Journal of Computational Physics 257 (2014) 1352–1372

Contents lists available at ScienceDirect

Journal of Computational Physics

www.elsevier.com/locate/jcp

The chain collocation method: A spectrally accurate calculus of forms

Dzhelil Rufat, Gemma Mason, Patrick Mullen, Mathieu Desbrun *

Applied Geometry Lab, Caltech, MS 305-16, Pasadena, CA 91125, United States

ARTICLE INFO

Article history:
Received 2 July 2012
Received in revised form 30 April 2013
Accepted 4 August 2013
Available online 20 August 2013

Keywords:
Dicrete Exterior Calculus
Chains
Cochains
Spectral accuracy
Hodge star
Wedge product

ABSTRACT

Preserving in the discrete realm the underlying geometric, topological, and algebraic structures at stake in partial differential equations has proven to be a fruitful guiding principle for numerical methods in a variety of fields such as elasticity, electromagnetism, or fluid mechanics. However, structure-preserving methods have traditionally used spaces of piecewise polynomial basis functions for differential forms. Yet, in many problems where solutions are smoothly varying in space, a spectral numerical treatment is called for. In an effort to provide structure-preserving numerical tools with spectral accuracy on logically rectangular grids over periodic or bounded domains, we present a spectral extension of the discrete exterior calculus (DEC), with resulting computational tools extending well-known collocation-based spectral methods. Its efficient implementation using fast Fourier transforms is provided as well.

1. Introduction

Recent years have seen the development of novel discretizations for a wide variety of systems of partial differential equations. In particular, preserving in the discrete realm the underlying geometric, topological, and algebraic structures at stake in differential equations has proven to be a fruitful guiding principle for discretization [1–4]. This geometric approach has led to numerical methods, analyzed in, e.g., [5,3], that inherit a variety of properties from the continuous world, and that surprisingly outperform their known theoretical guarantees [6]. However, geometric discretizations of elasticity, electromagnetism, or fluid mechanics have mostly been demonstrated using spaces of piecewise polynomial differential forms. Many problems where solutions are smoothly varying in space call for a spectral numerical treatment instead, as it produces low-error, exponentially converging approximations by leveraging fast implementations of transforms such as the Fast Fourier Transform. In an effort to provide structure-preserving numerical tools with spectral accuracy on logically rectangular grids over periodic or bounded domains, we present a spectral extension of the discrete exterior calculus described in [7–10]—and point out that the resulting computational tools extends well-known spectral collocation methods.

1.1. Review of previous work

Computational methods preserving geometric structures have become increasingly popular over the past few years, gaining acceptance among both engineers and mathematicians [11]. Computational electromagnetism [7,2], mimetic (or natural) discretizations [12,9], finite-dimensional exterior calculus (including Discrete Exterior Calculus (DEC, [13,8]), and Finite Element Exterior Calculus (FEEC, [1,3])) have all proposed discretizations that preserve vector calculus identities in order to

* Corresponding author.
E-mail address: mathieu@caltech.edu (M. Desbrun).

0021-9991/$ – see front matter
http://dx.doi.org/10.1016/j.jcp.2013.08.011

improve numerics. In particular, the relevance of exterior calculus (Cartan's calculus of differential forms [14]) and algebraic topology [15] to computations came to light.

Exterior calculus is a concise mathematical formalism to express differential and integral equations on smooth and curved spaces, while revealing the geometric structures at play and clarifying the nature of the physical quantities involved. At the heart of exterior calculus is the notion of differential forms, denoting antisymmetric tensors of arbitrary order. As integration of differential forms is an abstraction of the measurement process, this calculus of forms provides an intrinsic, coordinate-free approach particularly relevant to neatly describe a multitude of physical models making heavy use of line, surface and volume integrals [16–20]. Moreover, physical measurements, such as fluxes, represent local integrations over a small surface of the measuring instrument. Pointwise evaluation of such quantities does not have physical meaning; instead, one should manipulate these quantities only as geometrically-meaningful entities integrated over appropriate submanifolds.

Algebraic topology, specifically the notion of chains and cochains [21,15] has been used to provide a natural discretization of differential forms and to emulate exterior calculus on finite grids: a set of values on vertices, edges, faces, and cells are proper discrete versions of respectively pointwise functions, line integrals, surface integrals, and volume integrals [7]. This point of view is entirely compatible with the treatment of volume integrals in finite volume methods, or scalar functions in finite element methods; however, it also involves the "edge elements" and "facet elements" (as first introduced in computational electromagnetism) as special H_{div} and H_{curl} basis elements [22]. Equipped with such discrete forms of arbitrary degree, Stokes' theorem connecting differentiation and integration is automatically enforced if one thinks of differentiation as the dual of the boundary operator—a particularly simple operator on meshes. With these basic building blocks, important structures and invariants of the continuous setting directly carry over to the discrete world, culminating in a discrete Hodge theory [8,3]. As a consequence, such a discrete exterior calculus has already proven useful in many areas such as electromagnetism [7,2], fluid simulation [4,6], (re)meshing of surfaces [23,24], and graph theory [10] to mention a few. So far, only piecewise polynomial basis functions [1,25] have been employed in these applications, thus limiting their computational efficiency in terms of convergence rates.

1.2. Spectral methods

Spectral methods are a class of spatial discretizations of differential equations widely recognized as crucial in fluid mechanics, electromagnetics and other applications where solutions are expected to be smooth. Central to the efficiency of this large family of numerical methods is the fact that the approximation of a periodic C^∞ function by its trigonometric interpolation over evenly spaced points converges faster than any polynomial order of the step size. This is sometimes referred to as "spectral accuracy" or "super-convergence". In practice, spectral accuracy can be achieved for bounded domains through continuation methods [26] or using Gauss–Lobatto quadrature on Legendre or Chebyshev grids [27]. A larger number of spectral methods have been designed, varying in the mesh they consider (primal grids only, or staggered grids [28]), and the locations at which they enforce partial differential equations (PDE). Be it for Galerkin, Petrov–Galerkin, or collocation-based spectral schemes, it has however been noticed that besides constructing spectrally accurate approximations of the relevant fields and their derivatives involved in a PDE, numerically preserving conservation properties helps in obtaining stable and/or physically adequate results [27]. Yet, numerical schemes are often proven conservative a posteriori, as a formal approach to guarantee conservation properties by design remains elusive.

1.3. Motivations and contributions

Despite an increasingly large body of work on numerical approaches based on exterior calculus, developing a spectrally accurate calculus of discrete forms has received very little attention—with a few recent exceptions [29–32] that we will build upon. We present a discrete exterior calculus of differential forms on periodic or bounded domains, including wedge product, Hodge star, and exterior derivative, all of which converge spectrally under grid refinement while utilizing fast Fourier methods to remain computationally efficient. In order to construct a spectral representation of the operators on differential forms, we expand the conventional tools of spectral methods to give spectrally accurate approximations of fields for which integral values over specified domains are known—a process referred to as histopolation [29,25]. In this paper we construct a histopolation using trigonometric polynomials on periodic domains, and consider the extension to bounded domains using a Chebyshev grid, thereby allowing the use of the fast Fourier transform for efficient calculation.

Our work lays out a set of spectral, structure-preserving computational tools with the following distinguishing features:

- We leverage existing work in algebraic topology to discretize space through *chains* (linear combination of mesh elements) and differential forms through cochains (discrete forms). The resulting discrete de Rham complex, that by construction satisfies Stokes' theorem, offers a consistent, "structure-preserving" manipulation of integrals and differentials which respect important conservation laws. This approach, used mainly so far in non-spectral computations [1], was identified in [30,31] as a significant departure in the construction of conservative schemes from traditional spectral methods, since divergences, gradients, and curls are no longer computed through derivation but directly evaluated via metric-independent exterior derivative without having recourse to approximations—thus exactly enforcing the divergence theorem, Green's theorem, etc.

Table 1
Meaning of the basic notations used throughout this document.

Symbol	Meaning
d	Dimension of the domain
Ω	Spatial domain in $\mathbb{R}^d$
K	Computational grid of Ω
$\widetilde{K}$	Dual grid to K
σ^k	Primal k-dim element of grid K
$\tilde{\sigma}^k$	Dual k-dim element of grid $\widetilde{K}$
$*$	Duality operator ($*\sigma^k = \tilde{\sigma}^{n-k}$)
∂	Boundary operator on mesh elements
Λ^k	Space of differential k-forms ω^k over Ω
$\bar{\Lambda}^k$	Space of discrete k-forms $\bar{\omega}^k$ on K
D	Discrete exterior derivative operator **d**
W	Discrete wedge operator $\wedge$
H	Discrete Hodge star operator $\star$
$\mathcal{F}$	Discrete Fourier transform
$\hat{\cdot}$	Quantities expressed in Fourier space

- We extend the basic DEC operators to ensure spectral accuracy. Using a framework with some similarities to that of Bochev [9] and DEC [8], we construct operations on discrete differential forms by interpolating their finite-dimensional (point-sampled or integrated) representation, performing exterior calculus operations in continuous space, and point-sampling or integrating the results to project back to their finite-dimensional representation. We differ from previous spectral DEC approaches [32,30,31] in that we provide a natural Hodge star leveraging mesh duality, introduce a wedge product, satisfy consistency conditions between the various operators and basis functions, and provide their explicit expressions to facilitate implementation.
- Our spectrally accurate operators result in a *spectral chain collocation method* to solving partial differential equations: a differential equation is not enforced at a series of points, but at a series of mesh elements (points, edges, faces, etc.), extending point collocation methods to respect the natural geometric structure of the continuous equation. We also provide a discussion on the implementation of this spectral computational framework, and show results on Poisson problems with various boundary conditions.
- Finally, our approach is versatile: since it builds on the notion of integrated values (cochains to be precise), one can use the same computational tools on curved grids with only minor changes.

For simplicity of presentation, we will first treat regular and Chebyshev one-dimensional (1D) grids before extending our approach to grids of arbitrary dimension via tensor products. Other Chebyshev-like grids are easy to derive as well, since integrals of differential forms (the building blocks of our approach) are unchanged by pushforward or pullback. For the reader's convenience, Table 1 lists the main notations we will use throughout this paper.

2. Discrete differential forms and operators

Differential forms were pioneered by Cartan [14] in an effort to provide a unified approach to defining differentiation and integration over curves, surfaces, volumes, and higher-dimensional manifolds. Since then, forms have been shown crucial in elucidating the structures and invariants in physics [19]. Since most of our measurements of the world are of integral nature, forms have been also at the core of a number of numerical approaches targeting the definition of numerical counterparts of differential operators that are compatible with the geometric and topological structures underlying well-posed partial differential equations.

Most computational methods based on exterior calculus use the notion of simplicial k-cochain as a fundamental object that assigns a number to each simplex of dimension k of the simplicial complex. Cochains are, in a sense, a natural discretization of differential forms of degree k, which are objects to be integrated over k-dimensional domains to return a number; in fact, de Rham's theorem states that the (discretization) map from the de Rham complex to the simplicial cochain complex induces an isomorphism on cohomology. Conversely, one can go from cochains back to differential forms using Whitney forms [21,7] or higher-order piecewise polynomial finite element spaces [1,25]. These mimetic properties of finite-dimensional approximations to differential forms have led to the term *discrete differential forms* to describe cochains [8].

In this section, we review the principles and approaches used in discrete versions of exterior calculus in preparation for extending them to spectral computations. By construction, the calculus of discrete differential forms automatically preserves a number of important geometric structures, including integration by parts (with a proper treatment of boundaries), Stokes' theorem, the de Rham complex, Poincaré duality, Poincaré's lemma, and Hodge theory. Therefore, it provides a suitable foundation for a coordinate-free discretization of geometric field theories and associated computations, with applications including electromagnetism [2], incompressible Euler and complex fluids [4,6], and meshing algorithms [23,33]. The particular "flavor" of discrete differential forms and operators we will be leveraging is known as *discrete exterior calculus*, or *DEC*

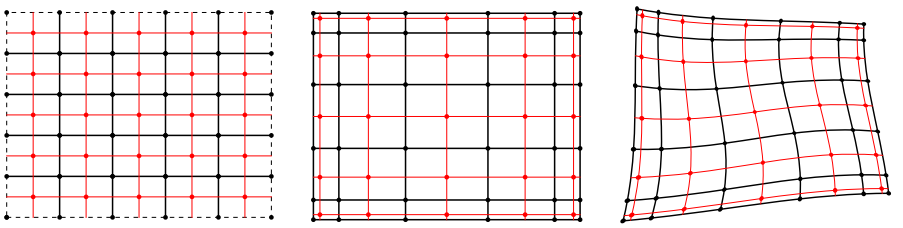

(a) *Periodic grid* (b) *Chebyshev grid* (c) *Logically rectangular grid*

Fig. 1. This paper discusses the implementation of a spectral exterior calculus on logically rectangular grids, such as (a) regular grids over periodic domains, (b) Chebyshev grids over bounded domains, or (c) an arbitrarily mapped rectangular grid. Primal elements are displayed in black, while their duals are in red. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

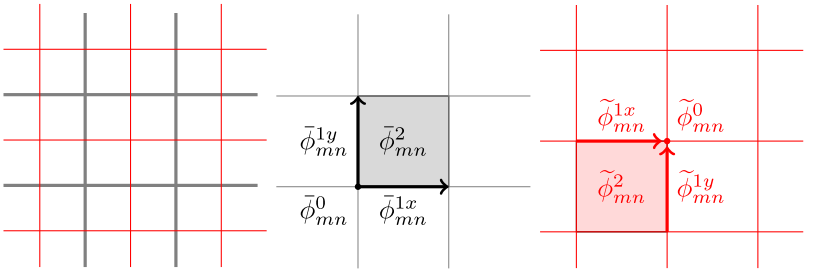

(a) *Primal/Dual grid* (b) *Primal grid* (c) *Dual grid*

Fig. 2. Illustration of a regular 2D grid (in black) along with its dual (in red), along with our mesh orientation and index convention. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

for short; see [34,35]. (For related work in this direction, see also [36] and [1]; DEC most significantly differs from these other discrete calculus approaches in the way the Hodge duality is defined through the use of a dual mesh.) We refer the reader to [8] for a brief introduction to DEC.

2.1. Mesh and dual mesh as spatial discretization

In our approach, the continuous domain (usually a smooth d-dimensional manifold Ω) is approximated by a mesh—more precisely, by a cell complex that is manifold, admits a metric, and is orientable. We will generally denote the complex by K, and a cell in the complex by σ. Superscripts will be used at times to denote the dimension of a cell; i.e., σ^0 represents a vertex of K, σ^1 represents an edge, etc. Moreover, we restrict our study to the case of logically rectangular grids over periodic and bounded domains, that include regular grids and Chebyshev grids in $\mathbb{R}^n$ (see Fig. 1); that is, each k-cell for $k \geqslant 1$ has $2k$ non-degenerate faces, and the mapping function between the rectangular grid of a rectangular domain and the actual spatial grid is assumed to be bijective and C^∞. Consequently, the mesh topology is easily encoded using indices as indicated in Fig. 2; we will also assume prescribed orientations as the figure indicates as well.

From a mesh K, one can construct a *dual mesh*—where duality is to be understood in the sense of graph duality. The dual mesh is defined via the location of its dual vertices, whereas its connectivity is directly induced by the primal connectivity. Once a primal mesh and its dual are defined, one can define the *duality operator* $*$ that maps each k-cell σ^k of the primal mesh K to a dual $(n-k)$-cell $*\sigma^k$ in its dual cell complex $*K$. Note that a refined definition of the dual mesh for bounded domains (where dual cells at the boundary are restricted ("clipped") to K to allow proper enforcement of boundary conditions) is assumed as well (see, for instance, [37,2]).

The canonical placement of primal and dual nodes for Chebyshev grids will be given in Section 3.2; but arbitrarily distorted grids will be straightforward to handle (see Section 6.1), owning to the fact that the exterior derivative and the pullback operator commute.

2.2. Discretization of differential forms

The fundamental objects of DEC are discrete differential forms. A discrete k-form $\bar{\omega}^k$, the discrete counterpart of a continuous k-form ω^k, assigns a real number $\bar{\omega}^k_i$ to each oriented k-dimensional cell σ^k_i in the mesh K—and zero for all other mesh elements. (The superscripts k are not actually required by the notation, but they are often useful as reminders of what dimension of form or cell we are dealing with.) The value of the form on σ^k_i is sometimes denoted by $\langle \omega^k, \sigma^k_i \rangle$, and represents the value of the continuous form ω^k integrated over the mesh element σ^k_i, i.e.,

$$\bar{\omega}_i^k \equiv \langle \omega^k, \sigma^k \rangle = \int_{\sigma^k} \omega^k.$$

For implementation purposes, discrete k-form can therefore be thought of as an array of values, one per k-dimensional cell in the mesh K. For example, discrete 0-forms sample values of an associated continuous 0-form at vertices, discrete 1-forms sample values of a continuous 1-form on all edges, etc. Note that we can even integrate these discrete forms over discrete paths by linearity: simply add the form's values on each cell in the path, taking care to flip the sign if the path is oriented opposite the cell. Formally, these "paths" of k-dimensional elements are called *chains*, and discrete differential forms are *cochains*, where $\langle \cdot, \cdot \rangle$ is the pairing between cochains and chains [15].

Discrete differential forms can be defined either on the mesh K or on its dual $*K$. We will refer to these as *primal forms* and *dual forms* respectively, and denote the space of primal discrete forms by $\bar{\Lambda}$ and the space of dual discrete forms by $\tilde{\Lambda}$. Note that there is a natural correspondence between primal k-forms $\bar{\omega}^k$ and dual $(n-k)$-forms $\tilde{\omega}^{n-k}$, corresponding to the mesh duality discussed in Section 2.1, for which each primal k-cell σ^k has an associated dual $(n-k)$-cell $*\sigma^k$ (Fig. 2). This important property will be used below to define the discrete Hodge star operator.

2.3. Reduction and reconstruction maps

The reduction and reconstruction maps provide a way to go back and forth between continuous forms and their discrete realizations.

Reduction. The reduction map (also called the de Rham map) is a linear operator $\mathcal{P}$ that projects a continuous form to its discrete realization on the grid through integration over mesh elements:

$$\begin{aligned} \mathcal{P}\colon \quad & \Lambda^k \to \bar{\Lambda}^k \\ & \omega^k \to \bar{\omega}^k \quad \text{with } \bar{\omega}_i^k = \left(\mathcal{P}\omega^k\right)_i \equiv \int_{\sigma_i^k} \omega^k. \end{aligned}$$

We will denote by $\widetilde{\mathcal{P}}$ the analogous operator mapping continuous forms to their *dual* discrete counterparts in a similar fashion:

$$\begin{aligned} \widetilde{\mathcal{P}}\colon \quad & \Lambda^k \to \tilde{\Lambda}^k \\ & \omega^k \to \tilde{\omega}^k \quad \text{with } \tilde{\omega}_i^k = \left(\widetilde{\mathcal{P}}\omega^k\right)_i \equiv \int_{\tilde{\sigma}_i^k} \omega^k. \end{aligned}$$

Note that this definition of reduction extends the notion of *point sampling*: while the reduction of a 0-form is found by simply point-sampling its value at each vertex of the grid, the reduction of a general k-form is its evaluation (i.e., integral) on all the k-dimensional elements (vertices for $k=0$, edges for $k=1$, faces for $k=2$, etc.) of the grid.

Reconstruction. Conversely, the reconstruction map ($\mathcal{R}$) is a map which reconstructs a continuous k-form from its discrete realization by interpolation for $k=0$, and by histopolation otherwise:

$$\begin{aligned} \mathcal{R}\colon \quad & \bar{\Lambda}^k \to \Lambda^k \\ & \bar{\omega}^k \to \omega^k. \end{aligned}$$

We will denote by $\widetilde{\mathcal{R}}$ the analogous operator mapping *dual* discrete forms to continuous forms in a similar fashion. Note that we will sometimes omit the dual sign $\tilde{\cdot}$ for clarity, as which reconstruction operator is meant is unambiguously implied by the (primal or dual) nature of the discrete form it is applied on.

If one has a set of basis functions $\{\phi_0^k(x), \phi_1^k(x), \ldots\}$ for k-forms that satisfy the property

$$\forall i, j, \quad \int_{\sigma_i^k} \phi_j^k = \delta_{ij}, \tag{1}$$

then, $\mathcal{R}$ can trivially be defined as

$$\omega^k = \mathcal{R}\bar{\omega}^k \equiv \sum_i \bar{\omega}_i^k \phi_i^k.$$

One can readily verify that this reduction map is a left inverse of the reconstruction map:

$$\mathcal{P}\mathcal{R} = \mathrm{Id}.$$

However, the converse is not true; the reconstruction map is only approximately the right inverse of the reduction map, with equality in the limit when the mesh element size h approaches 0:

$$\|\omega - \mathcal{R}\mathcal{P}\omega\| \xrightarrow[h\to 0]{} 0,$$

with a rate of convergence determined by the chosen norm on forms and the degree of the basis functions. While Whitney first introduced a one-sided inverse of the de Rham map using what amounts to piecewise linear basis functions [21], we will instead use global basis functions satisfying Eq. (1) to provide spectrally accurate reconstructions.

2.4. Operators on discrete forms

We now review the basic operators typically used in exterior calculus, along with their common discretizations. In particular, we list the properties that will naturally carry through our spectrally accurate discretization.

2.4.1. Exterior derivative $\mathbf{d}$

The continuous exterior derivative, $\mathbf{d} : \Lambda^k(\Omega) \to \Lambda^{k+1}(\Omega)$ extends the notion of derivative to differential forms, turning a k-form ω^k to a $(k+1)$-form $\mathbf{d}\omega^k$. This metric-independent operator satisfies linearity $\mathbf{d}(c_1\omega_1 + c_2\omega_2) = c_1\,\mathbf{d}\omega_1 + c_2\,\mathbf{d}\omega_2$, nilpotency $\mathbf{d}^2 = 0$, and locality $\mathbf{d}(\omega|_U) = (\mathbf{d}\omega)|_U$. The exterior derivative encompasses the usual linear operators in vector calculus (gradient, curl, and divergence); its nilpotency reflects the classical identities $\nabla\cdot(\nabla\times\vec{\mathbf{v}}) = 0$ and $\nabla\times(\nabla f) = 0$. Moreover, a number of theorems (the fundamental theorem of calculus, as well as Green's, divergence, and Gradient theorems) can be seen as special cases of the general *Stokes' theorem* which states that $\mathbf{d}$ and the boundary operator ∂ are dual operators, in the sense that:

$$\int_\sigma \mathbf{d}\omega = \int_{\partial\sigma} \omega.$$

Fortunately, one can use the notion of chains and cochains from algebraic topology to provide discrete realization of this derivative operator [21,15,7,8] that captures all its defining properties. Given its purely topological (metric-independent) definition, no special care is necessary to adapt the traditional discrete exterior derivative to our spectral framework (see Section 4.1).

2.4.2. Wedge product

The continuous wedge product is also a metric-independent operator, assembling a $(k+l)$-form from a k-form and an l-form; it hence generalizes the notion of cross product to differential forms. The wedge operator $\wedge : \Lambda^k(\Omega)\times\Lambda^l(\Omega) \to \Lambda^{k+l}(\Omega)$ is associative (i.e., $\alpha\wedge(\beta\wedge\gamma) = (\alpha\wedge\beta)\wedge\gamma$), bilinear ($(a\alpha + b\beta)\wedge\gamma = a(\alpha\wedge\beta) + b(\beta\wedge\gamma)$ and $\alpha\wedge(b\beta + c\gamma) = b(\alpha\wedge\beta) + c(\alpha\wedge\gamma)$), and anti-commutative ($\alpha^k\wedge\beta^l = (-1)^{kl}\beta^l\wedge\alpha^k$). The wedge product also satisfies a Leibniz rule with the exterior derivative:

$$\mathbf{d}(\alpha^k\wedge\beta^l) = \mathbf{d}\alpha^k\wedge\beta^l + (-1)^k\alpha^k\wedge\mathbf{d}\beta^l. \tag{2}$$

Various discretizations of this operator have been proposed over the years, from primal–dual versions [34], to metric-independent definitions [37,9]. While most properties of the wedge product are retained by these discrete counterparts, associativity is only satisfied for closed forms, or in the limit of mesh refinement [38].

2.4.3. Hodge star

The continuous Hodge star operator, denoted by $\star$, is a metric-dependent operator that maps k-forms on a d-dimensional manifold to $(d-k)$-forms:

$$\star: \quad \Lambda^k(\Omega) \to \Lambda^{d-k}(\Omega).$$

The Hodge star defines a dual in that if it is applied twice, the result is the identity up to a sign; more precisely, in a d-dimensional Euclidean space:

$$\star^k\star^{d-k} = (-1)^{k(d-k)}. \tag{3}$$

In dimension two, the Hodge star operator in coordinates satisfies:

$$\star 1 = \mathbf{d}x\wedge\mathbf{d}y, \qquad \star\mathbf{d}x = -\mathbf{d}y, \qquad \star\mathbf{d}y = \mathbf{d}x, \qquad \star(\mathbf{d}x\wedge\mathbf{d}y) = 1.$$

Along with the wedge product, the Hodge star defines a natural inner product $\langle .,. \rangle$ on k-forms:

$$\langle\alpha^k, \beta^k\rangle = \int_\Omega \alpha^k\wedge\star\beta^k = \int_\Omega (\alpha^k, \beta^k)\mu^d \tag{4}$$

where μ^d is the volume d-form.

The different methods to discretize the Hodge star operator fall mostly in two categories. A majority of approaches start from a discretization of the inner product of forms (Eq. (4)), then derive a discrete Hodge star from it. This leads to what is called the "*derived* discrete $\star$" in [9]. A few authors, instead, use a direct (called "*natural*" in [9]) definition of the Hodge star, which exploit the existence of primal/dual structures on meshes; i.e., the discrete Hodge star $\mathbf{H}$ now maps discrete *primal* forms to discrete *dual* forms, and vice versa through the inverse of the Hodge star. Discretizations of the Hodge star typically rely on local reconstructions and integration, the most simple one being what is often referred to as the *diagonal Hodge star* [7], corresponding to a piecewise constant reconstruction. Its implementation involves each primal value being multiplied by the ratio of the measure of the primal and dual mesh elements to obtain the dual value. It is hence represented by a diagonal matrix. Higher-order Galerkin Hodge stars have also been considered, resulting in better accuracy at the cost of denser matrices. As we will see, by choosing proper basis functions, we will be able to construct a spectrally accurate *natural* discrete Hodge star that is computationally efficient, as the star will become a Toeplitz matrix in the periodic case.

2.4.4. Other operators

There are important additional operators, including contraction by vector field, and Lie derivative (extending directional derivatives to differential forms). While these operators play a crucial role in certain applications such as computational fluids and other advection problems, their extension to our spectral framework brings additional difficulty due to their dependence on a vector field. They can, however, be handled in our framework, but *only through the use of coordinates*, whereas in the continuous world the application of a Lie derivative on a differential form is a coordinate-free metric-independent operation. A first step towards a geometric discretization of contractions and Lie derivatives was proposed in [39]; but a spectrally accurate discretization preserving key properties of these other operators is left for future work.

3. Basic spectral tools

Before delving into the design of spectrally accurate discrete operators, we must define a series of basic tools and conventions which will be particularly useful for our task. We start by defining proper periodic basis functions on regular grids, before describing the case of Chebyshev grids.

3.1. 1D periodic interpolator & histopolator functions

To build our spectral wedge and Hodge star operators to work on discrete forms of arbitrary degree, we will need not only spectral interpolating basis functions, but also spectral *histopolating* basis functions, i.e., basis functions which integrate to one over assigned intervals [30]. To this end, we consider a one-dimensional periodic domain of width 2π with N *regularly-spaced nodes*, and define over this canonical domain two scalar functions α_N and β_N—that we will respectively call *interpolator* and *histopolator*—as follows:

$$\alpha_N(x) = \frac{1}{N} \begin{cases} \cot\frac{x}{2} \sin\frac{Nx}{2} & \text{if } N \text{ even}, \\ \csc\frac{x}{2} \sin\frac{Nx}{2} & \text{if } N \text{ odd}, \end{cases} \tag{5}$$

and

$$\beta_N(x) = \begin{cases} \frac{1}{2\pi} - \frac{1}{4}\cos\frac{Nx}{2} + \frac{1}{N}\sum_{n=1}^{N/2} \frac{n\cos nx}{\sin\frac{n\pi}{N}} & \text{if } N \text{ even}, \\ \frac{1}{2\pi} + \frac{1}{N}\sum_{n=1}^{(N-1)/2} \frac{n\cos nx}{\sin\frac{n\pi}{N}} & \text{if } N \text{ odd}. \end{cases} \tag{6}$$

For convenience, let us denote any translation of the above functions using the notation below

$$\alpha_{N,n}(x) = \alpha_N(x - hn) \quad \text{and} \quad \beta_{N,n}(x) = \beta_N(x - hn)$$

where $h = 2\pi/N$ is the mesh's edge element size and the nodes are enumerated by $x_n = nh$. For any given number N of nodes, these two functions satisfy the following important properties mentioned in Eq. (1) (where δ_{mn} refers to the Kronecker delta):

$$\alpha_{N,n}(x_m) = \delta_{mn} \quad \text{and} \quad \int_{x_m - \frac{h}{2}}^{x_m + \frac{h}{2}} \beta_{N,n}(x)\, dx = \delta_{mn}.$$

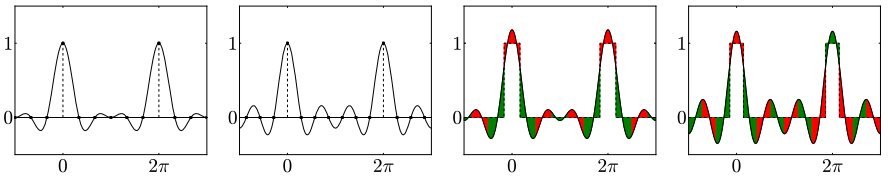

(a) $\alpha_6(x)$ (b) $\alpha_7(x)$ (c) $h\beta_6(x)$ (d) $h\beta_7(x)$

Fig. 3. Examples of periodic interpolators for $N = 6$ (a) and 7 (b), and corresponding periodic histopolators (c) and (d), scaled by $h = 2\pi/N$ for clarity. While the interpolator α_N satisfies $\alpha_N(nh \bmod 2\pi) = \delta_{0n}\ \forall n$, the histopolator β_N integrates to 1 over the dual cell straddling $x = 0$, and to 0 over other dual cells in the range $[0, 2\pi]$. Note that the alternating red and green colors are used to mark out dual cells, and to illustrate that the integral of β_N over each of these dual cells sums to zero or one. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

In other words, α_N provides a smooth interpolation of a discrete function with 1 at node x_0 and 0 at every other nodes, while β_N integrates to 1 over the interval $[x_0 - \frac{h}{2}, x_0 + \frac{h}{2}]$, and to 0 over all other intervals (see Fig. 3). Notice that these functions are the only Fourier series with N sinusoidal components satisfying the pointwise (resp., interval-wise) constraints. Note also that α_N and β_N are related to each other:

$$\int_{x-\frac{h}{2}}^{x+\frac{h}{2}} \beta_N'(\xi)\, d\xi = \alpha_N(x), \quad \text{or equivalently,} \quad \beta_N\left(x + \frac{h}{2}\right) - \beta_N\left(x - \frac{h}{2}\right) = \alpha_N'(x).$$

Next, we show that these two functions provide building blocks to construct basis functions for arbitrary forms on regular, periodic grids—from which we will derive basis functions for bounded domains via pushforward.

3.2. Spectral basis functions on regular grids

Basis functions for periodic grids in arbitrary dimensions can be easily built through tensor products of (translated) interpolators α_N and histopolators β_N, where the number of β_N used in the tensor product is equal to the degree of the form. To make this point clear, we will now introduce our notation for 1-, 2-, and 3-dimensional basis functions, with the extension to higher dimensions being straightforward. We will denote the basis functions as $\phi^{p[\text{comp}]}_{[\text{ind}]}$ with superscripts used to indicate the degree p, followed by the component of the form when appropriate (e.g., x, xy, etc.), and subscripts used for grid indices; for instance, $\phi^{1y}_{mnk}(.)$ is the function of $\mathbb{R}^3$ representing the $\mathbf{d}y$-component of the 1-form basis, located on the edge parallel to the y axis indexed by (m, n, k). See Fig. 2 for an illustration of the 2D notations.

Primal and dual nodes. On a 1D regular grid the N primal nodes are equally spaced, and the dual nodes are placed in the middle between two adjacent primal nodes. Consequently, the primal nodes are: $x_n = 2\pi n/N$, and the dual nodes are at: $\tilde{x}_n = 2\pi(n + \frac{1}{2})/N$.

Basis functions for forms in 1D. Because of their interpolation (resp., histopolation) property, translated versions of the functions α_N and β_N can directly be used as basis functions for 0- and 1-forms respectively as follows:

$$\phi^0_{N,n}(x) = \alpha_{N,n}(x) \quad \text{and} \quad \phi^1_{N,n}(x) = \beta_{N,n+\frac{1}{2}}(x)\, \mathbf{d}x,$$

where $n \in \{0, 1, \ldots, N-1\}$. Indeed, a discrete 0-form $\bar{\omega}^0$ (resp., a discrete 1-form $\bar{\omega}^1$) can be reconstructed as a smooth form through $\omega^0 = \sum_i \bar{\omega}^0_i \phi^0_i$ (resp., $\omega^1 = \sum_i \bar{\omega}^1_i \phi^1_i$); the reconstructed form then satisfies $\mathcal{P}\omega^0 = \bar{\omega}^0$ (resp., $\mathcal{P}\omega^1 = \bar{\omega}^1$). Similarly, dual basis functions are easily designed as well through:

$$\tilde{\phi}^0_{N,n}(x) = \alpha_{N,n+\frac{1}{2}}(x) \quad \text{and} \quad \tilde{\phi}^1_{N,n}(x) = \beta_{N,n}(x)\, \mathbf{d}x.$$

Basis functions for forms in 2D. In two dimensions, tensor products of the one-dimensional bases provide basis functions for 0-, 1-, and 2-forms on a regular $M \times N$ grid. These functions are expressed as follows:

$$\begin{aligned}
\phi^0_{MN,mn}(x, y) &= \phi^0_{M,m}(x) \wedge \phi^0_{N,n}(y),\\
\phi^{1x}_{MN,mn}(x, y) &= \phi^1_{M,m}(x) \wedge \phi^0_{N,n}(y),\\
\phi^{1y}_{MN,mn}(x, y) &= \phi^0_{M,m}(x) \wedge \phi^1_{N,n}(y),\\
\phi^2_{MN,mn}(x, y) &= \phi^1_{M,m}(x) \wedge \phi^1_{N,n}(y).
\end{aligned}$$

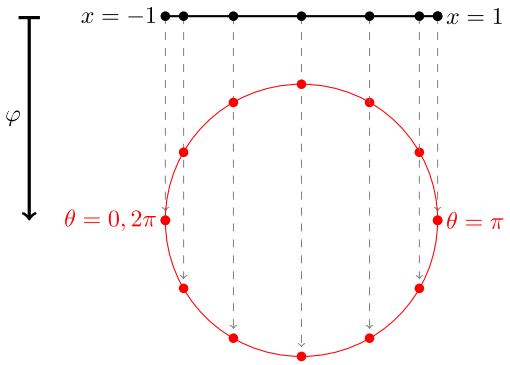

Fig. 4. The map φ mapping the canonical interval $[-1, 1]$ to the bottom unit hemicircle ($0 \leqslant \theta \leqslant \pi$).

One can easily check that these functions are 1 on their associated degree of freedom and 0 on all others (vertices for 0-forms, edges for 1-forms, and faces for 2-forms), thus offering a proper set of bases for smooth reconstructions of discrete forms. Formulas for the dual basis functions are strictly analogous, where $\tilde{\phi}$ is used in lieu of ϕ.

Basis functions for forms in 3D. The same construction can be used in three or higher dimensions. For completeness, we describe the three-dimensional basis functions for primal forms:

$$\begin{aligned}
\phi^{0}_{MNK,mnk}(x,y,z) &= \phi^{0}_{M,m}(x) \wedge \phi^{0}_{N,n}(y) \wedge \phi^{0}_{K,k}(z),\\
\phi^{1x}_{MNK,mnk}(x,y,z) &= \phi^{1}_{M,m}(x) \wedge \phi^{0}_{N,n}(y) \wedge \phi^{0}_{K,k}(z),\\
\phi^{1y}_{MNK,mnk}(x,y,z) &= \phi^{0}_{M,m}(x) \wedge \phi^{1}_{N,n}(y) \wedge \phi^{0}_{K,k}(z),\\
\phi^{1z}_{MNK,mnk}(x,y,z) &= \phi^{0}_{M,m}(x) \wedge \phi^{0}_{N,n}(y) \wedge \phi^{1}_{K,k}(z),\\
\phi^{2xy}_{MNK,mnk}(x,y,z) &= \phi^{1}_{M,m}(x) \wedge \phi^{1}_{N,n}(y) \wedge \phi^{0}_{K,k}(z),\\
\phi^{2xz}_{MNK,mnk}(x,y,z) &= \phi^{1}_{M,m}(x) \wedge \phi^{0}_{N,n}(y) \wedge \phi^{1}_{K,k}(z),\\
\phi^{2yz}_{MNK,mnk}(x,y,z) &= \phi^{0}_{M,m}(x) \wedge \phi^{1}_{N,n}(y) \wedge \phi^{1}_{K,k}(z),\\
\phi^{3}_{MNK,mnk}(x,y,z) &= \phi^{1}_{M,m}(x) \wedge \phi^{1}_{N,n}(y) \wedge \phi^{1}_{K,k}(z).
\end{aligned}$$

Note that since the wedge products above are the continuous (as opposed to discrete) ones, associativity holds, and no ambiguity is introduced by the omission of parentheses. Here again, it is easy to check that these functions provide smooth reconstructions from values of vertices, edges, faces, or volumes on a regular grid indexed by m, n, and k, in a periodic domain. Formulas for the dual basis functions are strictly analogous, where $\tilde{\phi}$ is used in lieu of ϕ.

3.3. Chebyshev grids over bounded domains

For bounded domains, a popular choice of spatial discretization in spectral methods is the use of Chebyshev computational grids [27]. It is well known that Chebyshev polynomials can be derived as the pullback of Fourier basis functions from a circle onto its diameter; see Fig. 4. We can, in fact, perform the same pullback of our regular-grid basis functions of forms to obtain new spectral bases of forms applicable for Chebyshev grids over non-periodic domains.

Denote by φ the map from the interval between -1 and 1 on the real line (with Cartesian coordinate x) to the unit semi-circle (with polar angle θ, see Fig. 4):

$$\begin{aligned}
\varphi\colon \quad [-1,1] \subset \mathbb{R} &\to [0,\pi] \subset \mathcal{S}^1\\
x &\mapsto \theta = \operatorname{acos}(-x).
\end{aligned}$$

Primal and dual nodes. Let the 1D grid consist of N points (including the boundaries) in the interval $[-1, 1]$. The corresponding unit circle will then have $2N - 2$ points (see Fig. 4). The primal nodes $\{x_n\}$ $(n = 0, \ldots, N-1)$ and dual nodes $\{\tilde{x}_n\}$ $(n = 0, \ldots, N-2)$ of the Chebyshev grid are thus

$$x_n = -\cos\frac{n\pi}{N-1} \quad \text{and} \quad \tilde{x}_n = -\cos\frac{(n+\frac{1}{2})\pi}{N-1}.$$

Basis functions. To design our basis functions, we first define functions on the semi-circle ($\theta \in [0, \pi]$) by mirroring/anti-mirroring the regular basis functions ϕ on the whole circle (see [40] for the usual case of primal 0-forms), in order to satisfy the interpolation/histopolation properties (Eq. (1)) on the semi-circle. The resulting functions κ are expressed as:

$$\kappa^0_{N,n} = \begin{cases} \phi^0_{2N-2,n}, & n = 0 \text{ or } n = N-1 \text{ (endpoints)}, \\ \phi^0_{2N-2,n} + \phi^0_{2N-2,2N-2-n}, & n \in \{1, \dots, N-2\} \text{ (midpoints)} \end{cases}$$

for primal 0-forms,

$$\kappa^1_{N,n} = \phi^1_{2N-2,n} - \phi^1_{2N-2,2N-3-n}, \quad n \in \{0, \dots, N-2\},$$

for primal 1-forms,

$$\tilde{\kappa}^0_{N,n} = \tilde{\phi}^0_{2N-2,n} + \tilde{\phi}^0_{2N-2,2N-3-n}, \quad n \in \{0, \dots, N-2\},$$

for dual 0-forms, and

$$\tilde{\kappa}^1_{N,n}(\theta) = \begin{cases} \delta_N(\theta), & n = 0, \\ [\tilde{\phi}^0_{2N-2,n} - \tilde{\phi}^0_{2N-2,2N-2-n} - \rho_{N,n}](\theta), & n \in \{1, \dots, N-2\}, \\ \delta_N(\pi - \theta), & n = N-1, \end{cases}$$

for dual 1-forms, where

$$\delta_N(\theta) = \left((N-1)^2 \alpha_{2N-2,0}(\theta) + \frac{1}{2} \cos\big((N-1)\theta\big) \right) \sin\theta \, \mathbf{d}\theta, \quad \text{and}$$

$$\rho_{N,n}(\theta) = 2\big(\gamma_{2N-2,n}\delta_N(\theta) + \gamma_{2N-2,N-n-1}\delta_N(\pi - \theta)\big)\, \mathbf{d}\theta, \quad \text{with:}$$

$$\gamma_{N,n} = \int_0^{\frac{\pi}{N}} \beta_{N,n}(\theta)\, d\theta$$

$$= \begin{cases} \frac{1-(-1)^k}{2N} + \frac{1}{N} \sum_{n=1}^{N/2} \frac{\sin(2kn\pi/N) - \sin((2k-1)n\pi/N)}{\sin\frac{n\pi}{N}} & \text{if } N \text{ even}, \\ \frac{1}{2N} + \frac{1}{N} \sum_{n=1}^{(N-1)/2} \frac{\sin(2kn\pi/N) - \sin((2k-1)n\pi/N)}{\sin\frac{n\pi}{N}} & \text{if } N \text{ odd}. \end{cases}$$

The function δ is used to deal with the special case of the two boundary dual (half-)edges, and the function ρ adds contributions to intermediate basis functions so that they integrate to zero at both boundary dual edges. Finally, we pull back the functions κ by φ to obtain the form basis functions ψ on the Chebyshev 1D grid ($\psi = \varphi^* \kappa$):

$$\psi^0_{N,n}(x) = \kappa^0_{N,n}\big(\operatorname{acos}(-x)\big),$$

$$\psi^1_{N,n}(x) = \kappa^1_{N,n}\big(\operatorname{acos}(-x)\big) \frac{\mathbf{d}x}{\sqrt{1-x^2}},$$

$$\tilde{\psi}^0_{N,n}(x) = \tilde{\kappa}^0_{N,n}\big(\operatorname{acos}(-x)\big),$$

$$\tilde{\psi}^1_{N,n}(x) = \tilde{\kappa}^1_{N,n}\big(\operatorname{acos}(-x)\big) \frac{\mathbf{d}x}{\sqrt{1-x^2}}.$$

One can easily check that the functions above satisfy the property in Eq. (1) required for basis functions, as the functions κ were designed to satisfy these properties, and the pullback φ^* commutes with integration.

Note that the basis functions for primal zero-forms turn out (unsurprisingly) to be the Lagrange polynomials of order N, i.e.,

$$\psi^0_{N,n}(x) = \prod_{\substack{m=0 \\ m \neq n}}^{N-1} \frac{x - x_m}{x_n - x_m}$$

where $\{x_n\}_{n=0,\dots,N-1}$ represent the coordinates of the primal points. All other basis functions of forms are *also* polynomials for any choice of N. Examples for the primal and dual basis functions for 0- and 1-forms are provided in Figs. 5 and 6 for $N = 7$.

Finally, basis functions in higher dimensions for arbitrary forms are derived using tensor products of these two primal and two dual 1D basis functions, just as we explained in Section 3.2.

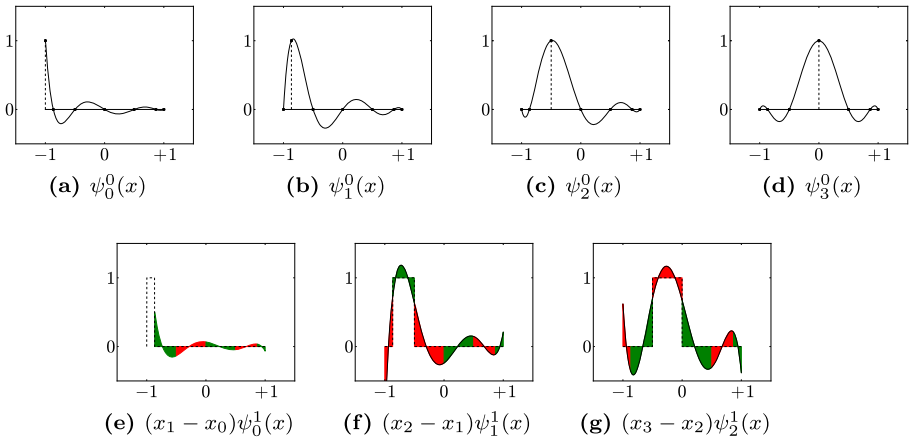

Fig. 5. Chebyshev primal basis functions for a grid with $N = 7$. We normalize the one-form basis functions by $x_n - x_{n-1}$ to have approximately the same scale in our visualizations.

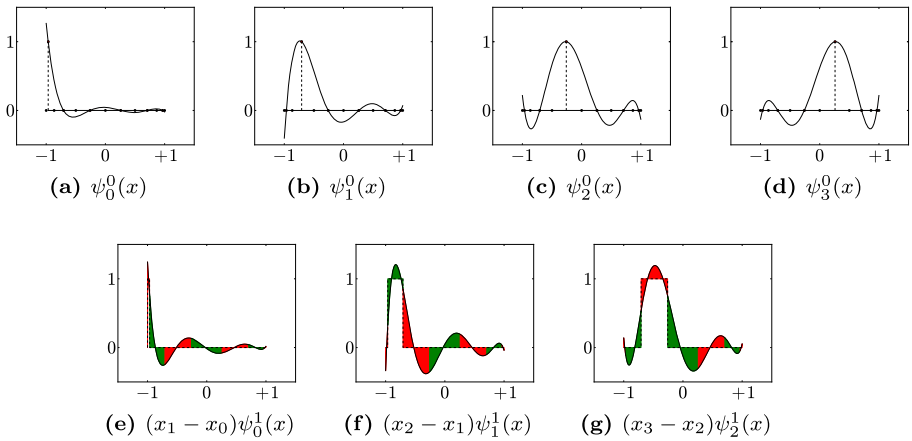

Fig. 6. Chebyshev dual basis functions for a grid with $N = 7$.

4. Spectrally accurate discrete operators

Equipped with these basis functions on regular and Chebyshev grids, we can now derive the discrete, spectrally accurate version of the exterior derivative $\mathbf{d}$, the wedge product $\wedge$, and the Hodge star $\star$. Emphasis is placed on the mathematical definitions; fast evaluations of these operators will be discussed subsequently in Section 5.

4.1. Discrete exterior derivative $\mathbf{D}$

The most common discrete realization of the exterior derivative based on algebraic topology [15] using chains and cochains is a linear operator $\mathbf{D}$:

$$\mathbf{D}: \quad \bar{\Lambda}^k \to \bar{\Lambda}^{k+1},$$

which is the (signed) incidence matrix between $(k+1)$ elements and k elements of the grid, with a sign determined by the relative orientation of the elements. In other words,

$$\boxed{\mathbf{D} = \partial^t,}$$

where ∂ refers to the *boundary operator* acting on chains (see [15,8]). Note that the operator $\mathbf{D}$ is thus implemented via a sparse matrix whose non-zero elements are only $+1$ and -1 values to indicate incidence between mesh cells. It also satisfies $\mathbf{D}^2 = 0$ like its continuous equivalent (since the boundary of a boundary is always the empty set), and that a *discrete Stokes' theorem* is also enforced on chains since the very definition of $\mathbf{D}$ is tantamount to enforcing

$$\int_\sigma \mathbf{d}\omega = \int_{\partial\sigma} \omega$$

for every mesh element σ. This discrete realization of the exterior derivative is exact, in the sense that the operator $\mathbf{D}$ commutes with the reduction operator:

$$\mathbf{d}\mathcal{P} = \mathcal{P}\mathbf{D}. \tag{7}$$

This implies that the following diagram fully commutes:

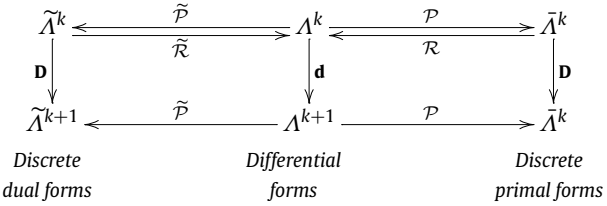

Therefore, as mentioned in Section 2.4.1, the classic discrete exterior derivative used in mimetic methods, DEC, or finite-dimensional exterior calculus needs no special treatment to obtain spectral accuracy. However, its use for spectral computations represents a clear departure from the conventional spectral methods [30], since all operators of classical vector calculus (divergence, gradient, and curl) can be achieved exactly via $\mathbf{D}$ [31]. In our framework, we also use the exterior derivative on dual forms:

$$\widetilde{\mathbf{D}}: \quad \widetilde{\Lambda}^k \to \widetilde{\Lambda}^{k+1}.$$

Its implementation and properties are no different from its primal version, since the adjacency of the dual mesh is directly derived from the adjacency on the primal:

$$\boxed{\widetilde{\mathbf{D}} = \tilde{\boldsymbol{\partial}}^t.}$$

We now turn to the definition of a discrete wedge product and discrete Hodge star operator.

4.2. Discrete wedge product $\mathbf{W}$

Various discrete definitions of the wedge product have been proposed, sometimes mixing primal and dual elements [34]. We instead follow Bochev and Hyman's treatment [9]: by utilizing the de Rham and reconstruction maps, one can define a discrete wedge product $\mathbf{W}(\bar{\alpha}, \bar{\beta})$ in a general way that applies to any pair of discrete (primal or dual) forms $\bar{\alpha}$ and $\bar{\beta}$ through:

$$\boxed{\mathbf{W}(\bar{\alpha}, \bar{\beta}) = \mathcal{P}(\mathcal{R}\bar{\alpha} \wedge \mathcal{R}\bar{\beta}).} \tag{8}$$

That is, we first reconstruct the two discrete forms with our trigonometric or Chebyshev basis functions, we then apply the wedge product to the two continuous forms that we have created, and finally, integrate (or, in the case of 0-forms, point-sample) the result to get the final discrete form $\mathbf{W}(\bar{\alpha}, \bar{\beta})$. The resulting operator satisfies a discrete Leibniz rule, equivalent to Eq. (2):

$$\mathbf{D}\mathbf{W}(\bar{\alpha}, \bar{\beta}) = \mathbf{W}\big(\mathbf{D}(\bar{\alpha}), \bar{\beta}\big) + (-1)^k \mathbf{W}\big(\bar{\alpha}, \mathbf{D}(\bar{\beta})\big),$$

where $\bar{\alpha}$ is a discrete k-form, and $\bar{\beta}$ is an arbitrary discrete form. Indeed, because the exterior derivative commutes with the reduction and reconstruction maps, this discrete Leibniz rule can be derived directly from the continuous Leibniz rule. Our spectrally accurate definition of the discrete wedge product is bilinear and anti-commutative, but it is *not* associative. It is, however, associative in the limit as the mesh size approaches zero. That is, the associator $\mathbf{W}(\bar{\alpha}, \mathbf{W}(\bar{\beta}, \bar{\gamma})) - \mathbf{W}(\mathbf{W}(\bar{\alpha}, \bar{\beta}), \bar{\gamma})$ (i.e., the measure of nonassociativity) will approach zero exponentially fast as the step size is reduced. The difficulty of creating an associative discrete wedge product has been previously noted by, e.g., Kotiuga [38], who refers to it as the "commutative cochain problem" and discusses some deeper topological reasons behind it. The spectral accuracy of our wedge operator will mitigate this lack of associativity exponentially fast under mesh refinement.

Note that the discrete wedge product of a (primal or dual) p-form and a (primal or dual) q-form can also be seen as a three tensor

$$\mathbf{W}_{ijk}^{p,q} = \int_{\sigma_i^{p+k}} \phi_j^p \wedge \phi_k^q,$$

since we can write the value of the wedge product on element σ_i^{p+q} as:

$$\big[\mathbf{W}(\bar{\alpha}, \bar{\beta})\big]_i = \sum_{j,k} \mathbf{W}_{ijk} \bar{\alpha}_j \bar{\beta}_k.$$

4.3. Discrete hodge star **H**

As explained earlier, our discrete Hodge star **H** exploits the notion of mesh duality in that the discrete Hodge star of a primal k-form is a dual $(d-k)$-form, and vice versa. The discrete Hodge star for a discrete form is realized conceptually by first reconstructing the continuous form with our primal (resp., dual) spectral bases, applying the continuous Hodge star to this form, and then projecting this form back to the dual (resp., primal) grid. In our notation, this can be written as

$$\mathbf{H}^k = \widetilde{\mathcal{P}} \star^k \mathcal{R}, \tag{9}$$

and is easily understood from the following diagram:

$$\begin{array}{ccc} \Lambda^k & \underset{\mathcal{R}}{\overset{\mathcal{P}}{\rightleftarrows}} & \bar{\Lambda}^k \\ \downarrow{\scriptstyle \star^k} & & \vdots{\scriptstyle \mathbf{H}^k} \\ \Lambda^{n-k} & \xrightarrow{\widetilde{\mathcal{P}}} & \widetilde{\Lambda}^{n-k} \end{array}$$

The operator **H** can be built as a matrix with easily precomputed entries. For instance, for k-forms, this matrix contains the terms

$$\mathbf{H}^k_{ij} = \int_{\tilde{\sigma}_i^{n-k}} \star \phi_j^k. \tag{10}$$

Similarly, the dual Hodge star operator is given by:

$$\widetilde{\mathbf{H}}^k_{ij} = \int_{\sigma_i^{n-k}} \star \tilde{\phi}_j^k. \tag{11}$$

Notice that if we require our discrete Hodge stars to satisfy the discrete equivalent of Eq. (3)

$$\mathbf{H}^k \widetilde{\mathbf{H}}^{n-k} = \widetilde{\mathbf{H}}^{n-k} \mathbf{H}^k = (-1)^{k(n-k)} \mathbf{Id}, \tag{12}$$

this imposes a constraint on dual basis functions. Indeed, for Eq. (12) to hold true, they must be a linear combination of the primal basis functions:

$$\tilde{\phi}_i^k = \sum_j \left[\left(\mathbf{H}^{n-k} \right)^{-1} \right]_{ji} \star \phi_j^{n-k}. \tag{13}$$

Both the regular (ϕ) and Chebyshev (ψ) basis functions that we defined above satisfy this constraint.

Finally, our Hodge matrices are dense; however, each matrix is *circulant* for a regular grid on periodic domains (due to the invariance of the grid under translation), and *centrosymmetric* for a Chebyshev grid (due to mirror symmetry around its center). While circulant matrices are special cases of Toeplitz matrices, for which multiplications by vectors can be done $\mathcal{O}(N \log N)$, we will show that *all* discrete Hodge stars for arbitrary dimensions can be efficiently implemented in Fourier space in $\mathcal{O}(N \log N)$ time complexity where N is the total number of spatial points.

5. Implementation in Fourier space

We now describe how to efficiently implement the operators described in the last section. We will see that the Hodge star and wedge operators can have very fast implementations as they can be expressed as the product of fast Fourier transforms and sparse matrices. On periodic, regular grids of size N^d, the basic DEC operators can be expressed in terms of the discrete equivalents of two elementary operators, namely a convolution with a box function over $[a, b]$ (denoted $\mathbf{I}^{a,b}$) and a translation by a (denoted $\mathbf{T}^a$). Both of these operations are expressed as diagonal matrices in Fourier space. The wedge product and Hodge star for Chebyshev grids require additional sparse operators that we describe next.

5.1. Sparse matrices for periodic 1D domains

We define our convolution and translation operators to act on an array of function values $\bar{f}_i$ at evenly spaced points x_i on a periodic grid by first constructing a trigonometric interpolation $\mathcal{R}\bar{f}$ as described in Section 2.3, then by transforming it as follows:

$$\mathbf{I}^{a,b} : \bar{f}_i \mapsto \int_{x_i+a}^{x_i+b} (\mathcal{R}\bar{f})(\xi)\, d\xi \quad \text{(integration)}, \tag{14}$$

$$\mathbf{T}^a : \bar{f}_i \mapsto (\mathcal{R}\bar{f})(x_i + a) \quad \text{(translation)}. \tag{15}$$

These two operators are implemented through Discrete Fourier Transforms as follows:

$$\mathbf{I}^{a,b} = \mathcal{F}^{-1}\hat{\mathbf{I}}^{a,b}\mathcal{F}, \qquad \mathbf{T}^{a,b} = \mathcal{F}^{-1}\widehat{\mathbf{T}}^{a,b}\mathcal{F}.$$

The operator $\mathcal{F}$ represents the shifted Discrete Fourier Transform and $\mathcal{F}^{-1}$ represents its inverse: rather than placing the zero-frequency term (the mean of the signal) first, it is shifted to the middle with the negative Nyquist frequency becoming the first element, and the positive Nyquist frequency becoming the last.

In Fourier space (denoted using a hat $\hat{\cdot}$), each operator is simply encoded by a diagonal matrix with the following coefficients:

$$\hat{\mathbf{I}}^{a,b}_{kk} = \frac{e^{ikb} - e^{ika}}{ik} \quad \text{and} \quad \widehat{\mathbf{T}}^{a}_{kk} = e^{ika},$$

where the special case $k = 0$ is treated as:

$$\lim_{k\to 0} \frac{e^{ikb} - e^{ika}}{ik} = b - a.$$

Not that since these two matrices are diagonal, the order in which they are applied does not matter.

5.2. Sparse matrices for 1D Chebyshev grids

For conciseness in later expressions, we also define a number of sparse matrices representing simple linear operators that will be particularly useful for the Hodge stars and wedge products on Chebyshev grids.

Mirroring. The mirroring matrices will allow us to map quantities from the Chebyshev grid to the unit circle and back (see Fig. 4). We will need four versions of mirroring: a mirroring $\mathbf{M}_0^+$ and anti-mirroring $\mathbf{M}_0^-$ operator for 0-forms mapping N nodal values of the Chebyshev grid to $2N-2$ values on the circle; and similarly, a mirroring $\mathbf{M}_1^+$ and anti-mirroring $\mathbf{M}_1^-$ operator mapping N nodal values of the Chebyshev grid to $2N$ values on the circle will be useful for 1-forms. Their expressions are:

$$\mathbf{M}_0^{\pm}\colon \begin{bmatrix} x_0 \\ x_1 \\ \vdots \\ x_{N-2} \\ x_{N-1} \end{bmatrix} \mapsto \begin{bmatrix} x_0 \\ x_1 \\ \vdots \\ x_{N-2} \\ x_{N-1} \\ \pm x_{N-2} \\ \vdots \\ \pm x_1 \end{bmatrix}, \qquad \mathbf{M}_1^{\pm}\colon \begin{bmatrix} x_0 \\ x_1 \\ \vdots \\ x_{N-2} \\ x_{N-1} \end{bmatrix} \mapsto \begin{bmatrix} x_0 \\ x_1 \\ \vdots \\ x_{N-2} \\ x_{N-1} \\ \pm x_{N-1} \\ \pm x_{N-2} \\ \vdots \\ \pm x_1 \\ \pm x_0 \end{bmatrix}.$$

The left inverses of the above operators, satisfying $\mathbf{M}_0^{\dagger}\mathbf{M}_0^{\pm} = \mathbf{Id}$ and $\mathbf{M}_1^{\dagger}\mathbf{M}_1^{\pm} = \mathbf{Id}$, are then given by

$$\mathbf{M}_0^{\dagger}\colon \begin{bmatrix} x_0 \\ x_1 \\ \vdots \\ x_{N-1} \\ x_N \\ \vdots \\ x_{2N-1} \end{bmatrix} \mapsto \begin{bmatrix} x_0 \\ x_1 \\ \vdots \\ x_{N-2} \\ x_{N-1} \end{bmatrix}, \qquad \mathbf{M}_1^{\dagger}\colon \begin{bmatrix} x_0 \\ x_1 \\ \vdots \\ x_{N-1} \\ x_N \\ \vdots \\ x_{2N} \end{bmatrix} \mapsto \begin{bmatrix} x_0 \\ x_1 \\ \vdots \\ x_{N-2} \\ x_{N-1} \end{bmatrix}.$$

Extensions. Extension matrices will allow us to resample forms at higher or lower resolutions, an operation needed in the wedge product as we will mention in Section 5.3. We define the upsampling extension matrix $\widehat{\mathbf{E}}^n$ that will act in Fourier space to add $2n$ samples as follows:

$$\forall n > 0, \quad \widehat{\mathbf{E}}^n: \quad \begin{bmatrix} x_0 \\ x_1 \\ \vdots \\ x_{N-2} \\ x_{N-1} \end{bmatrix} \mapsto \frac{N+2n}{N} \begin{bmatrix} \left.\begin{matrix} 0 \\ \vdots \\ 0 \end{matrix}\right\} n \\ x_0 \\ x_1 \\ \vdots \\ x_{N-2} \\ x_{N-1} \\ \left.\begin{matrix} 0 \\ \vdots \\ 0 \end{matrix}\right\} n \end{bmatrix}.$$

The left inverse of this extension matrix is a downsampling matrix $\widehat{\mathbf{E}}^{n\dagger}$ such that $\widehat{\mathbf{E}}^{n\dagger}\widehat{\mathbf{E}}^n = \mathbf{Id}$, and is expressed as:

$$\widehat{\mathbf{E}}^{n\dagger}: \quad \begin{bmatrix} \left.\begin{matrix} x_{-n} \\ \vdots \\ x_{-1} \end{matrix}\right\} n \\ x_0 \\ \vdots \\ x_{N-1} \\ \left.\begin{matrix} x_N \\ \vdots \\ x_{N+n-1} \end{matrix}\right\} n \end{bmatrix} \mapsto \frac{N-2n}{N} \begin{bmatrix} \left.\begin{matrix} x_0 + x_N \\ \vdots \\ x_{n-1} + x_{N+n-1} \end{matrix}\right\} n \\ x_n \\ \vdots \\ x_{N-n-1} \\ \left.\begin{matrix} x_{N-n} + x_{-n} \\ \vdots \\ x_{N-1} + x_{-1} \end{matrix}\right\} n \end{bmatrix}.$$

With these extension matrices, we can upsample a discrete form $\bar{\omega}$ from an N-point grid to an $(N+2n)$-point grid by first performing a Fourier transform $\mathcal{F}_N$, then applying $\widehat{\mathbf{E}}^n$, and applying an inverse Fourier transform $\mathcal{F}_{N+2n}^{-1}$ (see inset for an illustration of $\mathbf{E}^2$ and its left inverse):

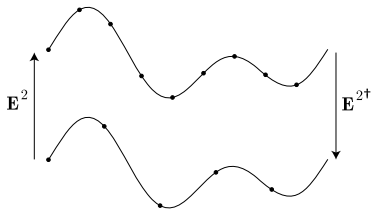

$$\mathbf{E}^n \bar{\omega} = \mathcal{F}_{N+2n}^{-1} \widehat{\mathbf{E}}^n \mathcal{F}_N \bar{\omega}.$$

Downsampling is done similarly through $\mathbf{E}^{n\dagger}$, with now $\mathcal{F}_N^{-1}$ and $\mathcal{F}_{N+2n}$.

Scaling. The metric-dependent matrix Ω^N is a diagonal matrix which has, as its diagonal elements, the weight by which each nodal value needs to be scaled by in order for the integration over its corresponding dual region on the unit circle to faithfully represent integration over its corresponding dual region on the original Chebyshev grid. Unlike previous matrices in this section, Ω^N is metric dependent in the sense that it does not only depend on the grid topology, but on the volume change between the grid on the circle and the Chebyshev grid:

$$\boldsymbol{\Omega}_{kk}^N = \sin\left(\frac{2\pi}{N}\left(k+\frac{1}{2}\right)\right) = \sqrt{1-\tilde{x}_k^2}.$$

Other matrices. Finally, we will need simple, additional matrices denoted $\mathbf{A}$, $\mathbf{B}$, $\mathbf{B}^\dagger$, and $\mathbf{C}$; they respectively lump every consecutive pairs of non-boundary dual edge values, evaluate the nodal values of the 1-form basis functions of the left and right boundary dual edges (i.e., $\delta_N(\theta)$ and $\delta_N(\pi-\theta)$ in Section 3.3), and remove the end node values. Their expressions are thus:

$$\mathbf{A} = \underbrace{\begin{pmatrix} 1 & & & & & & \\ & 1 & 1 & & & & \\ & & 1 & 1 & & & \\ & & & \ddots & \ddots & & \\ & & & & 1 & 1 & \\ & & & & & 1 & 1 \\ & & & & & & 1 \end{pmatrix}}_{N\times(2N-2)}, \qquad \mathbf{C} = \underbrace{\begin{pmatrix} 0 & & & & & \\ & 1 & & & & \\ & & 1 & & & \\ & & & \ddots & & \\ & & & & 1 & \\ & & & & & 0 \end{pmatrix}}_{N\times N},$$

$$\mathbf{B} = \underbrace{\begin{pmatrix} (N-1)^2+1/2 & & \\ -1/2 & & \\ 1/2 & & \\ -1/2 & \mathbf{0} & \\ 1/2 & & \\ \vdots & & \\ (-1)^{N+1}/2 & & \end{pmatrix}}_{N\times N}, \qquad \mathbf{B}^\dagger = \underbrace{\begin{pmatrix} & & (-1)^{N+1}/2 \\ & & \vdots \\ & & 1/2 \\ & \mathbf{0} & -1/2 \\ & & 1/2 \\ & & -1/2 \\ & & (N-1)^2+1/2 \end{pmatrix}}_{N\times N}.$$

5.3. *Wedge product* **W**

Let $\odot$ denote the element-wise multiplication (typically used in Fourier space) such that, for two vectors $\mathbf{a}$ and $\mathbf{b}$, $\mathbf{a}\odot\mathbf{b}$ is the vector with components

$$(\mathbf{a}\odot\mathbf{b})_i = \mathbf{a}_i\mathbf{b}_i.$$

Remember that we are (conceptually) working with bandlimited functions $f(x)$ whose frequencies are below the Nyquist rate for the spatial grid resolution; i.e., with functions that can be reconstructed exactly from their sample values $\{a_i = a(x_i)\}$ on the grid. However, when the element-wise product $\odot$ is applied to such functions, the result is no longer necessarily below the Nyquist rate, as the highest frequency can potentially be doubled by the multiplication. In order to compute wedge products *exactly within the Nyquist range*, we will have recourse to a temporary *upsampling* of the operands to a higher grid resolution, perform the multiplication at the higher resolution, and *downsample* back to the original grid resolution.

1D wedge. Based on our definition in Eq. (8), a wedge product of two 0-forms in 1D is expressed as

$$\bar{\omega}^0 = \mathbf{W}(\bar{\alpha}^0, \bar{\beta}^0) = \bar{\alpha}^0 \odot \bar{\beta}^0.$$

Note that there is no reason to use the up/downsampling operators here: while we obviously cannot capture the wedge product content past the Nyquist rate, all the frequencies below are exactly evaluated, thus leading to spectral accuracy. However, the wedge product of a 0-form and a 1-form differs, in the sense that to avoid losing accuracy, we must use the following definition:

$$\bar{\omega}^1 = \mathbf{W}(\bar{\alpha}^0, \bar{\beta}^1) = \mathbf{E}^{N\dagger}\mathbf{J}\big((\mathbf{E}^N\bar{\alpha}^0)\odot(\mathbf{E}^N\mathbf{J}^{-1}\bar{\beta}^1)\big),$$

where $\mathbf{J}$ maps from 0-forms to 1-forms (on the periodic domain) through:

$$\mathbf{J} = \mathbf{I}^{-\frac{h}{2},\frac{h}{2}}\mathbf{T}^{\frac{h}{2}}.$$

Indeed, $\mathbf{J}$ being an integration, the wedge without $\mathbf{E}^N$ would lose accuracy at the Nyquist rate. Should dual forms be used, they could easily be converted to primal forms via the translation operator $\mathbf{T}$ before the wedge product is applied. Therefore, wedge products in 1D can always be implemented with efficient evaluation through discrete fast Fourier transforms, both for periodic and Chebyshev grids since their expressions hold in each case.

2D and beyond. The various wedge products between forms can be expressed in terms of the following functions:

$$\mathbf{w}^{00}(\mathbf{a},\mathbf{b}) = \mathbf{a}\odot\mathbf{b},$$
$$\mathbf{w}^{01}\left(\mathbf{a}, \begin{pmatrix}\mathbf{b}_x\\ \mathbf{b}_y\end{pmatrix}\right) = \begin{pmatrix}\mathbf{a}\odot\mathbf{b}_x\\ \mathbf{a}\odot\mathbf{b}_y\end{pmatrix},$$
$$\mathbf{w}^{11}\left(\begin{pmatrix}\mathbf{a}_x\\ \mathbf{a}_y\end{pmatrix}, \begin{pmatrix}\mathbf{b}_x\\ \mathbf{b}_y\end{pmatrix}\right) = \mathbf{a}_x\odot\mathbf{b}_y - \mathbf{a}_y\odot\mathbf{b}_x$$

and similarly for higher dimensions.

Now, let:

$$\mathbf{J_x} \equiv \mathbf{J} \otimes \mathbf{Id}, \qquad \mathbf{J_y} \equiv \mathbf{Id} \otimes \mathbf{J}, \qquad \mathbf{J_{xy}} \equiv \mathbf{J} \otimes \mathbf{J},$$
$$\mathbf{E} \equiv \mathbf{E}^N \otimes \mathbf{E}^N, \qquad \mathbf{E}^\dagger \equiv \mathbf{E}^{N\dagger} \otimes \mathbf{E}^{N\dagger}.$$

Then we can define the wedge product in 2D for all possible pairs of forms as:

$$\mathbf{W}(\bar{\alpha}^0, \bar{\beta}^0) = \mathbf{w}^{00}(\bar{\alpha}^0, \bar{\beta}^0),$$
$$\mathbf{W}\left(\bar{\alpha}^0, \begin{pmatrix}\bar{\beta}^1_x \\ \bar{\beta}^1_y\end{pmatrix}\right) = \begin{pmatrix}\mathbf{E}^\dagger\mathbf{J_x} & \\ & \mathbf{E}^\dagger\mathbf{J_y}\end{pmatrix} \mathbf{w}^{01}\left(\mathbf{E}\bar{\alpha}^0, \begin{pmatrix}\mathbf{EJ_x}^{-1}\bar{\beta}^1_x \\ \mathbf{EJ_y}^{-1}\bar{\beta}^1_y\end{pmatrix}\right),$$
$$\mathbf{W}\left(\begin{pmatrix}\bar{\alpha}^1_x \\ \bar{\alpha}^1_y\end{pmatrix}, \begin{pmatrix}\bar{\beta}^1_x \\ \bar{\beta}^1_y\end{pmatrix}\right) = \mathbf{E}^\dagger\mathbf{J_{xy}}\mathbf{w}^{11}\left(\begin{pmatrix}\mathbf{EJ_x}^{-1}\bar{\alpha}^1_x \\ \mathbf{EJ_y}^{-1}\bar{\alpha}^1_y\end{pmatrix}, \begin{pmatrix}\mathbf{EJ_x}^{-1}\bar{\beta}^1_x \\ \mathbf{EJ_y}^{-1}\bar{\beta}^1_y\end{pmatrix}\right),$$
$$\mathbf{W}(\bar{\alpha}^0, \bar{\beta}^2) = \mathbf{E}^\dagger\mathbf{J_{xy}}\mathbf{w}^{00}(\mathbf{E}\bar{\alpha}^0, \mathbf{EJ_{xy}}^{-1}\bar{\beta}^2).$$

Here again, the use of up/downsampling is purely to ensure no loss of accuracy for frequencies at or below the Nyquist rate. Extension of the wedge product to dual forms or in higher dimensions is straightforward as the same principles apply. While the above derivation was for periodic domains, the wedge product is preserved under pullback; so the equations hold for any topologically equivalent space. For Chebyshev grids, $\mathbf{J}$ is no longer square, and there are only $N-1$ edges—but the formulas are basically unchanged.

5.4. Hodge star $\mathbf{H}$

We can also express our discrete Hodge stars in a way that lends itself to fast evaluations.

Periodic grid in 1D. The Hodge star for a periodic grid is very simple, as it is simply expressed as:

$$\mathbf{H}^0 = \widetilde{\mathbf{H}}^0 = \mathbf{I}^{-\frac{h}{2},\frac{h}{2}}, \qquad \mathbf{H}^1 = \widetilde{\mathbf{H}}^1 = \mathbf{I}^{-\frac{h}{2},\frac{h}{2}\,-1}. \tag{16}$$

Chebyshev grid in 1D. To define the Hodge star operator on a Chebyshev grid, we make use of the matrices we defined earlier in Section 5.2. In particular, we get:

$$\widetilde{\mathbf{H}}^0 = \mathbf{M}_1^\dagger\mathbf{I}^{-\frac{h}{2},\frac{h}{2}}\boldsymbol{\Omega}\mathbf{M}_1^+,$$
$$\mathbf{H}^1 = \mathbf{M}_1^\dagger\boldsymbol{\Omega}^{-1}\mathbf{I}^{-\frac{h}{2},\frac{h}{2}\,-1}\mathbf{M}_1^-.$$

The expressions for the other two stars are slightly more complicated because of the special treatment needed for the dual (half-)edges at the two ends of the domain:

$$\mathbf{H}^0 = \mathbf{A}\mathbf{M}_0^\dagger\mathbf{I}^{0,+\frac{h}{2}}\boldsymbol{\Omega}\mathbf{E}^{N-1}\mathbf{M}_0^+,$$
$$\widetilde{\mathbf{H}}^1 = \mathbf{M}_0^\dagger\big(\mathbf{T}^{-\frac{h}{2}}\boldsymbol{\Omega}^{-1}\mathbf{T}^{\frac{h}{2}} - \mathbf{B}\mathbf{I}^{0,\frac{h}{2}} - \mathbf{B}^\dagger\mathbf{I}^{-\frac{h}{2},0}\big)\mathbf{I}^{-\frac{h}{2},\frac{h}{2}\,-1}\mathbf{M}_0^-\mathbf{C} + \mathbf{B} + \mathbf{B}^\dagger.$$

2D and beyond. The extension of the fast 1D implementation to arbitrary dimension is rather direct: we construct Hodge stars on higher-dimensional grids by taking the Cartesian products of stars on one-dimensional grids. Using 2D for illustration purposes, we first introduce the following operators:

$$\mathbf{H_x} \equiv \mathbf{H} \otimes \mathbf{Id}, \qquad \mathbf{H_y} \equiv \mathbf{Id} \otimes \mathbf{H},$$

where $\mathbf{H_x}$ acts along the first index (x-direction), and $\mathbf{H_y}$ acts along the second index (y-direction). Then from these two 1D Hodge stars, the 2D Hodge star operators can be expressed as:

$$\mathbf{H^0} = \mathbf{H_x^0}\mathbf{H_y^0},$$
$$\mathbf{H^1} = \begin{pmatrix}\mathbf{0} & -\mathbf{H_x^0}\mathbf{H_y^1} \\ \mathbf{H_x^1}\mathbf{H_y^0} & \mathbf{0}\end{pmatrix},$$
$$\mathbf{H^2} = \mathbf{H_x^1}\mathbf{H_y^1}.$$

Equivalent equations hold for the dual Hodge stars. Note that we have $\widetilde{\mathbf{H}}^\mathbf{1}\mathbf{H}^\mathbf{1} = -\mathbf{Id}$ as desired.

For completeness, we also provide the 2D exterior derivative in terms of the Cartesian products of the 1D $\mathbf{D_x}$ and $\mathbf{D_y}$ operators:

$$\mathbf{D^0}\bar{\alpha}^0 = \begin{pmatrix}\mathbf{D_x}\bar{\alpha}^0 \\ \mathbf{D_y}\bar{\alpha}^0\end{pmatrix}, \qquad \mathbf{D^1}\begin{pmatrix}\bar{\alpha}^1_x \\ \bar{\alpha}^1_y\end{pmatrix} = -\mathbf{D_y}\bar{\alpha}^1_x + \mathbf{D_x}\bar{\alpha}^1_y.$$

6. Numerical tests

We now discuss how our spectral operators can be used for practical computations. In particular, we present numerical tests involving Poisson equations on 0- and 1-forms in 1D and 2D.

6.1. Solving differential equations

Our spectral framework can be used in various ways, either as is, or with minor modifications.

Chain collocation. Our computational framework lends itself naturally to a spectral method that we call *Chain Collocation Method*. Just as the point collocation method enforces a partial differential equation on all grid nodes, we now enforce partial differential equations on all appropriate chains. If the unknown is a scalar function, then the PDE is enforced at every node as usual; but for a k-form, the PDE will now be enforced on each k-chain. Since any differential equation can be rewritten with exterior forms, our operators can be used as is on *discrete* approximations of these forms, and the conversion of the continuous operators into discrete operators combining $\mathbf{D}, \mathbf{H}$, and $\mathbf{W}$ will enforce the differential equation on each appropriate mesh element. For instance, a 2D Poisson equation $\Delta \mathbf{v} = \mathbf{f}$ for a *vector field* $\mathbf{v}$ can be rewritten with, for instance, a 1-form ω as $\Delta\omega = \gamma$ where ω is now the circulation of the continuous field $\mathbf{v}$ on all edges, while γ is the circulation of $\mathbf{f}$—leading to a coordinate-free setup, drastically different from the usual coordinate-based treatment. Modulo boundary conditions, a direct discretization of the Laplace operator $\Delta = \mathbf{d} \star \mathbf{d} \star + \star \mathbf{d} \star \mathbf{d}$ will now enforce that a discrete form $\bar{\omega}$ satisfies the PDE in a discrete sense, i.e., on each edge. (Various 1D and 2D examples of this approach are shown at the end of this section.) Because of the algebraic topology properties of our framework, conservative differential equations will lead to conservative discrete equations, and structure-preservation will be enforced.

Other numerical approaches. Our framework can be easily altered to accommodate other numerical solves. First, different grids can be used if necessary, as we spelled out the constraints that one needs to enforce to get proper discretizations. Second, one can modify our Hodge stars by deriving them from the inner product instead: this is the approach spelled out in [32,31]. While the Hodge star operators are no longer converting primal to dual forms (and vice versa), the advantage of this approach is that the resulting operators inherit the symmetry of the inner product.

Logically rectangular grids. Finally, modifying our approach for arbitrary logically rectangular computational grids is simple. First, both $\mathbf{D}$ and $\mathbf{W}$ are unchanged since the integrals of forms are unchanged by the pullback φ^* and our numerical treatment of these operators is metric independent. The Hodge star operator needs, however, to be modified. In fact, only the matrices Ω_N require modification: one simply has to rescale the diagonal elements of Ω_N by how much the orthogonal Chebyshev grid is deformed into the computational grid. In this sense, our spectral framework is reminiscent of the regular DEC framework, as the diagonal Hodge stars are also subject to be changed based on the volume forms of the computational domains [8].

6.2. Convergence tests

In order to provide numerical evidence of spectral behavior, we tested our operators on a variety of examples. We show below a series of Poisson equations with Dirichlet and Neumann conditions in 1D and 2D.

1D Poisson. The traditional Poisson equation for a scalar function f reads:

$$\nabla^2 f = q,$$

which can be rewritten in exterior calculus as an equation on a 0-form f:

$$\star \mathbf{d} \star \mathbf{d} f = q.$$

Using a periodic domain (where a regular grid is placed at $[0, 2\pi]$) or a bounded domain (chosen to be $[-1, +1]$, discretized by a Chebyshev grid), we can directly discretize this last equation. The only choice we are left with is to decide whether to think of f as a primal ($\bar{f}$) or a dual ($\tilde{f}$) 0-form. Thus, the 1D Poisson equation is implemented as either:

$$\widetilde{\mathbf{H}}^1 \widetilde{\mathbf{D}}^{\mathbf{1}} \mathbf{H}^1 \mathbf{D}^{\mathbf{0}} \bar{f}^0 = \bar{q}^0,$$

or

$$\mathbf{H}^1 \mathbf{D}^{\mathbf{1}} \widetilde{\mathbf{H}}^1 \widetilde{\mathbf{D}}^{\mathbf{0}} \tilde{f}^0 = \tilde{q}^0.$$

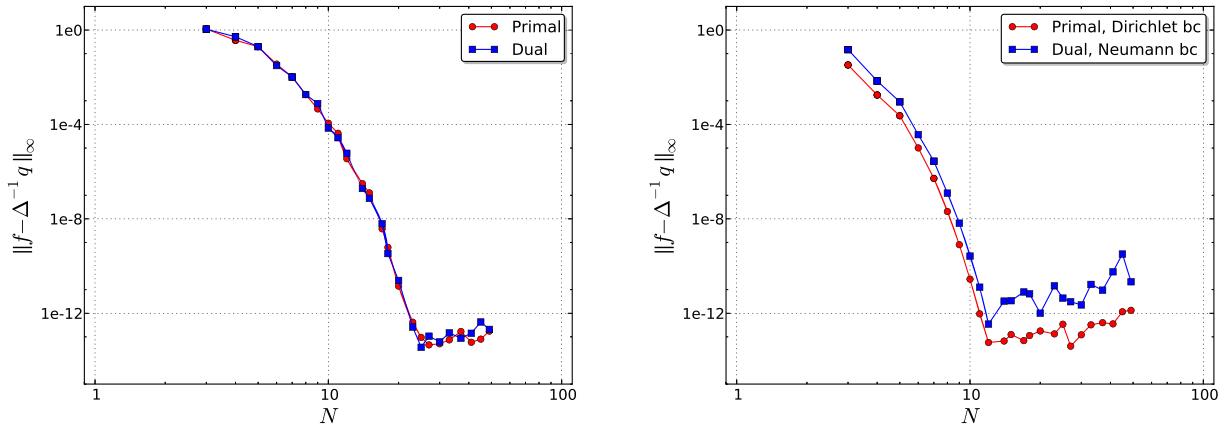

(a) *Periodic Domain* **(b)** *Bounded Domain (Chebyshev)*

Fig. 7. Convergence graphs for a 1D Poisson equation: (a) we solve $\Delta f = e^{\sin x}(\cos^2 x - \sin x)$ on a periodic domain for either a primal, or dual 0-form f; (b) we solve $\Delta f = e^x$ on a Chebyshev grid, for either a primal 0-form with Dirichlet boundary conditions $f(-1) = e^{-1}$ and $f(1) = e$, or for a dual 0-form with Neumann boundary conditions $f'(-1) = e^{-1}$ and $f'(+1) = e$. All of our results exhibit spectral convergence (measured through the L^∞ error $\|f - \Delta^{-1}q\|_\infty$), with the conventional plateau when we reach the limit of accuracy of the representation of floating point numbers.

Note that the boundary conditions (Dirichlet or Neumann) are simply incorporated into the $\mathbf{D}$ operators: prescribed primal values or dual values at the boundary will be used as is instead of being considered unknowns. This is substantially simpler than the traditional use of, e.g., generalized Lagrange polynomials for the treatment of boundary conditions. The convergence plots are shown in Fig. 7, proving spectral convergence.

2D Poisson. We also tested our approach on 2D problems. This time, we computed Poisson equations not only for scalar functions, but also for vector fields—or in exterior calculus jargon, we can compute the Laplacian $\Delta = \star\mathbf{d}\star\mathbf{d} + \mathbf{d}\star\mathbf{d}\star$ for 0-, 1-, and 2-forms. Given that these forms can also be performed on the primal or the dual, we now have six discrete versions of the Laplace operator:

$$\mathcal{P}(\star\mathbf{d}\star\mathbf{d} + \mathbf{d}\star\mathbf{d}\star)\mathcal{R} = \begin{cases} \widetilde{\mathbf{H}}^1\widetilde{\mathbf{D}}^1\mathbf{H}^1\mathbf{D}^0 & \text{for primal 0-forms,} \\ \mathbf{H}^1\mathbf{D}^1\widetilde{\mathbf{H}}^1\widetilde{\mathbf{D}}^0 & \text{for dual 0-forms,} \\ \widetilde{\mathbf{H}}^1\widetilde{\mathbf{D}}^0\mathbf{H}^2\mathbf{D}^1 + \mathbf{D}^0\widetilde{\mathbf{H}}^2\widetilde{\mathbf{D}}^1\mathbf{H}^1 & \text{for primal 1-forms,} \\ \mathbf{H}^1\mathbf{D}^0\widetilde{\mathbf{H}}^2\widetilde{\mathbf{D}}^1 + \widetilde{\mathbf{D}}^0\mathbf{H}^2\mathbf{D}^1\widetilde{\mathbf{H}}^1 & \text{for dual 1-forms,} \\ \mathbf{D}^1\widetilde{\mathbf{H}}^1\widetilde{\mathbf{D}}^0\mathbf{H}^2 & \text{for primal 2-forms,} \\ \widetilde{\mathbf{D}}^1\mathbf{H}^1\mathbf{D}^0\widetilde{\mathbf{H}}^2 & \text{for dual 2-forms.} \end{cases}$$

The convergence plots, shown in Fig. 8, demonstrate spectral convergence again.

Finally, to distinguish our approach from other conventional discretization techniques and reinforce the point that our spectral method is not solely limited to scalar functions, we depict a solution of the Poisson equation for primal 1-*forms* (i.e., the equivalent of a vectorial Poisson equation) on Fig. 9.

7. Conclusions and future work

In this paper, we laid out the foundations for a spectral instance of Discrete Exterior Calculus. Using chains and cochains as discretization for differential forms, we provide a discrete exterior derivative, a discrete Hodge star, and a discrete wedge product that are spectrally accurate. One can use the resulting calculus as a chain collocation method to solve differential equations.

While the three exterior operators we treated cover already a large range of differential operators (allowing a number of computational tasks including Hodge decomposition), contractions (exterior product) and Lie derivatives still need to be derived in a similar geometric, coordinate-free manner—currently, they can only be handled using coordinates. Note that defining the contraction operator would be enough, since one can then define the Lie derivative algebraically with Cartan's "magic formula" [37]. But a direct dynamic ("flowed-out") definition as presented in [39] may instead prove beneficial.

Mathematica and Python code that implements all our operators and produced our numerical results can be found at www.geometry.caltech.edu.

Acknowledgements

We thank Marc Gerritsma for discussions and references. We gratefully acknowledge partial funding from NSF grant CCF-1011944.

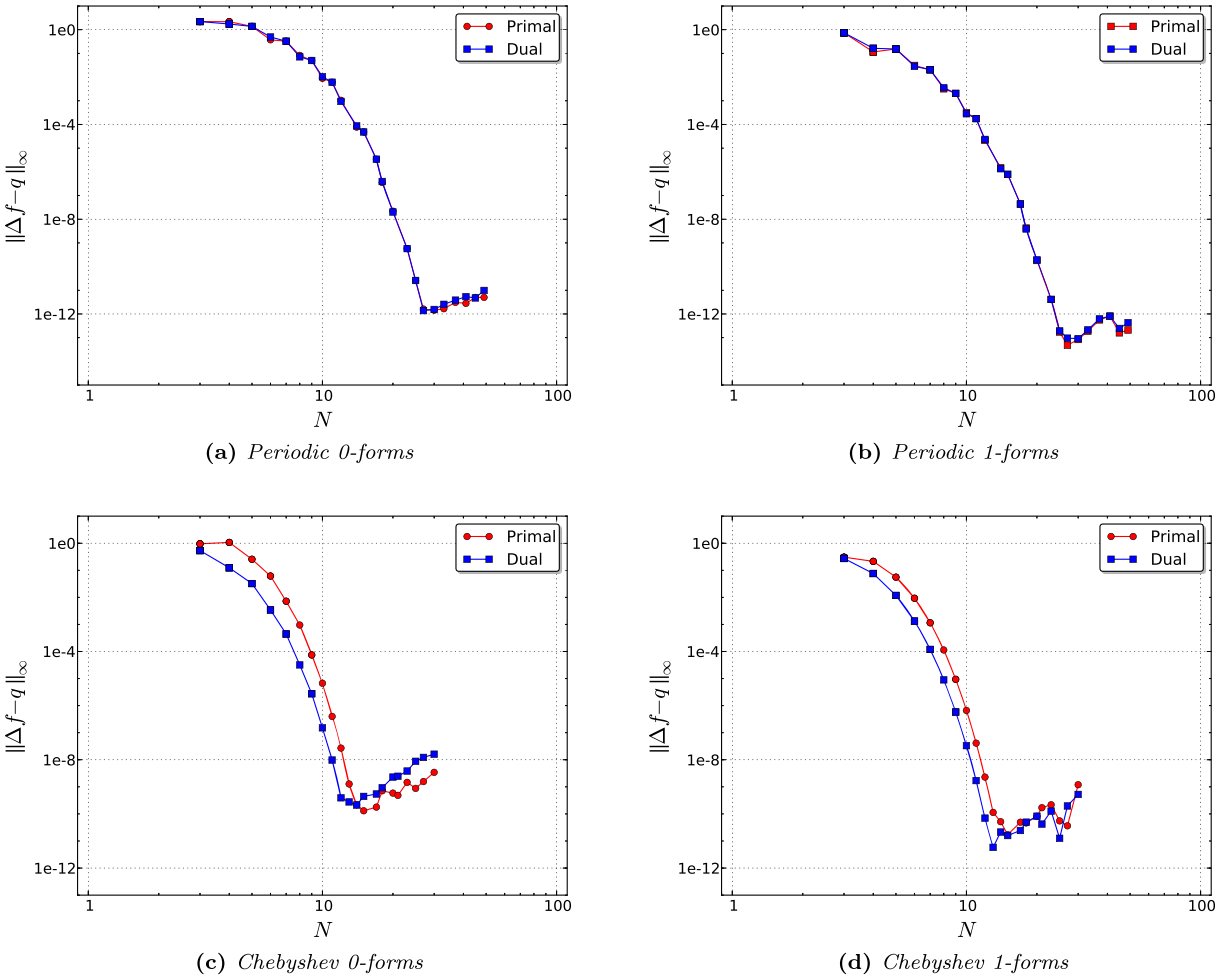

(a) *Periodic 0-forms* **(b)** *Periodic 1-forms*

(c) *Chebyshev 0-forms* **(d)** *Chebyshev 1-forms*

Fig. 8. Convergence graphs for a 2D Poisson equation; (a) we solve $\Delta f = e^{\sin x}(\cos^2 x - \sin x) + e^{\sin y}(\cos^2 y - \sin y)$ on a periodic domain for either a primal, or dual 0-form f; (b) now for $\Delta f = e^{\sin x}(\cos^2 x - \sin x)\,\mathbf{d}x + e^{\sin y}(\cos^2 y - \sin y)\,\mathbf{d}y$; (c) we solve $\Delta f = e^x + e^y$ on a Chebyshev grid, for either a primal 0-form with Dirichlet boundary conditions $f(x, y) = e^x + e^y$, or for a dual 0-form with Neumann boundary conditions $\nabla f(x) \cdot \mathbf{n} = (e^x e^y)^t \mathbf{n}$; (d) now for $\Delta f = e^x\,\mathbf{d}x + e^y\,\mathbf{d}y$. All of our results exhibit spectral convergence (measured through the L^∞ error $\|\Delta f - q\|_\infty$), with the conventional plateau in accuracy for fine meshes.

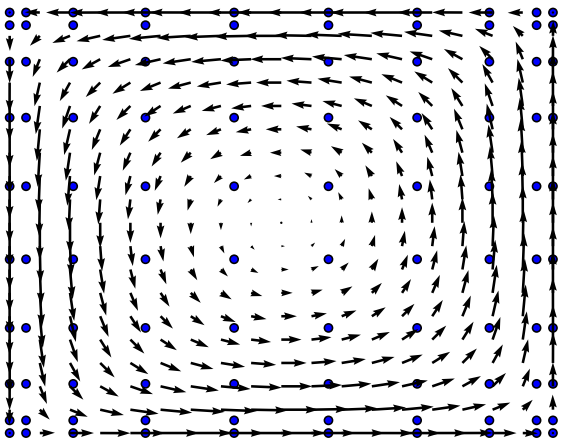

Fig. 9. Plot for the solution of $(\star\mathbf{d}\star\mathbf{d} + \mathbf{d}\star\mathbf{d}\star)f = 0$ with the boundary conditions of $\star f|_{\mathcal{B}} = 0$ (vector field is tangent to the boundary) and $\star\mathbf{d}\star f|_{\mathcal{B}} = 1$ (the curl at the boundary is equal to 1). The domain is $[-1, 1]^2$, discretized by a 10×10 Chebyshev grid.

References

[1] D.N. Arnold, R.S. Falk, R. Winther, Finite element exterior calculus, homological techniques, and applications, Acta Numer. 15 (2006) 1–155.
[2] A. Stern, Y. Tong, M. Desbrun, J.E. Marsden, Variational integrators for Maxwell's equations with sources, PIERS Online 4 (7) (2008) 711–715.
[3] D.N. Arnold, R.S. Falk, R. Winther, Finite element exterior calculus: from Hodge theory to numerical stability, Bull. Am. Math. Soc. 47 (2010) 281–354.
[4] D. Pavlov, P. Mullen, Y. Tong, E. Kanso, J.E. Marsden, M. Desbrun, Structure-preserving discretization of incompressible fluids, Physica D 240 (6) (2011) 443–458.
[5] M. Holst, A. Stern, Geometric variational crimes: Hilbert complexes, finite element exterior calculus, and problems on hypersurfaces, Found. Comput. Math. 12 (3) (2012) 263–293.
[6] E. Gawlik, P. Mullen, D. Pavlov, J. Marsden, M. Desbrun, Geometric, variational discretization of continuum theories, Physica D 240 (21) (2011) 1724–1760.
[7] A. Bossavit, Computational Electromagnetism, Academic Press, Boston, 1998.
[8] M. Desbrun, E. Kanso, Y. Tong, Discrete differential forms for computational modeling, in: A. Bobenko, P. Schröder (Eds.), Discrete Differential Geometry, Springer, 2007.
[9] P. Bochev, J. Hyman, Principles of mimetic discretizations of differential operators, in: IMA Vol. Math. Appl., vol. 142, 2006, pp. 89–119.
[10] L.J. Grady, J.R. Polimeni, Discrete Calculus: Applied Analysis on Graphs for Computational Science, Springer, 2010.
[11] D.N. Arnold, P.B. Bochev, R.B. Lehoucq, R.A. Nicolaides, M. Shashkov (Eds.), Compatible Spatial Discretizations, IMA Vol. Math. Appl., vol. 142, Springer, 2006.
[12] R.A. Nicolaides, X. Wu, Covolume solutions of three dimensional div-curl equations, SIAM J. Numer. Anal. 34 (1997) 2195.
[13] M. Desbrun, A.N. Hirani, J.E. Marsden, Discrete exterior calculus for variational problems in computer vision and graphics, in: Proc. CDC, vol. 42, 2003, pp. 533–538.
[14] É. Cartan, Les Systèmes Differentiels Exterieurs et leurs Applications Géometriques, Hermann, Paris, 1945.
[15] J.R. Munkres, Elements of Algebraic Topology, Addison–Wesley, Menlo Park, CA, 1984.
[16] W.L. Burke, Applied Differential Geometry, Cambridge University Press, 1985.
[17] R. Abraham, J.E. Marsden, T. Ratiu, Manifolds, Tensor Analysis, and Applications, Appl. Math. Sci., vol. 75, Springer, 1988.
[18] H. Flanders, Differential Forms and Applications to Physical Sciences, Dover Publications, 1990.
[19] T. Frankel, The Geometry of Physics, second edition, Cambridge University Press, United Kingdom, 2004.
[20] N. Robidoux, S. Steinberg, A discrete vector calculus in tensor grids, Comput. Methods Appl. Math. 11 (1) (2011) 23–66.
[21] H. Whitney, Geometric Integration Theory, Princeton University Press, 1957.
[22] J.-C. Nédélec, Mixed finite elements in $\mathbb{R}^3$, Numer. Math. 35 (1980) 315–341.
[23] Y. Tong, P. Alliez, D. Cohen-Steiner, M. Desbrun, Designing quadrangulations with discrete harmonic forms, in: A. Sheffer, K. Polthier (Eds.), Symposium on Geometry Processing, 2006, pp. 201–210.
[24] P. Mullen, F. de Goes, P. Memari, M. Desbrun, HOT: Hodge Optimized Triangulations, ACM Trans. Graph. 30 (3) (2011).
[25] K. Wang, Weiwei, Y. Tong, M. Desbrun, P. Schröder, Edge subdivision schemes and the construction of smooth vector fields, ACM Trans. Graph. 25 (3) (2006) 1041–1048.
[26] O.P. Bruno, Y. Han, M.M. Pohlman, Accurate, high-order representation of complex three-dimensional surfaces via Fourier continuation analysis, J. Comput. Phys. 227 (2) (2007) 1094–1125.
[27] C. Canuto, M. Hussaini, A. Quarteroni, T. Zang, Spectral Methods: Fundamentals in Single Domains, Sci. Comput. Ser., Springer, 2006.
[28] D.A. Kopriva, A staggered-grid multidomain spectral method for the compressible Navier–Stokes equations, J. Comput. Phys. 143 (1) (1998) 125–158.
[29] N. Robidoux, Polynomial histopolation, superconvergent degrees of freedom, and pseudospectral discrete Hodge operators, preprint, 2006.
[30] M. Bouman, A. Palha, J. Kreeft, M. Gerritsma, A conservative spectral element method for curvilinear domains, in: T.J. Barth, et al. (Eds.), Spectral and High Order Methods for Partial Differential Equations, in: Lect. Notes Comput. Sci. Eng., vol. 76, Springer, Berlin, Heidelberg, 2011, pp. 111–119.
[31] M. Gerritsma, Edge functions for spectral element methods, in: J.S. Hesthaven, et al. (Eds.), Spectral and High Order Methods for Partial Differential Equations, in: Lect. Notes Comput. Sci. Eng., vol. 76, Springer, Berlin, Heidelberg, 2011, pp. 199–207.
[32] A. Palha, M. Gerritsma, Spectral element approximation of the Hodge-$\star$ operator in curved elements, in: J.S. Hesthaven, et al. (Eds.), Spectral and High Order Methods for Partial Differential Equations, in: Lect. Notes Comput. Sci. Eng., vol. 76, Springer, Berlin, Heidelberg, 2011, pp. 283–291.
[33] K. Crane, M. Desbrun, P. Schröder, Trivial connections on discrete surfaces, Comput. Graph. Forum 29 (5) (2010) 1525–1533.
[34] A.N. Hirani, Discrete exterior calculus, PhD thesis, California Institute of Technology, Pasadena, CA, 2003.
[35] M. Leok, Foundations of computational geometric mechanics, PhD thesis, California Institute of Technology, Pasadena, CA, 2004.
[36] J. Harrison, Ravello lecture notes on geometric calculus—Part I, arXiv:math-ph/0501001, 2005.
[37] M. Desbrun, A.N. Hirani, M. Leok, J.E. Marsden, Discrete exterior calculus, arXiv:math/0508341v2.
[38] P.R. Kotiuga, Theoretical limitations of discrete exterior calculus in the context of computational electromagnetics, IEEE Trans. Magn. 44 (6) (2008) 1162–1165.
[39] P. Mullen, A. McKenzie, D. Pavlov, L. Durant, Y. Tong, E. Kanso, J.E. Marsden, M. Desbrun, Discrete Lie advection of differential forms, Found. Comput. Math. 11 (2) (2011) 131–149.
[40] L.N. Trefethen, Spectral Methods in MATLAB, SIAM, Philadelphia, 2000.