3D Hodge Decompositions of Edge- and Face-based Vector Fields

RUNDONG ZHAO, Michigan State University

MATHIEU DESBRUN, California Institute of Technology

GUO-WEI WEI and YIYING TONG, Michigan State University

Fig. 1. **Five-Component Vector Field Decomposition.** On a tetrahedral mesh of the kitten with a spherical cavity, a vector field is decomposed into a gradient field with zero potential on the boundary, a curl field with its vector potential orthogonal to the boundary, a pair of tangential and normal harmonic fields, and a harmonic field that is both a gradient and a curl field. Potential fields are shown in the corners of their corresponding components.

We present a compendium of Hodge decompositions of vector fields on tetrahedral meshes embedded in the 3D Euclidean space. After describing the foundations of the Hodge decomposition in the continuous setting, we describe how to implement a five-component orthogonal decomposition that generically splits, for a variety of boundary conditions, any given discrete vector field expressed as discrete differential forms into two potential fields, as well as three additional harmonic components that arise from the topology or boundary of the domain. The resulting decomposition is proper and mimetic, in the sense that the theoretical dualities on the kernel spaces of vector Laplacians valid in the continuous case (including correspondences to cohomology and homology groups) are exactly preserved in the discrete realm. Such a decomposition only involves simple linear algebra with symmetric matrices, and can thus serve as a basic computational tool for vector field analysis in graphics, electromagnetics, fluid dynamics and elasticity.

CCS Concepts: • **Mathematics of computing** → **Discretization**.

Additional Key Words and Phrases: Vector field decomposition, boundary conditions, discrete exterior calculus, homology, cohomology.

ACM Reference Format:

Rundong Zhao, Mathieu Desbrun, Guo-Wei Wei, and Yiying Tong. 2019. 3D Hodge Decompositions of Edge- and Face-based Vector Fields. *ACM Trans. Graph.* 38, 6, Article 181 (November 2019), 13 pages. <https://doi.org/10.1145/3355089.3356546>

1 INTRODUCTION

The existence of orthogonal decompositions of a given vector field into gradient and curl terms (that can be integrated into potentials) along with non-integrable parts (that are due to the topology of the domain) is a fundamental property leveraged in a variety of

static and dynamical problems — for instance, fluid simulation to enforce incompressibility. The mathematical foundations behind such decompositions were developed using the theory of differential forms for any finite-dimensional compact manifold without boundary early on [Hodge 1941], but were fully extended to manifolds with boundaries much more recently [Shonkwiler 2009].

In computer graphics, the analysis and processing of vector fields over *surfaces* have received plenty of attention in recent years. Consequently, the resulting computational tools needed to achieve a Hodge decomposition have been well documented and tested on various applications; see, e.g., recent surveys on surface vector field analysis [Vaxman et al. 2016; de Goes et al. 2016a]. For the case of vector fields over 3D bounded domains, discussions about the Hodge decomposition are significantly scarcer: while the usefulness of the Hodge decomposition is as prevalent as in 2D, the existing literature lacks a rigorous computational treatment of the full-blown decomposition over 3D domains of arbitrary topology. Our paper fills this void by offering both the theoretical foundations and a practical linear-algebra based implementation of a five-term Hodge decomposition of vector fields expressed as discrete forms for the most common boundary conditions used in computational science.

1.1 Related work

A variety of books present detailed expositions of the Hodge decomposition from a mathematical perspective (see [Abraham et al. 1988; Schwarz 1995] for two examples using a formulation based on differential forms), but provide no hints on computational approaches to implementing a discrete decomposition in the case of finite-dimensional vector field representations. Even more applied treatments (such as [Cantarella et al. 2002] which discusses the case of complicated topology at length) are often based on a Biot-Savart construction relying on volume integrals to prove the existence and uniqueness of the decomposition, but leave the computational aspects to realize such a decomposition mostly unaddressed. For the simpler case of a two-component decomposition (known as the Helmholtz decomposition), a number of papers describe how

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2019 Association for Computing Machinery.

0730-0301/2019/11-ART181 \$15.00

<https://doi.org/10.1145/3355089.3356546>

Fig. 2. **Helmholtz decomposition.** Vector field in a vase with a spherical cavity decomposed into a gradient and a curl field, but with a nonzero L_2 inner product between these two resulting components.

to compute the scalar and potential vector potentials [Amrouche et al. 1998], but no mention is made of the validity of the implied discretization of the cohomology and whether its dimensionality matches the continuous case based on the discrete choices of divergence, curl and gradient operators. Yet, the numerical issues generated by a failure to capture the proper cohomologies are well documented by now – see, e.g., the spurious (i.e., aphysical) modes in computational electromagnetism [Caorsi et al. 2001], or the typical checkerboard patterns in Poisson solves.

The most common use in computer graphics of a 3D vector field decomposition is arguably in incompressible fluid simulation; however, not all components are needed in this context since a simple pressure projection is typically used to remove all divergence [Stam 1999; Colin et al. 2006]. Most approaches in fluid dynamics discuss the case of ball-like topology, with the exception of methods using vorticity to reconstruct the velocity field, e.g., [Elcott et al. 2007; Chern 2017], which contain discussions on the treatment of domains with nonzero genus. Even in these cases, the decomposition is not comprehensive due to the specificities of typical boundary conditions in fluid animation. Finally, [Bhatia et al. 2013] provides a thorough survey of recent progress on 2D and 3D vector field decompositions in graphics and visualization, but also laments the lack of computational methods providing a five-component decomposition with proper discrete cohomology: until recently, 3D decompositions were mostly achieved for piecewise-constant vector fields on tetrahedral meshes as in [Tong et al. 2003], extending the 2D variational approach of [Polthier and Preuß 2000]; however, this decomposition overly inflates the size of the space of harmonic fields [de Goes et al. 2016a], leading to the wrong dimensionality of the cohomology. A cohomologically-correct five-component decomposition was very recently introduced for 2D surfaces in [Poelke and Polthier 2016; Razafindrazaka et al. 2018] under the name of “boundary-aware” Hodge decomposition; a corresponding 3D five-component decomposition was proposed in [Poelke 2017], extending the 2D decomposition from [Arnold and Falk 1989] and 3D decomposition from [Monk 1991]. However, it assumes piecewise-constant vector fields, making its extension to higher order basis functions unclear and its ability to handle mixed types of boundary condition (common in applications like fluid simulation) limited. Moreover, gauge conditions were not discussed, thus preventing efficient implementations purely based on symmetric positive definite matrices in 3D. Finally, an alternative way for visualization and analysis of 2D or 3D vector fields in bounded domains is to create a natural boundary condition for the gradient and curl components as suggested

in [Bhatia et al. 2014]. However, the lack of orthogonality between the resulting components limits its use in other applications.

1.2 Contributions at a glance

In this paper, we describe both the mathematical formulations and practical computations of a five-component decomposition of vector fields in $\mathbb{R}^3$. We begin with a review of Hodge theory expressed using differential forms, then provide its discretization using Discrete Exterior Calculus (DEC [Desbrun et al. 2008]). We offer:

- a practical procedure for five-component decompositions based on discrete vector fields provided as discrete 1-forms (edge values) or 2-forms (face values) on a tetrahedral mesh;
- a thorough discussion on the enforcement of boundary conditions using DEC discretization to ensure the correct cohomology (with the proper dimensionality of the topology-induced non-integrable parts of the vector field);
- and an effective method for solving the relevant Poisson equations with rank deficiency using only symmetric matrices.

Our exposition aims at serving both practitioners (as we spell out all the matrices involved and numerical treatments of their rank deficiency) and theoretically-minded researchers (as we carefully explain how the discrete setting mimics both boundary conditions and cohomologies). Thus, readers already well-versed in Hodge decomposition and DEC are invited to directly head to Sec 3.3, where our contributions are introduced.

2 BACKGROUND

Before delving into the actual discrete notion of Hodge decomposition, we present some background on the continuous notions that we wish to numerically emulate.

2.1 Helmholtz decomposition

In a bounded domain embedded in 3D Euclidean space, any vector field $\mathbf{v}$ can be expressed as the sum of the gradient of a scalar potential f and the curl of a vector potential $\mathbf{u}$, a two-component decomposition known as the *Helmholtz decomposition*, i.e.,

$$\mathbf{v} = \nabla f + \nabla \times \mathbf{u}.$$

The fields f and $\mathbf{u}$ can be constructed, for instance, using Green’s functions of the Laplacian operator through volume and boundary surface integrals. However, this decomposition is, in general, not an *orthogonal* decomposition, i.e., the L_2 -inner product between ∇f and $\nabla \times \mathbf{u}$ is not necessarily 0, and is not even unique without imposing proper boundary conditions (see Fig. 2). In practical problems, boundary conditions are often crucial, e.g., the slip wall (tangential) boundary conditions for fluid simulation, and the normal boundary condition for the electric field at an ideal conductor boundary. As the orthogonality between the gradient and curl parts are highly relevant for efficiency and accuracy in computational applications, a more general decomposition, called the *Helmholtz-Hodge decomposition* is called for; but it now involves components that are no longer integrable. Yet, these non-integrable parts are finite-dimensional and directly related to the topology of the domain through correspondences established by Poincaré, de Rham, and Hodge, as we briefly discuss next before spelling out the five-component decomposition.

	order 0	order 1	order 2	order 3
form	f^0	$v^1(a)$	$v^2(a, b)$	$f^3(a, b, c)$
d	f	$v \cdot a$	$v \cdot (a \times b)$	$f[(a \times b) \cdot c]$
$[dd = 0]$	df^0	dv^1	dv^2	df^3
$\star$	$(\nabla f)^1$	$(\nabla \times v)^2$	$(\nabla \cdot v)^3$	0
$[\star\star = 1]$	$\star f^0$	$\star v^1$	$\star v^2$	$\star f^3$
δ	f^3	v^2	v^1	f^0
$[\delta\delta = 0]$	δf^0	δv^1	δv^2	δf^3
$\wedge$	0	$(-\nabla \cdot v)^0$	$(\nabla \times v)^1$	$(-\nabla f)^2$
$[(\text{anti-})\text{commute}]$	$f^0 \wedge g^0$	$f^0 \wedge v^1$	$f^0 \wedge v^2, v^1 \wedge u^1$	$f^0 \wedge g^3, v^1 \wedge u^2$
	$(fg)^0$	$(fv)^1$	$(fv)^2, (v \times u)^2$	$(fg)^3, (v \cdot u)^3$

Fig. 3. **Exterior vs. traditional calculus:** odd rows show exterior calculus notations, and even rows give their more conventional expressions in 3D.

2.2 Vector fields through differential forms

Hodge theory is more conveniently and concisely described by differential k -forms and the exterior calculus based on these forms. While this notational formalism is more involved than the traditional vector notation, both are strictly equivalent, and exterior calculus more clearly identifies topological vs. metric operators; the reader unfamiliar with this equivalence is referred to tutorials [Desbrun et al. 2008; Lessig 2012]; we also provide a lookup table to peruse in Fig. 3 that summarize relevant equivalences (specific to 3D).

Forms as scalar or vector fields. A k -form ω^k is a pointwise multilinear mapping from k vectors to a scalar such that if two input vectors are swapped, the sign of the output is switched. Thus a 0- or 3-form in our $\mathbb{R}^3$ setting has only one degree of freedom (DoF) per point, and can be simply identified with a single-component field f (since they represent, respectively, a scalar field and a density field), while a 1- or 2-form has three DoF per point, and can be identified with a vector field v . We will use f^0, f^3, v^1 , and v^2 to denote f seen as a 0- or 3-form and v as a 1- or 2-form respectively.* How the DoFs are used in the antisymmetric linear map is listed in Fig. 3.

Operators on forms. Due to their antisymmetric tensorial nature, k -forms can be integrated on any k -submanifold. Additionally, the exterior derivative (or *differential*) d^k is an antisymmetrization of the partial derivatives of a k -form to produce a $(k+1)$ -form that satisfies the Stokes' theorem over any $(k+1)$ -submanifold $\mathcal{R}$ in M :

$$\int_{\mathcal{R}} d^k \omega^k = \int_{\partial \mathcal{R}} \omega^k.$$

Consequently, one can readily verify that $d^k d^{k-1} = 0$. Depending on the form degree that it is applied to, it encompasses the classical gradient, curl and divergence operators in one consistent type. In the remainder of this paper, we will often omit the superscript of d since it can be directly inferred from the type of its operand. Note that we conventionally call a form *closed* if its differential is zero. Additionally, the wedge product $\wedge$ is defined as an antisymmetrization of the tensor product of two mappings (a p -form and a q -form) to produce a $(p+q)$ -form: for $p+q > 3$, it is 0 since no degrees of freedom are left after antisymmetrization. Finally, the Hodge k -star $\star^k$ (or Hodge dual; we will also omit its superscript at times since the operand disambiguates its identity) is an isomorphism from a k -form ω^k to a $(3-k)$ -form $(\star \omega)^{3-k}$ by treating them as the same DoF used in mapping $3-k$ vectors (instead of k vectors in the same Euclidean coordinate system) to a scalar. Combinations of the basic

*This notation will allow us to keep the “musical” isomorphisms $\sharp$ and $\flat$ hidden to simplify expressions.

Fig. 4. **Helmholtz-Hodge decomposition.** On this example, a tangential field is decomposed into the orthogonal sum of a tangential gradient field, a tangential curl field, and a tangential harmonic field.

operators can be constructed. For instance, $\delta^k = (-1)^k \star^{k+1} d^k \star^k$ is usually called the *codifferential* operator (acting on a k -form and returning a $(k-1)$ -form).

Inner products of forms. On a compact manifold M , the space of k -forms $\Omega^k(M)$ is a Hilbert space when equipped with the inner product between two k -forms α and β defined as:

$$\langle \alpha, \beta \rangle = \int_M \alpha \wedge \star \beta = \int_M \beta \wedge \star \alpha.$$

In our setting, it corresponds to the L_2 -inner product between scalar fields for 0- or 3-forms, and to the L_2 -inner product between vector fields for 1- or 2-forms.

2.3 Hodge decomposition for boundaryless manifolds

Based on the linear map d^k on a boundaryless manifold, there exists an orthogonal decomposition of the space Ω^k written as

$$\Omega^k = \ker d^k \oplus \text{im } \delta^{k+1},$$

where $\ker$ denotes the kernel of an operator, $\oplus$ indicates an orthogonal sum of subspaces, and im denotes the image of an operator. This decomposition is simply a consequence of the fact that the kernel of a linear operator is the orthogonal complement of the range of its adjoint operator. Note that we have

$$\langle d\alpha, \beta \rangle = \langle \alpha, \delta\beta \rangle + \int_{\partial M} \alpha \wedge \star \beta, \quad (1)$$

which implies that δ is formally the adjoint of d only for boundaryless manifolds, i.e., when $\partial M = \emptyset$.

The kernel component can be further decomposed by noticing that $\ker d^k = \text{im } d^{k-1} + \mathcal{H}^k$, where $\mathcal{H}^k = \ker d^k / \text{im } d^{k-1}$ is the quotient space between the kernel of d^k and the image of d^{k-1} (also known as the de Rham cohomology). This property is simply a consequence of the important property $d^k d^{k-1} = 0$. As observed by Hodge, we can turn the direct sum into an *orthogonal* sum instead by picking one particular representative for each equivalence class in $\mathcal{H}^k$: the unique one that is orthogonal to $\text{im } d^{k-1}$. As $\Omega^k = \ker \delta^k \oplus \text{im } d^{k-1}$, one realizes that, in fact, $\mathcal{H}^k$ is isomorphic to $\ker d^k \cap \ker \delta^k$, i.e., to the space of *harmonic* k -forms that are both closed and coclosed—we will also denote it as $\mathcal{H}^k$ due to the natural isomorphism. Given this newfound orthogonality, we reach the Hodge decomposition theorem, which states:

$$\Omega^k = \text{im } d^{k-1} \oplus \text{im } \delta^{k+1} \oplus \mathcal{H}^k. \quad (2)$$

In other words, a k -form $\omega \in \Omega^k$ can be decomposed into the orthogonal sum of an exact form $d\alpha$, a coexact form $\delta\beta$, and a harmonic form $h \in \mathcal{H}^k$ (a form is “exact” if it is the differential of another form;

it is “coexact” if it is a codifferential instead). When $k = 1, 2$, this decomposition can be identified with its vector calculus equivalent, often referred to as the Helmholtz-Hodge decomposition in 3D:

$$\mathbf{v} = \nabla f + \nabla \times \mathbf{u} + \mathbf{h}. \quad (3)$$

According to de Rham’s theorem, $\mathcal{H}^k$ is isomorphic to a space called the singular cohomology $H^k(M)$, which is in turn isomorphic to the homology $H_{n-k}(M)$ by Poincaré duality, where H_k can be understood as the space of non-contractible k -dimensional closed manifolds. The dimensionality of $H_k(M)$ is a finite topological invariant, often referred to as the k -th Betti number $\beta_k = \dim H_k(M)$. Based on Eq. (1), $\mathcal{H}^k$ can also be equivalently defined through $\mathcal{H}^k = \{\alpha \in \Omega^k | \Delta \alpha = 0\}$, where the (de Rham) Laplacian Δ is defined as $\Delta = d\delta + \delta d$, with thus a finite-dimensional kernel. In fact, Δ is self-adjoint due to Eq. (1), and we can decompose k -forms as

$$\Omega^k = \text{im } \Delta^k \oplus \ker \Delta^k. \quad (4)$$

This fact suggests a simple computational approach to the Helmholtz-Hodge decomposition: if we can fix the finite rank deficiency of the Laplacian Δ , by projecting $\mathbf{v}$ to $\ker \Delta$ to get $\mathbf{h}$, and get $\mathbf{w}$ through $\Delta \mathbf{w} = \mathbf{v} - \mathbf{h}$, we can have the Hodge decomposition through $f = \nabla \cdot \mathbf{w}$ and $\mathbf{u} = \nabla \times \mathbf{w}$. However, in practice, our domain in 3D Euclidean space is always bounded and thus with a boundary—in which case, the boundary condition and orthogonality of the subspaces must be treated very carefully as we describe next.

2.4 Hodge decomposition for manifolds with boundary

To ensure adjointness of operators in the presence of boundary, there certainly are a variety of choices. A choice consistent with physical boundary conditions is to force the form α in the decomposition to be tangential to the boundary (we call a form α “tangential” or “parallel” if $\star \alpha$ is zero when applied to tangent vectors of the boundary), or normal to the boundary (we call a form α “normal” if α is zero when applied to tangent vectors of the boundary).

type	f^0	$\mathbf{v}^1$	$\mathbf{v}^2$	f^3
tangential	unrestricted	$\mathbf{v} \cdot \mathbf{n} = 0$	$\mathbf{v} \parallel \mathbf{n}$	$f _{\partial M} = 0$
normal	$f _{\partial M} = 0$	$\mathbf{v} \parallel \mathbf{n}$	$\mathbf{v} \cdot \mathbf{n} = 0$	unrestricted

Consequently, we can construct a Hodge decomposition as proposed in [Morrey 1956] through

$$\Omega^k = d\Omega_n^{k-1} \oplus \delta\Omega_t^{k+1} \oplus \mathcal{H}^k, \quad (5)$$

where Ω_t^{k+1} is the space of tangential forms (also known as Neumann forms since their normal components are fixed), Ω_n^{k-1} is the space of normal forms (or Dirichlet forms), and $\mathcal{H}^k = \ker d^k \cap \ker \delta^k$. Note that being both closed and coclosed is *stronger* than satisfying $\Delta \omega = 0$ when $\partial M \neq \emptyset$. (This point is important in our context, and we will come back to it in Sec. 4.) Nevertheless, one still has the orthogonality between the subspaces, using the adjointness of the operators with Dirichlet and Neumann boundary conditions for the potentials (sometimes called parallel and normal boundary conditions). However, $\mathcal{H}^k$ is infinite-dimensional in this case.

Complete decomposition. Friedrichs [1955] proposed two ways to decompose $\mathcal{H}^k$ orthogonally based on tangential or normal boundary conditions: $\mathcal{H}^k = \mathcal{H}_t^k \oplus (d\Omega_n^{k-1} \cap \mathcal{H}^k)$ as shown in Fig. 5, or $\mathcal{H}^k = \mathcal{H}_n^k \oplus (\delta\Omega_t^{k+1} \cap \mathcal{H}^k)$, which can be combined into the following five-component (Hodge-Morrey-Friedrichs) decomposition:

$$\Omega^k = d\Omega_n^{k-1} \oplus \delta\Omega_t^{k+1} \oplus (\mathcal{H}_t^k + \mathcal{H}_n^k) \oplus (d\Omega_n^{k-1} \cap \delta\Omega_t^{k+1}), \quad (6)$$

where the sum of the latter three terms spans the harmonic space $\mathcal{H}^k$, while $\mathcal{H}_t^k$ is the tangential harmonic space, $\mathcal{H}_n^k$ is the normal harmonic space, and the last term is *both* exact and coexact. Friedrichs also noted that $\mathcal{H}_t^k \cong H^k(M)$ and $\mathcal{H}_n^k \cong H^k(M, \partial M)$, i.e., these two special harmonic spaces are isomorphic to, respectively, the aforementioned (absolute) cohomology $H^k(M)$ now for a manifold with boundary, and the relative cohomology $H^k(M, \partial M)$ for which two k -forms differing only by a k -form on the boundary are treated as equivalent. In general, $\mathcal{H}_t^k$ and $\mathcal{H}_n^k$ are not L_2 -orthogonal with each other. However, they *are* orthogonal for domains in $\mathbb{R}^3$ according to [Shonkwiler 2009], as both the absolute and relative homologies are due to the boundary (we can always patch up the boundary to turn the domain into a ball). This indicates that we have a 5-component orthogonal decomposition, which is consistent with the work of [Cantarella et al. 2002]:

$$\omega^k = d\alpha_n^{k-1} \oplus \delta\beta_t^{k+1} \oplus h_t^k \oplus h_n^k \oplus \eta^k, \quad (7)$$

where η^k is the part that is both exact and coexact. Note that three components can be expressed through potentials under boundary conditions that ensure orthogonality (η being exact and coexact, it can be written as the differential or the codifferential of a potential), while the other two components h_t^k and h_n^k belong to finite-dimensional spaces spanned by harmonic basis fields determined by topology (resp., absolute and relative homologies). As a reminder, the dimensionality of the spaces of both the h_t^k and h_n^{3-k} components in our 3D setup is the Betti number β_k , with a rather intuitive topological meaning: it is the number of components for $k = 0$, tunnels for $k = 1$, cavities for $k = 2$, or simply 0 for $k = 3$.

Equations defining the potentials. Applying the codifferential δ to both sides of Eq. (7), one finds that the form potential α_n^{k-1} satisfying

$$\delta\omega^k = \delta d\alpha_n^{k-1}. \quad (8)$$

This equation can be highly underdetermined as $d\alpha_n^{k-1} = d(\alpha_n^{k-1} + d\gamma)$ for any $\gamma \in \Omega^{k-2}$. As $d\gamma$ does not influence the exact component, it is referred to as a **gauge** field and can be arbitrarily fixed through various gauge conditions. For example, we can enforce $\delta\alpha_n^{k-1} = 0$, turning the above equation into $\delta\omega = (\delta d + d\delta)\alpha_n^{k-1} = \Delta\alpha_n^{k-1}$. Since the rank deficiency of Δ restricted to the space of normal forms Ω_n^{k-1} is $\dim \mathcal{H}_n^{k-1}$, it is finite, so we can leverage this property to solve the corresponding linear system as we will see in Sec. 3.7 when we discuss discretizations and computations. Similarly, applying the differential d to both sides of Eq. (7) shows that

$$d\omega^k = \Delta\beta_t^{k+1}, \quad (9)$$

which has a rank deficiency of $\dim \mathcal{H}_t^{k+1}$ on tangential forms. We will also show in Sec. 3.8 that seeking a potential in $d\Omega \cap \delta\Omega$ for $k = 1, 2$ (i.e., for vector fields) can be achieved by solving a Laplace equation with Neumann boundary condition.

3 DISCRETE DECOMPOSITION OF VECTOR FIELDS

We now assume that our 3D domain M is discretized in the form of a tetrahedral mesh. We can then use the discrete exterior calculus (DEC [Desbrun et al. 2003, 2008; Crane et al. 2013]) as our primary tool to represent discrete differential forms, as DEC preserves the key identity $d\circ d = 0$ for simplicial meshes. We will show that our *discrete* five-component decomposition exhibits the desired orthogonality

Fig. 5. **Hodge-Morrey-Friedrichs decomposition.** For any 3D bounded domain, a vector field can always be decomposed into the orthogonal sum of the gradient of a scalar field vanishing on the boundary, the curl of a normal field, a tangential harmonic fields, and a harmonic gradient field.

defined by the L_2 -inner product between discrete differential forms, as well as the proper cohomology dimensionality for tangential and normal forms—thus exactly mimicking the continuous case and preventing the presence of spurious terms.

Given a tetrahedral mesh M , we denote the set of vertices, edges, faces and tets as $\mathcal{V}/\mathcal{E}/\mathcal{F}/\mathcal{T}$ respectively. We refer to the boundary (triangle) mesh as B , with the boundary vertex/edge/face sets as $\mathcal{V}_b/\mathcal{E}_b/\mathcal{F}_b$. Finally, we assume w.l.o.g that the domain is connected; otherwise one can treat each connected component separately.

3.1 Discrete forms as values on mesh elements

A continuous k -form can be discretized very naturally on a mesh: one can integrate it against every oriented k -simplex of the tet mesh. The resulting set of scalar values (one per oriented k -simplex) can then be seen as a discrete k -form; see [Desbrun et al. 2008; Wang et al. 2006] for details on how to reconstruct a continuous vector or scalar field from such discrete forms. In 3D, as we discussed in Sec. 2.2, vector fields can be interpreted as 1- or 2- forms, while scalar fields are 0- or 3- forms. So we will consider as an input vector field either a set of edge values (where each edge is given a fixed, but arbitrary orientation), or a set of face values (where, again, each face is given an arbitrary fixed orientation) encoding the associated discrete form. The whole discrete decomposition will then split the input discrete form into values on vertices, edges, faces, and/or tets, plus a few non-integrable components depending on the topology of the domain and the boundary condition.

Fig. 6. **Discrete de Rham cohomology.** The DEC linear operators provide a cohomology associated to the combinatorial operators D_k such that $D_{k+1}D_k = 0$ and the Hodge duality through the discrete Hodge stars S_k .

3.2 de Rham complex

Based on Stokes' theorem, the integral of $d\omega$ for a k -form ω on a $(k+1)$ -simplex is simply the signed sum of the integral of ω on the boundary faces of the simplex, where the sign is determined by the relative orientation between the simplex and a particular face. Thus, the exterior derivative d^k is simply encoded in the discrete setting as a matrix D_k which stores the signed incidence between $(k+1)$ -simplices and k -simplices [Desbrun et al. 2008]; it is thus

very sparse and completely combinatorial. The identity $d \circ d = 0$ can be easily verified to hold ($D_{k+1}D_k = 0$) with this discrete definition since the boundary of a boundary of an element is always empty. The Hodge star $\star^k$ is treated as a mapping from a discrete form ω^k (one value per k -simplex) to one value per corresponding $(n-k)$ -cells on a dual mesh—typically, the Voronoi dual structure of the tet mesh. The values on dual Voronoi $(n-k)$ -cells are treated as the integral of an $(n-k)$ -form stored on the dual mesh, and referred to as a *dual* discrete form. Thus, we will have two types of discrete forms (called primal forms $\in \Omega^k$ and dual forms $\in \bar{\Omega}^k$). Their isomorphism is through the Hodge duality $\star^k$, which, in the discrete setting, can be implemented as a *diagonal* matrix S_k , with diagonal entries representing the ratio between the $(n-k)$ -volume of the Voronoi cell and the k -volume of the corresponding primal k -simplex. Other more accurate Hodge star matrices can be used (such as the Galerkin Hodge star [Bossavit 2000]), but they must remain symmetric positive definite (SPD) to guarantee the correct dimensionality of the discrete cohomologies. We discuss how to construct sparse linear systems for non-diagonal Hodge stars after the exposition based on diagonal Hodge stars. We refer to the collection of discrete form spaces connected by the discrete counterparts of the d and $\star$ operators as the *discrete de Rham complex*, mimicking its continuous counterpart, see Fig. 6.

3.3 On the subtleties of boundary treatment

In order to provide a correct computational procedure to find the desired five-component decomposition, boundary values must be treated with caution: a naive derivation of operators without careful boundary treatment can lose the key adjoint properties that we seek to preserve. For instance, the general Laplacian operator for 1-forms is expressed in the continuous setting as $\Delta_1 = d\delta + \delta d$, the analog of the component-wise scalar Laplacian $\Delta = -\nabla \nabla \cdot + \nabla \times \nabla \times \equiv -\nabla^2$ of vector fields. Since we have discrete operators for d and δ , one could be tempted to directly define a discrete Laplacian L_1 as the symmetric matrix (corresponding to $\star \Delta_1$) through:

$$L_1 = D_1^T S_2 D_1 + S_1 D_0 S_0^{-1} D_0^T S_1.$$

This term-by-term conversion then corresponds to a discrete Dirichlet energy for discrete 1-forms $V \in \Omega_p^k$ defined as

$$\frac{1}{2} V^T L_1 V = \frac{1}{2} (D_1 V)^T S_2 (D_1 V) + \frac{1}{2} (S_0^{-1} D_0^T S_1 V)^T S_0 (S_0^{-1} D_0^T S_1 V).$$

However, one realizes that this energy mimics only the *non-boundary part* of the continuous identity (where $\mathbf{n}$ is the boundary normal):

$$\int_M \mathbf{v} \cdot \Delta \mathbf{v} = \int_M (\nabla \cdot \mathbf{v})^2 + \int_M (\nabla \times \mathbf{v})^2 - \int_{\partial M} [\mathbf{v} \nabla \cdot \mathbf{v} + \mathbf{v} \times \nabla \times \mathbf{v}] \cdot \mathbf{n}. \quad (10)$$

In other words, the continuous de Rham Laplacian (inset) implicitly contains boundary terms that are not zero except for very specific types of vector fields, thus adding spurious terms. In the following two sections, we describe how to construct discrete operators that properly treat the typical boundary conditions required in practical computations. We begin with the *tangential Laplace operator*, i.e., the de Rham Laplace operator for vector fields that are *tangent* to the domain boundary, before turning our attention to the *normal Laplace operator*.

3.4 Tangential vector Laplacian operator

Note that we have the choice between a primal 1-form or a primal 2-forms to represent a 3D vector field $\mathbf{u}$, which is equivalent to choosing a dual 2-form or a dual 1-form respectively. In order to assemble a proper discrete *tangential* vector Laplacian, we first discuss the case when our input is a primal 2-form $V \in \Omega_n^2 \subset \Omega^2$ (recall that $\mathbf{v}$ is tangential means that its corresponding 2-form $\mathbf{v}^2$ is normal), followed by the case of a primal 1-form.

2-form version. For an input primal 2-form $V \in \Omega_n^2$ to be normal (and thus corresponding to a tangential vector field $\mathbf{v}$) simply means that the boundary face values are zero, i.e., the flux through boundary triangle faces is null: $\forall \ell \in \mathcal{F}_b, V_\ell = 0$. Forcing the divergence calculation to consider the fluxes through boundary faces as zero is equivalent to simply removing the columns of D_2 that correspond to the boundary faces in order to keep only the interior face values. We denote the resulting matrix as $D_{2,\text{int}}$.

This idea is trivial to generalize: for $k = 0, 1, 2, 3$, we can create a version of D_k *restricted to the interior elements* using

$$D_{k,\text{int}} = P_{k+1} D_k P_k^T,$$

where P_k is the projection (or selection) matrix turning a discrete k -form to a *restricted* k -form containing only the values assigned to interior k -simplices. (Note while we use $D_{k,\text{int}}$ for conciseness, one can also implement this idea through a matrix $\tilde{D}_{k,\text{int}} = P_{k+1}^T P_{k+1} D_k P_k^T P_k$ (this time, without altering the size of the original matrix D_k), as it directly zeroes out the elements in the rows and columns corresponding to boundary elements.)

Similar to $D_{2,\text{int}}$ for divergence, one can show that $D_{1,\text{int}}$ provides the correct curl calculation for tangential vector fields. In addition

Fig. 7. **Absolute and relative homologies.** Homology generators and corresponding harmonic fields on a topological torus with a spherical cavity inside. The red loop (left) around tunnel represents the first homology, and the blue membrane is its dual in the second relative homology. The red curve (first relative homology generator, right) is a loop when the boundary is considered as a point, and the blue membrane is its dual in the second homology. Each harmonic field has the same circulation (resp., flux) on all loops (resp., membranes) that can deform into each other in the domain.

to removing the columns of D_1^T corresponding to boundary faces, it is necessary to remove the rows in D_1^T that correspond to the boundary edges: otherwise, the term corresponding to $\nabla \times \mathbf{v}$ would include a fictitious term assuming the tangential components along the boundary to be zero. More precisely, as shown in the inset, D_1^T sums up the integrals along the yellow dual polyline around the red boundary edge. Defining this term as the curl would amount to setting to zero the line integral along the dotted boundary path that forms a closed loop with the yellow polyline.

We also denote the discrete Hodge star for *interior* k -forms as

$$S_{k,\text{int}} = P_k S_k P_k^T.$$

Notice that $S_{3,\text{int}} = S_3$, since all tets are interior tets. Now the normal 2-form Laplacian can be expressed as

$$L_{2,n} = D_{2,\text{int}}^T S_3 D_{2,\text{int}} + S_{2,\text{int}} D_{1,\text{int}} S_{1,\text{int}}^{-1} D_{1,\text{int}}^T S_{2,\text{int}},$$

With this expression, we can verify that the harmonic forms defined by $h \in \ker L_{2,n}$ indeed correspond to the relative cohomology

$$H^2(M, \partial M) = \ker D_{2,\text{int}} / \text{im } D_{1,\text{int}}.$$

Indeed, we first note that both terms in $L_{2,n}$ are semi-positive definite. Thus $h^T L_{2,n} h = 0$ indicates $(D_{2,\text{int}} h)^T S_3 (D_{2,\text{int}} h) = 0$, which means $h \in \ker D_{2,\text{int}}$ as S_3 is positive definite. Similarly, $D_{1,\text{int}}^T S_{2,\text{int}} h = 0$, which implies $\forall V \in \text{im } D_{1,\text{int}}, V^T S_{2,\text{int}} h = W^T D_{1,\text{int}}^T S_{2,\text{int}} h = 0$ (where $V = D_{1,\text{int}} W$), thus h is orthogonal to $\text{im } D_{1,\text{int}}$. Consequently, h is the unique representative for its equivalence class in the quotient space, and we have the following theorem.

Discrete de Rham's Theorem for Normal 2-Forms. The space of discrete harmonic 2-forms normal to the boundary (i.e., our discrete counterpart to the de Rham cohomology $\mathcal{H}_{\text{dR},n}^2$) is isomorphic to the second (singular) relative cohomology group $\ker L_{2,n} \cong H^2(M, \partial M)$.

By Lefschetz duality, $H^2(M, \partial M) \cong H_1(M)$, the first homology group, which represents the independent “tunnels” of the shape M (see Fig. 7). So the dimension of $\ker L_{2,n}$ is $\beta_1 \equiv \dim H_1(M)$, exactly the sum of the genus for each connected component of the boundary ∂M . We can thus safely use a typical eigensolver to find the β_1 unit eigenvectors associated with the smallest eigenvalues of $L_{2,n} h = \lambda S_2 h$ (these eigenvalues will be 0 up to numerical accuracy): these are the basis of all harmonic 2-forms normal to the boundary. We assemble them into a (tall) matrix $\mathbb{H}_{2,n}$ of size $|\mathcal{F} \setminus \mathcal{F}_b| \times \beta_1$ as: $\mathbb{H}_{2,n} = [h_1 \dots h_{\beta_1}]$.

1-form version. If a 1-form discretization is used for input tangential vector fields $V \in \Omega_1^1$, a direct term-by-term discretization actually holds, i.e., the discrete tangential 1-form Laplacian is simply:

$$L_{1,t} = D_1^T S_2 D_1 + S_1 D_0 S_0^{-1} D_0^T S_1.$$

This is, in light of the previous case, not surprising: a full discretization of the vector field as a discrete one-form should also include one value per boundary Voronoi region (the intersection of the dark dual polyhedra dual to the red boundary vertex on the boundary surface as shown to the right), stored as U_b ; otherwise, the dual 2-form $S_1 U$ cannot be integrated over the boundary by lack of information. Consequently, the matrices D_1 and D_0^T should be augmented accordingly; but forcing the 1-form to be tangential

means that the extra rows/columns must be suppressed anyway just like in the previous case: the use of the original discrete exterior derivatives is thus justified. Since $S_0^{-1}D_0^T S_1$ is now the divergence operator, and at the boundary, the fluxes through interior dual faces are summed while getting no contribution from boundary (tangential condition), the resulting operator properly captures its continuous counterpart. We thus have a similar isomorphism theorem.

Discrete de Rham's Theorem for Tangential 1-Forms. The space of discrete harmonic 1-forms tangential to the boundary is isomorphic to the first cohomology group $\ker L_{1,t} \cong H^1(M)$.

Since the Hodge duality holds for singular cohomology ($H^1(M) \cong H^2(M, \partial M)$), $\ker L_{2,n}$ and $\ker L_{1,t}$ are isomorphic (see Fig. 7). The dimensionality of $\ker L_{1,t}$ is again $\dim H_2(M, \partial M) = \dim H_1(M) = \beta_1$. Solving for the first β_1 eigenvectors associated with the smallest eigenvalues of $L_{1,t}h = \lambda S_1 h$ (which will be 0 up to numerical accuracy) is thus also a viable approach to computing a basis. We assemble them into a (tall) matrix $\mathbb{H}_{1,t}$ of size $|\mathcal{E}| \times \beta_1$ as: $\mathbb{H}_{1,t} = [h_1 \dots h_{\beta_1}]$.

3.5 Normal vector Laplacian operator

The discrete expressions of the two normal Laplacian operators can be obtained by basically mirroring the arguments used earlier, as we now review for completeness.

1-form version. For an input primal 1-form $V \in \Omega_n^1$ to represent a normal vector field (i.e., a 1-form normal to the boundary), one must clearly have: $\forall e \in \mathcal{E}_B, V_e = 0$. Thus, modifying D_1 to become $D_{1,\text{int}}$ by removing the columns of the boundary edges as earlier is required. Moreover, the discrete dual divergence operator D_0^T must also be altered to become $D_{0,\text{int}}^T$ by removing the rows corresponding to boundary vertices: otherwise, a fictitious term in divergence $\nabla \cdot \mathbf{v}$ would (erroneously) assume the boundary fluxes to be zero. The discrete 1-form normal Laplacian is then

$$L_{1,n} = D_{1,\text{int}}^T S_{2,\text{int}} D_{1,\text{int}} + S_{1,\text{int}} D_{0,\text{int}} S_{0,\text{int}}^{-1} D_{0,\text{int}}^T S_{1,\text{int}}.$$

2-form version. Similar to the tangential 1-form $V \in \Omega_n^2$ case, we can augment the discrete tangential 2-form (corresponding to a normal vector field) with additional variables—this time, values of the line integral along each boundary dual edge to encode the tangential components of the 1-form. Setting these extra terms to 0 turns out to be equivalent to using the original D_k and S_k matrices to assemble the Laplacian, hence:

$$L_{2,t} = D_2^T S_3 D_2 + S_2 D_1 S_1^{-1} D_1^T S_2.$$

Like in the tangential case, the two related theorems follow.

Discrete de Rham's Theorem for Normal 1-Forms. The space of discrete harmonic 1-forms normal to the boundary is isomorphic to the first relative cohomology group $\ker L_{1,n} \cong H^1(M, \partial M)$.

Discrete de Rham's Theorem for Tangential 2-Forms. The space of discrete harmonic 2-forms tangential to the boundary is isomorphic to the second cohomology group $\ker L_{2,t} \cong H^2(M)$.

The dimension of the harmonic space is $\dim H^2(M) = \dim H_2(M) \equiv \beta_2$, which is the number of connected components of the boundary minus 1 (see Fig. 7). Solving for the first β_2 eigenvectors associated with the smallest eigenvalues of $L_{1,n}h = \lambda S_1 h$ and $L_{2,t}h = \lambda S_2 h$ (which will be 0 up to numerical accuracy) provides us with the basis of these tangential harmonic spaces. As earlier, we assemble them into two (tall) matrices $\mathbb{H}_{1,n}$ and $\mathbb{H}_{2,t}$ of size $|\mathcal{E} \setminus \mathcal{E}_B| \times \beta_2$ and $|\mathcal{F}| \times \beta_2$.

3.6 Normal and tangential scalar Laplacian operators

For the case of the Laplacian operator of scalar functions (represented as 0- or 3-forms), the exact same construction applies—but the expressions are simpler as only one part of the $d\delta + \delta d$ general expression is nonzero in these cases. We find:

$$L_{0,t} = D_0^T S_1 D_0 \quad L_{3,n} = S_{3,\text{int}} D_{2,\text{int}} S_{2,\text{int}}^{-1} D_{2,\text{int}}^T S_{3,\text{int}}$$

$$L_{0,n} = D_{0,\text{int}}^T S_{1,\text{int}} D_{0,\text{int}} \quad L_{3,t} = S_3 D_2 S_2^{-1} D_2^T S_3.$$

As in the continuous case, the rank deficiency of $L_{0,t}$ and $L_{3,n}$ is $\dim H^0(M) = \dim H^3(M, \partial M) = \dim H_0(M) \equiv \beta_0$, i.e., the number of connected components of the domain. The rank deficiency of $L_{0,t}$ and $L_{3,n}$ is, instead, $\dim H^3(M) = \dim H^0(M, \partial M) = \dim H_3(M) \equiv \beta_3 = 0$ as we cannot have a non-empty boundary of the 3D domain.

The reader may notice that just like for vector Laplacians, the normal Laplacians (where “normal” is meant in the differential form sense, i.e., with an n subscript) involve interior elements only, while the tangential Laplacians are assembled from full-blown differential and star operators. Thus the following formula can be used for any k (where terms containing an index < 0 or > 3 are considered null):

$$L_{k,t} = D_k^T S_{k+1} D_k + S_k D_{k-1} S_{k-1}^{-1} D_{k-1}^T S_k$$

$$L_{k,n} = D_{k,\text{int}}^T S_{k+1,\text{int}} D_{k,\text{int}} + S_{k,\text{int}} D_{k-1,\text{int}} S_{k-1,\text{int}}^{-1} D_{k-1,\text{int}}^T S_{k,\text{int}}.$$

3.7 Five-component decomposition

We are now ready to introduce our computational approach to evaluate the five-component decompositions, which depending on whether we start from a 1-form V^1 or 2-form V^2 input, reads

$$V^k = D_{k-1} \alpha^{k-1} + S_k^{-1} D_k^T S_{k+1} \beta^{k+1} + h_t^k + h_n^k + \eta^k \quad \text{for } k=1, 2.$$

They both correspond to the same vector field decomposition in vector calculus, i.e., $\mathbf{v} = \nabla f + \nabla \times \mathbf{u} + \mathbf{h}_t + \mathbf{h}_n + \nabla e$, where f is a scalar function that vanishes on ∂M (therefore, ∇f is orthogonal to the boundary), $\mathbf{u}$ is a vector potential that is orthogonal to ∂M (and thus $\nabla \times \mathbf{u}$ is a tangential vector field at the boundary), $\mathbf{h}_t$ is a tangential harmonic field, $\mathbf{h}_n$ is a normal harmonic field, and e is a harmonic scalar function (because of the exact and coexact nature of this last term, one can equivalently write it in vector calculus also as the curl $\nabla \times \mathbf{e}$ of a harmonic vector field $\mathbf{e}$).

Equations to solve for potentials. For the 1-form decomposition, one uses our preassembled Laplacian matrices to solve the two discrete form potentials α^0 (on vertices) and β^2 (on faces):

$$L_{0,n} \alpha^0 = D_{0,\text{int}}^T S_1 V^1, \quad L_{2,t} \beta^2 = S_2 D_1 V^1.$$

For the 2-form decomposition, we solve for α^1 (on edges) and β^3 (on tets) instead through:

$$L_{3,t} \beta^3 = S_3 D_2 V^2, \quad L_{1,n} \alpha^1 = D_{1,\text{int}}^T S_2 V^2.$$

Fig. 8. **Harmonic field basis.** Shown are $(\beta_1 = 1)$ tangential and $(\beta_2 = 3)$ normal harmonic basis fields spanning the corresponding harmonic spaces.

Fig. 9. **Resolving rank deficiency.** Randomly selected index sets to remove degeneracy of linear systems may result in very large inaccuracies in the solution of the linear system, unless our simple heuristic is used.

Topological parts. The next two parts are evaluated by just projecting the input form onto the eigenvectors of the vector Laplacian:

$$h_t^k = \mathbb{H}_{k,t} \mathbb{H}_{k,t}^T S_k V^k, \quad h_n^k = P_k^T \mathbb{H}_{k,n} \mathbb{H}_{k,n}^T P_k S_k V^k,$$

since $\mathbb{H}\mathbb{H}^T$ is a projection over the rows of $\mathbb{H}$.

Last term. The fifth element, i.e., the 1-form η^1 or 2-form η^2 , can finally be deduced from the other four components by subtracting them from the input. Note that η is completely determined by either its normal component at the boundary or its tangential components at the boundary. We will also introduce two alternative ways to directly compute it through either of its potentials in Sec. 3.8.

Resolving rank deficiency. The only technical issue in implementing the above linear solves is that some of the Laplacian matrices involved do not have full rank. Fortunately, we know *exactly* their rank deficiency, as well as a basis of their kernel (the associated harmonic forms). For instance, for an equation of the form $L_{2,t}x = y$, we know that the linear system is indefinite since $L_{2,t}$ has a rank deficiency of β_2 . One way to get a definite linear system is to add the constraint $\mathbb{H}_{2,t}^T S_2 x = 0$; but the system is now rectangular. One could instead solve $(L_{2,t} + S_2 \mathbb{H}_{2,t} \mathbb{H}_{2,t}^T S_2)x = y$ efficiently with an iterative solver since $S_2 \mathbb{H}_{2,t} \mathbb{H}_{2,t}^T S_2$ is dense but can be multiplied with a vector in $\mathcal{O}(n)$ time, where n is the number of faces.

Inspired by numerical strategies to pick a subset of columns in order to obtain an optimal condition number with high probability [Tropp 2009], we propose instead a simpler alternative since we already know the kernel and its topological origin. We randomly pick β_2 face indices of x . We assemble the small square sub-block of $\mathbb{H}_{2,t}$ corresponding to these indices, and check its condition number. After having tried $10\beta_2$ such randomly selected index sets, we pick the one with the lowest condition number among those with a determinant higher than the lowest 10% determinants — or we stop early if a condition number happens to be below a given safe threshold (we pick 5.0). This simple procedure has always performed reliably in all of our tests (see Fig. 9). Once we find a good set of indices, we remove these indices from x , along with the corresponding rows and columns of $L_{2,t}$, project the right hand side to $b - S_2 \mathbb{H}_{2,t} \mathbb{H}_{2,t}^T b$, and remove the β_2 indices of this resulting vector as well. This smaller (yet still symmetric) linear system will then have full rank (since we fixed the null space), and a solution of the original equation is the solution of this non-degenerate system where the few missing indices are set to 0. Note that we can finally project this solution to the space containing no harmonic fields using $\mathbb{H}_{2,t}$, if needed. Other rank deficiencies are fixed similarly.

Fig. 10. **Vector potential for tangential harmonic field.** For a tangential harmonic vector field (left) inside his kitten model forming a torus, we can compute its vector potential (right) whose curl is the original field.

3.8 Potentials for the harmonic components

While we proposed a simple eigen-based procedure to compute the tangential and normal harmonic spaces, we can exploit the fact that our domain is embedded in $\mathbb{R}^3$ to directly compute potentials that define the two topological and harmonic terms $\mathbf{h}_t$ and $\mathbf{h}_n^\dagger$. Depending on how the decomposition is used in practical applications, this alternative approach may be more efficient. For completeness, we also describe how to extract the potential (either as a gradient or a curl) of the fifth term η .

Tangential harmonic space. If one has already computed the generators for $H_2(M, \partial M)$, i.e., a set of independent surfaces in M that have their boundary loops in ∂M , we can construct one gradient field per generator that will be a *tangential* harmonic field. The gradient field is constructed by simulating a cut through the generator, allowing the potential to have two different (edge or face) values on the generator that differ exactly by 1 as done in global parameterization methods for quad meshing purposes [Tong et al. 2006]; we can then solve a Laplace equation with a single element fixed to remove the null space. Once these gradients are found, we run a Gram-Schmidt procedure to obtain an orthonormal basis for these tangential harmonic fields. Another simple strategy is to first restrict the computation to the nonzero genus boundary components. For each handle loop (that is, a non-contractible loop of the boundary which can be contracted inside the domain volume), one may build a tangential harmonic vector field $\mathbf{h}_t$ on the surface such that the circulation around the handle loop is 1; one can then extend $\mathbf{h}_t$ to the volume by solving a vector Laplace equation $\Delta \mathbf{h}_t = 0$ (i.e., either $L_{1,t} \mathbf{h}_t^1 = 0$ or $L_{2,n} \mathbf{h}_n^2 = 0$ depending on the decomposition being targeted) with Dirichlet boundary condition on all components of $\mathbf{h}_t$ at the boundary (for all other connected components of the boundary, it is set to 0). We will have β_1 such vector potentials, and they will span the entire space of tangential harmonic field, providing an alternative to the construction of $\mathbb{H}_{1,t}$ and $\mathbb{H}_{2,n}$. Moreover, the vector potential $\boldsymbol{\psi}_t$ of each basis element can be solved through $\Delta \boldsymbol{\psi}_t = \nabla \times \mathbf{h}_t$ with boundary conditions $(\nabla \times \boldsymbol{\psi}_t) \times \mathbf{n} = \mathbf{h}_t \times \mathbf{n}$ and by forcing the normal component of $\boldsymbol{\psi}_t$ to be 0 to impose the gauge condition,; e.g., for $\mathbf{h}_n^2$, we can solve for

[†] These potentials are not of the same nature as α and β : from Eq. (5), one can see that harmonic parts can *not* be written as the d of a normal form or the δ of a tangential form. But they are, however, in the range of d and the range of δ , so we can find potentials for them—just not with the same boundary conditions, hence the commonly-used term of “non-integrable” to describe these topological terms.

Fig. 11. **Potentials for the exact and coexact field.** In any 3D volume, the fifth vector component η in our decomposition (left) can be expressed both as a gradient field (middle) and a curl field (right).

the potential ψ_t^1 through $L_{1,t}\psi_t^1 = D_1^T S_2 h_n^2$, where the righthand side computes the tangential component of h_n^2 (i.e., it generates the tangential component of h_n^2 across the dual of boundary edges and produces 0 for the interior); we proceed for h_t^1 in a similar fashion, with $L_{2,n}\psi_n^2 = S_{2,\text{int}} D_{1,\text{int}} P_1 h_t^1$ this time (where $D_{1,\text{int}} P_1 h_t^1$ contains only the negative tangential component along boundary edges).

Normal harmonic space. We can similarly construct the elements of the kernels of the *normal* Laplacian matrices directly. Through the duality to $H^1(M, \partial M)$, we can represent these harmonic functions as combinations of simple gradient fields $\mathbf{h}_n = \nabla \phi$ (where ϕ be a discrete 3-form (resp., 0-form) when the input is a 2-form (resp., 1-form)), which are the solution of $\Delta \phi = 0$ with Dirichlet boundary conditions $\phi = 1$ on one of the connected component of the boundary mesh, and $\phi = 0$ on the rest of the boundary. This is a 3D extension of the procedure proposed in the Appendix of [Tong et al. 2006], and essentially corresponds to the problem of finding a static electric field with ideal conductor boundary for given potentials on the boundary. For 1-form inputs, it is solved on the graph of primal vertices and edges, while for 2-form inputs, it is solved on the dual graph for tets and faces. An additional Gram-Schmidt procedure is also necessary if one requires an orthonormal basis. If the potentials of these normal harmonic basis elements are needed, we can solve them in a mirrored way to the tangential case through: $L_{0,t}\phi_t^0 = D_0^T S_1 h_n^1$ and $L_{3,n}\phi_n^3 = S_3 D_{2,\text{int}} P_2 h_t^2$.

Potential(s) for fifth term. From η^1/η^2 , we can solve for their scalar potential e ($\eta^1 = D_2 e_t^0 \equiv de_t^0$ or $\eta^2 = S_{2,\text{int}}^{-1} D_{2,\text{int}}^T S_3 e_n^3 \equiv \delta e_n^3$) through:

$$L_{0,t}e_t^0 = D_0^T S_1 \eta^1, \quad \text{or} \quad L_{3,n}e_n^3 = S_3 D_{2,\text{int}} P_2 \eta^2,$$

where the right hand side only contains nonzero terms at the boundary (enforcing $\nabla e \cdot \mathbf{n} = \eta \cdot \mathbf{n}$). This is essentially the same setup as solving for potentials of normal harmonic fields. For such Neumann boundary conditions, we also need to fix $\beta_0 = \dim H_0(M)$ variables, since we can add one constant function to each connected component of the domain without changing the actual fifth component. Likewise for the vector potential $\mathbf{e}$ ($\eta^1 = S_{1,\text{int}}^{-1} D_{1,\text{int}}^T P_2 S_2 e_n^2 \equiv \delta e_n^2$ or $\eta^2 = D_{1,\text{int}} e_t^1 \equiv de_t^1$) by solving the same type of Laplace equation with boundary conditions $\nabla \times \mathbf{e} \times \mathbf{n} = \eta \times \mathbf{n}$ as for potentials of tangential harmonic fields, i.e.,

$$L_{2,n}e_n^2 = S_{2,\text{int}} D_{1,\text{int}} P_1 \eta^1 \quad \text{or} \quad L_{1,t}e_t^1 = D_{1,\text{int}}^T S_2 \eta^2.$$

Observe that directly applying δ to e_n^3 (resp., e_n^2) only provides correct values for η^2 (resp., η^1) on interior elements. Still, these potentials offer enough information for extrapolation to boundary elements through harmonicity of η^2 : if each tet contains at most one boundary face, η^2 on that face is the negated sum of the other three fluxes; likewise for η^1 if each boundary edge is incident to at least

one triangle with only one boundary edge. If the input mesh does not satisfy these conditions, local splits of offending tets and triangles can be applied. Alternatively, e_n^3 (resp., e_n^2) can be supplemented with one value per boundary face (resp., edge) for δ to generate the correct gradient (resp., curl) on each boundary element.

3.9 Counting argument

Both to further enhance the understanding of our discrete vector decomposition and to offer yet another approach to convince oneself that the counting is correct, we now review the number of degrees of freedom (DoFs) within each component in both representations. For the 1-form representation, $d\alpha^0$ contains $|\mathcal{V}| - |\mathcal{V}_B|$ DoFs, i.e. one value per interior vertex; $\delta\beta^2$ contains $|\mathcal{F}| - |\mathcal{T}| - \beta_2$ since we start with $|\mathcal{F}|$ values but need to get rid of $\dim \ker \delta_2 = \dim \text{im } \delta_3 + \dim H_2$. The non-integrable components h_t^1 and h_n^1 provide β_1 and β_2 DoFs respectively. Finally, η^1 provides $|\mathcal{V}_B| - \beta_2 - \beta_0$ DoFs, because $\beta_2 + \beta_0$ is the number of connected components of the boundary, and on each of them the total flux is 0. From the Poincaré-Euler formula

$$|\mathcal{V}| - |\mathcal{E}| + |\mathcal{F}| - |\mathcal{T}| = \beta_0 - \beta_1 + \beta_2 - \beta_3,$$

we can then verify that the number of values of the input 1-form ($|\mathcal{E}|$) is indeed the sum of the above DoFs (with $\beta_3 = 0$ in 3D). For the 2-form representation, following similar arguments, the DoFs for the five components are in the same order: $|\mathcal{E}| - |\mathcal{E}_B| + |\mathcal{V}| - |\mathcal{V}_B| - \beta_2$, $|\mathcal{T}|$, β_2 , β_1 , and $|\mathcal{F}_B| - \beta_2 - \beta_0$. Noting that $|\mathcal{V}_B| - |\mathcal{E}_B| + |\mathcal{F}_B| = 2(\beta_0 - \beta_1 + \beta_2)$ as it is the sum of the Euler characteristic $2 - 2g$ of each boundary component (one should not use the Euler characteristic of the volumetric domain!), we can again verify that they sum up to $|\mathcal{F}|$, as expected. We recap the various numbers of DoFs in Tab. 1.

	1-form DoFs	2-form DoFs
ω	$ \mathcal{E} $	$ \mathcal{F} $
$d\alpha$	$ \mathcal{V} - \mathcal{V}_B $	$ \mathcal{E} - \mathcal{E}_B + \mathcal{V} - \mathcal{V}_B - \beta_2$
$\delta\beta$	$ \mathcal{F} - \mathcal{T} - \beta_2$	$ \mathcal{T} $
h_t	β_1	β_2
h_n	β_2	β_1
η	$ \mathcal{V}_B - \beta_2 - \beta_0$	$ \mathcal{F}_B - \beta_2 - \beta_0$

Table 1. List of DoFs for 1-form and 2-form decompositions.

4 VARIATIONAL NATURE OF OUR DECOMPOSITION

Because we made sure our discrete treatment is closely mimicking the continuous five-component Hodge decomposition, it is directly related to variational approaches to vector decomposition based on L_2 energies. In particular, we point out that our discrete treatment can be understood as a particular enforcement of harmonicity with zero divergence and curl boundary conditions to enforce proper orthogonal projections.

Continuous notion of harmonicity. Because we are in $\mathbb{R}^3$, recall from Eq. (10) that the Laplace quadratic form satisfies:

$$\int_M \mathbf{v} \cdot \Delta \mathbf{v} = \int_M (\nabla \cdot \mathbf{v})^2 + (\nabla \times \mathbf{v})^2 \quad (11)$$

only if the boundary integral vanishes, i.e.:

$$\int_{\partial M} (\mathbf{v} \nabla \cdot \mathbf{v} + \mathbf{v} \times \nabla \times \mathbf{v}) \cdot \mathbf{n} = 0$$

Our choice of gauge in the decomposition proposed in Sec. 3.7, in fact, enforces the latter equality since it implies that we discretely enforce $\nabla \times \mathbf{v} = 0$ (with tangential $\mathbf{v}$) or $\nabla \cdot \mathbf{v} = 0$ (with normal $\mathbf{v}$) to make this boundary integrand identically zero. This is precisely why

our harmonic forms are not only harmonic ($\Delta v = 0$) in the interior of the domain, but have these boundary conditions enforced as well — explaining why we stated in Sec. 2.4 that forcing the forms in $\mathcal{H}^k$ to be closed and coclosed (hence, curl- and divergence-free) is *stronger* than just the notion of interior harmonicity. In the continuous setting, the consequence of these boundary conditions is that our construction can then be understood as forcing the tangential vector fields to satisfy one Dirichlet condition $\mathbf{v} \cdot \mathbf{n} = 0$ (tangentiality) and two Neumann conditions $\mathbf{n} \cdot \nabla_{\mathbf{t}_1} \mathbf{v} = 0$, $\mathbf{n} \cdot \nabla_{\mathbf{t}_2} \mathbf{v} = 0$ (where $\mathbf{t}_1$ and $\mathbf{t}_2$ are two local tangent direction forming a coordinate frame with the surface normal $\mathbf{n}$) to enforce a zero curl. On the other hand, the normal vector fields are constrained to satisfy two Dirichlet conditions $\mathbf{v} \cdot \mathbf{t}_1 = 0$ and $\mathbf{v} \cdot \mathbf{t}_2 = 0$, and one Neumann condition $\mathbf{n} \cdot \nabla \mathbf{v} = 0$ to enforce a zero divergence. These conditions are consistent with the formulation on generic manifold cases from [Demlow and Hirani 2014]. With these three conditions on the boundary added to the condition of harmonicity, the space of harmonic forms is finite-dimensional.

Dirichlet energy. In flat $\mathbb{R}^3$, the oft-used Dirichlet energy of vector fields can be converted to a volume integration using the Laplacian and a boundary term through integration by part:

$$\int_M |\nabla \mathbf{v}|^2 = \int_M \mathbf{v} \cdot \Delta \mathbf{v} + \int_{\partial M} \mathbf{v} \cdot \nabla \mathbf{n} \cdot \mathbf{v}. \quad (12)$$

Notice now that with the boundary conditions we enforced, by Eq. (11), we have three energies that match: the harmonic energy, the Laplace quadratic form, and the Dirichlet energy, i.e.,

$$\int_M (\nabla \cdot \mathbf{v})^2 + (\nabla \times \mathbf{v})^2 = \int_M \mathbf{v} \cdot \Delta \mathbf{v} = \int_M |\nabla \mathbf{v}|^2.$$

Variational nature. Due to its L_2 -orthogonality, one can conceive our decomposition as orthogonal projections onto subspaces — and thus as a variational problem. For instance the projection of V^1 onto $D_0 \alpha^0$ can be seen as minimization of

$\langle V^1 - D_0 \alpha^0, V^1 - D_0 \alpha^0 \rangle = \langle V^1, V^1 \rangle - 2 \langle V^1, D_0 \alpha^0 \rangle + \langle D_0 \alpha^0, D_0 \alpha^0 \rangle$. The scalar function α^0 is then entirely determined by adding a gauge to enforce that α is zero on the boundary. This type of variational arguments were already leveraged for a three-component 3D decomposition in [Tong et al. 2003] (extending the 2D work of [Polthier and Preuß 2000]); however, their choice of space of discrete vector fields (piecewise linear vector potential) did not lead to a cohomology-preserving discretization. Using DEC, instead, allows a discretization that captures the topological aspects correctly.

Our discrete Laplacian can really be seen as the counterpart of the continuous Laplacian, with particular boundary conditions (compatible with gauge conditions) added at the boundary to offer proper L_2 -orthogonality of the various terms of the Hodge decomposition.

5 EXTENSIONS AND SPECIALIZATIONS

While we described a discrete decomposition of vector fields given as 1- or 2-forms with a particularly canonical choice of gauges, we can extend our approach to different gauges in order to get different potentials, derive smaller, more specialized decompositions, or use non-diagonal Hodge star matrices without hindering efficiency.

5.1 Helmholtz decomposition

Our five-component decomposition can be trivially condensed into the two-component Helmholtz decomposition we described in Sec. 2.1:

∇f , $\mathbf{h}_n$, and η can all be expressed as a gradient field, and $\nabla \times \mathbf{u}$, $\mathbf{h}_t$, and η can be expressed as a curl field; no matter how we split the η term, we will get the expected decomposition of the type

$$\mathbf{v} = \nabla \phi + \nabla \times \psi.$$

However, if one wishes to ensure the orthogonality between the two components, we must put η entirely in the gradient part (resulting in a tangential curl field), or entirely in the curl part (resulting in a normal gradient field).

5.2 Specialized inputs

In some contexts, we can assume the input to be a tangential or normal vector field. In these cases, it is possible to specialize our decomposition and make it a three-component or even two-component decomposition instead. We provide a few examples to illustrate how this can be useful (and more efficient) in practice.

Tangential inputs. If we know that input vector field $\mathbf{v}$ is tangential, we can directly solve for a tangential vector field (called *Newtonian potential*) $\mathbf{w}$, such that the continuous decomposition becomes: $\mathbf{v} = \Delta \mathbf{w} + \mathbf{h}_t$. The discrete version is straightforward: one can solve for a 1-form w^1 or a 2-form w^2 based on the degree of the input form through:

$$L_{1,t} w_t^1 = S_1(V^1 - h_t^1) \quad \text{or} \quad L_{2,n} w_n^2 = P_2 S_2(V^2 - h_n^2),$$

since $\mathbf{v} - \mathbf{h}_t$ is in the image of Δ for tangential fields; note that the (tangential) harmonic part is computed directly by projection with the basis $\mathbb{H}$. This decomposition can be turned into a three-component decomposition as well through

$$\mathbf{v} = -\nabla(\nabla \cdot \mathbf{w}) + \nabla \times (\nabla \times \mathbf{w}) + \mathbf{h}_t.$$

We can further shift part of the curl field to the gradient field to make every component tangential: we can solve for the normal vector potential, and shift the rest to the gradient part.

Normal inputs. For normal vector fields, a similar approach leads to a two- or three-component decomposition:

$$\mathbf{v} = \Delta \mathbf{w} + \mathbf{h}_n \quad \text{or} \quad \mathbf{v} = -\nabla(\nabla \cdot \mathbf{w}) + \nabla \times (\nabla \times \mathbf{w}) + \mathbf{h}_n.$$

The discrete treatment to find the normal vector field $\mathbf{w}$ as a 1- or 2-form becomes then:

$$L_{1,n} w_n^1 = P_1 S_1(V^1 - h_n^1) \quad \text{or} \quad L_{2,t} w_t^2 = S_2(V^2 - h_n^2),$$

Again, one can shift part of the gradient field in the three-term decomposition to make the gradient field part normal to the boundary, which will make the curl field part normal to the boundary.

5.3 Mixed boundary conditions

Mixed and/or partial boundary conditions are sometimes required. Orthogonal decomposition into gradient and curl fields with boundary conditions and topology-determined finite-dimensional harmonic space can be established in the same fashion through relative homologies. In general, the boundary is the disjoint union of tangential, normal, and unconstrained regions: $\partial M = \partial M_t \cup \partial M_n \cup \partial M_u$. Sticking to the original full decomposition will lead to some components not satisfying the boundary conditions. One can make each component follow the given boundary conditions by replacing the original boundary conditions to enforce, instead:

$$\begin{aligned} \star d\alpha|_{\partial M_t} &= \star h_n|_{\partial M_t} = \star \eta|_{\partial M_t} = 0, & \delta \beta|_{\partial M_n} &= h_t|_{\partial M_n} = \eta|_{\partial M_n} = 0, \\ \alpha|_{\partial M_n \cup \partial M_u} &= h_n|_{\partial M_n \cup \partial M_u} = 0, & \star \beta|_{\partial M_t \cup \partial M_u} &= \star h_t|_{\partial M_t \cup \partial M_u} = 0. \end{aligned}$$

Fig. 12. **Mixed boundary conditions.** Our orthogonal decomposition extends naturally to mixed boundary conditions as well; in this example, no constraints are set on the blue regions, but tangential conditions are set on the rest of the boundary.

From a practical standpoint, we simply have to impose the typical conditions on α , β , h_n and h_t on the unconstrained regions to provide orthogonality and the uniqueness of the decomposition; one can add η to any combination of the other components to create the “natural” unconstrained behavior for chosen components as we described for the three-component decomposition. In order to solve for the various potentials, we can define an altered Laplacian $L_{k,A}$ for k -forms that are normal on a boundary region $A \subset \partial M$ and tangential on $\partial M \setminus A$ through the following matrix expression:

$$L_{k,A} = D_{k,A}^T S_{k+1,A} D_{k,A} + S_{k,A} D_{k-1,A} S_{k-1,A}^{-1} D_{k-1,A}^T S_{k,A},$$

where $S_{k,A} = P_{k,A} S_k P_{k,A}^T$, $D_{k,A} = P_{k+1,A} D_k P_{k,A}^T$, and $P_{k,A}$ is the projection of a full k -form on M to a k -form restricted to $M \setminus A$. With this definition, we can solve for α using $L_{k-1, \partial M_n \cup \partial M_u}$, for β using $L_{k+1, \partial M_n}$, for h_n using $L_{k, \partial M_n \cup \partial M_u}$, and for h_t using $L_{k, \partial M_n}$.

Note that h_t and h_n are no longer necessarily L_2 -orthogonal for general mixed boundary conditions as in the generic Hodge-Morrey-Friedrichs decomposition (Eq. (6)). Nevertheless, we can either combine the two components and create an orthonormal basis for the span of both low dimensional subspaces, and/or combine one of them to the other components. Let’s use the case $\partial M = \partial M_t \cup \partial M_u$ (i.e., where we only want to impose tangential forms on parts of the boundary) as an illustrative example. We propose the following decomposition by combining some of the components,

$$\omega^1 = d\alpha^0 + \delta\beta^2 + h_n^1$$

All three components are tangential on ∂M_t , α^0 is 0 on ∂M_u , β^2 is tangential on ∂M_t (i.e., the corresponding vector field is normal), and h_n^1 is normal on ∂M_u . Following the derivation of typical boundary conditions, we find that h_n^1 correspond to the relative homology $H_2(M, \partial M_t)$, or equivalently through its Hodge dual, to $H_1(M, \partial M_u)$. Thus, the space for h_n^1 is finite dimensional; Fig. 12 shows such an example where it is two dimensional.

Finally, the orthogonality of the various terms of the resulting five-component decomposition is properly enforced. E.g.,

$$\langle d\alpha^0, h_n^1 \rangle = \langle \alpha^0, \delta h_n^1 \rangle + \int_{\partial M} \alpha^0 \wedge \star h_n^1,$$

which is zero since $\delta h_n^1 = 0$, the boundary integral on ∂M_u is 0 as $\alpha^0 = 0$ there, and the leftover integral on ∂M_t vanishes due to the boundary condition on $\star h_n^1$. The exact same argument holds for the $\delta\beta^2$ term, while the orthogonality between $\delta\beta^2$ and h_n^1 is established through similar arguments. Note that when ∂M_u is replaced by ∂M_n , the case $\partial M = \partial M_t \cup \partial M_n$ is recovered.

5.4 Friedrichs decompositions

Finally, we note that η in the four-component decomposition can be merged with h_t to form a subspace that is both harmonic and a curl field, or with h_n to form a subspace that is both harmonic and a gradient field. Both four-component decompositions are sometimes called the Friedrichs decomposition.

5.5 Non-diagonal Hodge star

While the low-order “diagonal Hodge star” is often preferred in graphics due to its optimal sparsity [Desbrun et al. 2008], a variety of other discrete Hodge stars have been proposed [Mohamed et al. 2016]. Of particular interest are the Galerkin Hodge stars [de Goes et al. 2016b; Arnold 2018] which offer higher-order accuracy of approximation, at the price of requiring still sparse, but non-diagonal matrix representations. As they are symmetric positive definite, our decompositions apply without modification. However, S_k^{-1} can become a dense matrix, making the evaluation of the Laplace matrices much less efficient. We outline a procedure that only uses sparse matrices for the decomposition to be still strictly L_2 -orthogonal according to a non-diagonal Hodge star matrix S_k .

We first note that among the necessary discrete Laplacians for the decomposition of V^k ($k = 1, 2$), only $L_{k+1,t}$ involves S_k^{-1} . In other words, we can compute $D_{k-1}\alpha^{k-1}$, h_t^k , h_n^k , and $\eta = D_{k-1}\phi^{k-1}$ with the non-diagonal S_k . While it may be necessary to replace S_{k-1}^{-1} and S_{k-2}^{-1} by sparse substitute matrices $\tilde{S}_{k-1}$ and $\tilde{S}_{k-2}$ (e.g., identity matrices) to keep those systems sparse, it does not influence the actual accuracy of the decomposition: first, the L_2 -orthogonality in Ω_k for the components depends on S_k , which is not altered; second, the harmonic spaces remain the same since the kernels remain in $\ker d \cap \ker \delta$ under normal/tangential boundary conditions; third, the potential α may deviate from satisfying $\delta\alpha = 0$ exactly, but the error lies within the gauge field, so $D_{k-1}\alpha^{k-1}$ is still accurate.

For the final component, note that $\gamma \equiv \omega - d\alpha - h_t - h_n - \eta$ is in $\text{im } L_{k+1,t}$, so $S_{k+1}\beta = \gamma$ has a solution in $\Omega_{k+1,t}$. This means that

$$\left\langle D_k^T S_{k+1}\beta - S_k\gamma, D_k^T S_{k+1}\beta - S_k\gamma \right\rangle_{\tilde{S}_k}$$

can be minimized to exactly 0 in any weighted L_2 -inner product $\langle \cdot \rangle_{\tilde{S}}$, where $\tilde{S}_k$ is an arbitrary sparse SPD matrix. We can thus solve for the exact β without the inverse matrix S_k^{-1} through

$$(D_{k+1}^T S_{k+2} D_{k+1} + S_{k+1} D_k \tilde{S}_k D_k^T S_{k+1})\beta = S_{k+1} D_k \tilde{S}_k S_k \gamma.$$

If we take $\tilde{S}_k = S_k^{-1}$, the matrix is the same tangential Laplacian used for solving for β in our DEC decomposition; but we can now accommodate non-diagonal Hodge matrices as $\tilde{S}$ can be chosen arbitrarily:

we will still find the exact potential satisfying $S_k^{-1}D_k^T S_{k+1}\beta = \gamma$. This construction extends to arbitrary Hodge stars the approach described in [Bossavit and Kettunen 1999], where the authors realized that when the Galerkin Hodge star S_k (computed using Whitney forms, and thus non-diagonal) is multiplied by D_{k-1} on the right and D_{k-1}^T on the left, the result is no different than if the Galerkin Hodge star was replaced by a diagonal “lumped” matrix.

6 EXPERIMENTS

Decomposition zoo. In Fig. 1, we perform the full five-component vector field decomposition using a discrete 1-form representation. The connected domain contains one outside and one inside boundary components, with genus 1 and 0 respectively, thus $\beta_0 = 1, \beta_1 = 1, \beta_2 = 1, \beta_3 = 0$. We further evaluate the vector potential of the tangential harmonic component, the scalar potential of the normal harmonic component, and both potentials of the fifth (exact, coexact) component. We also numerically verified the L_2 -orthogonality of the five terms. In Fig. 8, we provide a depiction of all the harmonic field basis vectors for a model with a more complex topology (two spherical and one toroidal cavities).

To demonstrate the non-orthogonality when no boundary condition is imposed, we show in Fig. 2 a decomposition into the sum of a gradient field and a curl field, resulting from the five-component decomposition and after merging $d\alpha + h_n + \frac{1}{2}\eta$ and $\delta\beta + h_t + \frac{1}{2}\eta$ (or just summing up the potentials). Note that the L_2 -inner product between the two is then $\frac{1}{4}\langle\eta, \eta\rangle$.

When the input is a tangential field as in Fig. 4, its Helmholtz-Hodge decomposition contains three tangential fields, the gradient field $d\alpha + h_n + \eta$, the curl field $\delta\beta$, and the tangential harmonic field (non-integrable in the sense that it cannot be seen as the curl of a *normal* vector potential). It is also possible to obtain either one of the two four-component Hodge-Morrey-Friedrichs decompositions; e.g., in Fig. 5, we decompose the input into $d\alpha$, $\delta\beta$, h_t , and a harmonic gradient field $h_n + \eta$, which is harmonic with a scalar potential.

In Fig. 12 for mixed boundary conditions, we create a case mimicking fluids passing through a domain with two openings. As described in Sec. 5.3, the normal harmonic space (vector field normal to the unrestricted boundary and tangential to the tangential boundary condition region) is two dimensional, which can also be constructed through eigensolvers or through the corresponding relative homology. The gradient component can be constructed by solving a Poisson equation with the divergence of the input on the right hand side, and the same tangential boundary conditions. The rest can be expressed as the curl of a vector potential that is orthogonal to the boundary outside the openings.

Using the Galerkin Hodge star associated with Whitney basis functions [Bossavit 2000], the potential β for the $\delta\beta$ term in Fig. 13 is accurate with our approach. If the diagonal Hodge star S_D is used instead in the Laplacian to compute a different potential $\tilde{\beta}$, then $S_1^{-1}D_1^T \tilde{\beta}$ has a deviation from the γ term defined in Sec. 5.5 of around 48% (Fig. 13(top right)), but it still is orthogonal to the other components; if one tries $S_D^{-1}D_1^T \tilde{\beta}$ for consistency, then there is still a 1.5% deviation from γ (Fig. 13(bottom right)) and a 1.5% error in L_2 -inner product with the other components is now present.

Fig. 13. **Non-diagonal Hodge star.** Even for higher-order accurate Hodge stars, our decomposition still only requires sparse linear systems. Using a diagonal S_1 in the Laplacian produces inaccurate potentials (right), whether we use a curl operator with a diagonal (top) or non-diagonal (bottom) S_1 .

Our decomposition is also demonstrated on a simulated channel flow. The velocity field was generated with the OpenFOAM software, with a forced velocity on the round inlet and outlet, with free-slip and no-transfer boundary conditions on the interior walls. Our decomposition thus sets all regions away from the inlet and outlet with tangential conditions.

Performance and accuracy. For completeness, we also tested the assembly of the matrices on a laptop. For models with around 25K tets, we can perform the necessary solves using Conjugate Gradient in 2s even on a regular laptop with our unoptimized code. Note that if we prefactorize (through Cholesky decomposition) the Laplacians, we can much more efficiently perform the decomposition of arbitrary fields on the same domain through forward and backward substitutions, in less than a second. As shown in Fig. 9, our condition number based strategy to choose rows and columns to eliminate the null space of the Laplacian matrices is very effective in maintaining the accuracy of the linear system. Note that when working with non-diagonal Hodge stars, we can also either use Conjugate Gradient or precompute a Cholesky decomposition for the evaluation of the curl in (Fig. 13): for a 10K tet mesh, the iterative CG solve takes less than 1s, whereas the Cholesky factorization of the non-diagonal Hodge star takes 5s—but allows fast repeated evaluations.

Comparison to [Poelke 2017]. The only other existing 5-component 3D decomposition and our approach are based on very different discretization methods: Poelke represents a discrete vector field as piecewise constant per tetrahedron, while we use discrete 1- or 2-forms. In this sense, our work is complementary. Yet, and while cohomologies are preserved in both approaches, our representation also requires fewer DoFs as input, as the number of edges or faces is always smaller than three times the number of tets. Our approach also tackles the full 5-component decomposition using *only symmetric semi-positive definite matrices with smaller sizes*, resulting in higher computational efficiency: numerical experiments confirm that differential form based discretization leads to better accuracy, partially due to their exact line integral and flux sampling (i.e., linear precision vs piecewise-constant precision of the representation). Moreover, it is straightforward for us to formulate the relation of mixed boundary conditions to relative cohomologies, or to extend our construction using a higher-order L_2 -inner product. Finally, our eigensolver also produces L_2 -orthonormal basis for the cohomology more efficiently than the non- L_2 -orthonormal basis obtained in [Poelke 2017] through singular value decomposition of rectangular matrices.

7 CONCLUSION

We have detailed the construction of a five-component decomposition of vector fields on triangulated 3D domains for a large variety of typical boundary conditions. Our approach was shown to be consistent with the continuous theory on vector field analysis, and to capture the proper kernel spaces due to nontrivial homology and cohomologies, either from the topology or from the boundary of 3D domains. We showed that the numerical procedure based on discrete exterior calculus leads to straightforward constructions of the desired boundary conditions. We discussed how to assemble all the matrices involved, as well as how to handle their known rank deficiency (based on the topology of the domain) to ensure fast computations. We expect this straightforward numerical tool to benefit computational applications involving volumes in 3D Euclidean space, such as in geometric modeling, electromagnetism, fluid dynamics, elasticity and biomolecular science.

Fig. 14. **Decomposition of a channel flow simulation.** For a simulated channel flow (inlet and outlet in blue), the resulting vector field is decomposed into a curl field and a harmonic field, with the blue regions are set as unconstrained and all other boundary regions as tangential.

Limitations and future work. Our decomposition is restricted to domains in $\mathbb{R}^3$ with Euclidean metric. It can however be extended to any 3-manifold that can be embedded in $\mathbb{R}^3$: the orthogonality between $\mathcal{H}_t$ and $\mathcal{H}_n$ only depends on the topology. Moreover, it is possible to extend it to k -forms on any simplicial tessellation of compact n -manifolds with boundaries if we lift the restriction on the orthogonality between those two components, and compute the harmonic vector fields also through eigensolvers instead of our efficient alternatives through potentials designed for 3D domains. As a special case, 0-forms and n -forms on n -manifolds can always be orthogonally decomposed into the divergence of a tangential gradient field plus β_0 constant fields. Exploring spectral analysis of our tangential and normal Laplacian operators is also an interesting direction of research.

ACKNOWLEDGMENTS

Partial funding for this work was provided by NSF grants DMS-1721024 and IIS-1900473, and Pixar Animation Studios. Rundong wishes to thank Jin Huang at Zhejiang University for hosting him during the final editing phase of this paper.

REFERENCES

- Ralph Abraham, Jerrold E. Marsden, and Tudor Ratiu. 1988. *Manifolds, Tensor Analysis, and Applications*. Applied Mathematical Sciences, Vol. 75. Springer-Verlag.
- Cherif Amrouche, Christine Bernardi, Monique Dauge, and Vivette Girault. 1998. Vector potentials in three-dimensional non-smooth domains. *Mathematical Methods in the Applied Sciences* 21, 9 (1998), 823–864.
- Douglas Arnold. 2018. *Finite Element Exterior Calculus*. SIAM.
- Douglas N Arnold and Richard S Falk. 1989. A uniformly accurate finite element method for the Reissner–Mindlin plate. *SIAM J. Numer. Anal.* 26, 6 (1989), 1276–1290.

- Harsh Bhatia, Gregory Norgard, Valerio Pascucci, and Peer-Timo Bremer. 2013. The Helmholtz–Hodge decomposition—a survey. *IEEE Trans. Vis. Comp. Graph.* 19, 8 (2013), 1386–1404.
- Harsh Bhatia, Valerio Pascucci, and Peer-Timo Bremer. 2014. The natural Helmholtz–Hodge decomposition for open-boundary flow analysis. *IEEE Trans. Vis. Comp. Graph.* 20, 11 (2014), 1566–1578.
- Alain Bossavit. 2000. Computational electromagnetism and geometry. (5): The ‘Galerkin Hodge’. *J. Japan Soc. Appl. Electromagn. & Mech.* 8, 2 (2000), 203–209.
- Alain Bossavit and Lauri Kettunen. 1999. Yee-like schemes on a tetrahedral mesh, with diagonal lumping. *Int. J. Num. Model. Elec. Net. Dev. Fields* 12, 1–2 (1999), 129–142.
- Jason Cantarella, Dennis DeTurck, and Herman Gluck. 2002. Vector calculus and the topology of domains in 3-space. *American math. monthly* 109, 5 (2002), 409–442.
- Salvatore Caorsi, Paolo Fernandes, and Mirco Raffetto. 2001. On the Convergence of Galerkin Finite Element Approximations of Electromagnetic Eigenproblems. *SIAM J. Numerical Analysis* 38, 2 (2001), 580–607.
- Albert Chern. 2017. *Fluid dynamics with incompressible Schrödinger flow*. Ph.D. Dissertation. California Institute of Technology.
- Fabrice Colin, Richard Egli, and Feng Ying Lin. 2006. Computing a null divergence velocity field using smoothed particle hydrodynamics. *J. Comp. Phys.* 217, 2 (2006), 680–692.
- Keenan Crane, Fernando de Goes, Mathieu Desbrun, and Peter Schröder. 2013. Digital Geometry Processing with Discrete Exterior Calculus. In *SIGGRAPH Course #7*.
- Fernando de Goes, Mathieu Desbrun, Mark Meyer, and Tony DeRose. 2016b. Subdivision Exterior Calculus for Geometry Processing. *ACM Trans. Graph.* 35, 4, Article 133 (2016).
- Fernando de Goes, Mathieu Desbrun, and Yiyi Tong. 2016a. Vector Field Processing on Triangle Meshes. In *SIGGRAPH Course #27*.
- Alan Demlow and Anil N. Hirani. 2014. A posteriori error estimates for finite element exterior calculus: the de Rham complex. *Found. Comput. Math.* 14, 6 (2014), 1337–1371.
- Mathieu Desbrun, Anil N. Hirani, and Jerrold E. Marsden. 2003. Discrete Exterior Calculus for variational problems in computer vision and graphics. In *Inter. Conf. on Decision and Control*, Vol. 5. 4902–4907.
- Mathieu Desbrun, Eva Kanso, and Yiyi Tong. 2008. Discrete differential forms for computational modeling. In *Discrete Differential Geometry*, Alexander I. Bobenko et al. (Ed.). Birkhäuser Basel, 287–324.
- Sharif Elcott, Yiyi Tong, Eva Kanso, Peter Schröder, and Mathieu Desbrun. 2007. Stable, circulation-preserving, simplicial fluids. *ACM Trans. Graph.* 26, 1, Article 4 (2007).
- Kurt O. Friedrichs. 1955. Differential forms on Riemannian manifolds. *Communications on Pure and Applied Mathematics* 8, 4 (1955), 551–590.
- William V. D. Hodge. 1941. *The theory and applications of harmonic integrals*. Cambridge U. Press.
- Christian Lessig. 2012. A primer on differential forms. (2012). arXiv:1206.3323.
- Mamdouh S. Mohamed, Anil N. Hirani, and Ravi Samtaney. 2016. Comparison of discrete Hodge star operators for surfaces. *Comput. Aided Des.* 78 (2016), 118–125.
- Peter B Monk. 1991. A mixed method for approximating Maxwell’s equations. *SIAM J. Numer. Anal.* 28, 6 (1991), 1610–1634.
- Charles B. Morrey. 1956. A Variational Method in the Theory of Harmonic Integrals, II. *American Journal of Mathematics* 78, 1 (1956), 137–170.
- Konstantin Poelke. 2017. *Hodge-type decompositions for piecewise constant vector fields on simplicial surfaces and solids with boundary*. Ph.D. Dissertation. FU Berlin.
- Konstantin Poelke and Konrad Polthier. 2016. Boundary-aware hodge decompositions for piecewise constant vector fields. *Computer-Aided Design* 78 (2016), 126–136.
- Konrad Polthier and Eike Preuß. 2000. Variational approach to vector field decomposition. In *Data Visualization*. Springer, 147–155.
- Faniry Razafindrazaka, Pavlov Yevtushenko, Konstantin Poelke, Konrad Polthier, and Leonid Goubergrits. 2018. Hodge decomposition of wall shear stress vector fields characterizing biological flows. (2018).
- Günter Schwarz. 1995. *Hodge decomposition: a method for solving boundary value problems*. Springer-Verlag.
- Clayton Shonkwiler. 2009. Poincaré duality angles for Riemannian manifolds with boundary. *arXiv preprint arXiv:0909.1967* (2009).
- Jos Stam. 1999. Stable fluids. In *ACM SIGGRAPH proceedings*. 121–128.
- Yiyi Tong, Pierre Alliez, David Cohen-Steiner, and Mathieu Desbrun. 2006. Designing Quadrangulations with Discrete Harmonic Forms. In *Symp. Geo. Proc.* 201–210.
- Yiyi Tong, Santiago Lombeyda, Anil N. Hirani, and Mathieu Desbrun. 2003. Discrete Multiscale Vector Field Decomposition. *ACM Trans. Graph.* 22, 3 (2003), 445–452.
- Joel A. Tropp. 2009. Column Subset Selection, Matrix Factorization, and Eigenvalue Optimization. In *Symp. Disc. Algo.* 978–986.
- Amir Vaxman, Marcel Campen, Olga Diamanti, Daniele Panozzo, David Bommes, Klaus Hildebrandt, and Mirela Ben-Chen. 2016. Directional Field Synthesis, Design, and Processing (STAR). *Computer Graphics Forum* (2016).
- Ke Wang, Weiwei, Yiyi Tong, Mathieu Desbrun, and Peter Schröder. 2006. Edge subdivision schemes and the construction of smooth vector fields. *ACM Trans. Graph.* 25, 3 (2006), 1041–1048.