Chapter 6

Cohomology

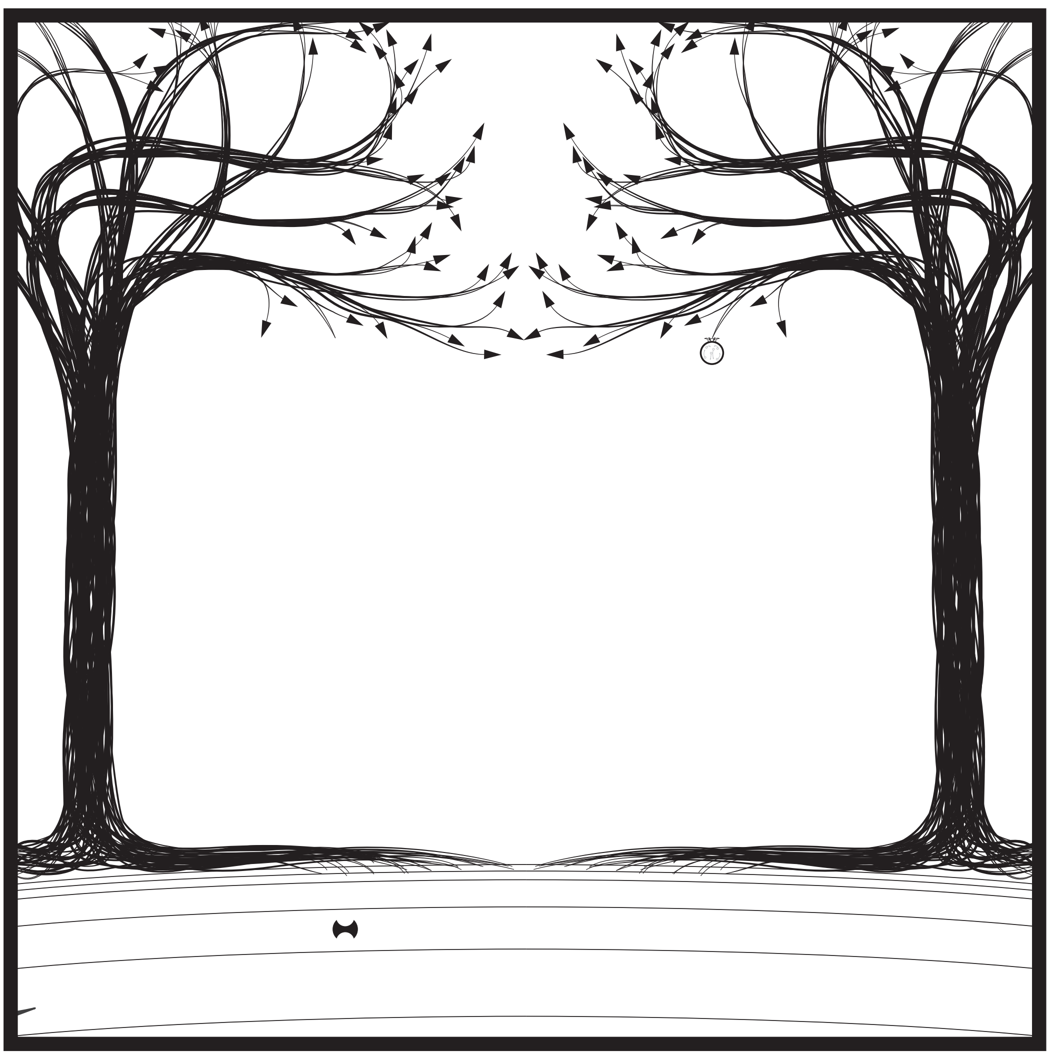

Cohomology is the mirror-image of homology, flipping geometric intuition for algebraic dexterity. The duality implicit in this theory is a subtle and easily underestimated tool in algebraic topology.

6.1 Duals

The first form of duality one encounters in Mathematics is combinatorial.[1]

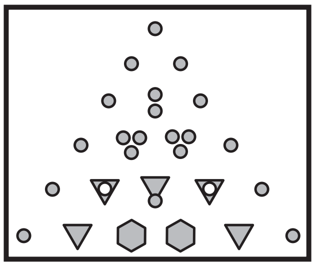

The number of ways to choose k items from $n \geq k$ is exactly the same as the number of ways to (not) choose $n-k$ items. This symmetry in counting manifests itself in numerous numerical miracles, from the reflection symmetry of Pascal's triangle, to the fact that $\chi(\mathbb{S}^{2n+1}) = 0$. The manner in which duality presents itself in Topology is best discovered through the familiar forms of linear algebra and calculus.

The **dual space** of a real vector space V is $V^\vee$, the vector space of all linear functionals $V \to \mathbb{R}$. The dual space satisfies $\dim(V^\vee) = \dim(V)$, and $(V^\vee)^\vee \cong V$ for V finite-dimensional. There is a corresponding notion of duality for linear transformations. If $f\colon V \to W$ is linear, the **dual map** or **adjoint** of f is $f^\vee\colon W^\vee \to V^\vee$ given by $(f^\vee(\eta))(v) = \eta(f(v))$. Note how the dual transformation reverses the direction of the map: it is the archetypal construct of this chapter.

Example 6.1 (Gradients) ⊙

Dual vector spaces play an important role in calculus on manifolds. The **cotangent space** to a manifold M at $p \in M$ is the vector space dual $T_p^*M = (T_pM)^\vee$ to the tangent space. The cotangent spaces, like their tangent space duals, fit together to form a bundle of vector spaces over M, the **cotangent bundle** T^*M. The analogue of a vector field is a **1-form**: a choice of T_p^*M continuous in p. For example, given a real-valued function $f\colon M \to \mathbb{R}$, the **gradient** of f is the 1-form df which, in local coordinates $\{x_i\}_1^n$, evaluates to

$$df = \sum_{i=1}^{n} \frac{\partial f}{\partial x_i} dx_i \,,$$

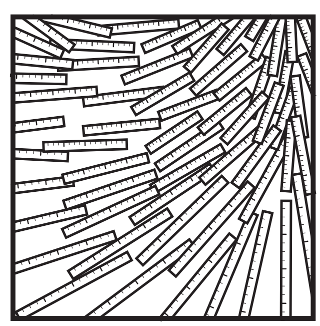

where dx_i is the dual to the x_i unit tangent vector. The *Chain Rule* implies that df is independent of the local coordinates used to express it. It is a common mistake to conflate the gradient of a function with a vector field ∇f: this is permissible in Euclidean space, but the pairing is not canonical. It is best to imagine a gradient df not as a vector field, but as a *ruler field* – a field of rulers with direction, orientation, and scale – along which tangent vectors are measured. ⊙

[1] That grammatical dualities cannot fail to precede even these is not a false statement.

6.2 Cochain complexes

A **cochain complex** is a sequence $\mathcal{C} = (C^\bullet, d)$ of **R**-modules C^k (**cochains**) and module homomorphisms $d^k \colon C^k \to C^{k+1}$ (**coboundary maps**) with the property that $d^{k+1} \circ d^k = 0$ for all k. The **cohomology** of a cochain complex is,

$$H^k(\mathcal{C}) = \ker d^k \big/ \operatorname{im} d^{k-1}.$$

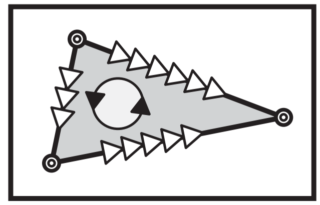

Cohomology classes are equivalence classes of **cocycles** in $\ker d$. Two cocylces are **cohomologous** if they differ by a **coboundary** in $\operatorname{im} d$. The simplest means of constructing cochain complexes is to dualize a chain complex $(C_\bullet, \partial)$. Given such a complex (with coefficients in, say, **R**), define $C^k = C_k^\vee$, the module of homomorphisms $C_k \to \mathbf{R}$. The coboundary d is the adjoint of the boundary ∂, so that

$$d \circ d = \partial^\vee \circ \partial^\vee = (\partial \circ \partial)^\vee = 0^\vee = 0.$$

The coboundary operator d is explicit: $(df)(\tau) = f(\partial\tau)$. For τ a k-cell, d implicates the **cofaces** – those $(k+1)$-cells having τ as a face.

Example 6.2 (Simplicial cochains) $\odot$

Examples in the simplicial category are illustrative. Consider a triangulated disc with a 1-cocycle on edges using $\mathbb{F}_2$ coefficients.

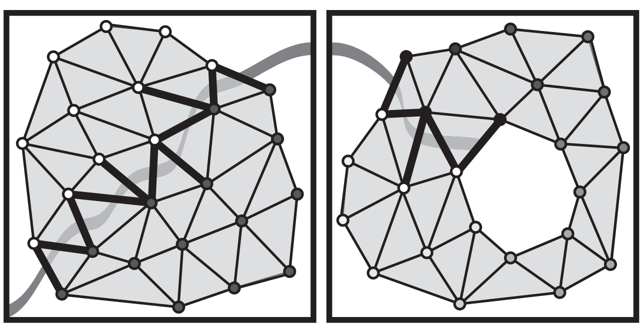

Any such 1-cocycle is the coboundary of a 0-cochain which labels vertices with 0 and 1 *on the left* and *on the right* of the 1-cocycle, so to speak: this is what a trivial class in $H^1(\mathbb{D}^2)$ looks like. On the other hand, if one considers a surface with some nontrivial H_1 – say, an annulus – then one can construct a similar 1-cocycle that is nonvanishing in H^1. The astute reader will notice the implicit relationship between such cocycles and gradients of a local *potential* over the vertices, with cohomology class in H^1 differentiating between those which are or are not globally expressible as a gradient of a potential. $\odot$

Example 6.3 (Integration) $\odot$

Consider a chain complex $\mathcal{C}$ with **R**-coefficients, freely generated by (oriented) simplices $\{\sigma\}$ in a simplicial complex X. The dual basis cochains can be thought of as characteristic functions $\{\mathbb{1}_\sigma\}$. The obvious pairing between $\mathcal{C}$ and its dual $\mathcal{C}^\vee$ in these bases permits an integral interpretation: for basis elements σ and τ, define $\int_\sigma \mathbb{1}_\tau$ to be the evaluation taking the value 1 if and only if $\tau = \sigma$ and 0 else.

This notation is illustrative. By definition of the coboundary d as the dual of the boundary ∂, one has $(d\mathbb{1}_\sigma)(\tau) = \mathbb{1}_\sigma(\partial\tau)$. Using linearity and the integral notation, one has for all cochains $\alpha \in C^p$ and chains $c \in C_{p+1}$,

$$\int_c d\alpha = \int_{\partial c} \alpha. \tag{6.1}$$

This notation should appeal to scientists and engineers who know from early education the utility of Stokes' Theorem. The hint of differentials is no accident: calculus and cohomology are integrally bound, see §6.9. ⊚

Example 6.4 (Kirchhoff's voltage rule) ⊚

Consider an electric circuit as a 1-dimensional cell complex, with circuit elements (resistors, capacitors, etc.) located in the interiors of edges. **Kirchhoff's voltage rule** states that the sum of the voltage potential differences across any loop in the circuit is zero. In the language of this chapter: *voltage is a 1-coboundary*. ⊚

Example 6.5 (Cuts and flows) ⊚

Consider a **directed graph** X: a 1-d cell complex with each edge oriented. A classical set of problems in combinatorial optimization concerns flows on X. Choose two nodes of X to be the source (s) and target (t) nodes, and assume that the graph is connected from s → t (respecting direction). A **flow** on X is an assignment of coefficients (in,

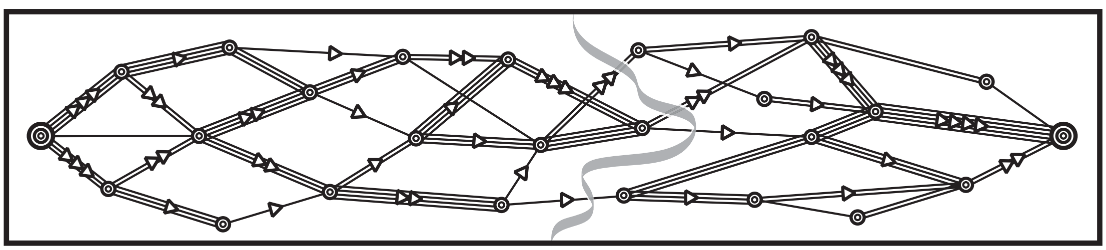

say, $\mathbb{N}$) to edges of X so that, for each node except s and t, the sum of the in-pointing edge flow values equals the sum of the out-pointing edge flow values. Motivated by problems in transportation and railway shipping, the classical **max flow problem** seeks to maximize the **flow value** (the net amount flowing from s or, equivalently, into t) over all flows. The problem is constrained in that each edge is assigned a **capacity** that dominates the possible flow value on that edge.

The **max-flow-min-cut theorem** states that the maximum possible flow value equals the minimal **cut value** as follows, where a **cut** is a subset of edges of X whose removal disconnects s and t (there are no longer directed paths from source to target) and its value is the sum of the edge capacities over the cut. This theorem is a minimax-type theorem in which a duality between flows and cuts is prominent.

In the language of this chapter, given X, s, and t, cuts and flows correspond to various chains and cochains. For clarity, assume that X has been augmented to have a single *feedback* edge from t to s with infinite capacity. In this case, any flow is an element of $H_1(X;\mathbb{Z})$, where the sign of the $\mathbb{Z}$-coefficients is consistent with the edge orientations: the condition for being a 1-cycle is *precisely* that of a conservative flow. A cut gives a 1-cocycle which is nontrivial in $H^1(X;\mathbb{Z})$, since, with the feedback edge, there is no way to separate vertices of X on either *side* of the cut. The reader rightly suspects that flow/cut duality is at heart a co/homological duality: see Example 10.25. ⊚

Example 6.6 (Arbitrage) ⊚

Certain simple examples from economics and social networks are expressible in the language of chains and cochains: the following is from [185]. Consider a static exchange market on a set of N commodities, $V = \{v_i\}_1^N$. Assume the existence of **exchange rates** – coefficients $\epsilon_{ij} > 0$ such that one unit of v_i is worth ϵ_{ij} units of v_j, with the natural symmetry that $\epsilon_{ji} = \epsilon_{ij}^{-1}$. Build a graph X, with nodes V and edge set equal to all $\{v_i, v_j\}$ such that exchange rates ϵ_{ij} are known.

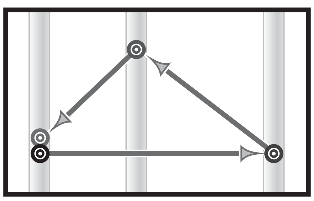

There is a natural $\mathbb{R}$-valued 1-cochain given by exchange rates as follows. Orient the 1-cells of X in an arbitrary but fixed fashion, and let the 1-cochain ξ evaluate on an oriented edge (v_i, v_j) to $\ln \epsilon_{ij}$. In a perfect exchange system, ξ is a cocycle: any cyclic sequence of trades from commodity v_i back to v_i returns the same amount (assuming zero transaction costs). The failure of ξ to be a cocycle indicates an arbitrage – a sequence of trades resulting in net gain. It is fascinating to consider what happens when the goods at issue are more complex financial items, whose restrictions on availability and convertibility yield richer topological objects both noisy (edge values are uncertain) and dynamic (time-varying). The reader may enjoy contemplating whether higher-order algebraic coefficients and/or higher-dimensional simplicial models (say, the flag complex) yield a useful approach to characterizing arbitrage in complex systems. ⊚

6.3 Cohomology

The definition of cohomology in terms of dualizing chain complexes, though less than intuitive, is efficient. By dualizing results of the previous two chapters, one immediately obtains the following:

1. Cohomology theories: cellular, singular, local, Čech, relative, reduced; all with arbitrary coefficients;
2. Functoriality and the homotopy invariance of cohomology;
3. Cohomological long exact sequences of pairs; the Mayer-Vietoris sequence; excision.

To correctly interpret and use these results *arrows must be reversed.* For example, the induced homomorphisms on cohomology twist composition: $H(f \circ g) = H(g) \circ H(f)$, in keeping with the way that duals of linear transformations behave. The Snake Lemma

(Lemma 5.5) holds, but when dualizing a short exact sequence of chains to cochains, the sequence flips. This manifests itself in, *e.g.*, the long exact sequence of a pair (X, A) as follows. The short exact sequence of cochain complexes,

$$0 \longrightarrow C^\bullet(X, A) \xrightarrow{j^\bullet} C^\bullet(X) \xrightarrow{i^\bullet} C^\bullet(A) \longrightarrow 0 \ ,$$

becomes the long exact sequence,

$$\longrightarrow H^{n-1}(A) \xrightarrow{\delta} H^n(X, A) \xrightarrow{H(j)} H^n(X) \xrightarrow{H(i)} H^n(A) \longrightarrow \ ,$$

where, in this text, the induced and connecting homomorphisms have the same notation in homology and cohomology and are distinguishable via context and direction. The novice reader should as an exercise write out the Mayer-Vietoris sequence on cohomology.

The beginner may be deflated at learning that cohomology seems to reveal no new data, at least as far as linear algebra can count.

Theorem 6.7 (Universal Coefficient Theorem). *For X a space with finite-dimensional homology over a field $\mathbb{F}$, $H^n(X; \mathbb{F}) \cong H_n(X; \mathbb{F})^\vee$.*

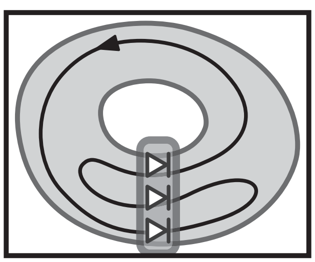

The situation is much more delicate in $\mathbb{Z}$-coefficients, but it is nevertheless true that cohomology is determined by knowing the homology and certain algebraic properties of the coefficient ring. Students may wonder why they should bother with this seemingly supererogatory cohomology, as it alike to homology in every respect except intuition. In the beginning, one should repeat the mantra that duality often simplifies algebraic difficulties.

Example 6.8 (Connectivity) ⊙

Note that the dimension of H^0, like that of H_0, yields connectivity data. The proof in the case of cohomology is *simpler* than for homology, since, by definition, $H^0 = \ker d^0$. There is nothing to quotient out, due to arrow reversal. In the cellular case, elements of $H^0(X)$ are functions on vertices whose oriented differences along edges (a 'discrete derivative') is everywhere zero – these equate to the singular interpretation of locally-constant functions on X. In either case, a suitable basis consists of characteristic functions of connected components of X. Note the differences. In homology, H_0 determines the number of path-connected components (homologous 0-cycles are connected by 1-chain paths) while cohomology H^0 measures connected components (as seen by functionals). ⊙

Example 6.9 (1-Cocycles) ⊙

One cartoon for understanding this distinction between local and global coboundaries is the popular optical illusion of the *impossible tribar*. When one looks at the tribar, the drawn perspective is locally realizable – one can construct a local depth function. However, a global depth function cannot be defined. The impossible tribar is a cartoon of a non-zero class in H^1 (properly speaking, $H^1(\mathbb{S}^1; \mathbb{R}^+)$, using the multiplicative reals for coordinates as a way to encode projective geometry [243]).

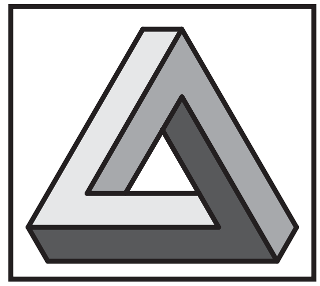

An even more cartoonish example that evokes nontrivial 1-cocycles is the popular game of *Rock, Paper, Scissors*, for which there are local but not global ranking functions. A local gradient of *rock-beats-scissors* does not extend to a global gradient. Perhaps this is why customers are asked to conduct rankings (*e.g.*, *Netflix* movie rankings or *Amazon* book rankings) as a 0-cochain (*"how many stars?"*), and not as a 1-cochain (*"which-of-these-two-is-better?"*): nontrivial H^1 is, in this setting, undesirable. The *Condorcet paradox* – that locally consistent comparative rankings can lead to global inconsistencies – is a favorite topic in voting theory. Its best explanation, cohomology, is less popular. ⦾

6.4 Poincaré duality

Homology and cohomology of manifolds express duality as a dimensional symmetry. Based on data from spheres, compact orientable surfaces, tori, and the Künneth formula in §4.2, one might guess that for M a compact orientable n-manifold and coefficients in a field, $\dim H_k M = \dim H_{n-k} M$, and likewise with cohomology. This is true and is a version of duality due to Poincaré.

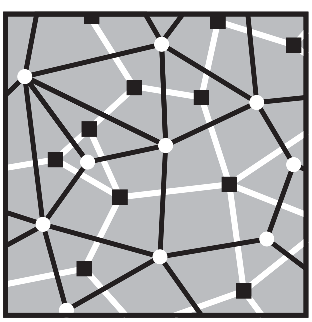

At the level of cellular homology, this duality has a geometric interpretation. Consider, *e.g.*, a compact surface with a polyhedral cell structure, and let $\mathcal{C}$ be the cellular chain complex with $\mathbb{F}_2$ coefficients. There is a dual polyhedral cell structure, yielding a chain complex $\overline{\mathcal{C}}$, where the dual cell structure places a vertex in the center of each original 2-cell, has 1-cells transverse to each original 1-cell, and, necessarily, has as its 2-cells neighborhoods of the original vertices. Each dual 2-cell is a polyhedral n-gon, where n is the degree of the original 0-cell dual. Note that these cell decompositions are truly dual and have the effect of reversing the dimensions of cells: k-cells generating C_k are in bijective correspondence with $(2-k)$-cells generating a modified cellular chain group $\overline{C}_{2-k}$. The dual complex consisting of $\overline{C}^k := \overline{C}_k^\vee$ and $\overline{d} = \overline{\partial}^\vee$ entwines with $C_\bullet$ in a diagram:

$$\begin{array}{ccccccccc} 0 & \longrightarrow & C_2 & \xrightarrow{\partial} & C_1 & \xrightarrow{\partial} & C_0 & \longrightarrow & 0 \\ & & \downarrow\cong & & \downarrow\cong & & \downarrow\cong & & \\ 0 & \longrightarrow & \overline{C}^0 & \xrightarrow{\overline{d}} & \overline{C}^1 & \xrightarrow{\overline{d}} & \overline{C}^2 & \longrightarrow & 0 \end{array}$$

The reader should check that the vertical maps are isomorphisms and, crucially, that the diagram is commutative. The equivalence of singular and cellular co/homology, along with Theorem 6.7, implies that, for a compact surface with $\mathbb{F}_2$ coefficients, $H_k \cong H^{2-k} \cong (H_{2-k})^\vee \cong H_{2-k}$. Though this style of proof generalizes to higher-dimensional manifolds, a better explanation for the symmetries present in co/homology is more algebraic, and proceeds in a manner that adapts to non-compact manifolds as well. To unwrap this, a modified cohomology theory is helpful.

Given a (singular, simplicial, cellular) cochain complex $\mathcal{C}^\bullet$ on a space X, consider the subcomplex $\mathcal{C}_c^\bullet$ of cochains which are compactly supported: each cochain is zero outside some compact subset of X. (In the simplicial or cellular setting, this is equivalent to building cochains from a finite number of basis cochains.) The coboundary map restricts to $d\colon C_c^k \to C_c^{k+1}$ and $d^2 = 0$, yielding a well-defined **cohomology with compact supports**, $H_c^\bullet(X)$. This cohomology satisfies the following:

1. $H_c^\bullet$ is not a homotopy invariant, but is a proper-homotopy (and hence a homeomorphism) invariant.
2. $H_c^k(\mathbb{R}^n) = 0$ for all k except $k = n$, in which case it is of rank one.
3. $H_c^\bullet(X) \cong H^\bullet(X)$ for X compact.
4. $H_c^\bullet(X) \cong H^\bullet(X, X{-}K)$ for K a sufficiently large compact set.

This can be subtle: in brief, one orders compact sets by inclusion and uses induced maps to take a limit. See, *e.g.*, [176, 3.3]. Compactly supported cohomology cleanly expresses manifold duality:

Theorem 6.10 (Poincaré Duality). *For M an n-manifold, there is a natural isomorphism* PD: $H_k(M;\mathbb{F}_2) \xrightarrow{\cong} H_c^{n-k}(M;\mathbb{F}_2)$.

For M a compact n-manifold, $H_k(M;\mathbb{F}_2) \cong H^{n-k}(M;\mathbb{F}_2)$. For orientable manifolds, the theorem holds with any field coefficients. Integer coefficients are more problematic: torsional elements lag in duality [46, 176]. With a slight change of perspective (in §6.9), a more precise form of PD will be given.

6.5 Alexander duality

Poincaré duality can be adapted to several related settings involving manifolds. Among the most useful is **Alexander duality**.

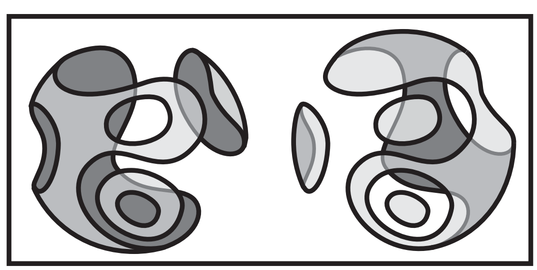

Theorem 6.11 (Alexander Duality). *Let $A \subset \mathbb{S}^n$ be compact, nonempty, proper, and locally-contractible. There is an isomorphism*

$$\mathrm{AD}\colon \tilde{H}_k(\mathbb{S}^n{-}A) \xrightarrow{\cong} \tilde{H}^{n-k-1}(A).$$

Proof. The condition on A is a form of tameness and is crucial, allowing one to choose a small open neighborhood U of A that deformation-retracts onto A. For $k > 0$,

$$\begin{aligned} H_k(\mathbb{S}^n - A) &\cong H_c^{n-k}(\mathbb{S}^n - A) && \text{[Poincaré duality]} \\ &\cong H^{n-k}\left(\mathbb{S}^n - A, (\mathbb{S}^n - A) - (\mathbb{S}^n - U)\right) && \text{[compact supports]} \\ &= H^{n-k}(\mathbb{S}^n - A, U - A) && \\ &\cong H^{n-k}(\mathbb{S}^n, U) && \text{[excision]} \\ &\cong \tilde{H}^{n-k-1}(U) && \text{[long exact sequence of pair]} \\ &\cong \tilde{H}^{n-k-1}(A). && \text{[deformation retraction]} \end{aligned}$$

Slight modifications are required for $k = 0$. The theorem holds for all coefficients. ⊙

6.6 Helly's Theorem

The following theorem is a classic result in convex geometry and geometric combinatorics. Using co/homology yields a transparent proof.

Theorem 6.12 (Helly's Theorem). *Let $\mathcal{U} = \{U_\alpha\}$ be a collection of $M > n+1$ compact convex subsets of $\mathbb{R}^n$ such that every $(n+1)$-tuple of distinct elements of $\mathcal{U}$ has a point in common. Then all elements of $\mathcal{U}$ have a point in common.*

Proof. Induct on M, beginning at $M = n+2$. Consider the nerve $\mathcal{N}(\mathcal{U})$ of $\mathcal{U}$. It is a subcomplex of the $(n+1)$-simplex which, by hypothesis, contains all faces. If the common intersection is empty, then $\mathcal{N}(\mathcal{U}) = \partial\Delta^{n+1} \simeq \mathbb{S}^n$. As the cover $\mathcal{U}$ is by convex sets, it is an acyclic cover (all nonempty intersections are homologically acyclic) and, via Theorem 2.4,

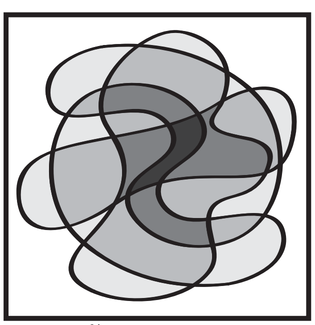

$$\check{H}_\bullet(\mathcal{U}) \cong H_\bullet(\mathcal{N}(\mathcal{U})) \cong H_\bullet\left(\cup_\alpha U_\alpha\right).$$

In other words, $\mathcal{N}(\mathcal{U}) \simeq \mathbb{S}^n$ has the homology of a subset of $\mathbb{R}^n$. However, it is impossible for a subset $A \subset \mathbb{R}^n$ to have the homology type of $\mathbb{S}^n$, thanks to Alexander duality, as follows. Note that the hypotheses for Theorem 6.11 are satisfied, and therefore,

$$H_n(A) \cong \tilde{H}^{n-n-1}(\mathbb{R}^n - A) = \tilde{H}^{-1}(\mathbb{R}^n - A) = 0,$$

since $\tilde{H}^{-1} = 0$ for all nonempty spaces. Thus, $\mathcal{N}(\mathcal{U}) \not\simeq \mathbb{S}^n$, and the cover $\mathcal{U}$ must have a common intersection point. The induction step is a simple modification of this proof. ⊙

The reader may have seen proofs of Helly's Theorem based on convex geometry or functional analysis. One benefit of the topological approach is that extensions to the non-compact and non-convex world are natural and easily discerned. The critical ingredient is that the resulting cover $\mathcal{U}$ is acyclic.

6.7 Numerical Euler integration

Alexander duality is a key step in developing a highly effective numerical method for performing integration with respect to Euler characteristic over a non-localized planar network. Recall from §3.7 that certain problems in data aggregation over a network are expressible as an integral with respect to $d\chi$ over a tame space X.

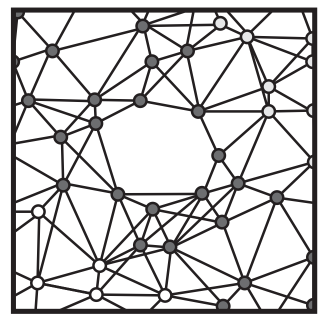

This assumption of integrating sensors over a continuous domain is highly unrealistic; however, expressing the true answer as an integral over a continuum provides a clue as to how to approximate the result over a discretely sampled domain. Given an integrand $h \in \mathsf{CF}(\mathbb{R}^n)$ sampled over a discrete set, computational formulae such as Equation (3.10) suggest that the estimation of the Euler characteristics of the upper excursion sets is an effective approach. However, if the sampling occurs over a network with communication links, then it is potentially difficult to approximate those Euler characteristics. Taking the flag complex of the network can lead to the existence of *fake holes* – higher-dimensional spheres (*cf.* §2.2) that ruin an Euler characteristic approximation.

Proposition 6.13 ([24]). *For $h\colon \mathbb{R}^2 \to \mathbb{N}$ constructible and upper semi-continuous,*

$$\int_{\mathbb{R}^2} h\, d\chi = \sum_{s=0}^{\infty} \left(\beta_0\{h > s\} - \beta_0\{h \leq s\} + 1\right), \tag{6.2}$$

where the zero$^{\text{th}}$ *Betti number $\beta_0 = \dim H_0$ is the number of connected components.*

Proof. Let A be a compact nonempty tame subset of $\mathbb{R}^2$. From the homological definition of the Euler characteristic and compactness of A,

$$\chi(A) = \sum_{k=0}^{\infty} (-1)^k \dim H_k(A),$$

where $H_\bullet$ denotes singular homology in field coefficients. Since $A \subset \mathbb{R}^2$, $H_k(A) = 0$ for all $k > 1$, and it suffices to compute $\chi(A) = \dim H_0(A) - \dim H_1(A)$. By a slight modification of Alexander duality,

$$\dim H_1(A) = \dim \tilde{H}^0(\mathbb{S}^2 - A) = \dim H_0(\mathbb{R}^2 - A) - 1,$$

where the last equality (if not obvious) follows from a long exact sequence on the pair $(\mathbb{S}^2 - A, \mathbb{R}^2 - A)$. Since h is upper semi-continuous, each of the upper excursion sets $A = \{h > s\}$ is compact. Noting that $\mathbb{R}^2 - A = \{h \leq s\}$, one has:

$$\int h\, d\chi = \sum_{s=0}^{\infty} \chi\{h > s\} = \sum_{s=0}^{\infty} \dim H_0(\{h > s\}) - (\dim H_0(\{h \leq s\}) - 1)\,.$$

$\odot$

Corollary 6.14. *A sufficient sampling condition to ensure exact computation of $\int h\, d\chi$ over a planar network is that the network correctly samples the connectivity of all the upper and lower excursion sets of h.*

This formulation is extremely important to numerical implementation of this integration theory to planar sensor networks, which, in practice, may be both non-localized and not dense enough to sufficiently cover regions [79].

6.8 Forms and Calculus

In the setting of smooth manifolds, cohomology flows naturally from multivariable calculus [46, 169]. One begins with multilinear algebra. Given a $\mathbb{R}$-vector space V, let $\mathbf{\Lambda}(V)$ denote the algebra of **forms** on V – alternating multilinear maps from products of V to $\mathbb{R}$. Given a basis $\{x_i\}_1^n$ for V, explicit generators for $\mathbf{\Lambda}(V)$ are given by the dual 1-forms $dx_i \in V^\vee$, where $dx_i \colon V \to \mathbb{R}$ returns the x_i^{th} coordinate of a vector in V. Scalar multiplication in $\mathbf{\Lambda}(V)$ is over $\mathbb{R}$ and the sum is induced by that on $V^\vee$. The product in the algebra $\mathbf{\Lambda}(V)$ is called the **wedge product** and denoted $\wedge$. It is alternating, meaning, in particular, that $dx_i \wedge dx_j = -dx_j \wedge dx_i$ for all i and j. The algebra $\mathbf{\Lambda}(V)$ is graded:

$$\mathbf{\Lambda}(V) := \bigoplus_{p=0}^{\infty} \mathbf{\Lambda}^p(V),$$

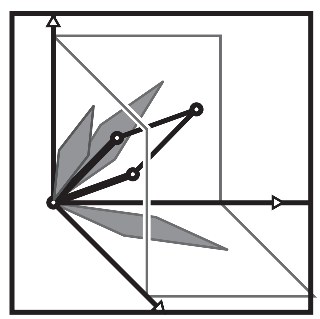

where $\mathbf{\Lambda}^p(V)$ is the vector space of p-forms, with basis $dx_{i_1} \wedge \cdots \wedge dx_{i_p}$ for $1 \leq i_1 < \cdots < i_p \leq n = \dim V$. A p-form takes as its argument an ordered p-tuple of vectors in V and returns a real number in a manner that is multilinear and alternating, *cf.* determinants. The uniqueness of the determinant implies that $\dim \mathbf{\Lambda}^n = 1$; the alternating property of $\wedge$ implies that $\mathbf{\Lambda}^p = 0$ for all $p > n = \dim V$.

These pointwise constructions pass to manifolds, yielding **differential forms**. Recall from §1.3 that for an n-manifold M, the tangent bundle T_*M is a collection of n-dimensional vector spaces, parameterized by points in M. A (smooth) vector field V is a choice of elements of T_xM varying (smoothly) in x, or, more precisely, a **section** taking $x \mapsto V(x) \in T_xM$. In like manner, the vector spaces of p-forms, $\mathbf{\Lambda}^p(T_xM)$, collectively form a "bundle" varying smoothly with $x \in M$, of which the cotangent bundle T^*M is the case $p = 1$. As with the gradient 1-forms of Example 6.1, a ***p*-form field** (shortened to *p-form* in practice) is a section $\alpha \colon M \to \mathbf{\Lambda}^p(T_*M)$ giving $\alpha_x \in \mathbf{\Lambda}^p(T_xM)$ varying smoothly in x. The space of all such sections – the p-form fields on M – is denoted $\Omega^p = \Omega^p(M)$. On all manifolds, $\Omega^0(M) = C^\infty(M;\mathbb{R})$, since $\dim \mathbf{\Lambda}^0 = 1$. Likewise, $\Omega^p(M) = 0$ for all $p > \dim M$. Algebraic operations on $\mathbf{\Lambda}$ pass to operations on $\Omega := \oplus_p \Omega^p$ operating pointwise. For example, the wedge product extends to $\wedge \colon \Omega^p \times \Omega^q \to \Omega^{p+q}$.

In the passage from $\mathbf{\Lambda}$ to Ω, form fields change from point-to-point: such changes are measured by a derivative. The appropriate differential operator for forms is the **exterior derivative** $d\colon \Omega^p \to \Omega^{p+1}$, which has the following properties:

1. d is linear with respect to addition of forms and scalar multiplication;
2. d satisfies a Leibniz rule: $d(\alpha \wedge \beta) = d\alpha \wedge \beta + d\beta \wedge \alpha$;
3. on 0-forms, d is the usual differential operator $d\colon f \mapsto df$; and
4. $d(df) = 0$ for all 0-forms f.

Example 6.15 (Vector calculus on $\mathbb{R}^3$) ⊚

On Euclidean $\mathbb{R}^3$, 1-forms $\alpha \in \Omega^1$ are representable as $\alpha = f_x dx + f_y dy + f_z dz$, where the coefficient functions f_i are smooth. A 2-form, $\beta \in \Omega^2$, is representable as $\beta = g_x dy \wedge dz + g_y dz \wedge dx + g_z dx \wedge dy$. Every 3-form on $\mathbb{R}^3$ is of the form $h\, dx \wedge dy \wedge dz$ for some h. The differential operator d is familiar to students of vector calculus. Recall from Equation (5.3) how the gradient, curl, and divergence operators tie together functions $C = C^\infty(\mathbb{R}^3)$ and vector fields $\mathfrak{X} = \mathfrak{X}(\mathbb{R}^3)$. This is conveyed via a commutative diagram,

$$\begin{array}{ccccccc} C & \xrightarrow{\nabla} & \mathfrak{X} & \xrightarrow{\nabla\times} & \mathfrak{X} & \xrightarrow{\nabla\cdot} & C \\ \cong\downarrow & & \cong\downarrow & & \cong\downarrow & & \cong\downarrow \\ \Omega^0 & \xrightarrow{d} & \Omega^1 & \xrightarrow{d} & \Omega^2 & \xrightarrow{d} & \Omega^3 \end{array} \quad . \tag{6.3}$$

On $\Omega^0(\mathbb{R}^3)$, d is the gradient operator, taking a function f not to its gradient vector field, but to the more natural gradient 1-form df. On $\Omega^1(\mathbb{R}^3)$, d is the curl operator; on $\Omega^2(\mathbb{R}^3)$, d acts as the divergence operator. Here, the vertical arrows identify 0- and 3-forms with functions, and identify vector fields with 1- and 2-forms in the obvious ways:

$$\begin{array}{lcl} \vec{F} = F_x\hat{\imath} + F_y\hat{\jmath} + F_z\hat{k} & \longrightarrow & F_x dx + F_y dy + F_z dz = \alpha_{\vec{F}} \\ \vec{F} = F_x\hat{\imath} + F_y\hat{\jmath} + F_z\hat{k} & \longrightarrow & F_x dy \wedge dz + F_y dz \wedge dx + F_z dx \wedge dy = \beta_{\vec{F}} \end{array}$$

The reader should return to simplicial examples of cochains and coboundaries and be convinced that what is measured on the algebraic level is indeed a discrete analogue of gradients, curls, and divergences.

Example 6.16 (Maxwell's equations) ⊚

The language of differential forms is ubiquitous in mathematical physics. Maxwell's equations admit a particularly simple interpretation. The calculus version of Maxwell's equations (on Euclidean $\mathbb{R}^3$, in a vacuum) are as follows:

$$\begin{array}{lcl} -\dfrac{1}{c}\dfrac{\partial \vec{B}}{\partial t} = \nabla \times \vec{E} & \qquad & \nabla \cdot \vec{B} = 0 \\ \dfrac{1}{c}\dfrac{\partial \vec{E}}{\partial t} = \nabla \times \vec{B} - \dfrac{4\pi}{c}\vec{J} & \qquad & \nabla \cdot \vec{E} = 4\pi\rho, \end{array}$$

where $\vec{E}$, $\vec{B}$, $\vec{J}$, ρ, and c are the electric field, magnetic field, current, charge density, and speed of light, respectively. Using the Euclidean structure to convert a field $\vec{F}$

into a 1-form $\alpha_{\vec{F}}$ or a 2-form $\beta_{\vec{F}}$ as in Example 6.15, one obtains:

$$d(c\alpha_{\vec{E}} \wedge dt + \beta_{\vec{B}}) = 0$$
$$d(c\alpha_{\vec{B}} \wedge dt - \beta_{\vec{E}}) = 4\pi\beta_{\vec{J}} \wedge dt - \rho\, dx \wedge dy \wedge dz.$$

These equations can be made more compact still and extended to arbitrary geometric manifolds using some of the constructions in §6.12. ⦾

Example 6.17 (Ideal fluids) ⦾

The equations of motion for a perfect inviscid fluid are likewise lifted to forms. Let $\vec{V}(t)$ be a time-dependent volume-preserving vector field representing the instantaneous motion of a fluid on Euclidean $\mathbb{R}^n$. Volume-preserving means that for the chosen volume form $\mu \in \Omega^n$ with $\mu \neq 0$ nowhere vanishing, the derivative $d\mu(\vec{V}, \cdot) = 0$ vanishes. Let $\alpha_{\vec{V}}$ be the (time-dependent) 1-form dual to $\vec{V}$ as per Example 6.15. Then, the Euler equations become

$$\frac{\partial \alpha_{\vec{V}}}{\partial t} + d\alpha_{\vec{V}}(\vec{V}, \cdot) = -dH, \tag{6.4}$$

where $H\colon M \to \mathbb{R}$ is a function (sometimes known as the *head* or *Bernoulli* function) combining pressure and kinetic energy terms. This formulation permits doing fluid dynamics on an arbitrary **Riemannian manifold** – a manifold M with a smoothly varying inner product $g(\cdot, \cdot)$ on tangent spaces. The dual 1-form $\alpha_{\vec{V}}$ to $\vec{V}$ is given by contraction into the metric: $\alpha_{\vec{V}} := g(\vec{V}, \cdot)$ [16]. The **vorticity** of such a fluid is usually presented as a vector field representing the curl of the velocity field: in fact, it is more properly defined to be the 2-form $d\alpha_{\vec{V}}$, the derivative of the 1-form dual to velocity. ⦾

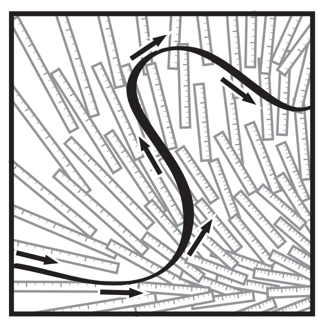

The language of differential forms is designed for integration: a p-form is perhaps best thought of as an object that can be integrated over a p-dimensional domain. Specifically, given an oriented p-dimensional submanifold with corners, S, there is an integral operator $\int_S\colon \Omega^p \to \mathbb{R}$ defined by evaluating the p-form pointwise on oriented p-tuples of tangent vectors to S and integrating on coordinate charts using the standard Lebesgue integral. As is the case in the more familiar setting of line integrals in vector calculus, the Chain Rule implies an invariance of the integral with respect to local orientation-preserving coordinate representations. Modulo details about induced orientations on boundaries, the fundamental theorem of calculus-with-forms is transparent:

Theorem 6.18 (Stokes' Theorem). *For α a compactly-supported p-form and S an oriented $(p+1)$-dimensional manifold (with boundary and/or corners),*

$$\int_S d\alpha = \int_{\partial S} \alpha.$$

6.9 De Rham cohomology

The antisymmetry property of forms is reminiscent of the cancelations via judicious choice of signs at the heart of all co/chain complexes. Thanks to this antisymmetry and the commutativity of mixed partial derivatives, $d^2 = 0$ always, as presaged by Example 6.15. This prompts the interpretation of Ω as a complex: the **de Rham complex** of a manifold M is the cochain complex ${}_{dR}\mathcal{C} = (\Omega^\bullet, d)$ of forms with coboundary the exterior derivative d. The **de Rham cohomology** of M, ${}_{dR}H^\bullet(M)$, is the cohomology of $\Omega^\bullet$. In this theory, it is traditional to call the cocycles **closed** forms and the coboundaries **exact** forms. *A de Rham cohomology class is an equivalence class of closed forms modulo exact forms.*

Example 6.19 (Winding numbers and de Rham cohomology) ⊙

The reduced de Rham cohomology ${}_{dR}\tilde{H}^\bullet(\mathbb{R}^n)$ of $\mathbb{R}^n$ is trivial for all n, thanks to the Fundamental Theorem of Integral Calculus: closed forms are exact. On the punctured plane $\mathbb{R}^2{-}0$, the closed 1-form,

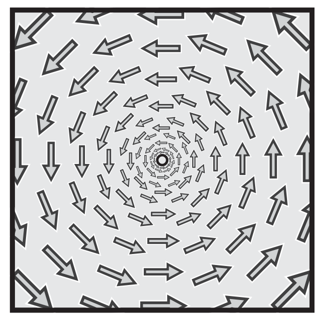

$$d\theta = \frac{x\,dy - y\,dx}{x^2 + y^2},$$

is *not* exact. Despite being denoted $d\theta$, there is no single-valued 0-form θ whose gradient 1-form is $d\theta$; thus, $d\theta$ defines a non-trivial cohomology class and generator of ${}_{dR}H^1(\mathbb{R}^2{-}0) \cong \mathbb{R}$. The integral of $d\theta$ over an oriented piecewise-smooth closed curve γ in the punctured plane yields an integer, and this is, precisely, the winding number of γ about 0: *cf.* Equation (3.4). ⊙

Example 6.20 (Wedge product) ⊙

In de Rham cohomology, the wedge product for forms descends to a product on cohomology. By defining $[\alpha] \wedge [\beta] := [\alpha \wedge \beta]$, one notes that since α and β are closed, $d(\alpha \wedge \beta) = d\alpha \wedge \beta + d\beta \wedge \alpha = 0$; furthermore, if α or β is exact, then so is $\alpha \wedge \beta$. Thus, $\wedge\colon {}_{dR}H^p \times {}_{dR}H^q \to {}_{dR}H^{p+q}$ turns ${}_{dR}H^\bullet$ into a ring. Since, locally, a basis Euclidean k-form measures oriented projected k-dimensional volumes, the wedge product inherits a volumetric interpretation. On the torus $\mathbb{T}^n$ with angular coordinates θ_i, the 1-forms $d\theta_i$, $i = 1 \ldots n$ generate the cohomology ring ${}_{dR}H^\bullet(\mathbb{T}^n)$. For example, the generator for ${}_{dR}H^n(\mathbb{T}^n) \cong \mathbb{R}$ is the volume form $d\theta_1 \wedge \cdots \wedge d\theta_n$. ⊙

It is no coincidence that the de Rham cohomology of $\mathbb{R}^n$ and $\mathbb{T}^n$ have the same dimensions as in the singular theory.

Theorem 6.21 (de Rham Isomorphism Theorem). *For M a manifold, ${}_{dR}H^\bullet(M) \cong H^\bullet(M;\mathbb{R})$.*

A little more effort yields a calculus-based version of Poincaré duality. Let $\Omega^\bullet_c(M)$ denote the complex of *compactly supported* forms on M, complete with differential

d. This yields a well-defined cohomology ${}_{dR}H^\bullet_c$ with compact supports, which, as an extension of Theorem 6.21, is isomorphic to the singular $H^\bullet_c$.

Theorem 6.22 (Poincaré Duality, de Rham version). *For M an oriented manifold of dimension n, wedge and integration of forms yield isomorphisms*

$$\int_M \blacksquare \wedge \square : {}_{dR}H^p(M) \xrightarrow{\cong} {}_{dR}H^{n-p}_c(M)^\vee$$

$$\int_{\blacksquare} \square : H_p(M;\mathbb{R}) \xrightarrow{\cong} {}_{dR}H^p(M)^\vee$$

It is an instructive exercise to show that integration descends to homology and cohomology. For $[\alpha] \in {}_{dR}H^p$, $\beta \in \Omega^{p-1}$, S a boundaryless p-dimensional submanifold and T a $(p+1)$-dimensional submanifold with boundary, then

$$\int_{S+\partial T} \alpha + d\beta = \int_S \alpha + \int_{\partial S} \beta + \int_T d\alpha + \int_{\partial T} d\beta = \int_S \alpha,$$

via Stokes' Theorem. Thus, only $[S]$ and $[\alpha]$ matter. The isomorphisms of Theorem 6.22 effect the Poincaré Duality isomorphism PD: $H_p(M;\mathbb{R}) \xrightarrow{\cong} {}_{dR}H^{n-p}_c(M)^\vee$ as in §6.4.

6.10 Cup products

The de Rham isomorphism of Theorem 6.21 allows one to import calculus-based intuition into cohomology theory. Several of the constructs that are natural and clear in the setting of manifolds lift to the more general (cellular, singular, algebraic) cohomology theories. This section explores the algebraic generalization of the wedge product. Recall from Example 6.20 that the wedge product on forms, $\wedge$, so implicitly familiar to students of multivariable calculus, descends to a product on de Rham cohomology classes, giving ${}_{dR}H^\bullet$ the structure of a ring. What is the singular analogue?

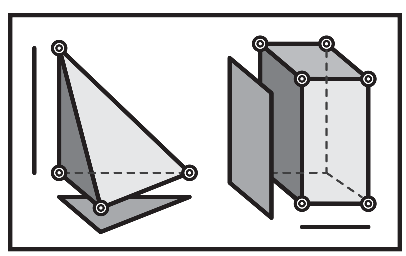

Just as one defines the wedge $\wedge$ on forms and then passes to cohomology, one defines a product on cochains. Let $\alpha \in C^p(X;\mathbf{R})$ be a p-cochain and $\beta \in C^q(X;\mathbf{R})$ a q-cochain. Define the **cup product**, $\alpha \smile \beta \in C^{p+q}(X;\mathbf{R})$ to be the cochain whose value on a singular $(p+q)$-simplex $\sigma\colon \Delta^{p+q} \to X$ is given by restriction of the canonical simplex $\Delta^{p+q} = [v_0, v_1, \ldots, v_{p+q}]$ to the 'first' p-simplex and the 'last' q-simplex:

$$(\alpha \smile \beta)\sigma := \alpha\left(\sigma|[v_0, v_1, \ldots, v_p]\right) \cdot \beta\left(\sigma|[v_p, v_{p+1}, \ldots, v_{p+q}]\right),$$

where the product is in the ring structure of the coefficients $\mathbf{R}$. One shows that for α and β cocycles, the cup product is a cocycle as well, therefore inducing an operation on $H^\bullet$ which gives it the structure of a ring. The cup product on $H^\bullet(X;\mathbf{R})$ for $\mathbf{R}$ a commutative ring is, like the wedge product, graded anti-commutative: $\beta \smile \alpha = (-1)^{pq}\alpha \smile \beta$.

If the reader finds this definition confusing, it is perhaps advisable to think of cubes rather than of simplices. For a cubical complex with α and β cochains on p-cubes and q-cubes respectively, then the product chain $\alpha \smile \beta$ evaluated on a *topological product* $(p+q)$-cube is the *algebraic product* of the cochain values on the factor p-cube and q-cube. In this setting, the parallel to the de Rham wedge product is clearest.

Example 6.23 (Projective space cohomology rings) ⊙

The ring structure on $\mathbb{P}^n$ in $\mathbb{F}_2$ coefficients is particularly satisfying: it is the ring of polynomials in one variable, x, modulo the ideal generated by x^{n+1}:

$$H^\bullet(\mathbb{P}^n;\mathbb{F}_2) \cong \mathbb{F}_2[x]/\left(x^{n+1}\right), \tag{6.5}$$

where $x \in H^1(\mathbb{P}^n;\mathbb{F}_2)$. This computation is not elementary (see, *e.g.*, [176]), but it has important consequences. For example, the ring structure reveals that $\mathbb{P}^3$ is not homotopic to $\mathbb{P}^2 \vee \mathbb{S}^3$, even though they have isomorphic $\mathbb{F}_2$ cohomology *groups*. Though both spaces have $H^k \cong \mathbb{F}_2$ for $k \leq 3$ and 0 otherwise, their ring structures differ since the generator $x \in H^1$ satisfies $x^3 = 0$ for $\mathbb{P}^2 \vee \mathbb{S}^3$ but $x^3 \neq 0$ for $\mathbb{P}^3$. ⊙

6.11 Currents

On smooth manifolds, calculus provides the convenient language of forms for cohomology. Duals of forms provide an extremely flexible interpolation between smooth and discrete homological structures on manifolds that allow one to talk about, *inter alia*, the homology class of a vector field. This section touches on analytic tools based on geometric measure theory [119, 134, 234]. To avoid the numerous technicalities involving regularity and rectifiability, let the reader assume (via restriction to the o-minimal structure of globally subanalytic sets) sufficient (piecewise) smoothness where needed.

Fix M an oriented manifold of dimension n. Let $\Omega_p(M) := \left(\Omega^p_c(M)\right)^\vee$ be the space of p-**currents** – real-valued functionals on compactly supported p-forms. Currents have a homological nature. Given any p-current $T \in \Omega_p$, the boundary of T, $\partial T \in \Omega_{p-1}$, is defined via the adjoint to the exterior derivative: $\partial T(\alpha) = T(d\alpha)$. A **cycle** is a current T with $\partial T = 0$. Clearly, $\partial^2 = 0$, and there is a resulting chain complex $(\Omega_\bullet(M), \partial)$ with ensuing **de Rham homology** ${}_{dR}H_\bullet(M)$. The analogue of Theorem 6.21 holds: ${}_{dR}H_\bullet(M) \cong H_\bullet(M;\mathbb{R})$.

The chief advantage in using currents is, as with all things homological, visualizability. An oriented p-dimensional submanifold in M is a p-current, since one can integrate a p-form over it: its de Rham homology class coincides with its singular homology class. For example, any (piecewise-smooth) oriented knot or link is a 1-current, since one can integrate a 1-form over oriented curves. Upon fixing a volume form on M, a piecewise-smooth vector field V is likewise a 1-current, since any 1-form

α pairs with V pointwise as $\alpha(V)$, which may then be integrated over M (to a finite value, thanks to compact support of α). The 2-currents on a manifold with volume form can range in shape from oriented surfaces to pairs of vector fields to a pair of tangent vectors at a single point.

Example 6.24 (Volume preserving links) ⊙

One beauty of the language of currents is that it allows one to compare both knots and vector fields on a manifold. Recall the definitions of links from Example 4.24 as disjoint embedded loops in $\mathbb{S}^3$. A vector field is an entirely different class of objects; one notes that the flowlines of a vector field have the potential to close up into periodic orbits – a link. The similarity seems to end there.

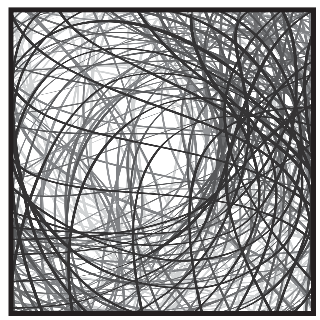

However, it follows from a result of Sullivan [284, Prop. II.25] that any closed nullhomologous 1-current, such as a volume-preserving vector field on an oriented manifold, can be realized as a limit of 1-currents supported on a compact 1-dimensional submanifold: an oriented link. This implies that any volume-preserving flow on $\mathbb{S}^3$ is the limit (in the sense of 1-currents) of a sequence of ever-lengthening, ever-coiling links. This suggests a reformulation of knot/link theory in terms of volume-preserving vector fields on $\mathbb{S}^3$. ⊙

Example 6.25 (Helicity and fluids) ⊙

It has been known for a long time what is the appropriate asymptotic analogue of linking number for volume-preserving vector fields on $\mathbb{S}^3$ [14, 16, 273]. The construction is as follows: given any two points $x, y \in \mathbb{S}^3$, evolve them forward under the flow of the vector field for times s and t respectively, until the flowlines come close to their starting points (that this happens for almost-every x and y infinitely often follows from the Poincaré Recurrence Theorem for volume-preserving flows [258]). Close these curves with short paths and compute the linking number (well-defined for almost all x and y). The limit of this linking number, normalized by st, converges as $s, t \to +\infty$ to a function $\ell k(x, y)$, which, when integrated over $\mathbb{S}^3 \times \mathbb{S}^3$ with the conserved volume form, yields the **asymptotic linking number** of the flow,

$$\ell k(V) := \int_{\mathbb{S}^3}\int_{\mathbb{S}^3} \ell k(x, y)\, \mathrm{dvol}_x\, \mathrm{dvol}_y.$$

The techniques of forms and currents makes the computation of this seemingly-intractable quantity elementary. A volume-preserving vector field V on $\mathbb{S}^3$ is closed and nullhomologous as a 1-current; this implies that the vector field contracted into the volume form μ yields an exact 2-form $\mu(V, \cdot, \cdot) = d\alpha$ for some α. The **helicity** of V is the integral of the wedge product of α with its derivative:

$$\mathcal{H}(V) := \int_{\mathbb{S}^3} \alpha \wedge d\alpha,$$

One shows well-definedness with respect to choice of α via Stokes' Theorem. Arnol'd [14] (following Moffat [232] (following Calŭgareanŭ [76])) showed that the helicity is the asymptotic linking number:

Theorem 6.26 (Helicity Theorem). $\mathcal{H}(V) = \ell k(V)$

As a corollary, the helicity is an invariant of V under the action of volume-preserving diffeomorphisms of $\mathbb{S}^3$, since linking numbers are unchanged by such. This is of great significance in fluid dynamics, since the velocity field of an ideal fluid evolves in time according to the Euler equations (Example 6.17), and the energy of the fluid (the integral of the norm of the velocity field) is bounded below by helicity – $\mathcal{H}$ is a topological measure of a fluid's inability to relax. The proof of the Helicity Theorem in [16] uses currents on $\mathbb{S}^3 \times \mathbb{S}^3$ to capture linking behavior. ⊚

Example 6.27 (Normal and conormal cycles) ⊚

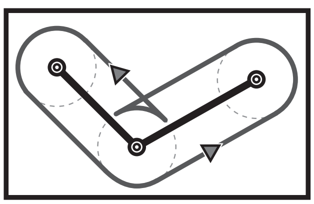

Many of the constructs of Chapter 3 concerning Euler characteristic and intrinsic volumes have a representation in terms of currents [234]. The **normal cycle** of a tame set $A \subset \mathbb{R}^n$ is a special $(n-1)$-current $\mathbf{N}^A$ on the unit cotangent bundle $UT^*\mathbb{R}^n \cong \mathbb{S}^{n-1} \times \mathbb{R}^n$. For $A \subset \mathbb{R}^n$ of positive codimension, the normal cycle is best visualized as having support on the set of points a *'unit'* distance from A.

The **conormal cycle** of a tame set $A \subset \mathbb{R}^n$ is a particular n-current $\mathbf{C}^A \in \Omega_n(T^*\mathbb{R}^n)$ on the cotangent bundle. It is, for lack of a better explanation in this text, the *cone* over the normal cycle. Each of the intrinsic volumes μ_k of §3.10, including Euler characteristic $\chi = \mu_0$, can be defined as the integral of a canonical form ($\alpha_k \in \Omega_c^{n-1}(UT^*\mathbb{R}^n)$ or $\omega_k \in \Omega_c^n(T^*\mathbb{R}^n)$) over the appropriate cycle:

$$\mu_k(A) = \int_{\mathrm{N}^A} \alpha_k = \int_{\mathrm{C}^A} \omega_k .$$

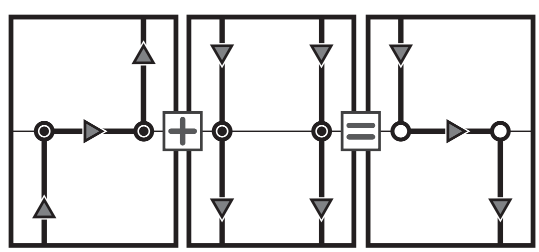

Additivity of the intrinsic volumes is expressed in terms of additivity of currents: *e.g.*, $\mathbf{N}^{A\cup B} = \mathbf{N}^A + \mathbf{N}^B - \mathbf{N}^{A\cap B}$ and likewise for $\mathbf{C}$. This means, *e.g.*, that one can see the difference between the Euler characteristic of a compact disc $\chi(\mathbb{D}^n) = 1$ and of its interior $\chi(\mathbb{D}^n - \partial\mathbb{D}^n) = (-1)^n$ as being a *reflection*. When subtracting the conormal cycle $\mathbf{C}^{\partial\mathbb{D}^n}$ from $\mathbf{C}^{\mathbb{D}^n}$, the support in $\mathbb{R}^n$ is the same, but the orientation in each axis is reversed. For n odd, this results in an orientation-reversal, reflected in the sign change. ⊚

6.12 Laplacians and Hodge Theory

With the addition of a geometric structure, there is another manifestation of Poincaré duality in cohomology for manifolds via partial differential equations. Recall that for a vector space V of dimension n, the algebra of forms, $\mathbf{\Lambda}(V)$, displays a combinatorial duality: $\dim \mathbf{\Lambda}^p(V) = \binom{n}{p} = \dim \mathbf{\Lambda}^{n-p}(V)$. Fix a geometry on V in the form of an inner product $\langle\cdot,\cdot\rangle$ and choose an orthonormal basis $\{x_i\}_1^n$. Fix also an orientation on V in the form of an equivalence class of orderings $[x_i]_1^n$ of basis elements up to even permutations. Define the **Hodge star** $\star\colon \mathbf{\Lambda}^p(V) \to \mathbf{\Lambda}^{n-p}(V)$ on basis elements as follows:

$$\star dx_{i_1} \wedge dx_{i_2} \wedge \cdots \wedge dx_{i_p} := dx_{i_{p+1}} \wedge \cdots \wedge dx_{i_{n-1}} \wedge dx_{i_n},$$

where the ordering $[x_{i_j}]_{j=1}^n$ respects orientation. Extend to all of $\mathbf{\Lambda}(V)$ via linearity. The Hodge star depends only on the inner product and the orientation, not on the basis itself. It satisfies a signed duality $\star\star = (-1)^{p(n-p)}\mathrm{Id}$.

If an oriented manifold M is Riemannian (see Example 6.17) then the Hodge star extends to $\star\colon \Omega^p \to \Omega^{n-p}$. For example, in Euclidean $\mathbb{R}^3$ with the standard basis as per Example 6.15, $\star\alpha_{\vec{F}} = \beta_{\vec{F}}$. Every oriented Riemannian manifold has a well-defined **volume form** $\mu \in \Omega^n$ which, in local orthonormal coordinates, is $dx_1 \wedge \cdots \wedge dx_n$ and which is given by $\mu = \star\mathbb{1}_M$. The Hodge star yields an inner product on each Ω^p via integration:

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

With this geometry in place, one may define a codifferential $\delta\colon \Omega^p \to \Omega^{p-1}$ given by the adjoint: $\langle\alpha, d\beta\rangle = \langle\delta\alpha,\beta\rangle$; more explicitly, $\delta := (-1)^{p(n-p)}\star d\star$. The **Laplacian** is the operator $\Delta\colon \Omega^\bullet \to \Omega^\bullet$ given by:

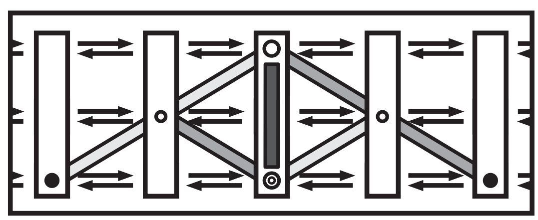

$$\Delta := (d+\delta)^2 = d\delta + \delta d.$$

Note that the Laplacian is degree zero, and for $p=0$ is the familiar second-order differential operator. The Laplacian blends analytic, geometric, and topological features. The **harmonic** forms are defined as ${}_\Delta H^\bullet(M) := \ker \Delta$, the kernel of the Laplacian.

Theorem 6.28 (Hodge Theorem). *For M a compact oriented Riemannian manifold, $\Omega^p(M)$ has an orthogonal decomposition:*

$$\Omega^p = d\Omega^{p-1} \oplus {}_\Delta H^p \oplus \delta\Omega^{p+1}.$$

Corollary 6.29. *For M a compact oriented Riemannian manifold, ${}_\Delta H^\bullet(M) \cong {}_{dR}H^\bullet(M)$.*

The bother of working with geometry has the following payoff: the Hodge star $\star$ is an incarnation of Poincaré duality. Let $\alpha \in {}_\Delta H^p$ be a harmonic form. Theorem 6.28 implies that $d\alpha = \delta\alpha = 0$. This implies that $\star\alpha$ is also harmonic, since

$$\Delta(\star\alpha) = (-1)^{p(n-p)}(d\star d\star + \star d\star d)\star\alpha = (d\star d + \star d\delta)\alpha = 0.$$

Thus, one may realize Poincaré duality as the isomorphism $\star\colon {}_{\Delta}H^k(M) \to {}_{\Delta}H^{n-k}(M)$; *cf.* Theorem 6.22.

One of the benefits of using differential-topological constructs is the ability to import and export ideas between smooth and discrete frameworks. There is a simple simplicial analogue of the Hodge theorem which has the advantage of requiring no forms, differentiability, or manifold structures, but merely an implicit geometry. Consider a cell complex X with cellular cochain complex $\mathcal{C} = (C^\bullet, d)$ in $\mathbb{R}$ coefficients. Choose an inner product $\langle\cdot,\cdot\rangle$ so that indicator functions over the cells of X are orthogonal. The implicit choice is in the **weight** of each simplex $\langle \mathbb{1}_\sigma, \mathbb{1}_\sigma\rangle = \omega_\sigma \in \mathbb{R}$. With this inner product structure, use the adjoint δ of d to define the **discrete Laplacian**: $\Delta = d\delta + \delta d$. As in the smooth theory, the harmonic cochains are ${}_{\Delta}H = \ker \Delta$. The following discrete Hodge theorem becomes a simple exercise.

Theorem 6.30 (Discrete Hodge Theorem). *For X a finite cell complex with choice of weights, $C^p = dC^{p-1} \oplus {}_{\Delta}H^p \oplus \delta C^{p+1}$.*

Example 6.31 (Graph Laplacian) ⊚

Discrete Laplacians have seen the greatest use in graph theory under the guise of the **graph Laplacian**. For X an undirected graph and $f\colon V(X) \to \mathbb{R}$ a function on the nodes, the graph Laplacian of f is defined as

$$(\Delta_X f)(v) = \sum_{w\colon (v,w)\in E(X)} f(w) - f(v).$$

The reader may verify that this agrees with the Laplacian on $C^0(X)$ under the obvious inner-product structure on chains. It is clear that a harmonic 0-chain is one in which each node's value is the average of its neighbors'. Graph Laplacians have found extensive use in algorithms [281], image processing [71] and much more. The principal (smallest nonzero) eigenvalue of the graph Laplacian controls the behavior of random walks on a graph and points to interesting generalizations in random simplicial complexes [171]. ⊚

Example 6.32 (Distributed homology computation) ⊚

The Laplacian is a local operator and, as such, is well-suited to distributed computation. The work of Tahbaz-Salehi and Jadbabaie [267] details the use of the simplicial Laplacian to distributed computation of the homological coverage criterion of §5.6. By Theorems 5.10 and 6.30, verified coverage in a sensor network in $\mathbb{R}^2$ follows from showing that $\ker \Delta = 0$ on discrete 1-forms of the flag complex F of the network. It is easy to see that the **heat equation**, $\frac{d}{dt}\alpha = \Delta\alpha$, has 0 as an asymptotically stable solution if and only if $\ker \Delta = 0$. Thus, by solving a heat equation with random initial conditions, one can safely (to the degree one trusts in random initial conditions) conclude coverage if the solution converges to zero. That this equation can be solved locally and in a distributed manner [235] should come as no surprise to the reader who has spent time with the heat equation, though the convergence to the solution can be slow. This can be improved by instead using a wave equation approach [265]. ⊚

6.13 Circular coordinates in data sets

In §5.14 the problem of determining the topology of a point cloud was addressed by means of persistent homology. A cohomological approach becomes the appropriate tool for addressing a related problem of coordinatizing a point cloud.

Assume for the present that $\mathcal{Q} \subset \mathbb{R}^n$ is a point-cloud whose topology is known or suspected to be sufficiently *circular* so as to merit outfitting circular coordinates $\Theta\colon \mathcal{Q} \to \mathbb{S}^1$ in a manner that respects the underlying topology of the space $X \subset \mathbb{R}^n$ (homotopic to $\mathbb{S}^1$) that $\mathcal{Q}$ is presumed to sample. Many of the existing algorithms for assigning circular coordinates to a point cloud [263, 288] have implicit convexity assumptions.

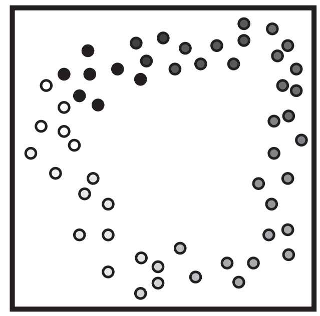

The solution of de Silva, Morozov, and Vejdemo-Johansson [89] is a slick application of algebraic topology that highlights the particular benefits of cohomology and the role of coefficients. The outline of their work is as follows.

1. One begins with the following result from homotopy theory: for any space X, the group of (basepoint-preserving) homotopy classes $[X, \mathbb{S}^1]$ of maps $X \to \mathbb{S}^1$ is isomorphic to $H^1(X; \mathbb{Z})$ (see §8.6). The coordinatization function $\Theta\colon X \to \mathbb{S}^1$ therefore is naturally approachable via cohomology.
2. To find a cohomology class for X based on a sampling of nodes $\mathcal{Q}$, compute the persistent cohomology of a sequence of Vietoris-Rips complexes, as in §5.13-5.14. For the Structure Theorem (Theorem 5.21), field coefficients are required; for numerical reasons (to avoid roundoff errors), coefficients in a *finite* field $\mathbb{F}_p$ are preferred.
3. A (persistent) class $[\alpha_p] \in H^1(X; \mathbb{F}_p)$ can be converted to a integral class $[\alpha] \in H^1(X; \mathbb{Z})$ by means of the following process. The short exact sequence of coefficients $0 \longrightarrow \mathbb{Z} \xrightarrow{\cdot p} \mathbb{Z} \longrightarrow \mathbb{F}_p \longrightarrow 0$ yields a short exact sequence of cochain complexes on X and, via Lemma 5.5, a long exact sequence on cohomology:

$$\longrightarrow H^1(X;\mathbb{Z}) \longrightarrow H^1(X;\mathbb{F}_p) \xrightarrow{\delta} H^2(X;\mathbb{Z}) \xrightarrow{\cdot p} H^2(X;\mathbb{Z}) \longrightarrow$$

   The kernel of $\cdot p\colon H^2(X;\mathbb{Z}) \to H^2(X;\mathbb{Z})$ consists of p-torsional cohomology classes: for $p > 2$ these would seem to be rare occurrences in *organic* spaces X living behind data sets. By exactness, $\ker(\cdot p) = \operatorname{im}\delta$; assuming this is zero, it implies that $H^1(X;\mathbb{Z}) \to H^1(X;\mathbb{F}_p)$ is surjective, and (persistent) classes in $\mathbb{F}_p$ coefficients therefore lift to integral classes.
4. The resulting integer cocycle α is perhaps a poor $\mathbb{S}^1$-coordinatization – all the circular motion may be concentrated over a small subset of X. To relax α to a smooth circular coordinate system, lift to $\mathbb{R}$ coefficients and find a cohomologous harmonic cocycle $\overline{\alpha} \in {}_{\Delta}H^1(X)$. Thanks to the local-averaging properties of the Laplacian, this 1-cocycle integrates to a well-regulated coordinate function $\Theta\colon X \to \mathbb{S}^1$.

For details on computational aspects and implementation, see [89]. This work illustrates well the utility of cohomology, while highlighting the delicate interplay between real, integral, and cyclic coefficients.

Notes

1. This chapter is woefully incomplete: a short, motivational text cannot do justice to cohomology theory. The interested reader should resolve to learn the theory properly. Hatcher [176] is, as ever, the best place to begin. For the de Rham theory, Bott and Tu [46] is the classic lucid source, and Fulton [135] is more elementary still.
2. The conflation of objects with duals is ubiquitous and insidious. Examples include confusing gradient 1-forms with vector fields and defining simplicial chains as functions from simplices to coefficients.
3. The idea of the impossible tribar as a 1-cocycle was suggested by Penrose [243]. One can imagine more interesting Escherian illusions based on H^2.
4. One of the many cohomology theories not covered in this text is related to configuration spaces. Given X a topological space and $\mathbf{R}$ a ring, let C^k be the set of functions (not necessarily maps!) $f\colon X^{k+1} \to \mathbf{R}$. This gives a complex $\mathcal{C}$ with differential d taking $f(x_0, \ldots, x_k)$ to the sum $\sum_i (-1)^i f(x_0, \ldots, x_{i-1}, x_{i+1}, \ldots, x_k)$. The set of all functions which vanish in a neighborhood of the grand diagonal $(x, \ldots, x)$ forms a subcomplex $\mathcal{C}^0$. The **Alexander-Spanier cohomology** of X is $H^\bullet(\mathcal{C}/\mathcal{C}^0)$. It is, for reasonable spaces, isomorphic to the singular $H^\bullet(X)$ [280]. This theory seems suspiciously relevant to applications in configuration spaces.
5. Helly's Theorem is important in a wide array of combinatorics and optimization problems [9]. That is has a purely topological proof is a testament to the power of topological methods. The homological proof was known to Helly and many topological generalizations exist. It is remarkable how many experts nevertheless consign Helly's Theorem to convex geometry.
6. There is, as one might suspect, a deeper form of duality, of which Poincaré, Alexander, and Lefschetz are emanations. **Verdier duality** for sheaves is perhaps the best encapsulation of the scope and power of duality theorems in co/homology. Chapter 9 will provide some of the requisite background for that theory.
7. The cup product is more important than may at first appear. It is good to visualize it using differential forms. Even better is its homological adjoint. Theorem 6.22 (and an illustration or two) hints at a relationship between $\smile$ and homology: the **cap product** $\frown\colon H_p(X) \times H^q(X) \to H_{p-q}(X)$ is defined on a chain σ and a cochain α via

   $$(\sigma \frown \alpha) := \alpha\left(\sigma|[v_0, v_1, \ldots, v_q]\right) \cdot \sigma|[v_q, v_{q+1}, \ldots, v_p].$$

   In field coefficients, $\frown$ and $\smile$ are related via $\beta(S \frown \alpha) = (\alpha \smile \beta)(S)$ for a cycle S and cocycles α, β.
8. The reader may wonder, given the utility of cohomology with compact supports, where is the corresponding homology with compact supports? This exists and goes under the name of **Borel-Moore homology**.
9. A very clean proof of the de Rham Theorem (6.21) uses a double complex [46] (see the notes to Chapter 5).
10. Hodge theory is merely a hint at how partial differential equations (in this case, Laplace's equation) on geometric manifolds can lead to topological invariants. Many subtler and deeper invariants come from other PDEs using auxiliary structures and have implications in string theory, algebraic geometry, knot theory, and more.

11. Simplicial/cellular Hodge theory is having significant impact in numerical analysis [12, 13, 32, 94] via **discrete exterior calculus**. The discretization of space and time can destroy auxiliary structures or symmetries within the underlying differential equations: *e.g.*, the symplectic structure implicit in celestial mechanics. By working with simplicial p-forms and discretizing the conservation laws and symmetries themselves, one is led to more accurate numerical solutions: see Example 10.23.