ $M_\Gamma^{-1} (C_\Gamma)^T \ M_{\mathcal{E}^*} C_\Gamma$

All operators are presented in their integrated versions when applicable.

circumscribing circles of each triangle and therefore differs from the rotated edge $e_{ik}^{\perp}$ used in FEM.

Differential operators. Two fundamental discrete operators are combined to create an entire suite of vector calculus: the exterior derivative d, taking k-forms into $(k+1)$-forms, and the Hodge star $\star$, taking primal k-forms into dual $2-k$ dual forms. For instance, the lumped $\star_1 : \mathcal{Z}_1 \to \mathcal{Z}_1^*$ is defined as $\star_{|ik,1} = \frac{|e_{ik}^*|}{|e_{ik}|}$. To streamline notation, we use M_1 to represent $\star_1$. $M_1 : |E| \times |E|$ is a diagonal matrix that contains the weights per edge. M_0 identifies with $M_{\mathcal{V}}$, as a diagonal matrix of Voronoi areas, and M_2 identifies with $M_{\mathcal{F}^*}$.

DEC operators also define a (de-Rham) sequence, as $d^2 = 0$ in the discrete setting. Therefore, DEC is also structure preserving. In the dual setting, we also work with the boundary operator $\partial = d^T$. Intuitively, ∂ sums up $(k+1)$-forms into k-forms of elements (chains) adjacent to them, with relation to the mutual orientation. The vector calculus operators are then interpreted as follows: the curl operator is simply d_1, where curl is a primal 2-form in DEC, and primal (weak) divergence is $(d_0)^T M_1$, producing a dual 0-form.

The DEC version of Hodge decomposition for 1-forms is such that for each $z_1 \in \mathcal{Z}_1$ there exist $z_0 \in \mathcal{Z}_0$ and $z_2 \in \mathcal{Z}_2$ such that

$$z_1 = d_0 z_0 + M_1^{-1} d_1^T M_2 z_2 + h_1, \tag{2}$$

where h_1 is a harmonic 1-form that is both closed and coclosed.

Between DEC and FEM. As linear discrete frameworks, DEC and FEM admit a similar power of expression, for instance $L_0 = L_{\mathcal{V}}$, the cotangent weights Laplacian. However, they are incompatible otherwise; $|\mathcal{Z}_1| = |E|$, whereas $|\mathcal{X}| = 2|F|$ (the ambient dimension in the raw representation is $3|F|$). As such, the differential operators are also different in dimensions.

Note that the commonly used diagonal M_1 is a lumped version of the "correct" (Galerkin) mass matrix for 1-forms, integrating over the interpolated linear Whitney forms [de Goes et al. 2014]. The lumped version results in diagonal matrices that are efficient to work with, especially with regard to solving equations. Moreover, interpolated closed (and, as a subset, exact) 1-forms are piecewise constant; in that case, the lumped M_1 is the correct inner product. This is the reason that FEM and DEC vertex Laplacians identify.

DEC has an advantage over FEM in the sense that it allows for a natural separation between the combinatorial differential operator d, and the metric encoded in the mass matrices, whereas PCVF spaces do not exhibit this separation. This distinction plays an important part in our definition of the subdivision operators.

4.3 Subdivision Exterior Calculus

Subdivision surfaces. A subdivision surface is a hierarchy of refined meshes, starting from a coarse *control mesh* and converging into a smooth fine mesh. We focus on approximative triangle-mesh schemes for both vertex-based and face-based functions. Extending notation from de Goes et al. [2016b] and Wang et al. [2006], we denote a subdivision operator as $S_{\mathcal{P}}^l$, where it subdivides an object of space $\mathcal{P}$ defined on a mesh in level l, denoted as $\mathcal{M}^l$, to an object on a mesh of the refined space in $\mathcal{M}^{l+1}$. For instance, $S_{\mathcal{E}^*}^5$ subdivides an unsigned integrated edge quantity in $\mathcal{E}^*$ from level 5 to level 6.

We denote the product of subdivision matrices from the coarsest level to a given level l as $\mathbb{S}_{\mathcal{P}}^l = \prod_{i=0}^{l-1} S_{\mathcal{P}}^i$. The columns of $\mathbb{S}_{\mathcal{P}}^l$ converge into refined basis functions $\Psi_{\mathcal{P}}^0$ defined on $\mathcal{M}^l$. These basis functions admit a nested refinable heirarchy:

$$\Psi_{\mathcal{P}}^0 \subset \Psi_{\mathcal{P}}^1 \subset \cdots \subset \Psi_{\mathcal{P}}^l, \tag{3}$$

where a function $\Psi_{\mathcal{P}}^k$ is a linear combination of basis functions at level $\Psi_{\mathcal{P}}^{k+1}$, encoded in the matrix $S_{\mathcal{P}}^k$: $\Psi_{\mathcal{P}}^k = \Psi_{\mathcal{P}}^{k+1} S_{\mathcal{P}}^k$. Note that $\Psi_{\mathcal{P}}^0 = \Psi_{\mathcal{P}}^l \mathbb{S}_{\mathcal{P}}^l$.

Structure-preserving subdivision. The essence of Subdivision Exterior Calculus (SEC) is the definition of stationary subdivision matrices for k-forms that commute with the differential as follows:

$$\begin{aligned} d_0 S_0 &= S_1 d_0, \\ d_1 S_1 &= S_2 d_1. \end{aligned} \tag{4}$$

This commutation subdivides exact 1-forms into exact 1-forms where the underlying 0-form is refined. Similarly, the curl of a fine 1-form is the subdivided curl of the coarse 1-form.

Restricted inner products. Choosing Loop subdivision [Biermann et al. 2000; Loop 1987] for S_0 and half-box spline subdivision [Prautzsch et al. 2002] for S_2 completely defines S_1, with some assumptions on the symmetry of the S_1 stencil. In de Goes et al. [2016b], the subdivision operator is mainly used for the purpose of defining mass matrices on the coarse mesh as *restricted* fine mass matrices:

$$\mathbb{M}^0_{\mathcal{P}} = \left(\mathbb{S}^l_{\mathcal{P}}\right)^T \cdot M^l_{\mathcal{P}} \cdot \mathbb{S}^l_{\mathcal{P}} \tag{5}$$

for the space $\mathcal{P}$ and associated subdivision matrix $\mathbb{S}^l_{\mathcal{P}}$ from level 0 to level l as earlier. The restricted mass matrix $\mathbb{M}^0_{\mathcal{P}}$ is exactly the product between subdivided $\mathcal{P}$-forms in the fine level l. The restricted mass matrices are in general no longer diagonal; however, they have a limited support (usually just two rings), derived from the support of the subdivision matrix. Working with restricted mass matrices provides SEC with a smooth, localized, and small function space on the limit surface in a structure-preserving manner that does not require special treatment for singularities, replacing the quadrature methods employed by IGA.

Divergence pollution. The relation of the SEC divergence to the fine DEC divergence reveals an interesting insight:

$$\begin{aligned}(d_0)^T \mathbb{M}^0_1 z^0_1 &= (d_0)^T \left(\mathbb{S}^l_1\right)^T \cdot M^l_1 \cdot \mathbb{S}^l_1 z^0_1 \\ &= \left(\mathbb{S}^l_0\right)^T (d_0)^T \cdot M^l_1 \cdot \left(\mathbb{S}^l_1 z^0_1\right) \\ &= \left(\mathbb{S}^l_0\right)^T \left((d_0)^T \cdot M^l_1 \cdot z^l_1\right).\end{aligned} \tag{6}$$

In words, the SEC divergence of a coarse 1-form z^0_1 subdivided into fine 1-form z^l_1 is not exactly the subdivided coarse divergence; it is rather equivalent only when tested against the *test functions* Ψ^0_0. Simply put, the divergence of the fine form might contain "high-frequency" components that are in $\ker(\mathbb{S}^l_1)^T$, where $(\mathbb{S}^l_1)^T$ acts effectively as a low-pass filter. We denote this as *divergence pollution.*

5 HALFEDGE FORMS

We aim to create a stationary subdivision scheme for PCVFs, inspired by SEC, achieving our goal of establishing a framework of hierarchical spaces for directional fields. For this, we need to first overcome the challenge of metric-free representation that allows for stationary commutation. We do so in the following by creating a halfedge representation for $\mathcal{X}$.

For each oriented edge e_{ik} adjacent to faces t_1 and t_2, we consider its *halfedges* $e_{ik,1}$ and $e_{ik,2}$ (with the notation of Figure 4). Note that they are both oriented in the same direction as e_{ik}; this departs from the usual doubly connected edge list convention [de Berg et al. 2008], where halfedges are of opposing orientations, and counterclockwise oriented with respect to their face normal. We choose to co-orient them with the edge, as it is a more natural convention for our differential operators.

We define Γ as the space of *null-sum* oriented scalar quantities on halfedges: for every face t with halfedges $e_{1,t}, e_{2,t}, e_{3,t}$, and with signs $s_{1,t}, s_{2,t}, s_{3,t}$ that encode the orientation of the respective halfedges with regard to t (see Section 4.1), we consider corresponding scalar quantities $\gamma_{1,t}, \gamma_{2,t}, \gamma_{3,t}$ that must satisfy

$$s_{1,t}\gamma_{1,t} + s_{2,t}\gamma_{2,t} + s_{3,t}\gamma_{3,t} = 0. \tag{7}$$

We denote $\gamma = \{\gamma_{e,t}\} \in \Gamma$ as a *halfedge form.*

Equivalence to $\mathcal{X}$. We represent the halfedges as row vectors $e_{\cdot,t}$. With that, we define the projection operator $P' : \mathcal{X} \to \Gamma$ as follows:

$$P'_{|t} = \begin{pmatrix} e_{1,t} \\ e_{2,t} \\ e_{3,t} \end{pmatrix}. \tag{8}$$

Note that $P'_{|t}$ has zero row sum, which is the sum of edges of a single triangle oriented with the proper signs; its null space is spanned by vectors along the normal of the triangle. For each $v \in \mathcal{X}$, the null sum of $\gamma = P'v$ is trivially satisfied. The operator P' is analogous to the "♭" operator that converts a vector field to a 1-form.

Conversely, for every $\gamma \in \Gamma$, which has null sum by definition, the system $P'v = \gamma$ has a single solution that is also a tangent vector (without normal components)—it can be reproduced by the Penrose-Moore pseudo-inverse $v = P'^{-1}\gamma$ (the analogue to the "♯" operator). This creates a bijection between the spaces Γ and $\mathcal{X}$, and they are therefore isomorphic. We are not aware of this construction made explicitly in the literature to represent the PCVF space $\mathcal{X}$; a similar construction is alluded to in Poelke and Polthier [2016].

Packed and unpacked representations. To naturally encode a null sum of $\gamma \in \Gamma$, in each face we only store the first two γ values: $\gamma_{1,t}$ and $\gamma_{2,t}$. The choice is made without loss of generality—the choice of edges can be arbitrary, except e_2 should follow e_1 in the counterclockwise order of the face. To reproduce all three when needed, we define an *unpacking* operator U as follows:

$$U_{|t}\begin{pmatrix}\gamma_{1,t}\\ \gamma_{2,t}\end{pmatrix} = \begin{pmatrix}\gamma_{1,t}\\ \gamma_{2,t}\\ s_3\left(-s_{1,t}\gamma_{1,t} - s_{2,t}\gamma_{2,t}\right)\end{pmatrix}. \tag{9}$$

The packed representation "costs" 2 scalars per triangle, which is exactly the true dimension of $\mathcal{X}$. The effect of the packing operator U^{-1} (in pseudo-inverse) is to simply throw away $\gamma_{3,t}$ if the null-sum condition is met: $U^{-1}_{|t} \cdot U_{|t} = I_{2\times 2}$. We get that $U_{|t} \cdot U^{-1}_{|t}$ is a 3×3 matrix that filters away any non-null sum while changing the γ values if they violate it; we always avoid using it in this capacity.

With this representation, we reduce P' to the operator we use in practice, P, where its pseudo-inverse P^{-1} is an actual inverse, and both are defined as

$$\begin{aligned}P_{|t} &= \begin{pmatrix} e_{1,t} \\ e_{2,t} \end{pmatrix} \\ P^{-1}_{|t} &= \frac{s_1 s_2}{2A_t}\begin{pmatrix} -e^{\perp}_{2,t} \\ e^{\perp}_{1,t} \end{pmatrix}^T.\end{aligned} \tag{10}$$

We use the convention $e^{\perp}_2 = J_{|t}e_2$. P and P^{-1} aggregate the preceding per-face matrices into global operators. Note that $(e_1) \cdot \left(-e^{\perp}_2\right) = N_t \cdot (e_1 \times e_2) = 2s_1 s_2 A_t$. As such, we have $P \cdot P^{-1} = I_{2\times 2}$ and $P^{-1} \cdot P$ is a 3×3 matrix that projects out the normal component from an ambient vector field in $\mathbb{R}^3$. We avoid the normal-component filtering capacity in our formulations here as well and provide a proof that $P^{-1} \cdot P$ is an identity for tangent vector fields in Appendix A.

5.1 Halfedge Differential Operators

We next redefine all differential operators for Γ with the underlying paradigm that they should be equivalent to the operators in

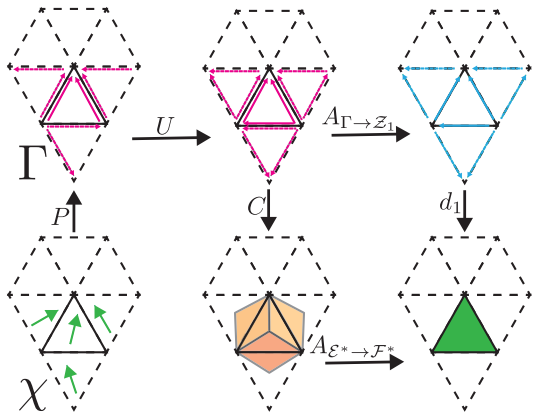

Fig. 5. The operators in the halfedge representation Γ.

$\mathcal{X}$, albeit formulated in Γ terms. We illustrate these operators in Figure 5 and provide the entire set of differential operators for $\mathcal{X}$ in the Γ setting in the rightmost columns of Table 2, comparing them to the analogous DEC and FEM operators.

Conforming gradient. Consider the assignment operator $A_{\mathcal{Z}_1\to\Gamma}$ that creates a halfedge form from a 1-form by copying the associated oriented scalar on an edge to its two halfedges. We then get

$$P^{-1}\cdot U^{-1}\cdot A_{\mathcal{Z}_1\to\Gamma}\cdot d_0 f = G_{\mathcal{V}} f. \tag{11}$$

The preceding relation demonstrates how DEC aligns with Γ where exact 1-forms, copied to halfedge forms, represent gradient fields—a fundamental parallel relation between DEC and $\mathcal{X}$. To avoid cumbersome notation, we denote $d_{0,\Gamma} = U^{-1}\cdot A_{\mathcal{Z}_1\to\Gamma}\cdot d_0$, which is the differential operator of dimensions $2|\mathcal{F}|\times|\mathcal{V}|$ in Γ space.

We extend the DEC d_1 to be the (oriented) sum operator $d_1 = \sum_{i=1}^{3} s_{i,t}\gamma_{i,t}$ (working similarly to DEC d_1, except with the halfedges of the face rather than 1-forms). To work with the packed form, we use $d_{1,\Gamma} = d_1\cdot U$. The null sum constraint is then encoded as the identity $d_{1,\Gamma}\cdot\gamma = 0$ for every $\gamma\in\Gamma$.

The transpose operator $(A_{\mathcal{Z}_1\to\Gamma})^T \equiv A_{\Gamma\to\mathcal{Z}_1}$ creates 1-forms from halfedge forms by summing up both halfedges scalars of each edge; we use it extensively in Section 5.3.

Curl. We consider again $\gamma_{ik,1}$ and $\gamma_{ik,2}$, the two halfedge forms restricted to the edge e_{ik} on the respective triangles t_1 and t_2. The curl operator $C:\Gamma\to\mathcal{E}^*$ is defined in Γ space as

$$C_{|ik} = \gamma_{ik,1} - \gamma_{ik,2}. \tag{12}$$

It is evident that curl-free fields in Γ (or the equivalent $\mathcal{X}$) are such that the halfedge forms are equal on both sides of the edge, which means they are isomorphic to 1-forms. As the null-sum constraint also dictates $d_{1,\Gamma}\gamma = 0$ by definition, we have have that a curl-free γ is isomorphic to a closed 1-form. However, a halfedge form that is not curl free is not compatible with any DEC quantity. Since we represent Γ with only two scalars per face, the complete definition for the curl operator is $C_\Gamma = C\cdot U$. Note that we have $C_\Gamma\cdot d_{0,\Gamma} = 0$, which preserves the discrete structure of $\mathcal{X}$.

5.2 Inner Product

The inner product between halfedge forms $\gamma_1,\gamma_2\in\Gamma$ is defined as

$$(P^{-1}\gamma_1)^T M_{\mathcal{X}}(P^{-1}\gamma_2) = \gamma_1^T(P^{-T}M_{\mathcal{X}}P^{-1})\gamma_2 = \gamma_1^T M_\Gamma\gamma_2. \tag{13}$$

$M_\Gamma: 2|\mathcal{F}|\times 2|\mathcal{F}|$ has the following simple structure:

$$M_{\Gamma|t} = \frac{1}{2}(U_{|t})^T\begin{pmatrix} cot(\alpha_1) & & \\ & cot(\alpha_2) & \\ & & cot(\alpha_3)\end{pmatrix}U_{|t}, \tag{14}$$

where α_j is the angle opposite edge j in face t and U is the unpacking operator as before. Simply put, we get a diagonal mass matrix for the unpacked null-summed $\gamma\in\Gamma$. We show the proof in Appendix B. Equipped with these basic operators, the divergence and Laplacian can be directly defined as in Table 2.

5.3 Mean-Curl Representation

Although the halfedge forms $\gamma\in\Gamma$ are equivalent to PCVFs in $\mathcal{X}$ through the projection operator P, we need an alternative and equivalent representation for them that reveals their differential properties, to be used in our subdivision schemes. Given the two halfedge forms $\gamma_{ik,1}$ and $\gamma_{ik,2}$ on both sides of edge ik adjacent to triangles t_1 and t_2 in our usual notation, we define the following:

$$\begin{aligned} z_{1|ik} &= \frac{\gamma_{ik,1}+\gamma_{ik,2}}{2} \Rightarrow z_1 = \frac{1}{2}A_{\Gamma\to\mathcal{Z}_1}\cdot U\cdot\gamma \\ \epsilon_{|ik} &= \frac{\gamma_{ik,1}-\gamma_{ik,2}}{2} \Rightarrow \epsilon = \frac{1}{2}C_\Gamma\cdot\gamma. \end{aligned} \tag{15}$$

In words, z_1 is the DEC 1-form that is the mean of the two halfedge forms and ϵ is half of the FEM curl of γ. This representation is trivially equivalent to that of the unpacked γ. We denote the conversion operator as W, defined as:

$$\begin{aligned} \begin{pmatrix} z_1 \\ \epsilon\end{pmatrix} &= W\cdot\gamma = \frac{1}{2}\begin{pmatrix} A_{\Gamma\to\mathcal{Z}_1} \\ C\end{pmatrix}U\cdot\gamma, \\ W^{-1} &= U^{-1}\begin{pmatrix} A^T_{\Gamma\to\mathcal{Z}_1} & C^T\end{pmatrix}. \end{aligned} \tag{16}$$

Note that $z_1\in\mathcal{Z}_1$ is a signed oriented quantity, whereas $\epsilon\in\mathcal{E}^*$ is an unsigned integrated quantity. We emphasize that the null-sum constraint $d_{1,\Gamma}\cdot\gamma = 0$ does not imply that z_1 is curl free in the DEC sense. In other words, we do not have $d_1 z_1 = 0$ in general.

Null sum constraint in mean curl. The mean-curl representation is not trivially equivalent to the packed Γ we use, since it has values for all edges, whereas γ is spanned by two halfedges within each triangle (hence the use of U^{-1} in W^{-1}). To get equivalence, we need to formulate the Γ null-sum requirement with (z_1,ϵ). This formulation has a surprisingly elegant form. Consider the face $t = ijk$ and the signs s for the coincident halfedge forms γ. Then,

$$\begin{aligned} d_{1,\Gamma|t}\cdot\gamma_{|t} = 0 = s_{ij}\gamma_{ij} + s_{jk}\gamma_{jk} + s_{ki}\gamma_{ki} = \\ s_{ij}z_{ij} + s_{jk}z_{jk} + s_{ki}z_{ki} - \epsilon_{ij} - \epsilon_{jk} - \epsilon_{ki} = \\ d_{1|t}z_{1|t} - A_{\mathcal{E}^*\to\mathcal{F}|t}\epsilon_{|t}, \end{aligned} \tag{17}$$

where $d_{1|t}$ is the DEC d_1 operator restricted to f and $A_{\mathcal{E}^*\to\mathcal{F}^*}$ is the summation operator $\epsilon_{ij}+\epsilon_{jk}+\epsilon_{ki}$ (analogous to $A_{\Gamma\to\mathcal{Z}_1}$). In global notation, the null-sum constraint reads as follows:

$$d_1 z_1 - A_{\mathcal{E}^*\to\mathcal{F}^*}\cdot\epsilon = 0. \tag{18}$$

Note again that when ϵ is 0, z_1 is a closed 1-form and we get the DEC identity $d_1 z_1 = 0$. More generally, as the DEC definition of

curl (see Table 2) is exactly $d_1 z_1$, the DEC face-based curl of the mean 1-form z_1 is then nothing but the face-summed edge-based FEM curl of the underlying field γ. We are not aware of this connection between DEC curl and FEM curl pointed out before.

The mean-curl representation reveals important ties between DEC and FEM more clearly:

- γ is FEM-exact if and only if z_1 is DEC-exact with the same function $f \in \mathcal{V}$ so that $d_0 f = z_1$, and where $\epsilon = 0$.
- γ is FEM-harmonic if and only if z_1 is DEC-harmonic. This is straightforward to see, as the DEC divergence operators $d_0^T M_1$ and D_Γ identify when $\epsilon = 0$.
- FEM-coexact γ does not correspond to coexact z_1; this is evident by the incompatible dimensions of the spaces. However, suppose that $\epsilon \in \mathcal{E}^*$ is the curl of γ; then, we have in this case a simple expression for the divergence of z_1: $d_0^T M_1 z_1 = D_\Gamma C^T \epsilon$.

Discussion: Refinable Hodge decomposition. Given the insights of the mean-curl representation, there is a subtle, yet important, distinction between the way DEC and FEM treat the Hodge decomposition, which we need to make to properly define the subdivision for PCVFs in $\mathcal{X}$. The DEC Hodge decomposition factors a 1-form $z_1 \in \mathcal{Z}_1$ into pointwise $z_0 \in \mathcal{V}$, harmonic part z_h, and *integrated* $z_2 \in \mathcal{F}^*$ (the equivalent of $\mathcal{Z}_2$). They further rely on refinable function spaces to perform subdivision (Section 4.3). For this, using integrated $\mathcal{F}^*$ is the correct choice, since $\mathcal{Z}_2$ admits a natural refinable hierarchy by triangle quadrisection. The pointwise dual 2-forms do not admit a refinable structure in this manner, and subdividing them directly would constitute as a "variational crime."

However, the FEM Hodge decomposition classically uses the pointwise elements in $\mathcal{E}$ to span its coexact part, which is, similarly to the dual 2-form space $\mathcal{Z}_2^*$, not a refinable space. Nevertheless, the Hodge decomposition can be defined in Γ analogously to DEC by using $f \in \mathcal{V}$, (half) curl $\epsilon \in \mathcal{E}^*$, and harmonic $h \in H_\Gamma$ as follows:

$$\forall \gamma \in \Gamma,\ \exists f \in \mathcal{V},\ \epsilon \in \mathcal{E}^*,\ h \in H_\Gamma :\ \gamma = d_{0,\Gamma} f + 2 M_\Gamma^{-1} C^T L_{\mathcal{E}}^{-1} \epsilon + h.$$

Other than just for revealing algebraic relations between FEM and DEC, we use the halfedge representation, mostly in its mean-curl representation, to establish PCVF subdivision schemes.

6 SUBDIVISION VECTOR FIELDS

Our purpose in constructing subdivision schemes for halfedge forms is the ability to work with PCVFs in a multiresolution structure-preserving manner. Specifically, we work with subdivided vector fields on very fine subdivided meshes, which are restricted to low-dimensional fields defined with the coarse control cage, for purposes of efficiency and robustness. We define $S_\mathcal{V}$ as the Loop subdivision matrix for vertex-based quantities and $S_{\mathcal{F}^*} = S_2$ as the half box spline face-based subdivision matrix, equivalent to S_0 and S_2 in the SEC scheme. For halfedge-based subdivision, we construct three distinct and interrelated operators for each subdivision level l:

- S_1^l, for 1-forms, of dimensions $|V^{l+1}| \times |V^l|$.
- $S_{\mathcal{E}^*}^l$, for unsigned integrated edge-based quantities (like curl), of dimensions $|E^{l+1}| \times |E^l|$.

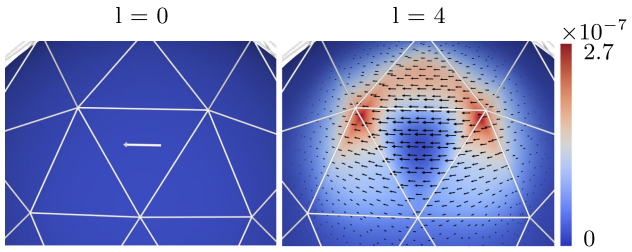

Fig. 6. A refined ($l = 4$) basis function from a single coarse ($l = 0$) unit-length vector. The color coding on the fine level depicts the per-face Hodge energy component $|C_\Gamma \gamma|^2 + |D_\Gamma \gamma|^2$ of the field, averaged to the vertices for visualization purposes. The glyph arrows on the fine level visualize directions and relative magnitudes.

- S_Γ^l, for halfedge forms composed of both. It is then of dimensions $2|F^{l+1}| \times 2|F^l|$.

S_1 is defined as in SEC (except our boundary modifications; see auxiliary material), so we need to define the latter two. For clarity, we often omit the level indicator l, as the operators are stationary, and the level can be understood from the context. We provide the full set of stencils in the auxiliary material.

To define structure-preserving operators on Γ, we require that S_Γ and $S_{\mathcal{E}^*}$ obey the following commutation rules:

$$\begin{aligned} d_{0,\Gamma} \cdot S_\mathcal{V} &= S_\Gamma \cdot d_{0,\Gamma} \\ C_\Gamma \cdot S_\Gamma &= S_{\mathcal{E}^*} \cdot C_\Gamma. \end{aligned} \tag{19}$$

In words, subdivided halfedge forms that represent gradient fields should result in gradient fields of the subdivided vertex-based scalar function, and the curl of a subdivided vector field should be equal to the subdivided curl of a vector field. To satisfy these conditions, our subdivision matrix for halfedge forms is defined directly on the mean-curl representation as follows:

$$S_\Gamma \cdot \gamma = W^{-1} \begin{pmatrix} S_1 & 0 \\ 0 & S_{\mathcal{E}^*} \end{pmatrix} \begin{pmatrix} z_1 \\ \epsilon \end{pmatrix} = W^{-1} \begin{pmatrix} S_1 & 0 \\ 0 & S_{\mathcal{E}^*} \end{pmatrix} W \cdot \gamma, \tag{20}$$

Since S_Γ is defined with the mean-curl representation that is in unpacked form, we need to verify that the null-sum requirement for the subdivided field $(z_1^{l+1}, \epsilon^{l+1})$ is satisfied before the application of W^{-1}, or otherwise W^{-1} will project the result unto the null-summed space Γ and the requirements in Equation (19) will not result in the promised structure-preserving properties. In other words, we require the following (as per Equation (18)):

$$d_1 z_1^{l+1} - A_{\mathcal{E}^* \to \mathcal{F}^*} \epsilon^{l+1} = 0.$$

As we inherit (albeit with some slight modifications) S_1 and $S_{\mathcal{F}^*}$ from SEC, our degrees of freedom for the requirements are in the definition of $S_{\mathcal{E}^*}$. To satisfy all requirements, we design it to adhere to the following additional commutation relation:

$$S_{\mathcal{F}^*} A_{\mathcal{E}^* \to \mathcal{F}^*} = A_{\mathcal{E}^* \to \mathcal{F}^*} S_{\mathcal{E}^*}. \tag{21}$$

In words, the face-based average of the subdivided curl should be equal to the subdivided face-based average of the coarse curl. This commutation elegantly preserves the null-sum requirement, as for level l, with mean z_1^l and half-curl ϵ^l, we get

$$\begin{aligned}
&\textit{(Level } l \textit{ null-sum constraint}\\
&\textit{(Equation(18)))} \quad d_1 z_1^l - A_{\mathcal{E}^* \to \mathcal{F}^*} \epsilon^l = 0 \Rightarrow\\
&\textit{(Subdivision)} \quad S_{\mathcal{F}^*} d_1 z_1^l - S_{\mathcal{F}^*} A_{\mathcal{E}^* \to \mathcal{F}^*} \epsilon^l\\
&\qquad = 0 \Rightarrow\\
&\textit{(Commutation)} \quad d_1 S_1 z_1^l - A_{\mathcal{E}^* \to \mathcal{F}^*} S_{\mathcal{E}^*} \epsilon^l\\
&\qquad = 0 \Rightarrow\\
&\textit{(Level } l+1 \textit{ null-sum constraint)} \quad d_1 z_1^{l+1} - A_{\mathcal{E}^* \to \mathcal{F}^*} \epsilon^{l+1} = 0.
\end{aligned} \tag{22}$$

Having secured the null-sum constraint, we can safely use W^{-1} to get the fine-level field γ^{l+1}, where all promised differential properties are guaranteed. We show an example of a basis function of the subdivision operator in Figure 6 and some examples of full subdivision vector fields in Figure 7.

6.1 Boundary Behavior

Our concepts of halfedges and the differential operators do not trivially extend to meshes with boundaries. Recall that our reasoning for subdivision is to commute with the gradient and the curl operators. However, the discrete curl operator on the boundary is not well defined for a single edge: consider a boundary face $t = ijk$ with boundary edge e_{ij} and the associated halfedge form $\gamma_{|ij}$. As studied in Poelke and Polthier [2016], the Hodge decomposition for meshes with boundaries admits several valid choices for decomposition, culminating in either Dirichlet or Neumann boundary conditions. We choose to assume that a function $f \in \mathcal{V}$ is defined everywhere, including the boundary, and that we commute with its gradient. Consequently, we assume that the boundary curl is zero by definition. In other words, on the boundary, we define $z_{1|ij} = \gamma_{ij}$ and $\epsilon_{|ij} = 0$. We adapt W and W^{-1} accordingly, noting that $z_{1|ij}$ is the only contribution to the field for boundary edge ij. Our subdivision matrices are designed to reflect this choice of decomposition, where $S_{\mathcal{E}^*}$ reproduces zero curl on the boundary, and S_1 is redefined to preserve the null-sum with this constrained $S_{\mathcal{E}^*}$. We show an illustration of boundary vector field basis functions on the boundary in Figure 8.

6.2 SHM Differential Operators

Following the reasoning of Section 4.3, we restrict M_Γ from a fine mesh back to a coarse mesh as follows:

$$\mathbb{M}_\Gamma^0 = \left(\mathbb{S}_\Gamma^l\right)^T \cdot M_\Gamma^l \cdot \mathbb{S}_\Gamma^l. \tag{23}$$

By this process of mass-matrix restriction, we process fine-level PCVFs that are spanned by the low-dimensional subdivided coarse-level PCVFs, directly on the coarse mesh. In analogy to SEC, we denote this technique as SHM.

By the commutation relations, the subdivided SHM curl of a field is equal to the fine curl, and when a field is SHM-exact on the coarse mesh, then it is also FEM-exact on the fine mesh,

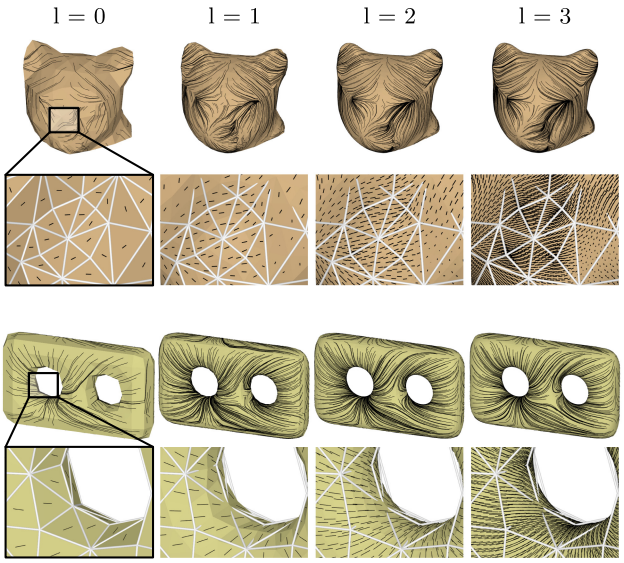

Fig. 7. Multiple levels of subdivision vector fields on the cathead (top, genus 0) and bitorus (bottom, genus 2) models. Note that the subdivsion preserves the features (sources, sinks, and vortices) of the fields.

where the fine function is the subdivision of the coarse one. Nevertheless, the SHM divergence behaves differently from the fine FEM divergence, as

$$\begin{aligned}
&\mathbb{D}_\Gamma^0 \gamma^0 = d_{0,\Gamma}^T \mathbb{M}_\Gamma^0 \gamma^0 =\\
&d_{0,\Gamma}^T \left(\mathbb{S}_\Gamma^l\right)^T \cdot M_\Gamma^l \cdot \mathbb{S}_\Gamma^l \gamma^0 =\\
&\left(\mathbb{S}_\mathcal{V}^l\right)^T d_{0,\Gamma}^T \cdot M_\Gamma^l \gamma^l = \left(\mathbb{S}_\mathcal{V}^l\right)^T D^l \gamma^l.
\end{aligned} \tag{24}$$

Note that we use $\mathbb{D}_\Gamma$ to denote the SHM divergence operator in line with other notation. In words, the divergence of a subdivided field is equal to the divergence of the resulting coarse field only through the restriction $(\mathbb{S}_\mathcal{V}^l)^T$. That essentially means that the divergence of the fine field might have "high frequency" components in $\ker(\mathbb{S}_\mathcal{V}^l)^T$ (Figure 9). This is an analogous phenomenon to the divergence pollution of SEC (Section 4.3). Note that the structure of SHM is preserved notwithstanding: SHM-exact fields are SHM curl free, and SHM-coexact fields are SHM divergence free.

The restricted mass matrices $\mathbb{M}$ are not diagonal anymore due to the two-ring support of any $\mathbb{S}$. Additionally, some operators are defined with inverse mass matrices, which are dense and nonlocal. In practice, we almost never need to compute the exact inverse, and we show how to circumvent this problem in the relevant applications.

Hodge decomposition. In Figure 9, we show a Hodge decomposition of a procedurally generated field with the SHM operators, subdivided to a fine-level $l = 3$. It is evident that the exact part subdivides as defined, but also that there is high-frequency divergence that pollutes the co-exact and the harmonic parts.

Hodge spectrum. The spectrum of the PCVF Hodge Laplacian $L_\mathcal{X}$ is studied in Brandt et al. [2017], where they show that the spectrum of $L_\mathcal{X}$ comprises harmonic fields (in its null space), gradients of eigenfunctions of $L_\mathcal{V}$, and cogradients of eigenfunctions

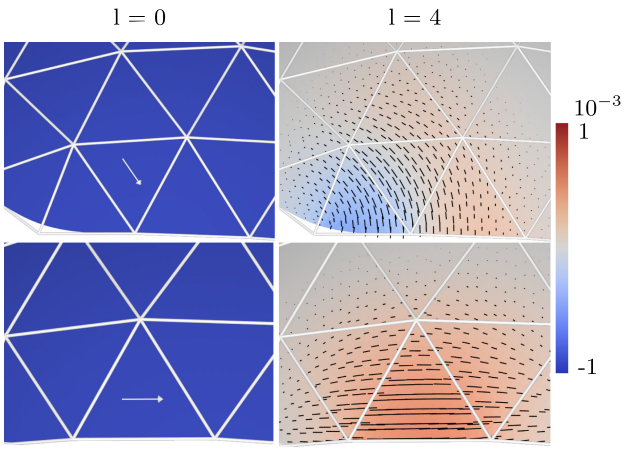

Fig. 8. Basis functions for coarse and fine levels near the boundary. The initial vector is of unit length. The color coding depicts the local face averaged curl $A_{\mathcal{E}^* \to \mathcal{F}^*} C \cdot \gamma$.

of $L_{\mathcal{E}}$. Using the SHM mass matrices, these relations still hold for the SHM Hodge Laplacian $\mathbb{L}_\Gamma$:

$$\begin{aligned} \forall \phi \in \mathcal{V},\ \lambda \in \mathbb{R},\ s.t.\ \mathbb{L}_{\mathcal{V}} \phi = \lambda \mathbb{M}_{\mathcal{V}} \phi \\ \Rightarrow \mathbb{L}_\Gamma \cdot d_{0,\Gamma} \cdot \phi = \lambda \mathbb{M}_\Gamma \cdot d_{0,\Gamma} \cdot \phi. \\ \forall \psi \in \mathcal{E},\ \mu \in \mathbb{R},\ s.t.\ \mathbb{L}_{\mathcal{E}} \psi = \mu \mathbb{M}_{\mathcal{E}} \psi \\ \Rightarrow \mathbb{L}_\Gamma \cdot \mathbb{M}_\Gamma^{-1} C_\Gamma^T \cdot \psi = \mu C_\Gamma^T \cdot \psi. \end{aligned} \quad (25)$$

Note that a term of $\mathbb{M}_\Gamma \cdot \mathbb{M}_\Gamma^{-1}$ was simplified in the right-hand side of the last equation. We used subdivision level $l = 3$ and computed the SHM Hodge eigenfunctions for several eigenvalues. We compare them against the ground-truth fine eigenfunctions in Figure 10. In Figures 11 and 12, we further analyze the relative error between the fine spectrum and the FEM and SHM spectra for the Hodge Laplacian. As can be seen, the SHM spectrum is a much better approximation of the fine Hodge spectrum than the coarse FEM one, for more than half of the full spectrum.

Errors and convergence. To study the behavior of our PCVF subdivision, we look at the behavior of the SHM Hodge Laplacian for the vector equation:

$$\mathbb{L}_\Gamma \cdot \gamma = b,$$

where $b \in \Gamma$ is some given field and $\mathbb{L}_\Gamma$ is the SHM Hodge Laplacian. We conduct two error and convergence tests as follows.

Projection error. We measure the error that is obtained by approximating the fine-level FEM with the low-dimensional SHM. For this, we choose the right-hand b^0 procedurally on some coarse mesh (level 0) and subdivide it several times to get b^l, where we consider the solution $\overline{\gamma}^l$ to the Hodge Laplacian system with this right hand as the ground-truth reference. For each level $0 \le k < l$, we solve for $\mathbb{L}_\Gamma^k \gamma^k = b^k$, where $\mathbb{L}_\Gamma^k$ is the SHM Hodge Laplacian at level k restricted from level l. We then subdivide γ^k to get $\gamma^{k \to l}$, and measure the L_2 and L_∞ error against the ground-truth solution $\overline{\gamma}^l$. For reference, we compare to a regular FEM solution at level k, computed as $L_\Gamma^k \gamma'^k = b^k$, also subdivided to level l and measured against the ground-truth solution. We show the results in Figure 13 and analyze convergence rates later in Table 3. It is

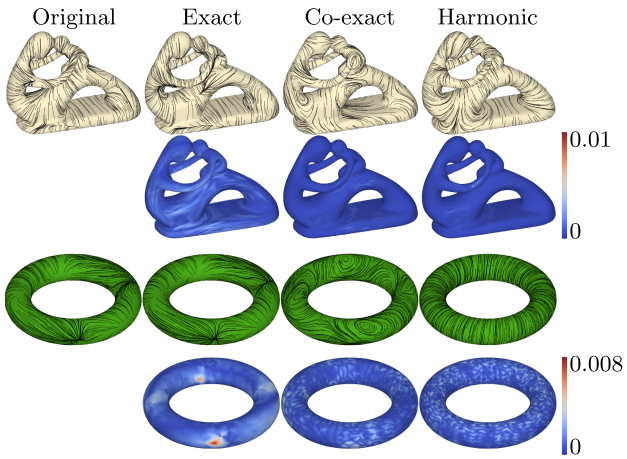

Fig. 9. SHM Hodge decomposition on models with genus 6 (first row) and 1 (third row), where a field is decomposed with SHM operators, subdivided, and shown in streamlines. The second and fourth row show the absolute value of the fine-level divergence, from which the high-frequency divergence pollution in the co-exact and harmonic part is evident. The fields are given by $\vec{v}_t(x, y, z) = \{\sin(\pi x)y, \sin(\pi xy)/r^2, \cos(\pi z) + x^2 + y^2\}$, where x, y, z are the coordinates of the face barycenter and $r^2 = x^2 + y^2 + z^2$. We add a random harmonic field from the null space of the Hodge Laplacian to the original field, and it is reproduced in the decomposition.

evident that the SHM solution has superior performance in terms of error when compared against the regular FEM solution, almost consistently with one to two order of magnitudes less error. Interestingly enough, the convergence rates are similar.

Operator error. We measure the error that is obtained on the coarse level $l = 0$ operator, by restricting the SHM operators only from a level $k < l$, rather than from the fine level l on which we wish to work. For instance, regular FEM operators are used when $k = 0$ and the full SHM when $k = l$. We show the result in Table 4. As is evident, the operator error diminishes quickly in the very coarse levels, but then it plateaus to a reasonable error. This suggests that a good approximation for processing on level l can be accomplished with a fairly low SHM level k; that can be explained by the rapid convergence of subdivision schemes [Dahmen 1986].

7 SUBDIVISION N-DIRECTIONAL FIELDS

We next extend our subdivision operators to N-directional subdivision, with the same structure-preserving guarantees. We do so by applying the local reduction of such fields into single-vector fields on branched cover spaces, which are introduced in Roy et al. [2018].

We work with N-directional fields that are elements of $\mathcal{X}^N$: in every face t, there are N indexed vectors $\{v_{t,1}, \ldots, v_{t,N}\}$, not necessarily symmetrically ordered. We assume that the field is equipped with a *matching*: a map between the vectors on a face t_1 to those in an adjacent face t_2, associated with the dual edge e between them. Furthermore, we assume the matching is (index) *order preserving*: the matching is parameterized by a per-edge index I_e, where a vector of index k on face t_1 is matched to vector of index $k + I_e$ (modulo N) on face t_2 (Figure 14). We denote the full matching as I_E.

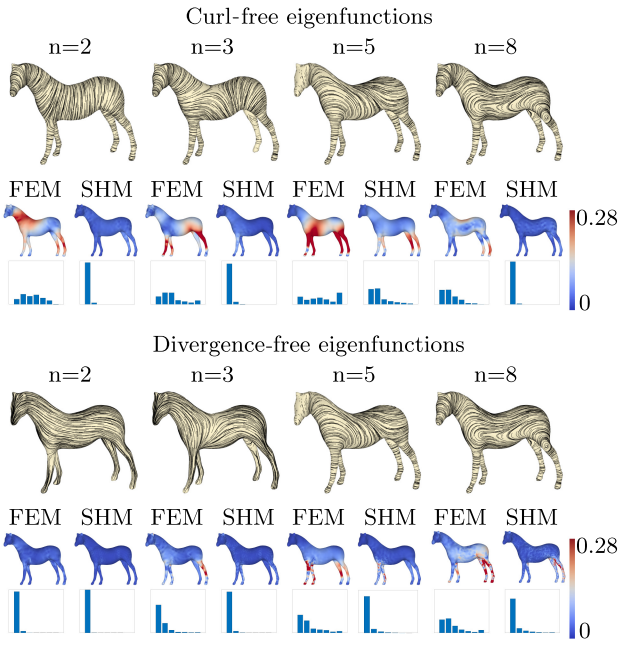

Fig. 10. Exact (top) and co-exact (bottom) eigenfunctions of the SHM ($l = 3$) Hodge Laplacian, subdivided and visualized at the fine level. The color coding denotes the norm difference $||P^{-1}\gamma_t' - P^{-1}\gamma_t^l||_2$ per triangle t, where γ_t' is the eigenfunction subdivided from the coarse-level (FEM and SHM) eigenfunction to the fine level (with normalization so that $\gamma_t'^T M_\Gamma^l \gamma_t' = 1$), and γ_t^l is the ground-truth eigenfunction of the fine-level Hodge Laplacian. The color scale depicts this pointwise error.

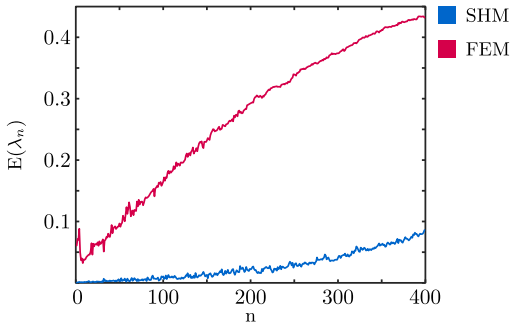

Fig. 11. Comparing eigenvalues between the fine-level conforming part of the Hodge Laplacian and the SHM or coarse-FEM eigenvalues, where the error is calculated for SHM as $E(\lambda_n) = |\lambda_{n,SHM} - \lambda_{n,\text{fine}}|/|\lambda_{n,\text{fine}}|$, and similarly for FEM. The same mesh as Figure 10 is used, with $|V| = 752$ at $l = 0$ and $|V| = 48{,}002$ at $l = 3$. We note that a concrete proof for this relation appears in Shoham et al. [2019].

The indices of the vertices are defined as $I_V = \frac{1}{N} d_0^T \cdot I_e$ [Crane et al. 2010], as d_0^T is the DEC boundary operator that encodes the dual cycle orientations around the vertex. A *regular* vertex v has $I_v = 0$, and otherwise it is called *singular*. The field on the 1-ring of a regular vertex can be *combed* (see Figure 14): it can be *locally* re-indexed in every face of the 1-ring such that $\forall e \in N(v)$, $I_e = 0$. With re-indexing, an N-field is locally reduced to N independent fields. A *fractional* singular vertex is defined by having $I_v \notin \mathbb{N}$,

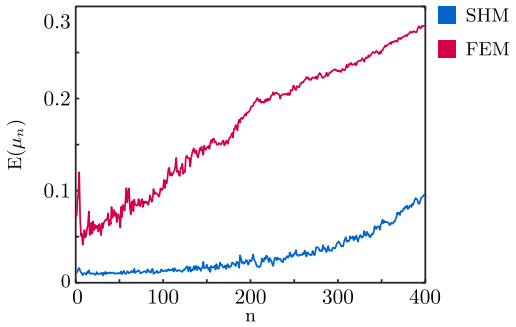

Fig. 12. Similar comparison as in Figure 11 for the nonconforming part of the Hodge Laplacian, where the error in SHM is $E(\mu_n) = |\mu_{n,SHM} - \mu_{n,\text{fine}}|/|\mu_{n,\text{fine}}|$, and similarly for coarse FEM. The same mesh has $|E| = 2{,}250$ at $l = 0$ and $|E| = 144{,}000$ at $l = 3$.

Table 3. Projection Error Convergence Rates for the Models in Figure 13

Error	Model		
	Star	Mannequin	Cone
SHM L_2	2.54	1.77	1.82
FEM L_2	2.00	1.91	2.00
SHM L_∞	1.90	0.978	1.10
FEM L_∞	1.80	1.06	1.07
SHM Curl L_2	1.95	1.13	1.59
FEM Curl L_2	1.87	1.24	1.54

The values correspond to the convergence factor β for the error hypothesis $a_0 h^{-\beta}$, where h is the mean edge length of the mesh. Fine levels l are at 6 for the Cone and Star experiment and 5 for the Mannequin experiment.

where such combing is not possible. Fields with fractional singularities cannot be *globally* combed. This is generally the case, as $\sum_{\forall v \in V} I_v = \chi(\mathcal{M})$, with $\chi(\mathcal{M})$ the Euler characteristic of the mesh. *Integral singularities* do not induce matching mismatches and therefore appear in single-vector fields as well, as sources, sinks, and vortices. They are basically sources of divergence and curl, and are irrelevant to our extension to N-directional fields.

7.1 Extending FEM Calculus

To be able to extend our subdivision scheme for N-directional fields, we need a concept of N-halfedge forms, N-scalar functions, and the entire suite of differential operators. For this, we next adapt existing notions from discrete calculus of branching coverings [Bommes et al. 2009; Diamanti et al. 2015; Kälberer et al. 2007]. Figure 15 presents an exemplification of the directional calculus presented here.

Seamless function spaces. Consider a vertex $v \in V$ with adjacent faces (in CCW order) $t_1, \ldots, t_d$ and associated corners $v_1, \ldots, v_d$. Further consider edges e_i between corners v_i and v_{i+1}. The function space $\mathcal{V}^N$ is parameterized by a vector $\mathbf{f}_{v_i}$ of N functions per corner i: $\mathbf{f}_{v_i} = (f_{v_1}, \ldots, f_{v_N})^T$. This amounts to $N \cdot d$ values for a single vertex (they are in fact spanned by a lower-dimensional parameter space, as we see in the following). The functions are matched across edges similarly to N-directional fields: consider two adjacent corners v_i and v_{i+1} across edge e_i with matching index I_{e_i}. We construct the permutation matrix $\pi(e_i)$ that represents

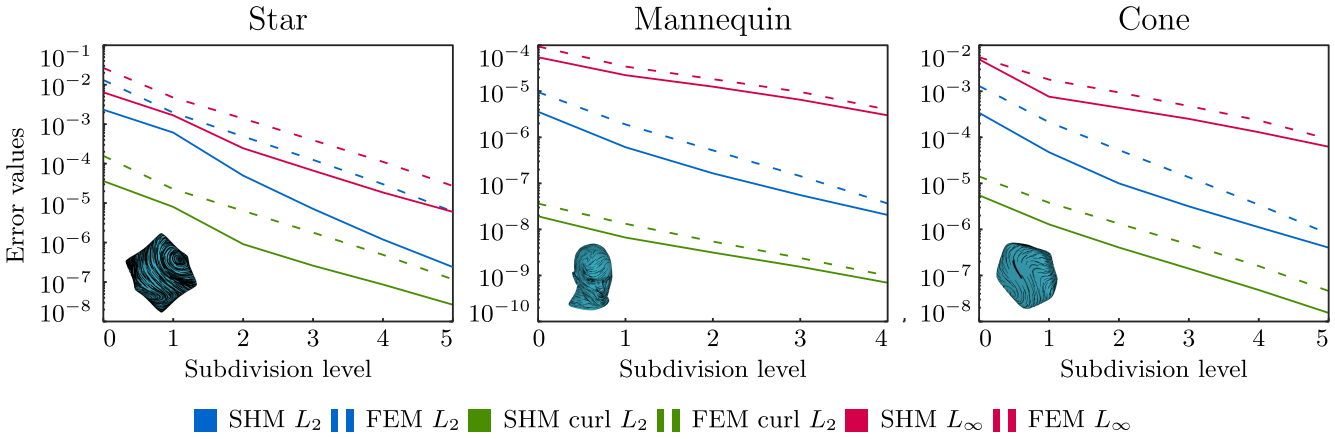

Fig. 13. L_2, L_∞, and the L_2 error of the curl, for the projection error of the vector Poisson equation on three models, as a function of subdivision level. We measure $L_\infty = \max |\delta\gamma|$ and $L_2 = \sqrt{(\delta\gamma)^T M_\Gamma \delta\gamma / \sum M_\Gamma}$, where $\delta\gamma = \bar{\gamma}^l - \gamma^l$, with $\bar{\gamma}^l$ as the ground-truth solution and γ^l as the subdivided solution. The curl L_2 error is measured as $\sqrt{(C\delta\gamma)^T M_{\mathcal{E}^*} C\delta\gamma / \sum M_{\mathcal{E}^*}}$. We use $l = 6$ for the Cone and Star models, and $l = 5$ for the Mannequin.

Table 4. Operator L_2 and L_∞ Errors for the Three Models

Level	Cone		Mannequin		Star	
	L_∞	L_2	L_∞	L_2	L_∞	L_2
0	10.8	7.22	0.229e-1	0.460e-4	2.02	0.618
1	4.88	0.326	0.145e-1	0.724e-5	0.578	0.496e-1
2	4.48	0.221	0.148e-1	0.631e-5	0.492	0.396e-1
3	4.46	0.222	0.148e-1	0.638e-5	0.468	0.409e-1
4	4.47	0.225	0.149e-1	0.644e-5	0.461	0.417e-1
5	4.47	0.227	—	—	0.459	0.420e-1

the map that the matching induces, to obtain the following:

$$\mathbf{f}_{v_{i+1}} = \pi(e_i) \cdot \mathbf{f}_{v_i}. \tag{26}$$

We always assume that within a single face, the corners have a trivial matching (so they are separate N functions); the only nontrivial matching is between corners across edges.

Combing. For regular vertices, and by successively applying Equation 26, we get that $\mathbf{f}_{v_1} = \pi(e_d) \cdot \mathbf{f}_{v_d}$. As such, we can *comb* the functions over regular vertices in the same way we do for directional fields: for a single 1-ring, we start from corner v_1 in face t_1 and transform every $\mathbf{f}_{v_i}$ into $\mathbf{f}_{v_1}$ by inverting Equation 26 recursively. We denote this linear transformation as $\Pi(v)$. Note that it means that there are only N independent functions in every regular vertex, parameterized by $\mathbf{f}_{v_1}$, which is expected.

Conforming operators. All conforming differential operators can be directly extended from the single-vector calculus around regular vertices, by conjugation with the combing (see Figure 15). For instance, we have that the divergence $D^N(v) : \mathcal{X}^N \to \mathcal{V}^{*N}$ is

$$D^N = \Pi^{-1}(v) \begin{pmatrix} d_0^T M_\mathcal{X} & & \\ & \dots & \\ & & d_0^T M_\mathcal{X} \end{pmatrix} \Pi(v). \tag{27}$$

In words, we comb a function and a field around a regular vertex, use the operators on every function in the vector f_{v_1} independently, and then comb back. The result is a vector of N scalars representing the independent divergences of the combed functions. Then, $\Pi^{-1}(v)$ combs the N scalars to corner-based values corresponding to original corner indexing. It is important to note that the identity of the "first" corner v_1 does not result in any loss of generality, due to the conjugation with $\Pi(v)$; the result per corner would be exactly the same regardless of which corner is first.

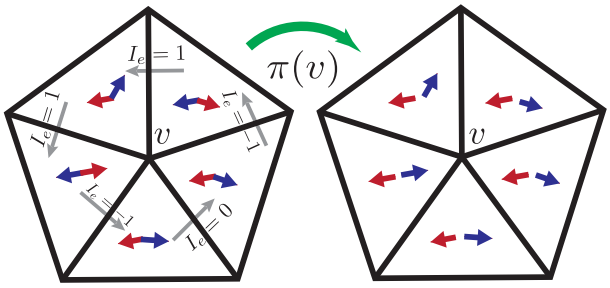

Fig. 14. Matching and combing. Left: A nonsingular matched 1-ring. Right: As the 1-ring is nonsingular, applying $\Pi(v)$ results in a combed separated field with a trivial matching.

The gradient operator G_V extends to $G^N_\mathcal{V} : \mathcal{V}^N \to \mathcal{X}^N$ by simply operating on the elements in the function values of the corners of the face independently, to produce N vectors. Therefore, it does not require combing; the corners of every single face are always trivially matched to each other.

Nonconforming operators. Nonconforming differential operators, namely the curl C^N, are easier to generalize: we only have to locally comb two faces sharing a single edge and then conjugate the curl operator independently for the N vectors in both faces with the combing operation. The result is a function in $\mathcal{E}^{*N}$. The rotated co-gradient $JG^N_\mathcal{E}$, exactly like $G^N_\mathcal{V}$, is defined per face and therefore does not require any matching or combing.

Structure-preserving calculus. It is easy to verify that directional-field calculus is structure preserving with relation to the sequence

around regular vertices. We have that $C^N \cdot G^N_{\mathcal{V}} = 0$, and that $D^N \cdot J \cdot G^N_{\mathcal{E}} = 0$ as well. The formal proof is straightforward, given the conjugation of combing and differential operators, and we omit it for brevity. Essentially, the existence of such sequences means that we can also define a directional Hodge decomposition, but we leave this line of research for future work.

Around singular vertices. For singular vertices, the product of $\pi(e)$ matrices leads to a nontrivial permutation matrix. In other words, "returning" to v_1 after applying Equation (26) successively, we get $\mathbf{f}_{v_1} \neq \pi(e_d) \cdot \mathbf{f}_{v_1}$. As such, conforming differential operators are not well defined for fractional singularities. To rationalize this, they can be interpreted as isolated boundary points in the field where there is not enough continuity by definition to allow for well-defined conforming operators. The nonconforming operators are well defined everywhere, as they only require two faces in every stencil.

7.2 Extending Γ^N

Calculus of halfedges is natural in the N-directional setting. We define $\gamma \in \Gamma^N$ as a vector of N scalars per halfedge. The operators $d^N_{0,\Gamma}$ and $d^N_{1,\Gamma}$ are trivially extended with respect to the matching of the corners. Note that we have a null-sum constraint for each element of γ independently. The same is done for per-face operators $P^N : \mathcal{X}^N \to \Gamma^N$ (and its inverse), unpacking operator U^N, and the summation operator $A_{\mathcal{E}^{*N} \to \mathcal{F}^{*N}}$.

The mean-curl representation and, consequently, the operator W^N are defined with the combing in the same manner as nonconforming differential operators like C^N: one of the halfedges in every edge is chosen arbitrarily as the "first," and then we define $A_{\Gamma^N \to \mathcal{Z}^N_1}$ to conjugate with the matching. As such, both the resulting mean z_1 and (half) curl ϵ are defined with relation to one of the halfedges, and this choice of "first halfedge" is well defined up to permutation.

7.3 Extending Subdivision Operators

Equipped with an extension of the Γ representation to Γ^N, we next extend our subdivision operators to work with directional fields and preserve their structure.

Branched Loop and half-box splines. For regular vertices, both the Loop $S_{\mathcal{V}}$ and the half-box spline $S^*_{\mathcal{F}}$ subdivision operators extend to the branched spaces $\mathcal{V}^N$ and $(\mathcal{F}^*)^N$ by conjugation with combing as well. For instance, for Loop subdivision we get the following:

$$(S_{\mathcal{V}})^N = \Pi^{-1}(v) \begin{pmatrix} S_{\mathcal{V}} & & \\ & \dots & \\ & & S_{\mathcal{V}} \end{pmatrix} \Pi(v). \tag{28}$$

The result creates new even and odd edges, where the permutation $\pi(e)$ for even edges is the same as the coarse edges they originate from, whereas $\pi(e)$ for odd edges is an identity, since they are created within coarse faces.

For singular vertices, we require a different definition of the subdivision operators. We do so by *unfolding the branch* (Figure 16): consider again a 1-ring with d faces, with singularity index $I_V = \frac{i}{N}$. We pick a single vector, follow its matching around the ring until we reach it again, and create a new ring just with this vector.

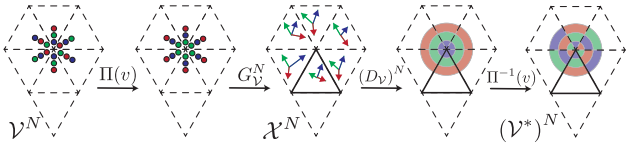

Fig. 15. Example of directional calculus with the N-Laplacian on a regular vertex. A corner-based vertex function is combed to a single vertex function, after which a gradient is applied to obtain a combed directional field. Applying the divergence then results in a combed integrated value on the vertex, which is combed back to the original labeling.

We then do so until all vectors are taken. That creates $\mathrm{GCD}(i, N)$ (greatest common divisor) new rings. We are always guaranteed to return to the original vector since $(\pi(v))^N = I$. We denote the unfolding operation as $\Phi(v)$. Then, we can conjugate $S_{\mathcal{V}}$ for singular vertices with the unfolding:

$$(S_{\mathcal{V}})^N = \Phi(v)^{-1} S_{\mathcal{V}} \Phi(v).$$

The unfolding $\Phi(v)$ is a generalization of the combing operator $\Pi(v)$ that allows us to extend all our subdivision operators without altering the original scalar subdivision stencils, as the commutation also works through the conjugation. For example, for a regular vertex, we just create N new rings each with the separated single vector field. As a result, we maintain all differential properties of the subdivision, and among them structure preserving of curl and exactness. We demonstrate this in Figures 17 and 18.

8 APPLICATIONS

In the following, we apply our SHM framework to several applications that use PCVF directional fields in their pipeline. We implemented the subdivision operators in C++ using the Directional library [Vaxman 2019] and the applications using MATLAB. Times are measured on a laptop with an Intel i7-7700HQ (2.8 GHz) CPU and 32 GB of RAM.

Vector field design. In Figure 19, we show an example of coarse-to-fine vector field design. Vectors are constrained on a small set of faces of a coarse mesh and interpolated to the rest of the mesh by minimizing the SHM (with level $l = 3$) Hodge energy

$$E_H(\gamma) = \mathbb{M}_{\Gamma^0} \left(|\mathbb{C}_\Gamma \gamma^0|^2 + |\mathbb{D}_\Gamma \gamma^0|^2 \right)$$

of a field γ^0 on the coarse mesh. This is done by solving $\mathbb{L}^3_\Gamma \gamma^0 = 0$ with fixed values for a subset of constrained faces. We then subdivide γ^0 to get γ^3 as our result. We get a fine smooth field efficiently designed with the coarse (restricted) degrees of freedom.

Earth mover's distance. We apply our subdivision to the optimal transport algorithm presented in Solomon et al. [2014]. For brevity, we do not consider meshes with boundary in this experiment. The formulation computes a geodesic vector field between two probability distributions $\mu_0, \mu_1 \in \mathcal{V}^*$ with $\sum_{v \in V} \mu_{0|v} = \sum_{v \in V} \mu_{1|v} = 1$, which are defined on the fine mesh of level l. These distributions are defined by densities $\rho_{0|1} = M^{-1}_{\mathcal{V}} \mu_{0|1} \in \mathcal{V}$. The geodesic field is computed to minimize (a simplification of) the 1-Wasserstein

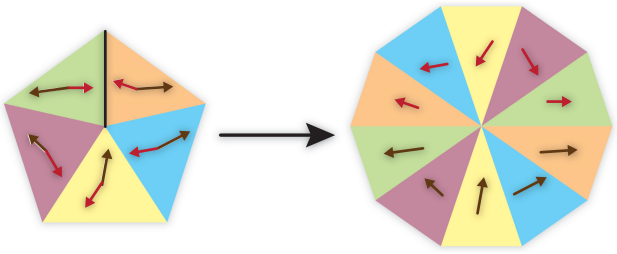

Fig. 16. The unfolding operator $\Phi(v)$, illustrated for a singular vertex in the space $\mathcal{X}^N$. We unfold a valence 5 vertex with singularity index $-\frac{1}{2}$ into a valence 10 ring with a single vector field. The vector field is then locally subdivided with S_Γ and then folded back.

distance $\zeta(\mu_0, \mu_1)$ between the probability measures as follows:

$$\zeta(\mu_0, \mu_1) = \inf_{g,h} \sum_{t \in F^l} A(t) |G_{\mathcal{V}|t} f + J G_{\mathcal{E}|t} g + h|_{L_2} \qquad (29)$$
$$s.t.\ L^l_{\mathcal{V}} f = M^l_{\mathcal{V}} (\rho_0 - \rho_1),$$

where $g \in \mathcal{E}^l$ and h is a harmonic field in H^l. $f \in \mathcal{V}^l$ is fully determined from the Laplacian constraint. To limit the solution space on a fine mesh, they use a spectral subspace for g from its Laplacian $L_\mathcal{E}$. We offer an alternative low-rank SHM approximation that uses coarse-mesh function values instead, which is more efficient due to the sparsity of the subdivision matrix. Here, we deviate from the multigrid V-cycle folding paradigm of SHM and solve the problem directly on the fine mesh. Nevertheless, we limit the solution space to subdivided coarse functions. To use the refinable conforming functions, we note that the underlying continuous norm is invariant to rotations. Therefore, we dualize the discretization of the problem: we consider mid-edge distributions $\rho'_0, \rho'_1 \in \mathcal{E}^*$, transform the problem to refinable $\gamma \in \Gamma$, and solve for

$$\zeta(\mu_0, \mu_1) = \inf_{f^0, h} \sum_{t \in F^k} \sqrt{\left(\gamma^l_{|t}\right)^T \cdot M^l_\Gamma \cdot \gamma^l_{|t}} \qquad (30)$$
$$s.t.\ \gamma^l = S_\Gamma (d_{0,\Gamma} f^0) + J G^l_\mathcal{E} g^l + S_\Gamma h^0$$
$$L^l_\mathcal{E} g^l = M^l_\mathcal{E} (\rho'_0 - \rho'_1).$$

In words, we solve for coarse f^0 so that its subdivided gradient γ^l creates the least-norm vector field with the Laplacian-computed coexact component g (we use a simply connected mesh with no harmonic component for simplicity). This is solved using the ADMM procedure described by Solomon et al. [2014]. Note that the coexact component is computed beforehand and therefore fixed after solving the Laplacian equation.

Our experiment is conducted as follows: we compute our SHM solution and compare the result to a spectral-subspace FEM solution with an increasing number of eigenbases. A similar accuracy (measured to the ground-truth solution in the fine level) is achieved with approximately 360 eigenvalues, at almost three times the computation time. We show our results in Figures 20 and 21.

Operator-based advection. Our framework can be used to modify the operator-based representation of PCVFs introduced in Azencot et al. [2013, 2015]. Their method constructs a discrete version of

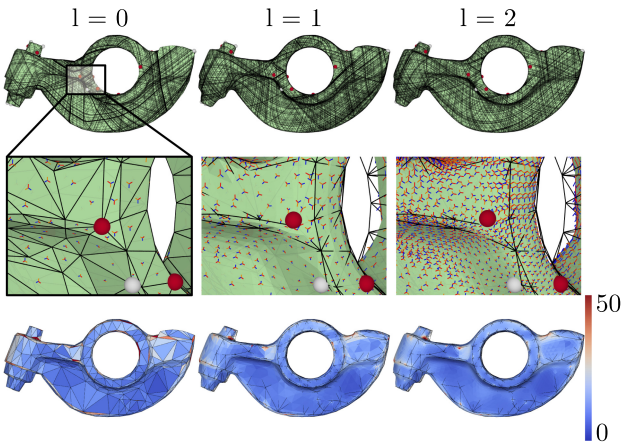

Fig. 17. A subdivided 3-directional field. The spheres mark the singularities (red = $-\frac{1}{3}$, white = $\frac{1}{3}$). Middle: Zooming in around a singularity. The color coding shows the vector norm of the curl per edge, averaged to faces and divided by face area. The curl is evidently refined under subdivision.

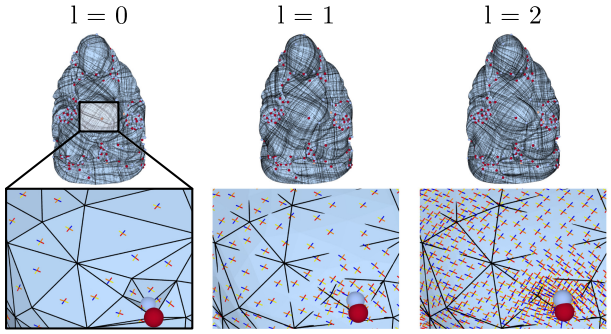

Fig. 18. An approximately curl-free 4-directional field subdivided. Bottom: zoom-in, with singularity colors red = $-\frac{1}{4}$, orange = $-\frac{2}{4}$, and white = $\frac{1}{4}$. The L_∞ norms of the curl per edge for the three levels are 5.05e-6, 9.50e-7, and 2.18e-7, respectively, which demonstrate that the (lack of) curl of the field is preserved under subdivision.

the classical representation of vector fields as derivations of scalar functions: for a field u and a scalar function f, the operator produces $\nabla_u f = \langle u, \nabla f \rangle$. Given a vector field $u \in \mathcal{X}$, their discrete operator is represented by a matrix $B_\mathcal{V} : |V| \times |V|$ on a mesh that is composed as follows:

$$B_\mathcal{V} = \frac{1}{3} (M_\mathcal{V})^{-1} A_{\mathcal{F} \to \mathcal{V}} M_\mathcal{X} B_\mathcal{F} G_\mathcal{V},$$

where $B_\mathcal{F} : |F| \times 3|F|$ is a matrix that performs the facewise dot-product of the face-based gradient with u, and $A_{\mathcal{F} \to \mathcal{V}}$ sums values from faces to adjacent vertices, in our usual notation. Essentially, the dot products are made per face and averaged to the vertices using the respective mass matrices of the mesh.

The operator representation makes it simple to advect a function $f \in \mathcal{V}$ on a surface: given time t, and the initial function value $f(0)$, they solve the advection equation in the weak sense, integrated over a spectrally reduced subspace. Consider Ψ as the matrix with n lowest Laplacian eigenvectors as columns. They work

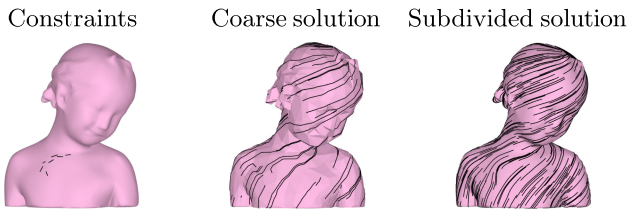

Fig. 19. SHM vector-field design. The local constraints (left) are interpolated to the rest of the coarse mesh (center) and subdivided to the fine mesh (right) at $l = 3$.

with a vector of coefficients α so that $f(t) = \Psi\alpha(t)$, and solve for

$$\Psi^T M_{\mathcal{V}} \cdot \frac{\partial\left(\Psi\alpha(t)\right)}{\partial t} = \Psi^T M_{\mathcal{V}} \cdot B_{\mathcal{V}} \Psi\alpha(t),$$

where $\Psi^T M_{\mathcal{V}} \Psi = Id$. We get $\alpha(t) = \exp(t \cdot B^{EIG}_{\mathcal{V}}) \cdot \alpha(0)$, where

$$B^{EIG}_{\mathcal{V}} = \Psi^T M_{\mathcal{V}} B_V \Psi.$$

We follow a similar construction, except we integrate over a subspace of refined subdivision basis functions, and our degrees of freedom are the coarse function f^0 so that $f(t) = S_{\mathcal{V}} f^0(t)$. Posing the system in the weak sense, we get the following:

$$S^T_{\mathcal{V}} M_{\mathcal{V}} \cdot \frac{\partial S_{\mathcal{V}} f^0(t)}{\partial t} = S^T_{\mathcal{V}} M_{\mathcal{V}} \cdot B_{\mathcal{V}} S_{\mathcal{V}} f^0(t),$$

where the solution is $f^0(t) = \exp(t \cdot B^{SHM}_{\mathcal{V}}) \cdot f^0(0)$, and

$$B^{SHM}_{\mathcal{V}} = \mathbb{M}^{-1}_{\mathcal{V}} S^T_{\mathcal{V}} M_{\mathcal{V}} B_{\mathcal{V}} S_{\mathcal{V}}.$$

The weak form is natural to the eigenfunction reduction, as the eigenfunctions are orthogonal w.r.t. $M_{\mathcal{V}}$, and $\psi^\dagger = \psi^T M_{\mathcal{V}}$. Nevertheless, we empirically witnessed that omitting $M_{\mathcal{V}}$ actually slightly improves the accuracy of SHM advection; we conjecture that this is since the subdivision bases are "more orthogonal" w.r.t. a uniform matrix but reserve the concrete analysis for future work.

We compute a ground-truth solution at the fine level $l = 5$, project it to the reduced basis, and compare the SHM advection (both uniform and weights with $M_{\mathcal{V}}$) against the spectral advection for a different number of eigenbases. We show the result in Figure 22 and the error in Figure 23. The spectral-subspace approximation has a comparable error profile between 50 and 100 eigens, but the computation is about 8 to 10 times as slow, where the eigenbasis extraction is the expensive part. Note that for both SHM and the spectral approximation, the error diverges with time, due to the high frequencies inevitably created by the advection equation.

Seamless parameterization. We next employ our structure-preserving subdivision for N-directional fields to compute coarse-to-fine curl-reduced fields. This allows us to compute fine-level rotationally seamless parameterizations (where the direction identifies across cuts but without perfect integer translations) with a very low integration error. We compute an N-RoSy (with Knöppel et al. [2013]) on the coarse mesh, optimize it to be (approximately) curl free with [Diamanti et al. 2015], and compute a coarse parameterization that consequently has a very small integration error. The subdivision preserves the small amount of curl, and the fine-level parameterization also has a small error as a result. We compare this process to performing the curl-free optimization on the fine mesh directly. Our coarse-to-fine optimization is faster in almost two orders of magnitude. We demonstrate this in Figure 24 and in the teaser (Figure 1).

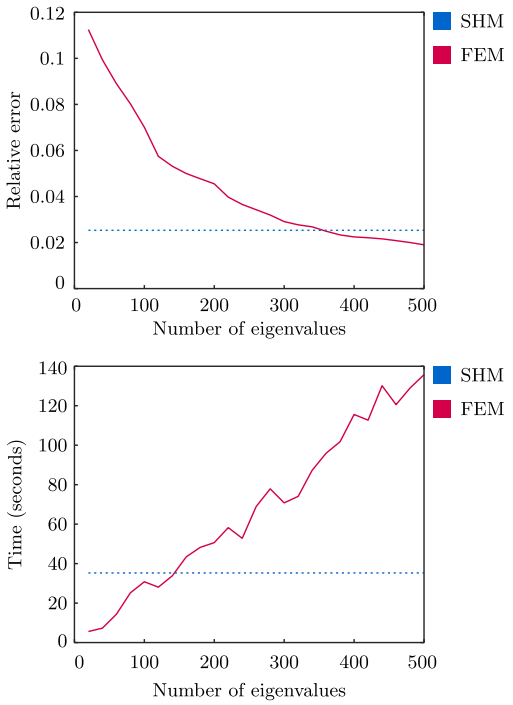

Fig. 20. Top: Relative error $|\zeta - \zeta_{GT}|/\zeta_{GT}$ of the computed distance value via SHM (horizontal dotted line, for reference) and the FEM spectral-subspace approximation with number of eigenvalues from 20 to 500 in steps of 20. Bottom: Total running time for SHM (horizontal dotted line, for reference) and the spectral approximation for the preceding specified number of eigenvalues.

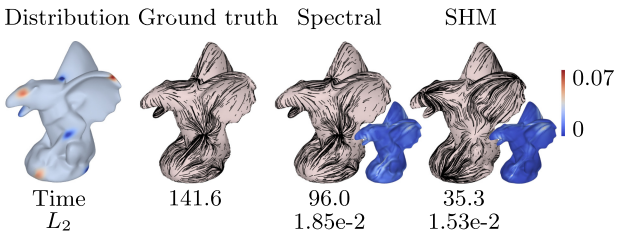

Fig. 21. A comparison of the EMD algorithm between the spectral (center-right) approximation with 360 eigenvalues and our SHM (right), with similar accuracy with relation to the ground truth (center-left). Left: The initial mass distribution. The big figures depict the geodesic fields γ^l resulting from the optimization, with insets depicting the vector norm per face of the difference between the fine-level solution and the chosen approximation. The total running times are given: for the spectral-subspace solution, 68.1 seconds were spent on the ADMM optimization, and 27.8 seconds on computing the basis and prefactoring. For SHM, this was 30.8 and 4.5 seconds, respectively. The L_2 error to the ground truth is provided.

9 DISCUSSION

Convergence and smoothness. As we discuss in the auxiliary material, our subdivision stencils for $S_{\mathcal{E}^*}$ and S_Γ have a few degrees

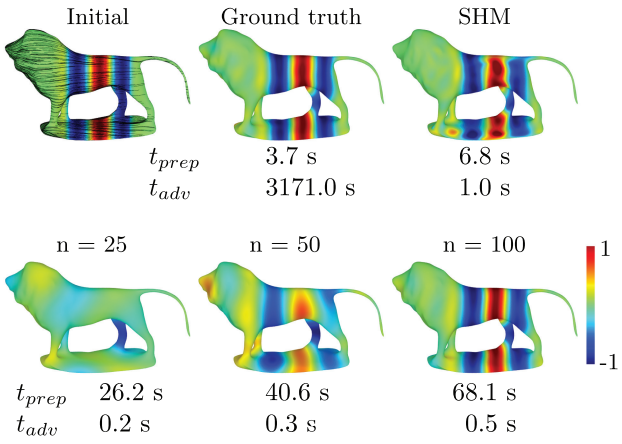

Fig. 22. Operator-based advection. The initial (color-coded) function f is projected to the reduced basis and advected with the given vector field v (streamlines), with time $t = 0.5$. The bottom row shows the approximation using the first n Laplacian eigenfunctions. We provide the preparation and advection times t_{prep} and t_{adv}. v is given by $v_t(x, y, z) = -x$ per triangle t, where x, y, z are the coordinates of the barycenter of t. f_0 (before projection) is given by $f_{0,v}(x, y, z) = \sin(\frac{4}{5}\pi x)$ for the x coordinate of vertex v. We localize the function below $x_{min} = \min(V_x) + \frac{1}{3}(\max(V_x) - \min(V_x))$ and above $x_{max} = \min(V_x) + \frac{2}{3}(\max(V_x) - \min(V_x))$ by multiplying $f_{0,v}$ with $\exp(-5|v_x - x_{min}|)$ below x_{min} and $\exp(-5|v_x - x_{max}|)$ above x_{max}. The model has $|V^0| = 435$ and $|V^5| = 449{,}532$.

of freedom (after counting the commutation constraints) that we use to optimize the spectrum of the subdivision stencil such that it converges in the limit since the subdominant eigenvalues are less than 1. We conjecture that since the fields are derivatives of smoothly subdivided functions, they are then one level of smoothness lower in the limit. However, we leave a formal theoretical analysis of convergence and smoothness to future work. We believe that a better design practice might be to allow $S_{\mathcal{V}}$ and $S_{\mathcal{F}^*}$ to vary entirely, where the smoothness of all subdivision operators is optimized concurrently (similarly to Huang and Schröder [2010]), rather than modify the existing schemes.

Dual formulation. Our Γ space uses $\langle v, e\rangle$ for the projection operator P. Nevertheless, the entire formulation can be made with the perpendicular $\langle v, e^{\perp}\rangle$, using the nonconforming divergence and conforming curl. This can be beneficial to fluid simulation.

Preconditioning and its disadvantages. The mass matrices of SHM are generally more strongly positive definite than those of the FEM in the coarse mesh. The reason is that the uniform and stationary subdivision operators average the mesh, and create better triangulations. Nevertheless, the fact that the subdivision does not commute with the fine mass matrix also creates the high-frequency divergence pollution in the subdivided fields. It is then worthwhile to try and explore alternatives that consider the mass matrices within the templates, to obtain precise fine Hodge decompositions.

Full multiresolution processing. Our article explores low-dimensional coarse-to-fine approximations. Moreover, the basis functions are not orthonormal as an eigenbasis, albeit considerably cheaper to obtain. Nevertheless, SHM can be augmented by

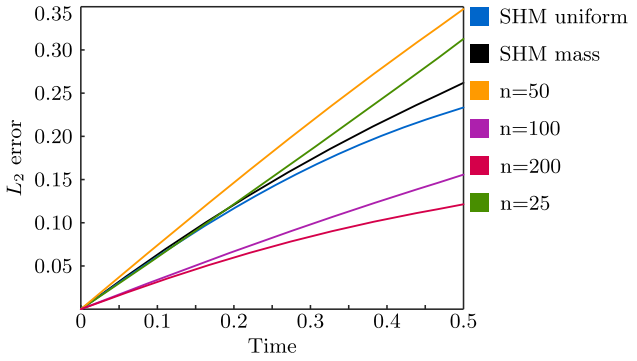

Fig. 23. Normalized L_2 error for the SHM (with and without mass matrix $M_{\mathcal{V}}$) and spectral approximations, depicted in Figure 22, compared to the ground truth for times $t = [0, 0.5]$, with step size of $\Delta t = \frac{1}{32}$. We define the ground truth $f(t)$ as the initial function, expanded in the respective basis and advected with the fine-level operator.

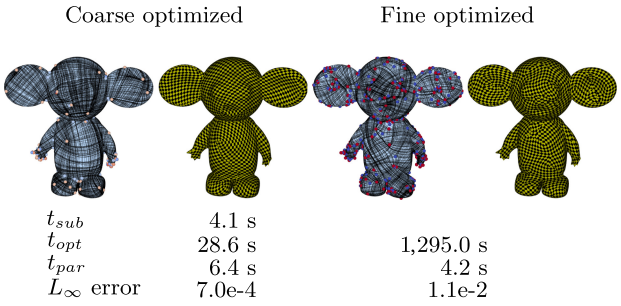

Fig. 24. Seamless parameterizations with coarse-to-fine curl-free fields. Left: Coarse curl optimization and subdivision, and fine-level parameterization. Right: Fine curl optimization and parameterization. The runtime is significantly reduced by optimizing in the coarse level, the integration error is lower, and the result is more appealing. The fine level is at $l = 2$. We use 250 iterations for the curl optimization. t_{sub} is the subdivision time, t_{opt} is the curl optimization time, and t_{par} is the parameterization time. The optimization time evidently dominates the total running time.

incorporating biorthogonal *subdivision wavelets* [Bertram 2004; Lounsbery et al. 1997], to obtain exact multiresolution representation of functions over the fine mesh, with the advantages of increasing locality—this could benefit applications such as solving diffusion problems.

Nontriangular meshes. The space $\mathcal{X}$ is not well defined for nonplanar polygonal meshes. Nevertheless, in the spirit of mimetic elements [Bossavit 1998], the space Γ, with its null-sum constraint, is well defined for any polygonal mesh, implicitly defining $\mathcal{X}$. As such, our framework can consider other subdivision operators (e.g., Catmull-Clark). We will explore this in future work.

General restriction operators. Finally, our setting is currently limited to subdivision surfaces. It could be beneficial to also allow for a multiresolution setting on general fine meshes using simplification operators (e.g., quadratic-error-based simplification [Garland and Heckbert 1997]) as the restriction operators. This should prove challenging as the vertex- and face-based restrictions have to be defined first but will allow a very general framework for directional-field processing on arbitrary triangle meshes.

APPENDIX

A P^{-1} AS INVERSE OF P

We next show that tangential vector fields are preserved under the operation $P^{-1} \cdot P$.

Let e_1, e_2, e_3 be the CCW oriented edges of a face in CCW order (i.e., $s_i = 1 \forall i \in \{1, 2, 3\}$), and let $\alpha_1, \alpha_2, \alpha_3$ be the angle opposite the corresponding edge. Here, we will interpret e_i as column vectors. Consider a tangential vector field, which is locally defined on a triangle t as $v_t = ae_1 + be_2$, without loss of generality. Then,

$$P_{|t} v_t = \begin{pmatrix} a|e_1|^2 + be_1 \cdot e_2 \\ ae_1 \cdot e_2 + b|e_2|^2 \end{pmatrix}. \tag{31}$$

Applying P^{-1}, we get

$$\begin{aligned} P_{|t}^{-1} P_{|t} v_t &= \frac{1}{2A} \begin{pmatrix} -e_2^{\perp,T} & e_1^{\perp,T} \end{pmatrix} \begin{pmatrix} a|e_1|^2 + be_1 \cdot e_2 \\ ae_1 \cdot e_2 + b|e_2|^2 \end{pmatrix} && (32) \\ &= \frac{1}{2A} \Big(-e_2^{\perp,T} \left(a|e_1|^2 + be_1 \cdot e_2 \right) && (33) \\ &\quad + e_1^{\perp,T} \left(b|e_2|^2 + ae_1 \cdot e_2 \right) \Big) && (34) \\ &= \frac{1}{2A} (X + Y), && (35) \end{aligned}$$

where we renamed the summed terms in the brackets for convenience. Note that $2A = |e_1||e_2| \sin \alpha_3$. We can express e_2 in terms of e_1 via orthogonal decomposition

$$e_2 = |e_2| \left(\cos(\pi - \alpha_3) \frac{e_1}{|e_1|} + \sin(\pi - \alpha_3) \frac{e_1^{\perp}}{|e_1|} \right), \tag{36}$$

which allows us to write

$$\begin{aligned} e_1^{\perp} &= \frac{|e_1|}{|e_2|} \frac{e_2}{\sin \alpha_3} + \cot \alpha_3 e_1 && (37) \\ &= \frac{|e_1|^2}{2A} e_2 + \cot \alpha_3 e_1 && (38) \\ &= (\cot \alpha_2 + \cot \alpha_3) e_2 + \cot \alpha_3 e_1, && (39) \end{aligned}$$

where we use $\frac{|e_1|^2}{2A} = \cot \alpha_2 + \cot \alpha_3$. Similarly,

$$\begin{aligned} -e_1 &= |e_1| \left(\cos \alpha_3 \frac{e_2}{|e_2|} + \sin \alpha_3 \frac{e_2^{\perp}}{|e_2|} \right) \iff && (40) \\ e_2^{\perp} &= -\frac{|e_2|}{|e_1|} \frac{e_1}{\sin \alpha_3} - \cot \alpha_3 e_2 && (41) \\ &= -(\cot \alpha_1 + \cot \alpha_3) e_1 - \cot \alpha_3 e_2. && (42) \end{aligned}$$

Working out $\frac{X}{2A}$, we get

$$\begin{aligned} \frac{X}{2A} &= ((\cot \alpha_1 + \cot \alpha_3) e_1 + \cot \alpha_3 e_2) \left(a \frac{|e_1|^2}{2A} + b \frac{e_1 \cdot e_2}{2A} \right). && (43) \\ &= ((\cot \alpha_1 + \cot \alpha_3) e_1 + \cot \alpha_3 e_2) \left(a(\cot \alpha_2 + \cot \alpha_3) \right. \\ &\quad \left. - b \cot \alpha_3 \right) && (44) \\ &= \cot^2 \alpha_3 (ae_1 + ae_2 - be_1 - be_2) + ae_1 && (45) \\ &\quad + a \cot \alpha_2 \cot \alpha_3 e_2 - b \cot \alpha_1 \cot \alpha_3 e_1, && (46) \end{aligned}$$

where we use $\frac{e_1 \cdot e_2}{2A} = -s_1 s_2 \cot \alpha_3$ and $\cot \alpha_1 \cot \alpha_2 + \cot \alpha_1 \cot \alpha_3 + \cot \alpha_2 \cot \alpha_3 = 1$ for the interior angles of a triangle. For $\frac{Y}{2A}$, we get

$$\begin{aligned} \frac{Y}{2A} &= ((\cot \alpha_2 + \cot \alpha_3) e_2 + \cot \alpha_3 e_1) \\ &\quad \times (-a \cot \alpha_3 + b(\cot \alpha_1 + \cot \alpha_3)) && (47) \\ &= \cot^2 \alpha_3 (-ae_2 - ae_1 + be_2 + be_1) + be_2 && (48) \\ &\quad -a \cot \alpha_2 \cot \alpha_3 e_2 + b \cot \alpha_1 \cot \alpha_3 e_1 && (49) \end{aligned}$$

and combining gives

$$\frac{1}{2A}(X + Y) = ae_1 + be_2 \tag{50}$$

as desired.

To generalize this result for arbitrarily oriented edges, we need to multiply $-e_2^{\perp}$ and $e_1^{\perp}$ by their appropriate sign, and sign the γ values to be correctly oriented. This amounts to

$$P_{|t}^{-1} = \frac{1}{2A} \begin{pmatrix} -s_2 e_2^{\perp,T} & s_1 e_1^{\perp,T} \end{pmatrix} \begin{pmatrix} s_1 & 0 \\ 0 & s_2 \end{pmatrix} = \frac{s_1 s_2}{2A} \begin{pmatrix} -e_2^{\perp} \\ e_1^{\perp} \end{pmatrix} \tag{51}$$

as stated before.

B INNER PRODUCT ON Γ

In the following, we develop the inner product mass matrix M_Γ to prove the formulation of Equation (14).

Consider a face $t = 123$ with three edges e_1, e_2, e_3 that are, without loss of generality, positively oriented toward the face. Further consider two halfedge forms $\gamma_x, \gamma_y \in \Gamma$ restricted to the face on these edges, $\gamma_{x|(1,2,3)}$ and $\gamma_{y|(1,2,3)}$, representing respective face-based vectors $v_{x|t}$ and $v_{y|t}$. We "pack" their representation to edges 1 and 2 alone, and by so trivially encoding the null-sum constraint $\gamma_{x|1} + \gamma_{x|2} + \gamma_{x|3} = 0$ (and respectively for γ_y). Following Equation (13), we have that the inner product M_Γ, restricted to the face, is given by

$$M_\Gamma = P^{-T} M_\chi P^{-1}.$$

This reproduces the inner product between v_y and v_x in the face t. We then get that

$$\begin{aligned} M_\Gamma = P^{-T} M_\chi P^{-1} &= \frac{1}{4A_t^2} \begin{pmatrix} -e_2^{\perp} \\ e_1^{\perp} \end{pmatrix} \begin{pmatrix} A_t & & \\ & A_t & \\ & & A_t \end{pmatrix} \begin{pmatrix} -e_2^{\perp} \\ e_1^{\perp} \end{pmatrix}^T \\ &= \frac{1}{4A_t} \begin{pmatrix} e_2^{\perp} \cdot e_2^{\perp} & -e_2^{\perp} \cdot e_1^{\perp} \\ -e_2^{\perp} \cdot e_1^{\perp} & e_1^{\perp} \cdot e_1^{\perp} \end{pmatrix} \end{aligned}$$

Consider the angles $\alpha_{1|2|3}$ opposite to edges $e_{1|2|3}$. Then, we use the identities

$$\begin{aligned} \frac{e_1^{\perp} \cdot e_1^{\perp}}{2A_t} &= cot\,(\alpha_2) + cot\,(\alpha_3) \\ \frac{-e_1^{\perp} \cdot e_2^{\perp}}{2A_t} &= cot\,(\alpha_3) \end{aligned}$$

for any cyclic shift of $(1, 2, 3)$. Then, we get

$$M_\Gamma = \frac{1}{2}\begin{pmatrix} cot\,(\alpha_1) + cot\,(\alpha_3) & -cot\,(\alpha_3) \\ -cot\,(\alpha_3) & cot\,(\alpha_2) + cot\,(\alpha_3) \end{pmatrix}$$
$$= \frac{1}{2}U^T\begin{pmatrix} cot(\alpha_1) & & \\ & cot(\alpha_2) & \\ & & cot(\alpha_3) \end{pmatrix}U,$$

where we use the unpacking operator $U = \begin{pmatrix} 1 & 0 \\ 0 & 1 \\ -1 & -1 \end{pmatrix}$.

ACKNOWLEDGMENTS

The authors would like to thank Bettina Speckmann for her support, and furthermore Fernando de Goes, Nilima Nigam, Mirela Ben-Chen, Justin Solomon, and Etienne Vouga for helpful discussions.

REFERENCES

Omri Azencot, Mirela Ben-Chen, Frédéric Chazal, and Maks Ovsjanikov. 2013. An operator approach to tangent vector field processing.*Computer Graphics Forum* 32, 5 (2013), 73–82.

Omri Azencot, Maks Ovsjanikov, Frédéric Chazal, and Mirela Ben-Chen. 2015. Discrete derivatives of vector fields on surfaces—An operator approach. *ACM Transactions on Graphics* 34, 3 (May 2015), Article 29.

Ivo Babuška and Manil Suri. 1994. The P and H-P versions of the finite element method, basic principles and properties. *SIAM Review* 36, 4 (Dec. 1994), 578–632. DOI: https://doi.org/10.1137/1036141

Mirela Ben-Chen, Adrian Butscher, Justin Solomon, and Leonidas Guibas. 2010. On discrete killing vector fields and patterns on surfaces. *Computer Graphics Forum* 29, 5 (2010), 1701–1711.

M. Bertram. 2004. Biorthogonal loop-subdivision wavelets. *Computing* 72, 1–2 (April 2004), 29–39. DOI: https://doi.org/10.1007/s00607-003-0044-0

Henning Biermann, Adi Levin, and Denis Zorin. 2000. Piecewise smooth subdivision surfaces with normal control. In *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH'00)*. ACM, New York, NY, 113–120. DOI: https://doi.org/10.1145/344779.344841

David Bommes, Henrik Zimmer, and Leif Kobbelt. 2009. Mixed-integer quadrangulation. *ACM Transactions on Graphics* 28, 3 (2009), Article 77.

Alain Bossavit. 1998. Introduction: Maxwell equations. In *Computational Electromagnetism*, A. Bossavit (Ed.). Academic Press, San Diego, CA, 1–30. DOI: https://doi.org/10.1016/B978-012118710-1/50002-7

Achi Brandt. 1977. Multi-level adaptive solutions to boundary-value problems. *Mathematics of Computation* 31, 138 (1977), 333–390.

Christopher Brandt, Leonardo Scandolo, Elmar Eisemann, and Klaus Hildebrandt. 2017. Spectral processing of tangential vector fields. *Computer Graphics Forum* 35, 6 (2016), 338–353.

Marcel Campen, David Bommes, and Leif Kobbelt. 2015. Quantized global parametrization. *ACM Transactions on Graphics* 34, 6 (2015), Article 192.

Fehmi Cirak, Michael J. Scott, Erik K. Antonsson, Michael Ortiz, and Peter Schröder. 2002. Integrated modeling, finite-element analysis, and engineering design for thin-shell structures using subdivision. *Computer-Aided Design* 34, 2 (2002), 137–148.

Keenan Crane, Fernando de Goes, Mathieu Desbrun, and Peter Schröder. 2013. Digital geometry processing with discrete exterior calculus. In *Proceedings of ACM SIGGRAPH 2013 Courses (SIGGRAPH'13)*.

Keenan Crane, Mathieu Desbrun, and Peter Schröder. 2010. Trivial connections on discrete surfaces. *Computer Graphics Forum* 29, 5 (2010), 1525–1533.

Michel Crouzeix and Pierre-Arnaud Raviart. 1973. Conforming and nonconforming finite element methods for solving the stationary Stokes equations I. *ESAIM: Mathematical Modelling and Numerical Analysis—Modélisation Mathématique et Analyse Numérique* 7, R3 (1973), 33–75. http://eudml.org/doc/193250.

Wolfgang Dahmen. 1986. Subdivision algorithms converge quadratically. *Journal of Computational and Applied Mathematics* 16, 2 (1986), 145–158. DOI: https://doi.org/10.1016/0377-0427(86)90088-9

Mark de Berg, Otfried Cheong, Marc van Kreveld, and Mark Overmars. 2008. *Computational Geometry: Algorithms and Applications* (3rd ed.). Springer-Verlag TELOS, Santa Clara, CA.

Fernando de Goes, Mathieu Desbrun, and Yiying Tong. 2016a. Vector field processing on triangle meshes. In *Proceedings of ACM SIGGRAPH 2016 Courses (SIGGRAPH'16)*. ACM, New York, NY, Article 27, 49 pages. DOI: https://doi.org/10.1145/2897826.2927303

Fernando de Goes, Mathieu Desbrun, Mark Meyer, and Tony DeRose. 2016b. Subdivision exterior calculus for geometry processing. *ACM Transactions on Graphics* 35, 4 (July 2016), Article 133, 11 pages. DOI: https://doi.org/10.1145/2897824.2925880

Fernando de Goes, Beibei Liu, Max Budninskiy, Yiying Tong, and Mathieu Desbrun. 2014. Discrete 2-tensor fields on triangulations. *Computer Graphics Forum* 33, 5 (2014), 1–12.

M. Desbrun, A. N. Hirani, M. Leok, and J. E. Marsden. 2005. Discrete exterior calculus. arXiv:math/0508341.

Olga Diamanti, Amir Vaxman, Daniele Panozzo, and Olga Sorkine-Hornung. 2014. Designing N-PolyVector fields with complex polynomials. *Computer Graphics Forum* 33, 5 (2014), 1–11.

Olga Diamanti, Amir Vaxman, Daniele Panozzo, and Olga Sorkine-Hornung. 2015. Integrable PolyVector fields. *ACM Transactions on Graphics* 34, 4 (2015), Article 38.

Michael Garland and Paul S. Heckbert. 1997. Surface simplification using quadric error metrics. In *Proceedings of of the 24th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH'97)*. 209–216.

Eitan Grinspun, Petr Krysl, and Peter Schröder. 2002. CHARMS: A simple framework for adaptive simulation. *ACM Transactions on Graphics* 21, 3 (July 2002), 281–290. DOI: https://doi.org/10.1145/566654.566578

Aaron Hertzmann and Denis Zorin. 2000. Illustrating smooth surfaces. In *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH'00)*. 517–526.

Anil N. Hirani. 2003. *Discrete Exterior Calculus*. Ph.D. Dissertation. California Institute of Technology.

Jinghao Huang and Peter Schröder. 2010. sqrt(3)-based 1-form subdivision. In *Curves and Surfaces: Proceedings of the 7th International Conference, Revised Selected Papers*. 351–368. DOI: https://doi.org/10.1007/978-3-642-27413-8_22

T. J. R. Hughes, J. A. Cottrell, and Y. Bazilevs. 2005. Isogeometric analysis: CAD, finite elements, NURBS, exact geometry and mesh refinement. *Computer Methods in Applied Mechanics and Engineering* 194, 39 (2005), 4135–4195. DOI: https://doi.org/10.1016/j.cma.2004.10.008

Bert Jüttler, Angelos Mantzaflaris, Ricardo Perl, and Martin Rumpf. 2016. On numerical integration in isogeometric subdivision methods for PDEs on surfaces. *Computer Methods in Applied Mechanics and Engineering* 302 (2016), 131–146. DOI: https://doi.org/10.1016/j.cma.2016.01.005

Felix Kälberer, Matthias Nieser, and Konrad Polthier. 2007. QuadCover—Surface parameterization using branched coverings. *Computer Graphics Forum* 26, 3 (2007), 375–384.

Felix Knöppel, Keenan Crane, Ulrich Pinkall, and Peter Schröder. 2013. Globally optimal direction fields. *ACM Transactions on Graphics* 32, 4 (2013), Article 59.

Bei-Bei Liu, Yiying Tong, Fernando de Goes, and Mathieu Desbrun. 2016. Discrete connection and covariant derivative for vector field analysis and design. *ACM Transactions on Graphics* 35, 3 (2016), Article 23.

Songrun Liu, Alec Jacobson, and Yotam Gingold. 2014. Skinning cubic Bézier splines and Catmull-Clark subdivision surfaces. *ACM Transactions on Graphics* 33, 6 (Nov. 2014), Article 190, 9 pages. DOI: https://doi.org/10.1145/2661229.2661270

Yang Liu, Helmut Pottmann, Johannes Wallner, Yong-Liang Yang, and Wenping Wang. 2006. Geometric modeling with conical meshes and developable surfaces. *ACM Transactions on Graphics* 25, 3 (2006), 1–9.

Charles Teorell Loop. 1987. *Smooth Subdivision Surfaces Based on Triangles*. Master's Thesis. Department of Mathematics, University of Utah.

Michael Lounsbery, Tony D. DeRose, and Joe Warren. 1997. Multiresolution analysis for surfaces of arbitrary topological type. *ACM Transactions on Graphics* 16, 1 (Jan. 1997), 34–73. DOI: https://doi.org/10.1145/237748.237750

P. Mullen, A. McKenzie, D. Pavlov, L. Durant, Y. Tong, E. Kanso, J. E. Marsden, and M. Desbrun. 2011. Discrete Lie advection of differential forms. *Foundations of Computational Mathematics* 11, 2 (2011).

Ashish Myles and Denis Zorin. 2012. Global parametrization by incremental flattening. *ACM Transactions on Graphics* 31, 4 (2012), 131–149.

Thien Nguyen, Keçstutis Karčiauskas, and Jörg Peters. 2014. A comparative study of several classical, discrete differential and isogeometric methods for solving Poisson's equation on the disk. *Axioms* 3, 2 (2014), 280–299. DOI: https://doi.org/10.3390/axioms3020280

Konstantin Poelke and Konrad Polthier. 2016. Boundary-aware Hodge decompositions for piecewise constant vector fields. *Computer-Aided Design* 78, C (Sept. 2016), 126–136. DOI: https://doi.org/10.1016/j.cad.2016.05.004

Konrad Polthier and Eike Preuß. 2003. Identifying vector field singularities using a discrete hodge decomposition. In *Visualization and Mathematics III*. Mathematics and Visualization. Springer, 113–134.

Hartmut Prautzsch, Wolfgang Boehm, and Marco Paluszny. 2002. *Bezier and B-Spline Techniques*. Springer-Verlag, Berlin, Germany.

Nicolas Ray, Bruno Vallet, Wan Chiu Li, and Bruno Lévy. 2008. N-symmetry direction field design. *ACM Transactions on Graphics* 27, 2 (2008), Article 10.

Lawrence Roy, Prashant Kumar, Sanaz Golbabaei, Yue Zhang, and Eugene Zhang. 2018. Interactive design and visualization of branched covering spaces. *IEEE*

Transactions on Visualization and Computer Graphics 24, 1 (Jan. 2018), 843–852. DOI: https://doi.org/10.1109/TVCG.2017.2744038

Andrew Sageman-Furnas, Albert Chern, Mirela Ben-Chen, and Amir Vaxman. 2019. Chebyshev nets from commuting PolyVector fields. *ACM Transactions on Graphics* 38, 6 (Nov. 2019), Article 172.

Marc Schober and Manfred Kasper. 2007. Comparison of hp-adaptive methods in finite element electromagnetic wave propagation. *COMPEL: The International Journal for Computation and Mathematics in Electrical and Electronic Engineering* 26, 2 (2007), 431–446. DOI: https://doi.org/10.1108/03321640710727782

Nicholas Sharp, Yousuf Soliman, and Keenan Crane. 2019. The vector heat method. *ACM Transactions on Graphics* 38, 3 (2019), Article 24.

Meged Shoham, Amir Vaxman, and Mirela Ben-Chen. 2019. Hierarchical functional maps between subdivision surfaces. *Computer Graphics Forum* 38, 5 (2019), 55–73.

Justin Solomon, Raif Rustamov, Leonidas Guibas, and Adrian Butscher. 2014. Earth mover's distances on discrete surfaces. *ACM Transactions on Graphics* 33, 4 (July 2014), Article 67, 12 pages. DOI: https://doi.org/10.1145/2601097.2601175

Jos Stam. 2003. Flows on surfaces of arbitrary topology. *ACM Transactions on Graphics* 22, 3 (July 2003), 724–731. DOI: https://doi.org/10.1145/882262.882338

Gilbert Strang and George Fix. 2008. *An Analysis of the Finite Element Method.* Wellesley-Cambridge Press.

Bernhard Thomaszewski, Markus Wacker, and Wolfgang Strasser. 2006. A consistent bending model for cloth simulation with corotational subdivision finite elements. In *Proceedings of the 2006 ACM SIGGRAPH/Eurographics Symposium on Computer Animation (SCA'06).* 107–116. http://dl.acm.org/citation.cfm?id=1218064.1218079.

Yiying Tong, Santiago Lombeyda, Anil N. Hirani, and Mathieu Desbrun. 2003. Discrete multiscale vector field decomposition. *ACM Transactions on Graphics* 22, 3 (2003), 445–452.

Amir Vaxman2019. Directional: A Library for Directional Field Synthesis, Design, and Processing. Retrieved January 6, 2020 from https://github.com/avaxman/Directional.

Amir Vaxman, Marcel Campen, Olga Diamanti, Daniele Panozzo, David Bommes, Klaus Hildebrandt, and Mirela Ben-Chen. 2016. Directional field synthesis, design, and processing. *Computer Graphics Forum* 35, 2 (2016), 545–572. DOI: https://doi.org/10.1111/cgf.12864

Ke Wang, Weiwei, Yiying Tong, Mathieu Desbrun, and Peter Schroder. 2006. Edge subdivision schemes and the construction of smooth vector fields. *ACM Transactions on Graphics* 25, 3 (2006), 1–8.

Max Wardetzky. 2006. *Discrete Differential Operators on Polyhedral Surfaces–Convergence and Approximation.* Ph.D. Dissertation. Freie Universität Berlin.

Mirko Zadravec, Alexander Schiftner, and Johannes Wallner. 2010. Designing quad-dominant meshes with planar faces. *Computer Graphics Forum* 29, 5 (2010), 1671–1679.

Eugene Zhang, Konstantin Mischaikow, and Greg Turk. 2006. Vector field design on surfaces. *ACM Transactions on Graphics* 25, 4 (2006), 1–33.

Received October 2018; revised October 2019; accepted November 2019