

Discrete Differential Forms for Computational Modeling

Mathieu Desbrun Eva Kanso* Yiyong Tong†

Applied Geometry Lab

Caltech‡

1 Motivation

The emergence of computers as an essential tool in scientific research has shaken the very foundations of differential modeling. Indeed, the deeply-rooted abstraction of smoothness, or *differentiability*, seems to inherently clash with a computer’s ability of storing only finite sets of numbers. While there has been a series of computational techniques that proposed discretizations of differential equations, the geometric structures they are simulating are often lost in the process.

1.1 The Role of Geometry in Science

Geometry is the study of space and of the properties of shapes in space. Dating back to Euclid, models of our surroundings have been formulated using simple, geometric descriptions, formalizing apparent *symmetries* and experimental *invariants*. Consequently, geometry is at the foundation of many current physical theories: general relativity, electromagnetism (E&M), gauge theory as well as solid and fluid mechanics all have strong underlying geometrical structures. Einstein’s theory for instance states that gravitational field strength is directly proportional to the *curvature of space-time*. In other words, the physics of relativity is *directly modelled* by the shape of our 4-dimensional world, just as the behavior of soap bubbles is modeled by their shapes. Differential geometry is thus, de facto, the mother tongue of numerous physical and mathematical theories.

Unfortunately, the inherent geometric nature of such theories is often obstructed by their formulation in vectorial or tensorial notations: the traditional use of a coordinate system, in which the defining equations are expressed, often obscures the underlying structures by an overwhelming usage of indices. Moreover, such complex expressions entangle the topological and geometrical content of the model.

1.2 Geometry-based Exterior Calculus

The geometric nature of these models is best expressed and elucidated through the use of the *Exterior Calculus of Differential Forms*, first introduced by Cartan [Cartan 1945]. This geometry-based calculus was further developed and refined over the twentieth century to become the foundation of modern differential geometry. The calculus of exterior forms allows one to express differential and integral equations on smooth and curved spaces in a consistent manner, while revealing the geometrical invariants at play. For example, the classical operations of gradient, divergence, and curl as well as the theorems of Green, Gauss and Stokes can all be expressed concisely in terms of differential forms and an operator on these forms called the exterior derivative—hinting at the generality of this approach.

Compared to classical tensorial calculus, this exterior calculus has several advantages. First, it is often difficult to recognize the

coordinate-independent nature of quantities written in tensorial notation: local and global invariants are hard to notice by just staring at the indices. On the other hand, invariants are easily discovered when expressed as differential forms by invoking either Stokes’ theorem, the Poincaré lemma, or by applying exterior differentiation. Note also that the exterior derivative of differential forms—the antisymmetric part of derivatives—is one of the most important parts of differentiation, since it is invariant under coordinate system change. In fact, Sharpe states in [Sharpe 1997] that every differential equation may be expressed in term of the exterior derivative of differential forms. As a consequence, several recent initiatives have been aimed at formulating physical laws in terms of differential forms. For recent work along these lines, the reader is invited to refer to [Burke 1985; Abraham et al. 1988; Lovelock and Rund 1993; Flanders 1990; Morita 2001; Carroll 2003; Frankel 2004] for books offering a theoretical treatment of various physical theories using differential forms.

1.3 Differential vs. Discrete Modeling

We have seen that a large amount of our scientific knowledge relies on a deeply-rooted differential (*i.e.*, smooth) comprehension of the world. This abstraction of differentiability allows researchers to model complex physical systems via concise equations. With the sudden advent of the digital age, it was therefore only natural to resort to computations based on such differential equations.

However, since digital computers can only manipulate finite sets of numbers, their capabilities seem to clash with the basic foundations of differential modeling. In order to overcome this hurdle, a first set of computational techniques (*e.g.*, finite difference or particle methods) focused on satisfying the continuous equations at a discrete set of spatial and temporal samples. Unfortunately, focusing on accurately discretizing the local laws often fails to respect important global structures and invariants. Later methods such as Finite Elements (FEM), drawing from developments in the calculus of variations, remedied this inadequacy to some extent by satisfying local conservation laws on average and preserving some important invariants. Coupled with a finer ability to deal with arbitrary boundaries, FEM became the de facto computational tool for engineers. Even with significant advances in error control, convergence, and stability of these finite approximations, the underlying structures of the simulated continuous systems are often destroyed: a moving rigid body may gain or loose momentum; or a cavity may exhibit fictitious eigenmodes in an electromagnetism (E&M) simulation. Such examples illustrate some of the loss of fidelity that can follow from a standard discretization process, failing to preserve some fundamental geometric and topological structures of the underlying continuous models.

The cultural gap between theoretical and applied science communities may be partially responsible for the current lack of proper discrete, computational modeling that could mirror and leverage the rich developments of its differential counterpart. In particular, it is striking that the calculus of differential forms has not yet had an impact on the mainstream computational fields, despite excellent initial results in E&M [Bossavit 1998] or Lagrangian mechanics [Marsden and West 2001]. It should also be noticed that

*Now at the University of Southern California.

†Now at Michigan State University.

‡E-mail: {mathieu|eva|yiyong}@caltech.edu

some basic tools necessary for the definition of a discrete calculus already exist, probably initiated by Poincaré when he defined his cell decomposition of smooth manifolds. The study of the structure of ordered sets or simplices now belongs to the well-studied branch of mathematics known as *Combinatorial Differential Topology and Geometry*, which is still an active area of research (see, e.g., [Forman 2003] and [Björner and Welker 1995] and references therein).

1.4 Calculus ex Geometrica

Given the overwhelming geometric nature of the most fundamental and successful calculus of these last few centuries, it seems relevant to *approach computations from a geometric standpoint*.

One of the key insights that percolated down from the theory of differential forms is rather simple and intuitive: one needs to recognize that different physical quantities have different properties, and must be treated accordingly. Fluid mechanics or electromagnetism, for instance, make heavy use of line integrals, as well as surface and volume integrals; even physical measurements are performed as specific local integrations or averages (think flux for magnetic field, or current for electricity, or pressure for atoms' collisions). Pointwise evaluations or approximations for such quantities are not the appropriate discrete analogs, since the defining geometric properties of their physical meaning cannot be enforced naturally. Instead, *one should store and manipulate those quantities at their geometrically-meaningful location*: in other words, we should consider values on vertices, edges, faces, and tetrahedra as proper discrete versions of respectively pointwise functions, line integrals, surface integrals, and volume integrals: only then will we be able to manipulate those values without violating the symmetries that the differential modeling tried to exploit for predictive purposes.

1.5 Similar Endeavors

The need for improved numerics have recently sprung a (still limited) number of interesting related developments in various fields. Although we will not try to be exhaustive, we wish to point the reader to a few of the most successful investigations with the same “flavor” as our discrete geometry-based calculus, albeit their approaches are rarely similar to ours. First, the field of *Mimetic Discretizations of Continuum Mechanics*, led by Shashkov, Steinberg, and Hyman [Hyman and Shashkov 1997], started on the premise that spurious solutions obtained from finite element or finite difference methods often originate from inconsistent discretizations of the operators div , curl , and grad , and that addressing this inconsistency pays off numerically. Similarly, *Computational Electromagnetism* has also identified the issue of field discretization as the main reason for spurious modes in numerical results. An excellent treatment of the discretization of the Maxwell's equations resulted [Bossavit 1998], with a clear relationship to the differential case. Finally, recent developments in *Discrete Lagrangian Mechanics* have demonstrated the efficacy of a proper discretization of the Lagrangian of a dynamical system, rather than the discretization of its derived Euler-Lagrange equations: with a discrete Lagrangian, one can ensure that the integration scheme satisfies an exact discrete least-action principle, preserving all the momenta directly for arbitrary orders of accuracy [Marsden and West 2001]. Respecting the defining geometric properties of both the fields and the governing equations is a common link between all these recent approaches.

1.6 Advantages of Discrete Differential Modeling

The reader will have most probably understood our bias by now: we believe that the systematic construction, inspired by Exterior Calculus, of **differential, yet readily discretizable computational foundations** is a crucial ingredient for numerical fidelity. Because

many of the standard tools used in differential geometry have discrete combinatorial analogs, the *discrete versions of forms or manifolds* will be formally identical to (and should partake of the same properties as) the continuum models. Additionally, such an approach should clearly maintain the separation of the topological (*metric-independent*) and geometrical (*metric-dependent*) components of the quantities involved, keeping the geometric picture (*i.e.*, intrinsic structure) intact.

A *discrete differential modeling approach to computations* will also be often much simpler to define and develop than its continuous counterpart. For example, the discrete notion of a differential form will be implemented simply as values on mesh elements. Likewise, the discrete notion of orientation will be more straightforward than its continuous counterpart: while the differential definition of orientation uses the notion of equivalence class of atlases determined by the sign of the Jacobian, the orientation of a mesh edge will be one of two directions; a triangle will be oriented clockwise or counterclockwise; a volume will have a direction as a right-handed helix or a left-handed one; no notion of atlas (a collection of consistent coordinate charts on a manifold) will be required.




Figure 1: Typical 2D and 3D meshes: although the David head appears smooth, its surface is made of a triangle mesh; tetrahedral meshes (such as this mechanical part, with a cutaway view) are some typical examples of irregular meshes on which computations are performed. David's head mesh is courtesy of Marc Levoy, Stanford.

1.7 Goal of This Chapter

Given these premises, this chapter was written with several purposes in mind. First, we wish to demonstrate that the foundations on which powerful methods of computations can be built are quite approachable—and are not as abstract as the reader may fear: the ideas involved are very intuitive as a side effect of the simplicity of the underlying geometric principles.

Second, we wish to help bridge the gap between applied fields and theoretical fields: we have tried to render the theoretical bases of our exposition accessible to computer scientists, and the concrete implementation insights understandable by non-specialists. For this very reason, the reader should not consider this introductory exposition as a definite source of knowledge: it should instead be considered as a portal to better, more focused work on related subjects. We only hope that we will ease our readers into foundational concepts that can be undoubtedly and fruitfully applied to all sorts of computations—be it for graphics or simulation.

With these goals in mind, we will describe the background needed to develop a principled, geometry-based approach to computational modeling that gets around the apparent mismatch between differential and discrete modeling.

2 Relevance of Forms for Integration

The evaluation of differential quantities on a discrete space (mesh) is a nontrivial problem. For instance, consider a piecewise-linear

2-dimensional surface embedded in a three-dimensional Euclidean space, *i.e.*, a triangle mesh. Celebrated quantities such as the Gaussian and mean curvatures are delicate to define on it. More precisely, the Gaussian curvature can be easily proven to be zero everywhere *except* on vertices, where it is a Dirac delta function. Likewise, the mean curvature can only be defined in the distributional sense, as a Dirac delta function on edges. However, through local *integrations*, one can easily manipulate these quantities numerically: if a careful choice of non-overlapping regions is made, the delta functions can be properly integrated, rendering the computations relatively simple as shown, for example, in [Meyer et al. 2003; Hildebrandt and Polthier 2004]. Note that the process of integration to suppress discontinuity is, in spirit, equivalent to the idea of weak form used in the Finite Element method.

This idea of integrated value has predated in some cases the equivalent differential statements: for instance, it was long known that the genus of a surface can be calculated through a cell decomposition of the surface via the Euler characteristic. The actual Gauss-Bonnet theorem was, however, derived later on. Now, if one tries to discretize the Gaussian curvature of a piecewise-linear surface in an arbitrary way, it is not likely that its integral over the surface equals the desired Euler characteristic, while its discrete version, defined on vertices (or, more precisely, on the dual of each vertex), naturally preserves this topological invariant.

2.1 From Integration to Differential Forms

Integration is obviously a linear operation, since for any disjoint sets A and B ,

$$\int_{A \cup B} = \int_A + \int_B.$$

Moreover, the integration of a smooth function over a subset of measure zero is always zero; for example, an area integral of (a lower dimensional object such as) a curve or a point is equal to zero. Finally, integration is *objective* (*i.e.*, relevant) only if its evaluation is invariant under change of coordinate systems. These three properties combined directly imply that the integrand (*i.e.*, the whole expression after the integral sign) has to be *antisymmetric*. That is, the basic building blocks of any type of integration are **differential forms**. Chances are, the reader is already very well acquainted with forms, maybe without even knowing it.

2.1.1 An Intuitive Definition

A differential form (also denoted as exterior¹ differential form) is, informally, an integrand, *i.e.*, a quantity that can be integrated. It is the dx in $\int dx$ and the $dx dy$ in $\iint dx dy$. More precisely, consider a smooth function $F(x)$ over an interval in \mathbb{R} . Now, define $f(x)$ to be its derivative, that is,

$$f(x) = \frac{dF}{dx},$$

Rewriting this last equation (with slight abuse of notation for simplicity) yields $dF = f(x)dx$, which leads to:

$$\int_a^b dF = \int_a^b f(x)dx = F(b) - F(a). \quad (1)$$


This last equation is known as the Newton-Leibnitz formula, or the first fundamental theorem of calculus. The integrand $f(x)dx$

¹The word “exterior” is used as the exterior algebra is basically built out of an *outer* product.

is called a *1-form*, because it can only be integrated over any 1-dimensional (1D) real interval. Similarly, for a function $G(x, y, z)$, we have:

$$dG = \frac{\partial G}{\partial x} dx + \frac{\partial G}{\partial y} dy + \frac{\partial G}{\partial z} dz,$$

which can be integrated over any 1D curve in \mathbb{R}^3 , and is also a 1-form. More generally, a *k-form* can be described as an entity ready (or designed, if you prefer) to be integrated on a kD (sub)region. Note that forms are valued zero on (sub)regions that are of higher or lower order dimension than the original space; for example, 4-forms are zero on \mathbb{R}^3 . These differential forms are extensively used in mathematics, physics and engineering, as we already hinted at the fact in Section 1.4 that most of our measurements of the world are of integral nature: even digital pictures are made out of local area integrals of the incident light over each of the sensors of a camera to provide a set of values at each pixel on the final image (see inset). The importance of this notion of forms in science is also evidenced by the fact that operations like gradient, divergence, and curl can all be expressed in terms of forms only, as well as fundamental theorems like Green’s or Stokes.



2.1.2 A Formal Definition

For concreteness, consider the n -dimensional Euclidean space \mathbb{R}^n , $n \in \mathbb{N}$ and let \mathcal{M} be an open region $\mathcal{M} \subset \mathbb{R}^n$; \mathcal{M} is also called an *n-manifold*. The vector space $T_x\mathcal{M}$ consists of all the (tangent) vectors at a point $x \in \mathcal{M}$ and can be identified with \mathbb{R}^n itself. A k -form ω^k is a rank- k , anti-symmetric, tensor field over \mathcal{M} . That is, at each point $x \in \mathcal{M}$, it is a multi-linear map that takes k tangent vectors as input and returns a real number:

$$\omega^k : T_x\mathcal{M} \times \dots \times T_x\mathcal{M} \longrightarrow \mathbb{R}$$

which *changes sign for odd permutations of the variables* (hence the term antisymmetric). Any k -form naturally induces a k -form on a submanifold, through restriction of the linear map to the domain that is the product of tangent spaces of the submanifold.

Comments on the Notion of Pseudo-forms There is a closely related concept named pseudo-form. Pseudo-forms change sign when we change the orientation of coordinate systems, just like pseudo-vectors. As a result, the integration of a pseudo-form does not change sign when the orientation of the manifold is changed. Unlike k -forms, a pseudo- k -form induces a pseudo- k -form on a submanifold *only* if a transverse direction is given. For example, fluid flux is sometimes called a pseudo-2-form: indeed, given a transverse direction, we know how much flux is going through a piece of surface; it does not depend on the orientation of the surface itself. Vorticity is, however, a true 2-form: given an orientation of the surface, the integration gives us the circulation around that surface boundary induced by the surface orientation. It does *not* depend on the transverse direction of the surface. But if we have an orientation of the ambient space, we can always associate transverse direction with internal orientation of the submanifold. Thus, in our case, we may treat pseudo-forms simply as forms because we can consistently choose a representative from the equivalence class.

2.2 The Differential Structure

Differential forms are the building blocks of a whole calculus. To manipulate these basic blocks, Exterior Calculus defines seven operators:

- ◇ d : the exterior derivative, that extends the notion of the differential of a function to differential forms;

- ◇ \star : the Hodge star, that transforms k -forms into $(n-k)$ -forms;
- ◇ \wedge : the wedge product, that extends the notion of exterior product to forms;
- ◇ \sharp and \flat : the sharp and flat operators, that, given a metric, transform a 1-form into a vector and vice-versa;
- ◇ i_X : the interior product with respect to a vector field X (also called contraction operator), a concept dual to the exterior product;
- ◇ \mathcal{L}_X : the Lie derivative with respect to a vector field X , that extends the notion of directional derivative.

In this chapter, we will restrict our discussions to the first three operators, to provide the most basic tools necessary in computational modeling.

2.3 A Taste of Exterior Calculus in \mathbb{R}^3

To give the reader a taste of the relative simplicity of Exterior Calculus, we provide a list of equivalences (in the continuous world!) between traditional operations and their Exterior Calculus counterpart in the special case of \mathbb{R}^3 . We will suppose that we have the usual Euclidean metric. Then, forms are actually quite simple to conceive:

- 0-form \Leftrightarrow scalar field
- 1-form \Leftrightarrow vector field
- 2-form \Leftrightarrow vector field
- 3-form \Leftrightarrow scalar field

To be clear, we will add a superscript on the forms to indicate their rank. Then applying forms to vector fields amounts to:

- 1-form: $u^1(v) \Leftrightarrow u \cdot v$.
- 2-form: $u^2(v, w) \Leftrightarrow u \cdot (v \times w)$.
- 3-form: $f^3(u, v, w) \Leftrightarrow fu \cdot (v \times w)$.

Furthermore, the usual operations like gradient, curl, divergence and cross product can all be expressed in terms of the basic exterior calculus operators. For example:

$$\begin{aligned} d^0 f &= \nabla f, \quad d^1 u = \nabla \times u, \quad d^2 u = \nabla \cdot u; \\ \star^0 f &= f, \quad \star^1 u = u, \quad \star^2 u = u, \quad \star^3 f = f; \\ \star^0 d^2 \star^1 u^1 &= \nabla \cdot u, \quad \star^1 d^1 \star^2 u^2 = \nabla \times u, \quad \star^2 d^0 \star^3 f = \nabla f; \\ f^0 \wedge u &= fu, \quad u^1 \wedge v^1 = u \times v, \quad u^1 \wedge v^2 = u^2 \wedge v^1 = u \cdot v; \\ i_v u^1 &= u \cdot v, \quad i_v u^2 = u \times v, \quad i_v f^3 = fv. \end{aligned}$$

Now that we have established the relevance of differential forms even in the most basic vector operations, time has come to turn our attention to make this concept of forms readily usable for computational purposes.

3 Discrete Differential Forms

Finding a discrete counterpart to the notion of differential forms is a delicate matter. If one was to represent differential forms using their coordinate values and approximate the exterior derivative using finite differences, basic theorems such as Stokes' theorem would not hold numerically. The main objective of this section is therefore to present a proper discretization of the forms on what are known as simplicial complexes. We will show how this discrete geometric structure, well suited for computational purposes, is designed to preserve all the fundamental differential properties. For simplicity, we restrict the discussion to forms on 2D surfaces or 3D regions embedded in \mathbb{R}^3 , but the construction is applicable to general manifolds in arbitrary spaces. In fact, the only necessary assumption is that the embedding space must be a vector space, a natural condition in practice.

3.1 Simplicial Complexes and Discrete Manifolds

For the interested reader, the notions we introduce in this section are defined formally in much more details (for the general case of k -dimensional spaces) in references such as [Munkres 1984] or [Hatcher 2004].




Figure 2: A 1-simplex is a line segment, the convex hull of two points. A 2-simplex is a triangle, i.e., the convex hull of three distinct points. A 3-simplex is a tetrahedron, as it is the convex hull of four points.

3.1.1 Notion of Simplex

A k -simplex is the generic term to describe the simplest mesh element of dimension k —hence the name. By way of motivation, consider a three-dimensional mesh in space. This mesh is made of a series of adjacent tetrahedra (denoted *tets* for simplicity throughout). The vertices of the tets are called 0-simplices. Similarly, the line segments or edges form 1-simplices, the triangles or faces form 2-simplices, and the tets form 3-simplices. Note that we can define these simplices in a top-down manner too: faces (2-simplex) can be thought of as boundaries of tets (3-simplices), edges (1-simplices) as boundaries of faces, and vertices (0-simplices) as boundaries of edges.

The definition of a simplex can be made more abstract as a series of k -tuples (referring to the vertices they are built upon). However, for the type of applications that we are targeting in this chapter, we will often not make any distinction between an abstract simplex and its topological realization (connectivity) or geometrical realization (positions in space).

Formally, a k -simplex σ_k is the non-degenerate convex hull of $k+1$ geometrically distinct points $v_0, \dots, v_k \in \mathbb{R}^n$ with $n \geq k$. In other words, it is the intersection of all convex sets containing (v_0, \dots, v_k) ; namely:

$$\sigma_k = \left\{ x \in \mathbb{R}^n \mid x = \sum_{i=0}^k \alpha^i v_i \text{ with } \alpha^i \geq 0 \text{ and } \sum_{i=0}^k \alpha^i = 1 \right\}.$$

The entities v_0, \dots, v_k are called the *vertices* and k is called the dimension of the k -simplex., which we will denote as:

$$\sigma_k = \{v_0 v_1 \dots v_k\}.$$

3.1.2 Orientation of a Simplex

Note that all orderings of the $k+1$ vertices of a k -simplex can be divided into two equivalent classes, i.e., two orderings differing by an even permutation. Such a class of orderings is called an *orientation*. In the present work, we always assume that local orientations are *given* for each simplex; that is, each element of the mesh has been given a particular orientation. For example, an edge $\sigma_1 = \{v_0 v_1\}$ in Figure 2 has an arrow indicating its default orientation. If the opposite orientation is needed, we will denote it as $\{v_1 v_0\}$, or, equivalently, by $-\{v_0 v_1\}$. For more details and examples, the reader is referred to [Munkres 1984; Hirani 2003].

3.1.3 Boundary of a Simplex

Any $(k-1)$ -simplex spanned by a subset of $\{v_0, \dots, v_k\}$ is called a $(k-1)$ -face of σ_k . That is, a $(k-1)$ -face is simply a $(k-1)$ -simplex

whose k vertices are all from the $k+1$ vertices of the k -simplex. The union of the $(k-1)$ -faces is what is called the *boundary* of the k -simplex. One should be careful here: because of the default orientation of the simplices, the formal *signed* sum of the $(k-1)$ -faces defines the boundary of the k -simplex. Therefore, the boundary operator takes a k -simplex and gives the sum of all its $(k-1)$ -faces with 1 or -1 as coefficients depending on whether their respective orientations match or not, see Figure 4.




Figure 3: The boundary operator ∂ applied to a triangle (a 2-simplex) is equal to the signed sum of the edges (i.e., the 1-faces of the 2-simplex).

To remove possible mistakes in orientation, we can define the *boundary operator* as follows:

$$\partial\{v_0v_1 \cdots v_k\} = \sum_{j=0}^k (-1)^j \{v_0, \dots, \widehat{v}_j, \dots, v_k\}, \quad (2)$$

where \widehat{v}_j indicates that v_j is missing from the sequence, see Figure 3. Clearly, each k -simplex has $k+1$ $(k-1)$ -faces. For this statement to be valid even for $k = 0$, the empty set \emptyset is usually defined as a (-1) -simplex face of every 0-simplex. The reader is invited to verify this definition on the triangle $\{v_0, v_1, v_2\}$ in Figure 3:

$$\partial\{v_0, v_1, v_2\} = \{v_1, v_2\} - \{v_0, v_2\} + \{v_0, v_1\}.$$





Figure 4: Boundary operator applied to a triangle (left), and a tetrahedron (right). Orientations of the simplices are indicated with arrows.

3.1.4 Simplicial Complex

A simplicial complex is a collection \mathcal{K} of simplices, which satisfies the following two simple conditions:

- ◇ every face of each simplex in \mathcal{K} is in \mathcal{K} ;
- ◇ the intersection of any two simplices in \mathcal{K} is either empty, or an entire common face.



Computer graphics makes heavy use of what is called *realizations* of simplicial complexes. Loosely speaking, a realization of a simplicial complex is an embedding of this complex into the underlying space \mathbb{R}^n . Triangle meshes in 2D and tet meshes in 3D are examples of such simplicial complexes (see Figure 1). Notice that polygonal meshes can be easily triangulated, thus can be easily turned into simplicial complexes. One can also use the notion of *cell complex* by allowing the elements of \mathcal{K} to be non-simplicial; we will restrict our explanations to the case of simplicial complexes for simplicity.

3.1.5 Discrete Manifolds

An n -dimensional discrete manifold \mathcal{M} is an n -dimensional simplicial complex that satisfies the following condition: for each simplex, the union of all the incident n -simplices forms an n -dimensional ball (i.e., a disk in 2D, a ball in 3D, etc), or half a ball if the simplex is on the boundary. As a consequence, each $(n-1)$ -simplex has exactly two adjacent n -simplices—or only one if it is on a boundary.

Basically, the notion of discrete manifold corresponds to the usual Computer Graphics acceptance of “manifold mesh”. For example in 2D, discrete manifolds cannot have isolated edges (also called sticks or hanging edges) or isolated vertices, and each of their edges is adjacent to 2 triangles (except for the boundary; in that case, the edge is adjacent to only one triangle). A surface mesh in 3D cannot have a “fin”, i.e., an edge with more than two adjacent triangles. To put it differently, infinitesimally-small, imaginary *inhabitants* of a n -dimensional discrete manifolds would consider themselves living in \mathbb{R}^n as any small neighborhood of this manifold is isomorphic to \mathbb{R}^n .




Figure 5: (a) A simplicial complex consisting of all vertices $\{v_0, v_1, v_2, v_3\}$ and edges $\{e_0, e_1, e_2, e_3, e_4\}$. This simplicial complex is not a discrete manifold because the neighborhoods of the vertices v_1 and v_2 are not 1D balls. (b) If we add the triangles f_0 and f_1 to the simplicial complex, it becomes a 2-manifold with one boundary.

3.2 Notion of Chains

We have already encountered the notion of chain, without mentioning it. Recall that the boundary operator takes each k -simplex and gives the *signed* sum of all its $(k-1)$ -faces. We say that the boundary of a k -simplex produces a $(k-1)$ -chain. The following definition is more precise and general.

3.2.1 Definition

A k -**chain** of an oriented simplicial complex \mathcal{K} is a set of values, one for *each* k -simplex of \mathcal{K} . That is, a k -chain c can then be thought of as a linear combination of all the k -simplices in \mathcal{K} :

$$c = \sum_{\sigma \in \mathcal{K}} c(\sigma) \cdot \sigma, \quad (3)$$

where $c(\sigma) \in \mathbb{R}$. We will denote the group of all k -chains as \mathcal{C}_k .

3.2.2 Implementation of Chains

Let the set of all k -simplices in \mathcal{K} be denoted \mathcal{K}^k , and let its cardinality be denoted as $|\mathcal{K}^k|$. A k -chain can simply be stored as a *vector* (or array) of dimension $|\mathcal{K}^k|$, i.e., one number for each k -simplex $\sigma_k \in \mathcal{K}^k$.

3.2.3 Boundary Operator on Chains

We mentioned that the boundary operator ∂ was returning a particular type of chain, namely, a chain with coefficients equal to either 0,

1, or -1 . Therefore, it should not be surprising that we can extend the notion of boundary to act also on k -chains, simply by linearity:

$$\partial \sum_k c_k \sigma_k = \sum_k c_k \partial \sigma_k.$$

That is, from one set of values assigned to all simplices of a com-




Figure 6: (a) An example of 1-chain being the boundary of a face (2-simplex); (b) a second example of 1-chain with 4 nonzero coefficients.

plex, one can deduce another set of values derived by weighting the boundaries of each simplex by the original value stored on it. This operation is very natural, and can thus be implemented easily as explained next.

3.2.4 Implementation of the Boundary Operator

Since the boundary operator is a linear mapping from the space of k -simplices to the space of $(k-1)$ -simplices, it can simply be represented by a *matrix* of dimension $|\mathcal{K}^{k-1}| \times |\mathcal{K}^k|$. The reader can convince herself that this matrix is sparse, as only immediate neighbors are involved in the boundary operator. Similarly, this matrix contains only the values 0, 1, and -1 . Notice that in 3D, there are three non-trivial boundary operators ∂_k (∂_1 is the boundary operator on edges, ∂_2 on triangles, ∂_3 on tets). However, the operator needed for a particular operation is obvious from the type of the argument: if the boundary of a tet is needed, the operator ∂_3 is the only one that makes sense to apply; in other words, the boundary of a k -simplex σ_k is found by invoking $\partial_k \sigma_k$. Thanks to this context-dependence, we can simplify the notation and remove the subscript when there is no ambiguity.

3.3 Notion of Cochains

A k -cochain ω is the *dual* of a k -chain, that is to say, ω is a linear mapping that takes k -chains to \mathbb{R} . One writes:

$$\begin{aligned} \omega : \mathcal{C}_k &\rightarrow \mathbb{R}, \\ c &\mapsto \omega(c), \end{aligned}$$

which reads as: a k -cochain ω operates on a k -chain c to give a scalar in \mathbb{R} . Since a chain is a linear combination of simplices, a cochain returns a linear combination of the values of that cochain on each simplex involved.

Clearly, a cochain also corresponds to one value per simplex (since all the k -simplices form a basis for the vector space \mathcal{C}_k , and we only need to know the mapping of vectors in this basis to determine a linear mapping), and hence the notion of duality of chains and cochains is appropriate. But contrary to a chain, a k -cochain is *evaluated* on each simplex of the dimension k . In other words, a k -cochain can be thought of as a *field* that can be evaluated on each k -simplex of an oriented simplicial complex \mathcal{K} .

3.3.1 Implementation of Cochains

The numerical representation of cochains follows from that of chains by duality. Recall that a k -chain can be represented as a vector c_k of length equal to the number of k -simplices in \mathcal{M} . Similarly, one may represent ω by a vector ω^k of the same size as c_k .

Now, remember that ω operates on c to give a scalar in \mathbb{R} . The linear operation $\omega(c)$ translates into an inner product $\omega^k \cdot c_k$. More specifically, one may continue to think of c_k as a *column vector* so that the \mathbb{R} -valued linear mapping ω can be represented by a *row vector* $(\omega^k)^t$, and $\omega(c)$ becomes simply the matrix multiplication of the row vector $(\omega^k)^t$ with the column vector c_k . The evaluation of a cochain is therefore trivial to implement.

3.4 Discrete Forms as Co-Chains

The attentive reader will have noticed by now: *k-cochains are discrete analogs to differential forms*. Indeed, a continuous k -form was defined as a linear mapping from k -dimensional sets to \mathbb{R} , as we can only integrate a k -form on a k -(sub)manifold. Note now that a k D set, when one has only a mesh to work with, is simply a *chain*. And a linear mapping from a chain to a real number is what we called a cochain: *a cochain is therefore a natural discrete counterpart of a form*.

For instance a 0-form can be evaluated at each point, a 1-form can be evaluated on each curve, a 2-form can be evaluated on each surface, *etc.* Now if we *restrict* integration to take place only on the k -submanifold which is the sum of the k -simplices in the triangulation, we get a k -cochain; thus k -cochains are a discretization of k -forms. One can further map a continuous k -form to a k -cochain. To do this, first integrate the k -form on each k -simplex and assign the resulting value to that simplex to obtain a k -cochain on the k -simplicial complex. This k -cochain is a discrete representation of the original k -form.

3.4.1 Evaluation of a Form on a Chain

We can now naturally extend the notion of evaluation of a differential form ω on an arbitrary chain simply by linearity:

$$\int_{\sum_i c_i \sigma_i} \omega = \sum_i c_i \int_{\sigma_i} \omega. \quad (4)$$

As mentioned above, the integration of ω on each k -simplex σ_k provides a discretization of ω or, in other words, a mapping from the k -form ω to a k -cochain represented by:

$$\omega[\hat{i}] = \int_{\sigma_i} \omega.$$

However convenient this chain/cochain standpoint is, in practical applications, one often needs a point-wise value for a k -form or to evaluate the integration on a particular k -submanifold. How do we get these values from a k -cochain? We will cover this issue of *form interpolation* in Section 6.

4 Operations on Chains and Cochains

4.1 Discrete Exterior Derivative

In the present discrete setting where the discrete differential forms are defined as cochains, defining a discrete exterior derivative can be done very elegantly: Stokes' theorem, mentioned early on in Section 2, can be used to *define* the exterior derivative d . Traditionally, this theorem states a vector identity equivalent to the well-known curl, divergence, Green's, and Ostrogradsky's theorems. Written in terms of forms, the identity becomes quite simple: it states that d applied to an arbitrary form ω is evaluated on an arbitrary simplex σ as follows:

$$\int_{\sigma} d\omega = \int_{\partial\sigma} \omega. \quad (5)$$

You surely recognize the usual property that an integral over a k -dimensional set is turned into a boundary integral (*i.e.*, over a set of dimension $k-1$). With this simple equation relating the evaluation of $d\omega$ on a simplex σ to the evaluation of ω on the boundary of this simplex, the exterior derivative is *readily defined*: each time you encounter an exterior derivative of a form, replace any evaluation over a simplex σ by a direct evaluation of the form itself over the boundary of σ . Obviously, Stokes' theorem will be enforced by construction!

4.1.1 Coboundary Operator

The operator d is called the *adjoint* of the boundary operator ∂ : if we denote the integral sign as a pairing, *i.e.*, with the convention that $\int_{\sigma} \omega = [\omega, \sigma]$, then applying d on the left hand side of this operator is equivalent to applying ∂ on the right hand: $[d\omega, \sigma] = [\omega, \partial\sigma]$. For this very reason, d is sometimes called the coboundary operator.

Finally, by linearity of integration, we can write a more general form of Stokes' theorem, now extended to arbitrary chains as follows:

$$\int_{\sum_i c_i \sigma_i} d\omega = \int_{\partial(\sum_i c_i \sigma_i)} \omega = \int_{\sum_i c_i \partial\sigma_i} \omega = \sum_i c_i \int_{\partial\sigma_i} \omega.$$

Consider the example shown in Figure 7. The discrete exterior derivative of the 1-form, defined as numbers on edges, is a 2-form represented by numbers on oriented faces. The orientation of the 1-forms may be opposite to that induced on the edges by the orientation of the faces. In this case, the values on the edges change sign. For instance, the 2-form associated with the d of the 1-forms surrounding the oriented shaded triangle takes the value $\omega = 2 - 1 - 0.75 = 0.25$.




Figure 7: Given a 1-form as numbers on oriented edges, its discrete exterior derivative is a 2-form. In particular, this 2-form is valued 0.25 on the oriented shaded triangle.

4.1.2 Implementation of Exterior Derivative

Since we use vectors of dimension $|\mathcal{K}^k|$ to represent a k -cochain, the operator d can be represented by a matrix of dimension $|\mathcal{K}^{k+1}| \times |\mathcal{K}^k|$. Furthermore, this matrix has a trivial expression. Indeed, using the matrix notation introduced earlier, we have:

$$\int_{\partial c} \omega = \omega^t (\partial c) = (\omega^t \partial) c = (\partial^t \omega)^t c = \int_c d\omega.$$

Thus, the matrix d is simply equal to ∂^t . This should not come as a surprise, since we previously discussed that d is simply the adjoint of ∂ . Note that care should be used when boundaries are present. However, and without digging too much into the details, it turns out that even for discrete manifolds *with* boundaries, the previous statement is valid. Implementing the exterior derivative while preserving Stokes' theorem is therefore a trivial matter in practice. Notice that just like for the boundary operator, there is actually more

than one matrix for the exterior derivative operator: there is one per simplex dimension. But again, the context is sufficient to actually know which matrix is needed. A brute force approach that gets rid of these multiple matrices is to use a notion of super-chain, *i.e.*, a vector storing *all* simplices, ordered from dimension 0 to the dimension of the space: in this case, the exterior derivative can be defined as a single, large sparse matrix that contains these previous matrices as blocks along the diagonal. We will not use this approach, as it makes the exposition less intuitive in general.

4.2 Exact/Closed Forms and Poincaré Lemma

A k -form ω is called exact if there is a $(k-1)$ -form α such that $\omega = d\alpha$, and it is called closed if $d\omega = 0$.




Figure 8: (a) The 2-form on the oriented shaded triangles defined by the exterior derivative d of the 1-form on the oriented edges is called an exact 2-form; (b) The 1-form on the oriented edges whose derivative d is identically zero is called a closed 1-form.

It is worth noting here that every exact form is closed, as will be seen in Section 4.3. Moreover, it is well-known in the continuous setting that a closed form on a smooth contractible (sub)-manifold is locally exact (to be more accurate: exact over any disc-like region). This result is called the *Poincaré lemma*. The discrete analogue to this lemma can be stated as follows: given a closed k -cochain ω on a star-shaped complex, that is to say, $d\omega = 0$, there exists a $(k-1)$ -cochain α such that $\omega = d\alpha$. For a formal statement and proof of this discrete version, see [Desbrun et al. 2004].

4.3 Introducing the deRham Complex

The boundary of a boundary is the empty set. That is, the boundary operator applied twice to a k -simplex is zero. Indeed, it is easy to verify that $\partial \partial \sigma_k = 0$, since each $(k-2)$ -simplex will appear exactly twice in this chain with different signs and, hence, cancel out (try it at home!). From the linearity of ∂ , one can readily conclude that the property $\partial \partial = 0$ is true for all k -chains since the k -simplices form a basis. Similarly, one has that the discrete exterior derivative satisfies $d d = \partial^t \partial^t = (\partial \partial)^t = 0$, analogously to the exterior derivative of differential forms (notice that this last equality corresponds to the equality of mixed partial derivatives, which in turn is responsible for identities like $\nabla \times \nabla = 0$ and $\nabla \cdot \nabla \times = 0$ in \mathbb{R}^3).




Figure 9: The chain complex of a tetrahedron with the boundary operator: from the tet, to its triangles, to their edges, and to their vertices.

4.3.1 Chain Complex

In general, a *chain complex* is a sequence of linear spaces, connected with a linear operator D that satisfies the property $D D = 0$. Hence, the boundary operator ∂ (resp., the coboundary operator d) makes the spaces of chains (resp., cochains) into a chain complex, as shown in Figures 9 and 13.

When the spaces involved are the spaces of **differential forms**, and the operator is the exterior derivative d , this chain complex is called

the *deRham complex*. By analogy, the chain complex for the spaces of **discrete** forms and for the coboundary operator is called the discrete deRham complex (or sometimes, the cochain complex).

4.3.2 Examples

Consider the 2D simplicial complex in Figure 10(a) and choose the oriented basis of the i -dimensional simplices ($i = 0$ for vertices, $i = 1$ for edges and $i = 2$ for the face) as suggested by the ordering in the figure.




Figure 10: Three examples of simplicial complexes. The first one is not manifold. The two others are.

One gets $\partial(f_0) = e_0 - e_4 - e_3$; this can be identified with the vector $(1, 0, 0, -1, -1)$, representing the coefficient in front of each simplex. By repeating similar calculations for all simplices, one can readily conclude that the boundary operator ∂ is given by:

$$\partial_2 = \begin{pmatrix} 1 \\ 0 \\ 0 \\ -1 \\ -1 \end{pmatrix}, \quad \partial_1 = \begin{pmatrix} -1 & 0 & 0 & -1 & 0 \\ 1 & -1 & 0 & 0 & 1 \\ 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & -1 & 1 & -1 \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix},$$

That is, the chain complex under the boundary operator ∂ can be written as:

$$0 \longrightarrow C_2 \xrightarrow{\partial_2} C_1 \xrightarrow{\partial_1} C_0 \longrightarrow 0$$

where C_i , $i = 0, 1, 2$, denote the spaces of i -chains.

Consider now the domain to be the mesh shown in Figure 10(b). The exterior derivative operator, or the coboundary operator, can be expressed as:

$$d^0 = \begin{pmatrix} -1 & 1 & 0 & 0 \\ 0 & -1 & 1 & 0 \\ 1 & 0 & 0 & -1 \\ 0 & 1 & 0 & -1 \\ 0 & -1 & 0 & 1 \end{pmatrix}, \quad d^1 = \begin{pmatrix} 1 & 0 & 0 & 1 & 1 \\ 0 & 1 & 1 & 0 & -1 \end{pmatrix}.$$

It is worth noting that, since d is adjoint to ∂ by definition, the coboundary operator d induces a cochain complex:

$$0 \longleftarrow C^2 \xleftarrow{d^1} C^1 \xleftarrow{d^0} C^0 \longleftarrow 0$$

where C^i , $i = 0, 1, 2$, denote the spaces of i -cochains.

Finally, suppose the domain is the tetrahedron in Figure 10(c), then the exterior derivative operators are:

$$d^0 = \begin{pmatrix} -1 & 1 & 0 & 0 \\ 0 & -1 & 1 & 0 \\ -1 & 0 & 0 & -1 \\ 0 & 1 & 0 & -1 \\ 0 & 0 & -1 & 1 \end{pmatrix}, \quad d^1 = \begin{pmatrix} 1 & 1 & -1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 1 & -1 & 0 \\ 0 & 1 & 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 1 & 0 & 1 \end{pmatrix}, \quad d^2 = (-1 \ 1 \ 1 \ -1).$$

4.4 Notion of Homology and Cohomology

Homology is a concept dating back to Poincaré that focuses on studying the topological properties of a space. Loosely speaking, homology does so by counting the number of holes. In our case, since we assume that our space is a simplicial complex (*i.e.*, triangulated), we will only deal with *simplicial homology*, a simpler, more straightforward type of homology that can be seen as a discrete version of the continuous definition (in other words, it is equivalent to the continuous one if the domain is triangulated). As we are about to see, the notion of discrete forms is intimately linked with these topological notions. In fact, we will see that (co)homology is the study of the relationship between *closed* and *exact* (co)chains.

4.4.1 Simplicial Homology

A fundamental problem in topology is that of determining, for two spaces, whether they are topologically equivalent. That is, we wish to know if one space can be morphed into the other without having to puncture it. For instance, a sphere-shaped tet mesh is not topologically equivalent to a torus-shaped tet mesh as one cannot alter the sphere-shaped mesh (*i.e.*, deform, refine, or coarsen it locally) to make it look like a torus.

The key idea of homology is to define *invariants* (*i.e.*, quantities that cannot change by continuous deformation) that characterize topological spaces. The simplest invariant is the number of connected components that a simplicial complex has: obviously, two simplicial complexes with different numbers of pieces cannot be continuously deformed into each other! Roughly speaking, homology groups are an extension of this idea to define more subtle invariants than the number of connected components. In general, one can say that homology is a way to *define* the notion of holes/voids/tunnels/components of an object in any dimension.

Cycles and their Equivalence Classes Generalizing the previous example to other invariants is elegantly done using the notion of *cycles*. A cycle is simply a *closed k -chain*; that is, a linear combination of k -simplices so that the boundary of this chain (see Section 3.2) is the empty set. Any set of vertices is a closed chain; any set of 1D loops are too; etc. Equivalently, a k -cycle is any k -chain that belongs to $\text{Ker } \partial_k$, by definition.

On this set of all k -cycles, one can define *equivalence classes*. We will say that a k -cycle is *homologous* to another k -cycle (*i.e.*, in the same equivalence class than the other) when these two chains differ by a boundary of a $(k+1)$ -chain (*i.e.*, by an *exact chain*). Notice that this exact chain is, by definition (see Section 4.2), in the image of ∂_{k+1} , *i.e.*, $\text{Im } \partial_{k+1}$. To get a better understanding of this notion of equivalence class, the reader is invited to look at Figure 11: the 1-chains L_1 and L_3 are part of the same equivalence class as their difference is indeed the boundary of a well-defined 2D chain—a rubber-band shape in this case. Notice that as a consequence, L_1 can be deformed into L_3 without having to tear the loop apart. However, L_2 is not of this class, and thus cannot be deformed into L_3 ; there's no 2-chain that corresponds to their difference.

4.4.2 Homology Groups

Let us now use these definition in the simple case of the 0^{th} homology group \mathcal{H}_0 .

Homology Group \mathcal{H}_0 The boundary of any vertex is \emptyset . Thus, any linear combination of vertices is a 0-cycle by definition. Now if two vertices v_0 and v_1 are connected by an edge, $v_1 - v_0$ (*i.e.*, the difference of two cycles) is the boundary of this edge. Thus, by our previous definition, two vertices linked by an edge are *homologous* as their difference is the boundary of this edge. By the same reasoning, any two vertices taken from the *same* connected component are, also, homologous, since there exists a chain of edges in between. Consequently, we can pick only one vertex per connected component to form a basis of this homology group. Its dimension, β_0 , is therefore simply the number of connected components. The basis elements of that group are called *generators*, since they generate the whole homology group.

Homology Group \mathcal{H}_1 Let us proceed similarly for the 1^{st} homology class: we now have to consider 1-cycles (linear combinations of 1D loops). Again, one can easily conceive that there are different *types* of such cycles, and it is therefore possible to separate all possible cycles into different equivalence classes. For instance,

the loop L_1 in Figure 11 is topologically distinct from the curve L_2 : one is around a hole and the other is not, so the difference between the two is *not* the boundary of a 2-chain. Conversely, L_1 is in the same class as curve L_3 since they differ by one connected area. Thus, in this figure, the 1st homology group is a 1-dimensional group, and L_1 (or L_3 , equivalently) is its unique generator. The reader is invited to apply this simple idea on a triangulated torus, to find two loops as generators of \mathcal{H}_1 .




Figure 11: Example of Homology Classes: the cycles L_1 and L_2 are topologically distinct as one encloses a hole while the other does not; L_1 and L_3 are however in the same equivalence class.

Formal Definition of Homology Groups We are now ready to generalize this construction to all homology groups. Remember that we have a series of k -chain spaces:

$$C_n \xrightarrow{\partial_n} C_{n-1} \cdots \xrightarrow{\partial_2} C_1 \xrightarrow{\partial_1} C_0$$

with the property that $\partial \partial$ is the empty set. This directly implies that the image of C_j through ∂_j is always in the kernel of ∂_{j-1} —such a series is called a *chain complex*. Now, the homology groups $\{\mathcal{H}_k\}_{k=0..n}$ of a chain complex based on ∂ are defined as the following quotient spaces:

$$\mathcal{H}_k = \text{Ker } \partial_k / \text{Im } \partial_{k+1}.$$

The reader is invited to check that this definition is *exactly* what we did for the 0th and 1st homology groups—and it is now valid for any order: indeed, we use the fact that closed chains (belonging to $\text{Ker } \partial$) are homologous *iff* their difference is in $\text{Im } \partial$, and this is exactly what this quotient vector space is.

Example Consider the example in Figure 10(a). Geometrically, \mathcal{H}_0 is nontrivial because the simplicial complex σ is disconnected (it is easy to see $\{v_0, v_4\}$ form a basis for \mathcal{H}_0), while \mathcal{H}_1 is nontrivial since the cycle $(e_1 - e_2 + e_4)$ is not the boundary of any 2-chain of σ ($\{(e_1 - e_2 + e_4)\}$ is indeed a basis for this 1D space \mathcal{H}_1).

Link to Betti Numbers The dimension of the k -th cohomology group is called k -th Betti number; $\beta_k = \dim \mathcal{H}_k$. For a 3D simplicial complex embedded in \mathbb{R}^3 , these numbers have very straightforward meanings. β_0 is the number of connected components, β_1 is the number of tunnels, β_2 is the number of voids, while β_3 is the number of 4D holes, which is 0 in the Euclidean (flat 3D) case. Finally, note that $\sum_{k=0..n} (-1)^k \beta_k$, where β_k is the k -th Betti number, gives us the well-known Euler characteristic.

4.4.3 Cohomology Groups

The definition of homology groups is much more general than what we just reviewed. In fact, the reader can take the formal definition in the previous section, replace all occurrences of chain by cochain, of ∂ by d , and reverse the direction of the operator between spaces (see Section 4.3.2): this will also define equivalence classes. Because cochains are dual of chains, and d is the adjoint of ∂ , these equivalence classes define what are actually denoted as *cohomology groups*: the cohomology groups of the deRham complex for the coboundary operator are simply the quotient spaces $\text{Ker } d / \text{Im } d$. Finally, note that the homology and cohomology groups are not only dual notions, but they are also isomorphic; therefore, the cardinalities of their bases are equal.

4.4.4 Calculation of the Cohomology Basis

One usual way to calculate a cohomology basis is to calculate a Smith Normal Form to obtain the homology basis first (possibly using progressive meshes [Gu and Yau 2003]), with a worst case complexity of $O(n^3)$, and then find the corresponding cohomology basis derived from this homology basis. We provide an alternative method here with worst case complexity also equal to $O(n^3)$. The advantage of our method is that it directly calculates the cohomology basis.

Our algorithm is a modified version of an algorithm in [Edelsbrunner et al. 2000], although they did not use it for the same purpose². We will use *teatrmmrow#(.)* to refer to the row number of the *last nonzero coefficient* in a particular column.

The procedure is as follows:

1. Transform d^k (size $|\mathcal{K}^{k+1}| \times |\mathcal{K}^k|$) in the following manner:

```

// For each column of  $d^k$ 
for( $i = 0; i < |\mathcal{K}^k|; i++$ )
  // Reduce column  $i$ 
  repeat
     $p \leftarrow \text{row\#}(d^k[i])$ 
    find  $j < i$  such that  $p = \text{row\#}(d^k[j])$ 
    make  $d^k[i][p]$  zero by adding to  $d^k[i]$  a multiple of  $d^k[j]$ 
  until  $j\_not\_found$  or column  $i$  is all zeros

```

At the end of this procedure, we get $D^k = d^k N^k$, whose non-zero column vectors are linearly independent of each other and with different $\text{row\#}(\cdot)$, and N^k is a non-singular upper triangular matrix.

2. Construct $K^k = \{N_i^k \mid D_i^k = 0\}$ (where N_i^k and D_i^k are column vectors of matrices N^k and D^k respectively). K^k is a basis for kernel of d^k .
3. Construct $I^k = \{N_i^k \mid \exists j \text{ such that } i = \text{row\#}(D_j^{k-1})\}$
4. Construct $P^k = K^k - I^k$
 P^k is a basis of the cohomology.

Short proof of correctness: First, notice that the N_i^k 's are all linearly independent because N^k is nonsingular. For any nonzero linear combination of vectors in P^k , $\text{row\#}(\cdot)$ of it (say i) equals the max of $\text{row\#}(\cdot)$ of vectors with non-zero coefficients. But i is not $\text{row\#}(\cdot)$ of any $D_j^{(k-1)}$ (and thus any linear combination of them) by definition of P^k . Therefore, we know that the linear combination is not in the image space of d^{k-1} (since the range of d^{k-1} is the same as D^{k-1} , by construction). Thus, P^k spans a subspace of $\text{Ker}(d^k) / \text{Im}(d^{k-1})$ of dimension $\text{Card}(P^k)$.

One can also prove that I^k is a subset of K^k . Pick an N_i^k with $i = \text{row\#}(D_j^{(k-1)})$. We have: $d^k D_j^{(k-1)} = 0$ (since $d^k \circ d^{k-1} = 0$). Now $\text{row\#}(\tau \equiv (N^k)^{-1} d^{(k-1)}_j) = i$ (the inverse of an upper triangular matrix is also an upper triangular matrix). So consequently, $0 = d^k d^{(k-1)}_j = D^k (N^k)^{-1} d^{(k-1)}_j = D^k \tau$ means that $D_i^k = 0$ because the columns of D^k are linearly independent or 0. Therefore, $\text{Card}(P^k) = \text{Card}(K^k) - \text{Card}(I^k) = \dim(\text{Ker}(d^k)) - \dim(\text{Im}(d^{(k-1)}))$, and we conclude that, P^k spans $\text{Ker}(d^k) / \text{Im}(d^{k-1})$ as expected. ■

²Thanks to David Cohen-Steiner for pointing us to the similarities

4.4.5 Example

Consider the 2D simplicial complex in Figure 10(a) again. We will show an example of running the same procedure described above to compute a homology basis. The only difference with the previous algorithm is that we use ∂ instead of d , since we compute the homology basis instead of the cohomology basis.

1. Compute the $D^K = \partial_k N^{k,s}$ and $N^{k,s}$: D^2 is trivial, as it is the same as ∂_2 .

$$D^1 = \begin{pmatrix} -1 & 0 & 0 & 0 & 0 \\ 1 & -1 & 0 & 0 & 0 \\ 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & -1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix}, \quad N^1 = \begin{pmatrix} 1 & 0 & 0 & -1 & 0 \\ 0 & 1 & 0 & -1 & 1 \\ 0 & 0 & 1 & 1 & -1 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{pmatrix},$$

2. Construct the K 's:

$$K^0 = \left\{ \begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 1 \\ 0 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 0 \\ 1 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 0 \\ 0 \\ 1 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 1 \end{pmatrix} \right\} \\ = \{v_0, v_1, v_2, v_3, v_4\}$$

(N^0 is the identity)

$$K^1 = \left\{ \begin{pmatrix} -1 \\ -1 \\ 1 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 1 \\ -1 \\ 0 \\ 1 \end{pmatrix} \right\} = \{(-e_0 - e_1 + e_2 + e_3), (e_1 - e_2 + e_4)\}$$

3. Construct the I 's:

$$I^0 = \{v_1 \ (1 = \text{row}\#(D_0^1)), \\ v_2 \ (2 = \text{row}\#(D_1^1)), \\ v_3 \ (3 = \text{row}\#(D_2^1))\} \\ I^1 = \{(e_1 - e_2 + e_4) \ (4 = \text{row}\#(D_0^2))\}$$

4. Consequently, the homology basis is:

$$P^0 = \{v_0, v_1, v_2, v_3, v_4\} - \{v_1, v_2, v_3\} = \{v_0, v_4\} \\ P^1 = \{(-e_0 - e_1 + e_2 + e_3)\}$$

This result confirms the basis we gave in the example of Section 4.4.2. (Note that $-(-e_0 - e_1 + e_2 + e_3) - (e_1 - e_2 + e_4) = e_0 - e_4 - e_3 = \partial f_0$, thus $(-e_0 - e_1 + e_2 + e_3)$ spans the same homology space as $(e_1 - e_2 + e_4)$).

4.5 Dual Mesh and its Exterior Derivative

Let us introduce the notion of *dual mesh* of triangulated manifolds, as we will see that it is one of the key components of our discrete calculus. The main idea is to associate to each *primal* k -simplex a *dual* $(n-k)$ -cell. For example, consider the tetrahedral mesh in Figure 13, we associate a dual 3-cell to each primal vertex (0-simplex), a dual polygon (2-cell) to each primal edge (1-simplex), a dual edge (1-cell) to each primal face (2-simplex), and a dual vertex (0-cell) to the primal tet (3-simplex). By construction, the number of dual $(n-k)$ -cells is equal to that of primal k -simplices. The collection of dual cells is called a *cell complex*, which need not be a simplicial complex in general.

Yet, this dual complex inherits several properties and operations from the primal simplicial complex. Most important is the notion of *incidence*. For instance, if two primal edges are on the same primal face, then the corresponding dual faces are incident, that is, they share a common dual edge (which is the dual of the primal common face). As a result of this incidence property, one may easily derive

a boundary operator on the dual cell complex and, consequently, a discrete exterior derivative! The reader is invited to verify that this exterior derivative on the dual mesh can be simply written as the opposite of a primal one *transposed*:

$$d_{Dual}^{n-k} = (-1)^k (d_{Primal}^{k-1})^t. \quad (6)$$

The added negative sign appears as the orientation on the dual is induced from the primal orientation, and must therefore be properly accounted for. Once again, an implementation can overload the definition of this operator d when used on dual forms using this previous equation. In the remainder of our chapter, we will be using d as a contextual operator to keep the notations as simple as possible. Because we have defined a proper exterior derivative on the dual mesh (still satisfying $d \circ d = 0$), this dual cell complex also carries the structure of a chain complex. The structure on the dual complex may be linked to that of the primal complex using the Hodge star (a metric-dependent operator), as we will discuss in Section 5.




Figure 12: A 2-dimensional example of primal and dual mesh elements. On the top row, we see the primal mesh (a triangle) with a representative of each simplicial complex being highlighted. The bottom row shows the corresponding circumcentric dual cells (restricted to the triangle).

4.5.1 Dualization: The $*$ Operator

For simplicity, we use circumcentric (or Voronoi) duality to construct the dual cell complex. The circumcenter of a k -simplex is defined as the center of the k -circumsphere, which is the unique k -sphere that has all $k+1$ vertices of the k -simplex on its surface. In Figure 12, we show examples of circumcentric dual cells of a 2D mesh. The dual 0-cell associated with the triangular face is the circumcenter of the triangle. The dual 1-cell associated with one of the primal edges is the line segment that joins the circumcenter of the triangle to the circumcenter of that edge, while the dual 2-cell associated with a primal vertex is corner wedge made of the convex hull of the circumcenter of the triangle, the two centers of the adjacent edges, and the vertex itself (see Figure 12, bottom left). Thereafter, we will denote as $*$ the operation of *duality*; that is, a primal simplex σ will have its dual called $*\sigma$ with the orientation induced by the primal orientation and the manifold orientation. For a formal definition, we refer the reader to [Hirani 2003] for instance. It is also worth noting that other notions of duality such as the barycentric duality may be employed. For further details on dual cell (or “block”) decompositions, see [Munkres 1984].

4.5.2 Wedge Product

In the continuous setting, the wedge product \wedge is an operation used to construct higher degree forms from lower degree ones; it is the antisymmetric part of the tensor product. For example, let α and β be 1-forms on a subset $\mathcal{R} \subset \mathbb{R}^3$, their wedge product $\alpha \wedge \beta$ is a 2-form on \mathcal{R} . In this case, one can relate the wedge product to the cross product of vector fields on \mathcal{R} . Indeed, if one considers the vector representations of α and β , the vector proxy to $\alpha \wedge \beta$ is the cross product of the two vectors. Similarly, the wedge product of a

1-form γ with the 2-form $\omega = \alpha \wedge \beta$ is a 3-form $\mu = \gamma \wedge \omega$ (also called volume-form) on \mathcal{R} which is analogous to the scalar triple product of three vectors.

A discrete treatment of the wedge operator can be found in [Hirani 2003]. Here, we only need to introduce the notion of a discrete *primal-dual wedge product*: given a primal k -cochain γ and a dual $(n-k)$ -cochain ω , the discrete wedge product $\gamma \wedge \omega$ is an n -form (or a volume-form). For instance, in the example depicted in the inset, the wedge product of the primal 1-cochain with the dual 1-cochain is a 2-form associated with the diamond region defined by the convex hull of the union between the primal *and* dual edge (see inset).



5 Metric-Dependent Operators on Forms

Notice that up to now, we did *not* assume that a metric was available, *i.e.*, we never required anything to be *measured*. However, such a metric is necessary for many purposes. For instance, simulating the behavior of objects around us requires measurements of various parameters in order to be able to model laws of motion, and compare the numerical results of simulations. Consequently, a certain number of operations on forms can only be defined once a metric is known, as we shall see in this section.

5.1 Notion of Metric and Inner Product

A *metric* is, roughly speaking, a nonnegative function that describes the “distance” between neighboring points of a given space. For example, the Euclidean metric assigns to any two points in the Euclidean space \mathbb{R}^3 , say $\mathbf{X} = (x_1, x_2, x_3)$ and $\mathbf{Y} = (y_1, y_2, y_3)$, the number:

$$d(\mathbf{X}, \mathbf{Y}) = \|\mathbf{X} - \mathbf{Y}\| = \sqrt{(x_1 - y_1)^2 + (x_2 - y_2)^2 + (x_3 - y_3)^2}$$

defining the “standard” distance between any two points in \mathbb{R}^3 . This metric then allows one to measure length, area, and volume. The Euclidean metric can be expressed as the following quadratic form:

$$g^{\text{Euclid}} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}.$$

Indeed, the reader can readily verify that this matrix g satisfies: $d^2(\mathbf{X}, \mathbf{Y}) = (\mathbf{X} - \mathbf{Y})^t g (\mathbf{X} - \mathbf{Y})$. Notice also that this metric *induces an inner product* of vectors. Indeed, for two vectors \mathbf{u} and \mathbf{v} , we can use the matrix g to define:

$$\mathbf{u} \cdot \mathbf{v} = \mathbf{u}^t g \mathbf{v}.$$

Once again, the reader is invited to verify that this equality does correspond to the traditional dot product when g is the Euclidean metric. Notice that on a non-flat manifold, subtraction of two points is only possible for points infinitesimally close to each other, thus the metric is actually defined pointwise for the tangent space at each point: it does not have to be constant. Finally, notice that a volume form can be induced from a metric by defining $\mu^n = \sqrt{\det(g)} dx^1 \wedge \dots \wedge dx^n$.

5.2 Discrete Metric

In the discrete setting presented in this paper, we only need to measure length, area, and volume of the simplices and dual cells (note

these different notions of sizes depending on dimension will be denoted “intrinsic volumes” for generality). We therefore do not have a full-blown notion of a metric, only a *discrete metric*. Obviously, if one were to use a finer mesh, more information on the metric would be available: having more values of length, area, and volume in a neighborhood provides a better approximation of the real, continuous metric.

5.3 The Differential Hodge Star

Let us go back for a minute to the differential case to explain a new concept. Recall that the metric defines an inner product for vectors. This notion also extends to forms: given a metric, one can define the product of two k -forms $\in \Omega^k(\mathcal{M})$ which will measure, in a way, the projection of one onto the other. A formal definition can be found in [Abraham et al. 1988]. Given this inner product denoted $\langle \cdot, \cdot \rangle$, we can introduce an operator \star , called the *Hodge star*, that maps a k -form to a complementary $(n-k)$ -form:

$$\star : \Omega^k(\mathcal{M}) \rightarrow \Omega^{n-k}(\mathcal{M}),$$

and is defined to satisfy the following equality:

$$\alpha \wedge \star \beta = \langle \alpha, \beta \rangle \mu^n$$

for any pair of k -forms α and β (recall that μ^n is the volume form induced by the metric g). However, notice that the wedge product is very special here: it is the product of k -form and a $(n-k)$ -form, two complementary forms. This fact will drastically simplify the discrete counterpart of the Hodge star, as we now cover.




Figure 13: On the first line, the ‘primal’ chain complex is depicted and on the second line we see the dual chain complex (*i.e.*, cells, faces, edges and vertices of the Voronoi cells of each vertex of the primal mesh).

5.4 Discrete Hodge Star

In the discrete setting, the Hodge star becomes easier: we only need to define how to go from a *primal* k -cochain to a *dual* $(n-k)$ -cochain, and vice-versa. By definition of the dual mesh, k -chains and dual $(n-k)$ -chains are represented by vectors of the same dimension. Similarly to the discrete exterior derivative (coboundary) operator, we may use a matrix (this time of size $|\mathcal{K}^k| \times |\mathcal{K}^{n-k}|$) to represent the Hodge star. Now the question is: what should the coefficients of this matrix be?

For numerical purposes we want it to be symmetric, positive definite, and sometimes, even diagonal for faster computations. One such diagonal Hodge star can be defined with the diagonal elements as the ratio of intrinsic volumes of a k -simplex and its dual $(n-k)$ -simplex. In other words, we can define the discrete Hodge star through the following simple rule:

$$\frac{1}{|\sigma_k|} \int_{\sigma_k} \omega = \frac{1}{|\star \sigma_k|} \int_{\star \sigma_k} \star \omega \quad (7)$$

Therefore, any primal value of a k -form can be easily *transferred* to the dual mesh through proper scaling—and vice-versa; to be precise, we have:

$$\star_k \star_{n-k} = (-1)^{k(n-k)} \text{Id}, \quad (8)$$

which means that \star on the dual mesh is the inverse of the \star on the primal *up to a sign* (the result of antisymmetry of the wedge product, which happens to be positive for any k -form when $n = 3$).

So we must use the inverse of the Hodge star to go from a dual $(n-k)$ -cochain to a k -cochain. We will, however, use \star to indistinguishably mean either the star or its inverse, as there is no ambiguity once we know whether the operator is applied to a primal or a dual form: this is also a context-dependent operator.

Implementation Based on Eq. (7), the inner product of forms α^k and β^k at the diamond-shaped region formed by each k -simplex and its dual $(n-k)$ -simplex is simply the product of the value of α at that k -simplex and value of $\star\beta$ at that dual $(n-k)$ -simplex. Therefore, the sum over the whole space gives the following inner product (which involves only linear algebra matrix and vector multiplications)

$$\langle \alpha^k, \beta^k \rangle = \alpha^t \star \beta. \quad (9)$$

where the Hodge star matrix has, as its only nonzero coefficients, the following diagonal terms:

$$(\star_k)_{qq} = |(\star\sigma)_q|/|(\sigma_q)|.$$

Notice that this definition of the inner product, when $\alpha = \beta$, induces the definition of the norm of k -forms.

Again, there are three different Hodge stars in \mathbb{R}^3 , one for each simplex dimension. But as we discussed for all the other operators, the dimension of the form on which this operator is applied disambiguates which star is meant. So we will not encumber our notation with unnecessary indices, and will only use the symbol \star for any of the three stars implied.

The development of an accurate, yet fast to compute, Hodge star is still an active research topic. However, this topic is beyond the scope of the current chapter.

5.5 Discrete Codifferential Operator δ

We already have a linear operator d which maps a k -form to a $k+1$ -form, but we do not have a linear operator which maps a k -form to a $(k-1)$ -form. Having defined a discrete Hodge star, one can now create such an *adjoint* operator δ for the discrete exterior derivative d . Here, adjoint is meant with respect to the inner product of forms; that is, this operator δ satisfies:

$$\langle d\alpha, \beta \rangle = \langle \alpha, \delta\beta \rangle \quad \forall \alpha \in \Omega^{k-1}(M), \beta \in \Omega^k(M)$$

For a smooth, compact manifold without boundary, one can prove that the operator $(-1)^{n(k-1)+1} \star d\star$ satisfies the above condition [Abraham et al. 1988]. Let us try to use the same definition in the discrete setting; *i.e.*, we wish to define the discrete δ applied to k -forms by the relation:

$$\delta \equiv (-1)^{n(k-1)+1} \star d\star, \quad (10)$$

Beware that we use the notation d to mean the context-dependent exterior derivative. If you apply δ to a primal k -form, then the exterior derivative will be applied to a *dual* $(n-k)$ -form, and thus,

Equation 6 should be used. Once this is well understood, it is quite straightforward to verify the following series of equalities:

$$\begin{aligned} \langle d\alpha, \beta \rangle &\stackrel{\text{Eq. (9)}}{=} (d\alpha)^t \star \beta = \alpha^t d^t \star \beta \\ &\stackrel{\text{Eq. (8)} \ w/ \ k \leftrightarrow k-1}{=} \alpha^t (-1)^{(k-1)(n-(k-1))} \star \star d^t \star \beta \\ &= \alpha^t (-1)^{n(k-1)+1} \star \star (-1)^k d^t \star \beta \\ &\stackrel{\text{Eq. (10)}}{=} \langle \alpha, \delta\beta \rangle \end{aligned}$$

holds on our discrete manifold. So indeed, the discrete d and δ are also adjoint, in a similar fashion in the discrete setting as they were in the continuous sense. For this reason, δ is called the *codifferential operator*.

Implementation of the Codifferential Operator Thanks to this easily-proven adjointness, the implementation of the discrete codifferential operator is a trivial matter: it is simply the product of three matrices, mimicking exactly the differential definition mentioned in Eq. (10).

5.6 Exercise: Laplacian Operator

At this point, the reader is invited to perform a little exercise. Let us first state that the Laplacian Δ of a form is defined as: $\Delta = \delta d + d\delta$. Now, applied to a 0-form, notice that the latter term disappears. Question: in 2D, what *is* the Laplacian of a function f at a vertex i ? The answer is actually known: it is the now famous *cotangent formula* [Pinkall and Polthier 1993], since the ratio of primal and dual edge lengths leads to such a trigonometric equality. Applied to a 1-form, however, the expression does have both terms as explicitly given in [Fisher et al. 2007].

6 Interpolation of Discrete Forms

In Section 3.4, we argued that k -cochains are discretizations of k -forms. This representation of discrete forms on chains, although very convenient in many applications, is not sufficient to fulfill certain demands such as obtaining a pointwise value of the k -form. As a remedy, one can use an interpolation of these chains to the rest of space. For simplicity, these interpolation functions can be taken to be linear (by linear, we mean with respect to the coordinates of the vertices).

6.1 Interpolating 0-forms

It is quite obvious how to linearly interpolate discrete 0-forms (as 0-cochains) to the whole space: we can use the usual vertex-based linear interpolation basis, often referred to as the *hat function* in the Finite Element literature. This basis function will be denoted as φ_i for each vertex v_i . By definition, φ_i satisfies:

$$\varphi_i = 1 \quad \text{at } v_i, \quad \varphi_i = 0 \quad \text{at } v_j \neq v_i$$

while φ_i linearly goes to zero in the one-ring neighborhood of v_i . The reader may be aware that these functions are, within each simplex, *barycentric coordinates*, introduced by Möbius in 1827 as mass points to define a *coordinate-free geometry*.

With these basis functions, one can easily check that if we denote a vertex v_j by σ_j , we have:

$$\int_{v_j} \varphi_{v_i} = \int_{\sigma_j} \varphi_{\sigma_i} = \int_{\sigma_j} \varphi_i = \begin{cases} 1 & \text{if } i = j, \\ 0 & \text{if } i \neq j. \end{cases}$$

Therefore, these interpolating functions represent a basis of 0-cochains, that exactly corresponds to the dual of the natural basis of 0-chains.

6.2 Interpolating 1-forms

We would like to be able to extend the previous interpolation technique to 1-forms now. Fortunately, there is an existing method to do just that: the *Whitney* 1-form (used first in [Whitney 1957]) associated with an edge σ_{ij} between v_i and v_j is defined as:

$$\varphi_{\sigma_{ij}} = \varphi_i d\varphi_j - \varphi_j d\varphi_i.$$

A direct computation can verify that:

$$\int_{\sigma_{kl}} \varphi_{\sigma_{ij}} = \begin{cases} 1 & \text{if } i = k \text{ and } j = l, \\ -1 & \text{if } i = l \text{ and } j = k, \\ 0 & \text{otherwise.} \end{cases}$$

Indeed, it is easy to see that the integral is 0 when we are not integrating it on edge e_{ij} , because at least one of the vertices (say, i) is not on the edge, thus, $\varphi_i = 0$ and $d\varphi_i = 0$ on the edge. However, along the edge σ_{ij} , we have $\varphi_i + \varphi_j = 1$, therefore:

$$\int_{\sigma_{ij}} \varphi_{\sigma_{ij}} = \int_{\varphi_i=1}^{\varphi_i=0} (\varphi_i d(1-\varphi_i) - (1-\varphi_i) d\varphi_i) = \int_{\varphi_i=1}^{\varphi_i=0} (-d\varphi_i) = 1.$$

We thus have defined a correct basis for 1-cochains.

6.3 Interpolating with Whitney k -Forms

One can extend these 1-form basis functions to arbitrary k -simplices. In fact, Whitney k -forms are defined similarly:

$$\varphi_{\sigma_{i_0, i_1, \dots, i_k}} = k! \sum_{j=0 \dots k} (-1)^j \varphi_{i_j} d\varphi_{i_0} \wedge \dots \wedge \widehat{d\varphi_{i_j}} \dots \wedge d\varphi_{i_k}$$

where $\widehat{d\varphi_{i_p}}$ means that $d\varphi_{i_p}$ is excluded from the product. Notice how this definition exactly matches the case of vertex and edge bases, and extends easily to higher dimensional simplices.

Remark If a metric is defined (for instance, the Euclidean metric), we can simply identify $d\varphi$ with $\nabla\varphi$ for the real calculation. This corresponds to the notion of *sharp* (\sharp), but we will not develop this point other than for pointing out the following remark: the traditional gradient of a linear function f in 2D, known to be constant per triangle, can indeed be re-written à la Whitney:

$$\nabla f = \sum_i f_i \nabla \varphi_i \quad \varphi_i + \varphi_j + \varphi_k = 1 \quad \sum_{i,j,i \neq j} (f_j - f_i)(\varphi_i \nabla \varphi_j - \varphi_j \nabla \varphi_i).$$

The values $(f_j - f_i)$ are the edge values associated with the gradient, i.e., the values of the one-form df .




Figure 14: $\nabla\varphi$ for the vertex on top

Basis of Forms The integration of the Whitney form φ_{σ_k} associated with the k -simplex σ_k will be 1 on that particular simplex, and 0 on all others. Indeed, it is a simple exercise to see that the integration of φ_{σ_k} is 0 on a different k -simplex, because there is *at least one* vertex of this simplex v_j that does not belong to σ_k , so its hat function φ_j is valued 0 everywhere on σ_k . Since φ_j or $d\varphi_j$ appears in every term, the integral of φ_{σ_k} is 0. To see that the integral is 1 on the simplex itself, we can use Stokes' theorem (as our discrete forms satisfy it exactly on simplices): first, suppose $k < n$, and pick a $(k+1)$ -simplex, such that the k -simplex σ_k is a *face* of it. Since it is 0 on other faces, the integral of the Whitney form is equal to the integral of $d\varphi_{\sigma_k} = (k+1)! d\varphi_{i_0} \wedge \dots \wedge d\varphi_{i_k}$ on the $(k+1)$ -simplex, if we use φ_{i_j} as a local reference frame for the integration, $\int_{\sigma_{k+1}} d\varphi_{i_0} \wedge \dots \wedge d\varphi_{i_k}$ is simply the volume of a standard simplex, which is $\frac{1}{(k+1)!}$, thus the integral is 1. The case when $k = n$ is essentially the same as $k = n - 1$.

This means that these Whitney forms are forming a basis of their respective form spaces. In a way, these bases are an extension of the Finite Element bases defined on nodes, or of the Finite Volume elements that are constant per tet.

Note finally that the Whitney forms are not continuous; however, they *are* continuous *along the direction of the k -simplex* (i.e., tangential continuity for 1-forms, and normal continuity for 2-forms); this is the only condition needed to make the integration well defined. In a way, this property is the *least* we can ask them to be. We would lose generality if we were to add any other condition! The interested reader is referred to [Bossavit 1998] for a more thorough discussion on these Whitney bases and their relations to the notion of weak form used in the Finite Element Method.

7 Application to Hodge Decomposition

We now go through a first application of the discrete exterior calculus we have defined up to now. As we will see, the discrete case is often much simpler than its continuous counterpart; yet it captures the same properties.

7.1 Introducing the Hodge Decomposition

It is convenient in some applications to use the Helmholtz-Hodge decomposition theorem to decompose a given continuous vector field or differential form (defined on a smooth manifold \mathcal{M}) into components that are mutually orthogonal (in the \mathcal{L}^2 sense), and easier to compute (see [Abraham et al. 1988] for details). In fluid mechanics for example, the velocity field is generally decomposed into a part that is the gradient of a potential function and a part that is the curl of a stream vector potential (see Section 8.3 for further details), as the latter one is the incompressible part of the flow. When applied to k -forms, this decomposition is known as the *Hodge decomposition for forms* and can be stated as follows:

Given a manifold \mathcal{M} and a k -form ω^k on \mathcal{M} with appropriate boundary conditions, ω^k can be decomposed into the sum of the exterior derivative of a $(k-1)$ -form α^{k-1} , the codifferential of a $(k+1)$ -form β^{k+1} , and a harmonic k -form h^k :

$$\omega^k = d\alpha^{k-1} + \delta\beta^{k+1} + h^k.$$

Here, we use the term *harmonic* to mean that h^k satisfies the equation $\Delta h^k = 0$, where Δ is the Laplacian operator defined as $\Delta = dd + \delta d$. The proof of this theorem is mathematically involved and requires the use of elliptic operator theory and similar tools, as well as a careful study of the boundary conditions to en-

sure uniqueness. The discrete analog that we propose has a very simple and straightforward proof as shown below.

7.2 Discrete Hodge Decomposition

In the discrete setting, the discrete operators such as the exterior derivative and the codifferential can be expressed using matrix representation. This allows one to easily manipulate these operators using tools from linear algebra. In particular, the discrete version of the Hodge decomposition theorem becomes a simple exercise in linear algebra. Note that we will assume a boundaryless domain for simplicity (the generalization to domains with boundary is conceptually as simple).

Theorem 7.1 *Let \mathcal{K} be a discrete manifold and let $\Omega^k(\mathcal{K})$ be the space of discrete Whitney k -forms on \mathcal{K} . Consider the linear operator $d^k : \Omega^k(\mathcal{K}) \rightarrow \Omega^{k+1}(\mathcal{K})$, such that $d^{k+1} \circ d^k = 0$, and a discrete Hodge star which is represented as a symmetric, positive definite matrix. Furthermore, define the codifferential (the adjoint of the operator d) as done in Section 5.5; namely, let $\delta^{k+1} = (-1)^{n(k-1)+1} (\star^k)^{-1} (d^k)^t \star^{k+1}$. In this case, the following orthogonal decomposition holds for all k :*

$$\Omega^k(\mathcal{K}) = d\Omega^{k-1}(\mathcal{K}) \oplus \delta\Omega^{k+1}(\mathcal{K}) \oplus \mathcal{H}^k(\mathcal{K})$$

where \oplus means orthogonal sum, and $\mathcal{H}^k(\mathcal{K})$ is the space of harmonic k -forms on \mathcal{K} , that is, $\mathcal{H}^k(\mathcal{K}) = \{h \mid \Delta^k h = 0\}$.

Proof For notational convenience, we will omit the superscript of the operators when the rank is obvious. We first prove that the three component spaces are orthogonal. Clearly, using the facts that the Laplacian operator Δ is equal to $d\delta + \delta d$ and that d and δ are adjoint operators, one has that $\forall h \in \mathcal{H}^k$:

$$\begin{aligned} \langle \Delta h, h \rangle = 0 &\Rightarrow \langle d\delta h, h \rangle + \langle \delta d h, h \rangle = \langle dh, dh \rangle + \langle \delta h, \delta h \rangle = 0 \\ &\Rightarrow dh = 0 \text{ and } \delta h = 0 \end{aligned}$$

Also, for all $\alpha \in \Omega^{k-1}(\mathcal{K})$ and $\beta \in \Omega^{k+1}(\mathcal{K})$, one has:

$$\langle d\alpha, \delta\beta \rangle = \langle dd\alpha, \beta \rangle = 0$$

and

$$\langle d\alpha, h \rangle = \langle \alpha, \delta h \rangle = 0 \quad \langle h, \delta\beta \rangle = \langle dh, \beta \rangle = 0$$

Now, any k -form that is perpendicular to $d\Omega^{k-1}(\mathcal{K})$ and $\delta\Omega^{k+1}(\mathcal{K})$ must be in $\mathcal{H}^k(\mathcal{K})$, because this means $dh = 0$ and $\delta h = 0$, so $\Delta h = d\delta h + \delta d h = 0$.

Alternatively, we can prove that:

$$\Omega^k(\mathcal{K}) = \Delta\Omega^k(\mathcal{K}) \oplus \mathcal{H}^k(\mathcal{K}).$$

By analogy to the previous argument, it is easy to show that $\Delta\Omega^k$ is orthogonal to \mathcal{H}^k . Additionally, the dimension of these two spaces sum up to the dimension of Ω^k , which means the decomposition is complete. ■

Note that the reader can find a similar proof given in Appendix B of [Frankel 2004], where it is used for Kirchhoff's Circuit Laws. There, Frankel does not mention that we can actually use cochains as the discretization of forms, and his operations using a "metric" of cochains can be interpreted as a Hodge star.

Implementation of the Discrete Hodge Decomposition

Before we discuss how to numerically implement the discrete Hodge decomposition, we prove a useful result (that has a continuous analog).

Lemma 7.2 *In the discrete setting, one can find exactly one harmonic cochain from each cohomology equivalence class.*

Proof It can be readily shown that the bases of harmonic cochains and the cohomology groups both have the dimension equal to $\dim(\text{Ker } d^k) - \dim(\text{Im } d^{k-1})$. To this end, recall that a cohomology basis is defined as is $\text{Ker}(d^k)/\text{Im}(d^{k-1})$ and has dimension $\dim(\text{Ker } d^k) - \dim(\text{Im } d^{k-1})$. Now, in order to see that the space of harmonic cochains has this same dimension, simply note that: $\text{Ker}(d^k) = d\Omega^{k-1} \oplus \mathcal{H}^k$.

Now, the equation $\delta(\omega + df) = 0$ has a solution for each ω in one cohomology equivalence class. We know that the cochains forming different cohomology groups are linearly independent, hence, we conclude that these harmonic cochains span \mathcal{H}^k . ■

By virtue of the above lemma, the implementation of the Hodge decomposition is simply recursive in the rank of the form (*i.e.*, cochain). The case of 0-forms is trivial: fix one vertex to a constant, and solve the Poisson equation for 0-forms. Now suppose that we have a decomposition working for $(k-1)$ -forms, and we look for the decomposition of k -forms. Our approach is to get the harmonic component h^k first, so that we only need to solve a Poisson equation for the rest:

$$\Delta\omega^k = f^k - h^k \quad (11)$$

One is left with the problem of finding a basis of harmonic forms. Since we are given a Hodge star operator, we will use it to define the metric on the space of cochains. This metric allows us to define a basis for harmonic k -form (the dimension of this harmonic space is generally small, since it is the k -th Betti number β_k). First, one needs to calculate the cohomology basis $\{P_i\}$ based on the algorithm in Section 4.4.4. Once we have $\{P_i\}$, we solve one special decomposition of $(k-1)$ -forms by first computing the forms f_i satisfying:

$$\Delta f_i = -\delta P_i \quad (12)$$

Now $H^k = P_i + df_i$ gives us the forms in basis for harmonic k -form space. After normalization, we have the basis to calculate the projection $h^k = HH^t f^k$, where we assemble all H^k into a matrix H . This completes the procedure of calculating the decomposition.

A nonsingular matrix is often preferable when it comes to solve a linear system efficiently; we can change the Laplacian matrix slightly to make the Poisson equation satisfy this requirement. First, we can get an orthonormal basis for harmonic form space (the dimension is β^k). Now for basis e^j (column vector with j -th element equal to 1, and 0 everywhere else), take the distance of e^j to the harmonic space $|e^j - HH^t e^j|$; notice that this can be done in constant time. Now take out the j -th column and j -th row of Δ if e^j has the smallest distance from harmonic space, and repeat the step for β^k times. We are left with a nonsingular matrix, and the solution to the new linear system is a solution to the original Poisson equation.

8 Others Applications

8.1 Form-based Proof of Tutte's Theorem

The notion of forms as convenient, intrinsic substitutes for vector fields has been used to provide a concise proof of the celebrated *Tutte's Embedding Theorem*. This important result in graph theory

states that if one fixes the boundary of a 3-connected graph (*i.e.*, a typical polygonal mesh) to a convex domain in the plane and ensures that every non-boundary vertex is a *strict convex combination* of its neighbors, then one obtains a planar straight-line embedding of the graph. In other words, this embedding procedure will not result in fold-overs. A significantly shorter alternative to the original proof of this theorem was proposed by Gortler, Gotsman, and Thurston [Gortler et al. 2006], using discrete 1-forms on edges. We now present a sketch of their approach, using a formulation more in line with the terms we used in this paper.

A Tutte embedding assigns to each vertex v_i of a graph G some 2D coordinates $\mathbf{X}(v_i) = (x(v_i), y(v_i))$. By definition, each interior vertex v_i satisfies a linear condition on its coordinates of the form: $\mathbf{X}(v_i) = \sum_{v_j \in \mathcal{N}(i)} w_{ij} \mathbf{X}(v_j)$, where $\mathcal{N}(i)$ is the set of 1-ring neighbors of vertex v_i . These coefficients w_{ij} are all nonnegative due to the condition of strict convex combination mentioned above. Now, for a given Tutte embedding, one can construct a 0-form $z(v) = \alpha x(v) + \beta y(v)$ for any pair of positive coefficients α and β . Notice that this 0-form satisfies the same convex combination condition: $z(v_i) = \sum_{v_j \in \mathcal{N}(i)} w_{ij} z(v_j)$. As they are non negative, one can identify these coefficients w_{ij} with the diagonal Hodge star of primal 1-forms (see Section 7) defined by a particular metric. Therefore, the relationship $0 = \sum_{v_j \in \mathcal{N}(i)} w_{ij} (z(v_j) - z(v_i))$ is equivalent to: $d \star dz = 0$. There are two immediate conclusions:

- ◊ the 1-form $\omega = dz$ is closed (since it is the exterior derivative of a 0-form), and
- ◊ it is also co-closed since $\delta\omega = (\star d \star) dz = \star(d \star dz) = 0$.

To use the previously defined 1-form ω to prove Tutte’s theorem, Gortler *et al.* then invoke the usual definition of index of vector fields, *i.e.*, the number of revolutions that the direction of the vector fields does along any small curve around this vertex. This concept is one of the oldest in Algebraic Topology, initially stated by Poincaré and then developed by Hopf and Morse in the continuous case. Its discrete counterpart was first proposed by Banchoff, and used for instance in [Lazarus and Verroust 1999]. A discrete Poincaré-Hopf index theorem also holds, stating that the sum of all indices must be equal to 2 for a genus-0 patch. The final argument uses the link between (co)closed forms and their indices. Indeed, because we found a closed *and* coclosed form ω , it can be easily shown that these two properties induce that the index of each face must be less or equal to zero, as well as the index of each vertex. Because the boundary of the patch is convex, only two vertices on the boundary have index 1. Since all the indices must sum to 2 and each interior index must be less than zero, we can conclude that *each interior index is zero*. Because this argument is valid for every positive pair (α, β) , one can easily deduce that each interior face is convex and each vertex is a “wheel”; thus, injectivity can be guaranteed.

This rather elegant proof demonstrates how discrete forms and their obvious links to Algebraic Topology can be quite powerful in a variety of applications. We also point the interested reader to other papers, such as [Mercat 2001; Gu and Yau 2003], for which special discrete Hodge stars are defined to satisfy a discrete definition of conformality: there are also very interesting research on this particular topic, once again using the calculus of exterior forms.

8.2 Electromagnetism with Forms

Electromagnetism can be formulated very elegantly using differential forms. For a detailed exposition of the geometric structure in E&M, we refer the reader to [Bossavit 1998] and [Warnick et al. 1997]. In this approach, the electric field E is represented by a 1-form as the integral of E along a path traced by a test charge q ,

and is equal to the electromotive force experienced by that charge. The electric displacement L as well as the current density J are represented by 2-forms. The charge distribution ρ is a 3-form. The magnetic field B is represented by a 2-form since it is measured as a flux, whereas the magnetic field intensity H is a 1-form.

With these conventions, Maxwell’s equations can be rewritten as follows:

$$\partial_t B + dE = 0, \quad -\partial_t L + dH = J, \quad dL = \rho, \quad (13)$$

subject to the *constitutive* equations:

$$L = \epsilon E, \quad H = \mu B, \quad (14)$$

where ϵ is the permittivity, and μ is the permeability. The constitutive relations (14) are very similar to the Hodge star operator that transforms a k -form to an $(n-k)$ -form. Here, ϵ operates on the electric field E (1-form) to yield the electric displacement L (2-form) while μ transforms the magnetic field B (2-form) into the magnetic field intensity H (1-form). To this end, one may think of both ϵ and μ as Hodge star operators induced from appropriately chosen metrics. Note that the balance laws in (13) are metric-independent.

As the reader can guess, one can readily discretize this representation of the physical quantities E, L, \dots and the associated system of equations (13-14) using the tools presented in this chapter. The resulting numerical algorithms preserve *exactly* the geometric structure of the system [Bossavit 1998], even in the presence of charge and/or for irregular meshes [Stern et al. 2008].

8.3 Fluids


The geometric structure of Fluid Mechanics, specifically Euler’s equations for inviscid fluids, has been investigated (see [Marsden and Weinstein 1983] and references therein). In this geometric framework, vorticity is represented as a two-form (an area-form) and Euler’s equations can be written as vorticity advection. Roughly speaking, vorticity measures the rotation of a fluid parcel; we say the fluid parcel has vorticity when it spins as it moves along its path. Vorticity advection means that the vorticity (as a two-form) moves dynamically as if it is pushed forward by the fluid flow. The integral of the vorticity on a given bounded domain is equal, by Stokes’ theorem, to the circulation around the loop enclosing the domain. This quantity, as the loop is advected by the fluid, is conserved in the absence of external forcing, as well as the total energy of the fluid. Inspired by this geometric viewpoint and in light of the present development of Discrete Exterior Calculus, one can develop a discrete differential approach to fluid mechanics and an integration scheme that satisfy conservation of circulation, see [Elcott et al. 2007] for further details.

8.4 Developments in Geometry Processing

Discrete forms, with their geometric nature we described up to now, lend themselves naturally to geometric applications. From the design of barycentric coordinates over arbitrary polytopes (see, *e.g.*, [Warren et al. 2007]) usable as dual form basis, to the design of smooth, higher-order Whitney forms [Wang et al. 2006], discrete forms have been shown particularly relevant in several geometry processing tasks. Recent applications include point set reconstruction [Alliez et al. 2007], Eulerian treatment of interfaces [Mullen et al. 2007], as well as conformal parameterization and its use for quadrangle meshing [Tong et al. 2006]—note the recent non-linear version that offers a complete notion of conformal equivalence for triangle meshes [Springborn et al. 2008], including a discrete metric.

9 Conclusions

In this chapter, we have provided an introduction to discrete differential forms and explained how they can be extremely useful in computational science. A convenient Discrete Exterior Calculus solely based on values stored on a discrete manifold has been derived. In the common 3D case, this calculus for scalar and vector fields can be summarized by the following schematic graph:



We have also given a discrete version of the Hodge decomposition, useful for a number of computations in various fields. Despite numerous recent developments this geometric approach to computations is still nascent, and many details need to be explored and proven superior to current approaches. In order to work towards this goal, more work needs to be done to further demonstrate that this idea of forms as fundamental readily-discretizable elements of differential equations can be successfully used in various other contexts where predictive power is crucial.

Acknowledgments

The authors wish to first thank Jerrold E. Marsden for his tremendous support along the way. Peter Schröder, Herbert Edelsbrunner, and John M. Sullivan helped us with a thorough proofreading of this document. Bugs were also reported by Henry Adams and F. Fan. We finally wish to acknowledge the support of Alain Bossavit, Anil Hirani, Melvin Leok, David Cohen-Steiner, Sharif Elcott, Pierre Alliez, and Eitan Grinspun.

References

ABRAHAM, R., MARSDEN, J., AND RATIU, T., Eds. 1988. *Manifolds, Tensor Analysis, and Applications*. Applied Mathematical Sciences Vol. 75, Springer.

ALLIEZ, P., COHEN-STEINER, D., TONG, Y., AND DESBRUN, M. 2007. Voronoi-based variational reconstruction of unoriented point sets. In *Symposium on Geometry Processing*, 39–48.

BJÖRNER, A., AND WELKER, V. 1995. The homology of “k-equal” manifolds and related partition lattices. *Advances in Math.* 110, 277–313.

BOBENKO, A., AND SEILER, R., Eds. 1999. *Discrete Integrable Geometry and Physics*. Clarendon Press.

BOCHEV, P. B., AND HYMAN, J. M. 2005. Principles of mimetic discretizations of differential operators. Preprint: LA-UR-05-4244.

BOSSAVIT, A. 1998. *Computational Electromagnetism*. Academic Press, Boston.

BURKE, W. L. 1985. *Applied Differential Geometry*. Cambridge University Press.

CARROLL, S. 2003. *Spacetime and Geometry: An Introduction to General Relativity*. Pearson Education.

CARTAN, É. 1945. *Les Systèmes Différentiels Extérieurs et leurs Applications Géométriques*. Hermann, Paris.

DESBRUN, M., LEOK, M., AND MARSDEN, J. E. 2004. Discrete Poincaré Lemma. *Appl. Num. Math.*

DIMAKIS, A., AND MÜLLER-HOISSEN, F. 1994. Discrete Differential Calculus, Graphs, Topologies, and Gauge Theory. *Journal of Mathematical Physics* 35, 6703–6735.

DORAN, C., AND LASENBY, A., Eds. 2003. *Geometric Algebra for Physicists*. Cambridge University Press.

EDELSBRUNNER, H., LETSCHER, D., AND ZOMORODIAN, A. 2000. Topological persistence and simplification. In *IEEE Symposium on Foundations of Computer Science*, 454–463.

ELCOTT, S., TONG, Y., KANSO, E., DESBRUN, M., AND SCHRÖDER, P. 2007. Discrete, Circulation-preserving and Stable Simplicial Fluid. *ACM Trans. on Graphics* 26(1) (Jan.).

FISHER, M., SCHRÖDER, P., DESBRUN, M., AND HOPPE, H. 2007. Design of tangent vector fields. *ACM Transactions on Graphics* 26, 3 (July), 56:1–56:9.

FLANDERS, H. 1990. *Differential Forms and Applications to Physical Sciences*. Dover Publications.

FLANDERS, H., Ed. 2001. *Geometric Methods for Computational Electromagnetics*. EMW Publishing, Cambridge Mass.

FORMAN, R. 2003. Bochner’s Method for Cell Complexes and Combinatorial Ricci Curvature. *J. of Discrete and Computational Geom.* 29, 3, 323–374.

FRANKEL, T. 2004. *The Geometry of Physics*. Second Edition. Cambridge University Press, United Kingdom.

GORTLER, S., GOTSMAN, C., AND THURSTON, D. 2006. One-Forms on Meshes and Applications to 3D Mesh Parameterization. *Computer Aided Geometric Design* 23, 2, 83–112.

GROSS, P. W., AND KOTIUGA, R. 2004. *Electromagnetic Theory and Computation: A Topological Approach*. Cambridge University Press.

GU, X., AND YAU, S.-T. 2003. Global conformal surface parameterization. In *Proceedings of the Eurographics/ACM SIGGRAPH Symposium on Geometry Processing*, Eurographics Association, 127–137.

HARRISON, J. 2005. Ravello Lecture Notes on Geometric Calculus – Part I. Tech. rep., UC Berkeley.

HATCHER, A. 2004. *Algebraic Topology*. Cambridge University Press.

HILDEBRANDT, K., AND POLTHIER, K. 2004. Anisotropic filtering of non-linear surface features. In *Computer Graphics Forum*, M.-P. Cani and M. Slater, Eds., vol. 23. Proc. Eurographics 2004.

HIRANI, A. N. 2003. *Discrete Exterior Calculus*. PhD thesis, Caltech.

HYMAN, J. M., AND SHASHKOV, M. 1997. Natural Discretizations for the Divergence, Gradient, and Curl. *International Journal of Computers and Mathematics with Applications* 33.

KANSO, E., ARROYO, M., TONG, Y., YAVARI, A., MARSDEN, J. E., AND DESBRUN, M. 2007. On the Geometric Character of Continuum Mechanics. *Z. angew. Math Phys.* 58.

LAZARUS, F., AND VERROUST, A. 1999. Level Set Diagrams of Polyhedral Objects. In *Proceedings of the 5th ACM Symposium on Solid Modeling and Applications*, 130–140.

- LOVELOCK, D., AND RUND, H. 1993. *Tensors, Differential Forms, and Variational Principles*. Dover Publications.
- MARSDEN, J. E., AND WEINSTEIN, A. 1983. Coadjoint orbits, vortices and Clebsch variables for incompressible fluids. *Physica D* 7, 305–323.
- MARSDEN, J. E., AND WEST, M. 2001. Discrete Mechanics and Variational Integrators. *Acta Numerica* 10, 357–514.
- MERCAT, C. 2001. Discrete Riemann Surfaces and the Ising Model. *Commun. Math. Phys.* 218, 1, 177–216.
- MEYER, M., DESBRUN, M., SCHRÖDER, P., AND BARR, A. H. 2003. Discrete Differential-Geometry Operators for Triangulated 2-Manifolds. In *Proceedings of Visualization and Mathematics III*, 35–57.
- MORITA, S. 2001. *Geometry of Differential Forms*. Translations of Mathematical Monographs, Vol. 201. Am. Math. Soc.
- MULLEN, P., MCKENZIE, A., TONG, Y., AND DESBRUN, M. 2007. A variational approach to eulerian geometry processing. *ACM Trans. on Graphics (SIGGRAPH)* (July), 66.
- MUNKRES, J. R. 1984. *Elements of Algebraic Topology*. Addison-Wesley, Menlo Park, CA.
- PINKALL, U., AND POLTHIER, K. 1993. Computing Discrete Minimal Surfaces. *Experimental Mathematics* 2, 1, 15–36.
- RAMASWAMY, V., AND SHAPIRO, V. 2004. Combinatorial Laws for Physically Meaningful Design. *J. of Computing and Information Science in Engineering* 4, 1, 3–10.
- SCHREIBER, U. 2003. On Superstrings in Schrödinger Representation. *Preprint*.
- SHARPE, R. W. 1997. *Differential Geometry:Cartan's Generalization of Klein's Erlangen Programm*. Springer-Verlag, NY.
- SPRINGBORN, B., SCHRÖDER, P., AND PINKALL, U. 2008. Conformal equivalence of triangle meshes. *ACM Trans. Graph.* 27, 3, 1–11.
- STERN, A., TONG, Y., DESBRUN, M., AND MARSDEN, J. E. 2008. Variational integrators for maxwell's equations with sources. *PIERS* 4, 7, 711–715.
- TONG, Y., ALLIEZ, P., COHEN-STEINER, D., AND DESBRUN, M. 2006. Designing quadrangulations with discrete harmonic forms. In *ACM/EG Symposium on Geometry Processing*, 201–210.
- WANG, K., WEIWEI, TONG, Y., DESBRUN, M., AND SCHRÖDER, P. 2006. Edge subdivision schemes and the construction of smooth vector fields. *ACM Trans. on Graphics (SIGGRAPH)* 25, 3 (July), 1041–1048.
- WARNICK, K. F., SELFRIDGE, R. H., AND ARNOLD, D. V. 1997. Teaching Electromagnetic Field Theory Using Differential Forms. *IEEE Trans. on Education* 40, 1, 53–68.
- WARREN, J., SCHAEFER, S., HIRANI, A., AND DESBRUN, M. 2007. Barycentric coordinates for convex sets. *Advances in Computational Mathematics* 27, 3, 319–338.
- WHITNEY, H. 1957. *Geometric Integration Theory*. Princeton Press, Princeton.
- ZAPATRIN, R. 1996. Polyhedral Representations of Discrete Differential Manifolds. *Preprint*.

Further Reading

Despite a large number of theoretical books, we are aware of only a few books with a truly “applied flavor” in line with this chapter. For applications based on this exterior calculus or other geometric algebras, see [Bossavit 1998; Flanders 2001; Bobenko and Seiler 1999; Doran and Lasenby 2003; Gross and Kotiuga 2004; Ramaswamy and Shapiro 2004]. The reader interested in the application of differential forms to E&M is further referred to [Warnick et al. 1997], for applications in fluid mechanics see [Marsden and Weinstein 1983], and in elasticity see [Kanso et al. 2007] and [Frankel 2004]. The reader is also invited to check out current developments of variants of DEC, for instance, in [Dimakis and Müller-Hoissen 1994; Schreiber 2003; Zapatrin 1996; Harrison 2005], as well as the nice, thorough review from Bochev and Hyman [Bochev and Hyman 2005].

Finally, the interested reader can find additional material on the following websites:

Graphics and Applied Geometry at Caltech:

<http://multires.caltech.edu/pubs/>

<http://www.geometry.caltech.edu/>

Computational E&M (Alain Bossavit):

<http://www.lgep.supelec.fr/mse/perso/ab/bossavit.html>

Discrete Vector Fields and Combinatorial Topology (R. Forman):

<http://math.rice.edu/~forman/>

Discrete Mechanics at Caltech (Jerrold E. Marsden):

<http://www.cds.caltech.edu/~marsden/>