FROM HOMOLOGY TO DE RHAM COHOMOLOGY

An Expository Journey through the Topology of Smooth Manifolds

Authors

Anshveer Bindra, Nathan Lee, Isabelle Wu

August 23, 2025

Contents

1 Introduction

Mathematics has a curious tendency: as soon as one language for describing a phenomenon is mastered, another emerges—offering a radically different viewpoint, yet somehow speaking about the same thing. In geometry and topology, one such shift occurs when we move from the *local* differential properties of objects to *global* invariants that persist under continuous deformations. This is the transition from the world of smooth functions, curvature formulas, and differential equations to the world of *homology* and *cohomology* theories.

Up to this point in one's study, it is easy to become immersed in geodesics, curvature tensors, and integrations over manifolds, extracting precise geometric information from the smooth structure. The Gauss–Bonnet theorem, for instance, offers a tantalizing glimpse of something deeper: integrating curvature over a closed surface yields an integer, the Euler characteristic, which remains unchanged by bending or stretching the surface. A purely *analytic* quantity (the integral of curvature) mysteriously encodes a purely *topological* number.

This is not an isolated miracle. It belongs to a vast network of ideas whose central question is:

> *How can we translate questions about shape and connectivity into algebraic language, so that they can be computed, compared, and understood in new ways?*

The modern answer begins with the theory of *homology* and extends, in elegant and far-reaching ways, to *cohomology*.

1.1 A Historical Detour: From Shapes to Algebra

To appreciate why cohomology was ever conceived, we must step back and examine the intellectual landscape of the 18th through early 20th centuries. During this period, geometry underwent a transformation: it ceased to be merely the study of lines and figures drawn on paper and became the study of *space* itself.

The seeds of this transformation can be traced to the work of Leonhard Euler in the mid-18th century. In his analysis of polyhedra, Euler noticed a remarkable invariant: for any convex polyhedron, the quantity

$$V - E + F$$

(where V is the number of vertices, E the number of edges, and F the number of faces) is always 2. This number, now recognized as the *Euler characteristic*, was one of the first examples of a topological invariant—unchanged under continuous deformation, so long as the polyhedron was not torn or glued.

The 19th century saw an acceleration in the study of such invariants. Johann Benedict Listing coined the term "Topologie" in 1847, marking a conceptual break from traditional geometry. August Möbius and Camille Jordan developed early methods to describe and classify surfaces beyond their metric properties. Yet these developments were often ad hoc—ingenious, but lacking a unified framework. A Möbius strip could be described, a Klein bottle imagined, but there was no general theory for cataloging such spaces by their essential features.

The unifying vision came from Henri Poincaré in the 1890s. In his monumental series of papers, *Analysis Situs*, Poincaré introduced a systematic way to detect and classify the "holes" in a space. He defined numbers—later called *Betti numbers* in honor of Enrico Betti—that, in rough terms, count the number of independent cycles: closed loops, closed surfaces, and their higher-dimensional analogues that cannot be shrunk to a point. Poincaré's insight was revolutionary: one could now assign to each space a sequence of algebraic invariants that faithfully recorded its large-scale structure.

By the early 20th century, Poincaré's program was placed on a rigorous algebraic footing, largely through the work of Emmy Noether and her contemporaries. Intuitive surfaces and loops were replaced by formal *chains*, *boundaries*, and *quotient groups*, giving birth to modern *homology theory*. Homology provided a universal framework for distinguishing spaces: a sphere and a torus, though locally indistinguishable to a differential geometer, now had entirely different algebraic signatures.

1.2 From Counting Holes to Algebraic Invariants

The genius of homology lies in its shift from local to global. Locally, a sphere and a torus each resemble the Euclidean plane; any sufficiently small patch can be flattened without distortion. Yet globally, their structures are profoundly different. A sphere has no 1-dimensional holes, while a torus has one in each independent direction of its loop.

Topology asks: what properties of a shape survive stretching and bending, but not tearing or gluing? Algebraic topology answers by associating to each space algebraic objects—numbers, groups, rings—that remain unchanged under homeomorphisms. Homology was the first grand embodiment of this philosophy, systematically cataloging the "holes" in every dimension.

1.3 Why Cohomology?

Homology, powerful as it is, tells only half the story. Very soon after its introduction, mathematicians realized that there is a natural *dual* point of view: instead of examining the chains that make up a space, one could examine the algebraic data—functions, differential forms, or more general cochains—defined *on* those chains.

This gave rise to *cohomology*. Initially introduced as the algebraic dual of homology, cohomology quickly revealed itself to be richer. Cohomology groups naturally form a ring via the cup product, encoding not only which cycles exist but how they interact with one another. This algebraic structure gave cohomology a unifying role in fields as diverse as algebraic geometry, number theory, and physics.

1.4 Enter de Rham Cohomology

The bridge between cohomology and differential geometry was built in the 1930s by Georges de Rham. His guiding question was simple to state: given a closed differential form on a smooth manifold, when is it exact—that is, the exterior derivative of another form? If it is not exact, what topological feature of the manifold obstructs it?

De Rham's answer was the creation of *de Rham cohomology*, whose groups measure precisely these obstructions. His theorem—now a cornerstone of modern geometry—states that for a smooth manifold, the de Rham cohomology groups (defined analytically) are isomorphic to the singular cohomology groups with real coefficients (defined topologically). In one stroke, de Rham linked the analytic world of smooth forms with the algebraic world of topological invariants.

1.5 From Gauss–Bonnet to a General Principle

The Gauss–Bonnet theorem becomes almost inevitable from this perspective. The Gaussian curvature K is not merely a geometric measure; it defines a closed 2-form on the manifold. Integrating this form over the entire surface yields the Euler characteristic, a purely topological number. This is a special case of a far more general principle:

> *Integrals of certain closed differential forms compute topological invariants.*

De Rham cohomology is the natural framework for understanding and generalizing such results.

1.6 The Plan of This Exposition

In what follows, we will:

1. Motivate and define *homology* through geometric intuition and illustrative examples, introducing cycles, boundaries, and the construction of homology groups.

2. Shift to *cohomology*, first in the abstract algebraic setting, then concretely via differential forms, the exterior derivative, and the generalized Stokes' theorem.

3. Develop *de Rham cohomology* in detail, proving key results and computing examples on familiar manifolds such as spheres and tori.

4. Situate these concepts within the broader mathematical landscape, emphasizing their unifying role in geometry, topology, and analysis.

By the end, we will see that cohomology—and de Rham cohomology in particular—is not merely a rephrasing of topology in analytic language, but a profound framework in which classical results like Gauss–Bonnet find their natural home, and from which new insights continually emerge.

2 Review of Prerequisites

Before we embark on the study of homology and cohomology, it is essential to recall certain foundational concepts from differential geometry, multivariable calculus, and manifold theory.

2.1 Smooth Manifolds and Tangent Spaces

We begin by recalling the notion of a *smooth manifold.* A smooth manifold M of dimension n is, informally, a space that *locally* resembles $\mathbb{R}^n$ and has a globally consistent system of smooth coordinate charts. This means that for every point $p \in M$, there exists a neighborhood $U \subset M$ and a homeomorphism (called a *chart*)

$$\varphi : U \to \varphi(U) \subset \mathbb{R}^n$$

such that transition maps between overlapping charts are smooth functions.

The *tangent space* T_pM at a point p can be introduced in several equivalent ways: as equivalence classes of curves through p, as derivations acting on smooth functions, or as the span of coordinate partial derivatives. Intuitively, T_pM contains all possible velocity vectors of curves in M passing through p. For surfaces embedded in $\mathbb{R}^3$, one may visualize T_pM as the plane that just touches the surface at p without intersecting it locally.

2.2 Orientation

An *orientation* on a manifold is, loosely speaking, a consistent choice of "direction" for volume measurement throughout the manifold. In $\mathbb{R}^n$, the standard orientation is given by the ordered basis $(\mathbf{e}_1, \ldots, \mathbf{e}_n)$. For a surface $S \subset \mathbb{R}^3$, an orientation can be induced by a smooth choice of unit normal vector field $\mathbf{n}$ satisfying certain compatibility conditions with the parametrization of S.

This notion becomes crucial when formulating integral theorems: the sign of an integral over an oriented manifold depends on the chosen orientation, and changing orientation generally changes the sign of the result.

2.3 Line and Surface Integrals in $\mathbb{R}^3$

In the classical setting of vector calculus, we frequently integrate vector fields along curves or across surfaces.

Line integrals: Let $\mathbf{F} : U \subset \mathbb{R}^3 \to \mathbb{R}^3$ be a vector field and let C be a smooth curve parametrized by $\mathbf{r}(t)$, $a \leq t \leq b$. The *line integral* of $\mathbf{F}$ along C is

$$\int_C \mathbf{F} \cdot d\mathbf{r} = \int_a^b \mathbf{F}(\mathbf{r}(t)) \cdot \mathbf{r}'(t)\, dt.$$

Physically, this can represent the work done by the force field $\mathbf{F}$ in moving a particle along C.

Surface integrals: Let S be a smooth oriented surface parametrized by $\mathbf{r}(u, v)$, $(u, v) \in D \subset \mathbb{R}^2$, with unit normal vector $\mathbf{n}$ consistent with the given orientation. The *surface integral* of $\mathbf{F}$ over S is given by

$$\iint_S \mathbf{F} \cdot d\mathbf{S} = \iint_D \mathbf{F}(\mathbf{r}(u, v)) \cdot \left(\frac{\partial \mathbf{r}}{\partial u} \times \frac{\partial \mathbf{r}}{\partial v} \right) du\, dv.$$

Here $d\mathbf{S} = \mathbf{n}\, dS$ denotes the oriented surface element.

2.4 The Classical Stokes' Theorem

The classical version of Stokes' theorem relates a surface integral of the curl of a vector field to a line integral of the field around the boundary of the surface.

Theorem 2.1 (Classical Stokes' Theorem)**.** Let S be an oriented smooth surface in $\mathbb{R}^3$ with a smooth, simple, closed, positively oriented boundary curve ∂S. Let $\mathbf{F} : U \to \mathbb{R}^3$ be a smooth vector field defined on an open set U containing S. Then

$$\iint_S (\nabla \times \mathbf{F}) \cdot d\mathbf{S} = \oint_{\partial S} \mathbf{F} \cdot d\mathbf{r}.$$

This theorem is a unifying statement: the classical fundamental theorem of calculus, Green's theorem in the plane, and the divergence theorem can all be seen as special cases of Stokes' theorem under suitable choices of S and $\mathbf{F}$.

In intuitive terms, Stokes' theorem tells us that the total "circulation" of a vector field through a surface can be measured equally well by examining the behavior of the field along the surface's boundary. This principle reflects the deep interplay between local differential properties (curl) and global integral properties (line integrals).

We assume familiarity with these results and focus on recalling them only as needed for our generalization.

3 Homology Theory

3.1 Simplices and Simplicial Complexes

When faced with a complicated geometric object—such as a manifold—one natural strategy is to replace it by a simpler object that retains enough of its essential structure for our purposes. This is a recurring theme in mathematics: we trade the full complexity of an object for a tractable approximation that is easier to analyse. In topology, one particularly effective approach is to build such approximations out of the simplest possible building blocks.

Let us begin in low dimensions. In 0 dimensions, the simplest geometric object is just a point. In 1 dimension, the simplest object is a line segment, determined entirely by its two endpoints. In 2 dimensions, the simplest object is a filled-in triangle: once we fix its three vertices, its shape is completely determined. In 3 dimensions, the corresponding object is the solid tetrahedron determined by four vertices. Observing this pattern, we may ask:

Is there a single notion of "simplest possible shape" that makes sense in every dimension?

Indeed, such an object exists, and it is called a *simplex*.

Definition 3.1. An *n-simplex* Δ^n is the convex hull of a set of $n+1$ affinely independent points in $\mathbb{R}^n$. Equivalently, it is the simplest possible n-dimensional geometric figure determined by $n+1$ vertices. Topologically, Δ^n is homeomorphic to the n-dimensional closed ball D^n.

For example:

- Δ^0 is a point.
- Δ^1 is a closed line segment.
- Δ^2 is a filled-in triangle.

Note that the n-simplex is topologically equivalent to D^n, the n-ball. Observe that every simplex contains lower-dimensional simplices on its boundary.

Definition 3.2. An *m-face* of an n-simplex ($m \leq n$) is the convex hull of any subset of $m+1$ vertices of the simplex. Faces with dimension strictly less than n are called *proper faces*.

Two simplices are said to be *properly situated* if their intersection is either empty or a face of both simplices (i.e., a simplex itself).

At this point, it is natural to imagine combining these building blocks. If two triangles share an entire edge, they fit together neatly along that edge; if two tetrahedra share an entire triangular face, they join seamlessly along that face. This leads to the idea of a *complex* built from simplices, with the rule that they should only meet in shared faces.

Definition 3.3. A *simplicial complex* K is a finite collection of simplices satisfying:

1. Every face of a simplex in K is also a simplex in K.
2. The intersection of any two simplices in K is either empty or a face of both, i.e, if $A, B \in K$, then A and B are both properly situated.

The *dimension* of K is the maximum of the dimensions of its simplices.

There is a deep reason simplicial complexes are important: every reasonable topological space (in particular, every compact smooth manifold) can be represented as the geometric realization of some simplicial complex.

Theorem 3.1 (Triangulation Theorem)**.** Every compact smooth manifold admits a triangulation; that is, it is homeomorphic to the geometric realization of a simplicial complex.

While we omit the proof, the idea is that even the most intricate surface can be "tiled" with sufficiently small simplices, just as a curved line can be approximated by short straight segments.

One can also consider this in a more abstract way.

Definition 3.4. An **abstract simplicial complex** is a finite set of vertices $\{v_0, \ldots, v_k\}$ together with a collection of finite subsets, called *abstract simplices*, such that if σ is an abstract simplex in the collection, then every subset of σ is also an abstract simplex.

This definition is entirely combinatorial: it does not tell us *where* the vertices lie in space or even whether they do so at all. It simply tells us which sets of vertices are allowed to form simplices. From this combinatorial data, we can obtain a genuine geometric object by *realizing* the complex.

Definition 3.5. A **geometric realization** of an abstract simplicial complex R is obtained by mapping each vertex v_i to a distinct point in some Euclidean space $\mathbb{R}^n$ and interpreting each abstract simplex $\{v_{i_0}, \ldots, v_{i_m}\}$ as the convex hull of its images in $\mathbb{R}^n$.

A particularly natural choice is the *standard realization*: if R has $k+1$ vertices, we may take $n = k+1$ and assign

$$v_0 \mapsto e_1, \quad v_1 \mapsto e_2, \quad \ldots, \quad v_k \mapsto e_{k+1},$$

where $\{e_1, \ldots, e_{k+1}\}$ are the standard basis vectors of $\mathbb{R}^{k+1}$. Each simplex then becomes a standard coordinate simplex.

Of course, we are free to place the vertices in different positions in $\mathbb{R}^n$, obtaining geometrically different-looking realizations. However, the combinatorial structure ensures that all realizations of the same abstract simplicial complex are topologically equivalent.

Theorem 3.2. If R is an abstract simplicial complex, any two geometric realizations of R are homeomorphic.

The intuition here is that the abstract complex encodes exactly the connectivity and incidence relations between simplices, and these are precisely the features that a homeomorphism preserves.

With this notion in hand, we can now think of shapes both as purely combinatorial objects and as embedded geometric ones. This dual viewpoint will be crucial: when defining homology, we will want to pass from the geometry of a space to an algebraic object that depends only on its combinatorial structure. Before doing so, we will need to introduce a few more definitions that let us manipulate simplices algebraically, keeping track of orientation and boundaries in a way that will let us detect and measure "holes."

3.2 Homology Groups

Up to this point, we have discussed simplices and simplicial complexes as the combinatorial scaffolding on which we can approximate and study topological spaces. Our goal in homology theory is to extract algebraic invariants from these structures—quantities that remain unchanged under continuous deformations and thus reveal intrinsic features of the space.

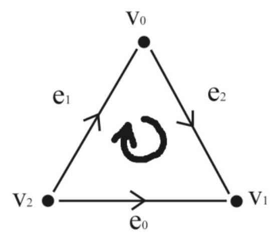

Figure 1: An oriented 2-simplex

To do so, we will need to be precise about how simplices are assembled, and in particular, how their orientations interact when forming boundaries. This leads us naturally to the notion of *oriented simplices*.

Definition 3.6. Let S be the set of vertices of a simplex. An **orientation** of the simplex is obtained by choosing an ordering of the vertices in S. If two such orderings differ by an *even permutation*, they are said to represent the same orientation; if they differ by an *odd permutation*, they represent the *opposite* orientation. Thus, any simplex admits exactly two possible orientations.

Intuitively, the orientation specifies a "direction" or "handedness" for traversing the vertices of the simplex. In the case of a 1-simplex (an edge), orientation corresponds to choosing a starting vertex and an ending vertex. For a 2-simplex (a triangle), it is equivalent to deciding whether we traverse its vertices clockwise or counterclockwise.

The choice of orientation on an n-simplex determines orientations on each of its $(n-1)$-dimensional faces in a systematic way. This is essential for defining boundaries consistently across a simplicial complex: the induced orientations on shared faces must cancel appropriately when forming chains.

Example 3.1. Consider the 2-simplex in Figure 1 with vertices (v_0, v_1, v_2), oriented in that order. The orientation induced on its edges (1-faces) is given by:

$$e_2 = (v_0, v_1), \quad e_0 = (v_1, v_2), \quad e_1 = (v_2, v_0).$$

Each edge inherits its direction from the omission of one vertex of the triangle, with a sign determined by the position of the omitted vertex.

Definition 3.7. Let $A^n = (v_0, v_1, \ldots, v_n)$ be an oriented n-simplex. The orientation of its $(n-1)$-dimensional face obtained by removing the i-th vertex, i.e. the face with vertex set

$$\{v_0, \ldots, v_{i-1}, v_{i+1}, \ldots, v_n\},$$

is given by:

$$F_i = (-1)^i (v_0, \ldots, v_{i-1}, v_{i+1}, \ldots, v_n).$$

Here, the factor $(-1)^i$ accounts for whether the removal of v_i from the ordering results in an even or odd permutation relative to the induced ordering of the face. This sign convention ensures that when we later define boundary operators, the boundaries of boundaries will cancel, a property essential for the consistency of homology.

Up to this point, our attention has been on individual simplices and their orientations. However, in a simplicial complex we rarely care about a single simplex in isolation. Shapes in topology are built from many simplices glued together, and we often want to describe entire collections of them at once.

How can we record such a collection in a way that remembers not only which simplices are present, but also how many times each appears and with which orientation? One natural idea is to take a "formal sum" of oriented simplices, allowing positive coefficients to indicate agreement with the chosen orientation and negative coefficients to indicate reversal. This leads us to the notion of an *n-chain*, an algebraic object that encodes both the combinatorial structure and the orientation data of a simplicial complex.

Definition 3.8. Given a set $A_1^n, \ldots, A_k^n$ of arbitrarily oriented n-simplices in a complex K and an abelian group G, we define an *n-chain with coefficients in G* as a formal sum

$$x = g_1 A_1^n + g_2 A_2^n + \cdots + g_k A_k^n,$$

where $g_i \in G$.

Up to this point, we have been talking about individual oriented simplices. But what if we want to treat several simplices together — perhaps even repeat them or take certain ones with a "negative" orientation? To do this rigorously, we attach a *coefficient* from G to each simplex and add them as a formal sum. These objects are called *n-chains*.

From now on, we will take $G = \mathbb{Z}$. This means every oriented simplex is assigned an integer coefficient — positive, negative, or zero.

The set of all n-chains forms an abelian group under addition: if

$$x = \sum_{i=1}^{k} g_i A_i^n \quad \text{and} \quad y = \sum_{i=1}^{k} h_i A_i^n,$$

then

$$x + y = \sum_{i=1}^{k} (g_i + h_i) A_i^n.$$

We denote this group by L_n.

Before we can talk about how chains interact, we need a way to capture the intuitive idea of the *boundary* of a simplex. Take a triangle as an example: it is completely determined by its three oriented edges. Similarly, a tetrahedron is determined by its four oriented triangular faces. In each case, the lower-dimensional simplices that "frame" the original simplex form what we naturally think of as its boundary.

From a combinatorial viewpoint, this boundary should be a sum of all those faces, each with an orientation induced from the simplex itself. The signs here are crucial: they ensure that when two simplices share a face, those shared faces appear with opposite orientations in the two boundaries and hence cancel when we sum over a complex — just as one would expect for a well-behaved boundary operator.

Thus, the boundary of an oriented n-simplex should be an $(n-1)$-chain consisting of all its $(n-1)$-faces, each taken with the correct induced orientation. We now formalize this notion:

Definition 3.9. Let A^n be an oriented n-simplex in a complex K. The *boundary* of A^n is the $(n-1)$-chain of K over $\mathbb{Z}$ given by

$$\partial(A^n) = A_0^{n-1} + A_1^{n-1} + \cdots + A_n^{n-1},$$

where A_i^{n-1} is the ith oriented $(n-1)$-face of A^n. If $n = 0$, we define $\partial(\Delta^0) = 0$.

Here we are extending the geometric idea of a boundary into algebraic form. Each oriented simplex has faces, and the orientations of these faces are not arbitrary — they are determined by the orientation of the original simplex. The boundary operator ∂ records these faces, preserving their induced orientations.

We can apply ∂ not just to a single simplex but to any chain, by extending it linearly:

$$\partial\left(\sum_{i=1}^{k} g_i A_i^n\right) = \sum_{i=1}^{k} g_i\,\partial(A_i^n).$$

That is, $\partial : L_n \to L_{n-1}$ is a group homomorphism. The definition of the boundary operator naturally raises a question: what happens if we take the boundary *twice*? Geometrically, the boundary of a boundary should be empty — after all, if you start with a triangle, its boundary is a loop of edges, and each edge's boundary consists of its two vertices, which always cancel out in pairs. Let us verify this idea explicitly for a 2-simplex.

Example 3.2. Let us compute $\partial(\partial(\Delta^2))$, where Δ^2 is the oriented 2-simplex shown in Figure 1. The boundary of Δ^2 is the sum of its three oriented edges:

$$\partial(\Delta^2) = e_1 + e_2 + e_3$$

where $e_1 = (v_0, v_1)$, $e_2 = (v_1, v_2)$, and $e_3 = (v_2, v_0)$.

Applying ∂ again, we use the definition of the boundary for 1-simplices:

$$\partial(\partial(\Delta^2)) = \partial(v_0, v_1) + \partial(v_1, v_2) + \partial(v_2, v_0)$$

$$= [(v_1) - (v_0)] + [(v_2) - (v_1)] + [(v_0) - (v_2)].$$

Now, because L_0 (the group of 0-chains) is abelian, the terms cancel: (v_1) appears once positively and once negatively, the same for (v_0) and (v_2). No vertex survives in the sum, and we obtain:

$$\partial(\partial(\Delta^2)) = 0.$$

This computation is not a coincidence: it is a concrete example of the general principle $\partial \circ \partial = 0$, which says that the boundary of a boundary is always empty. The cancellation here is exactly the algebraic manifestation of the geometric fact that once you have taken the full boundary of a simplex, there is nothing left to "border" those boundary pieces.

As one can observe, taking the boundary twice does indeed give 0. This is no accident. The same argument works in higher dimensions: since ∂ is linear and every oriented simplex has faces that cancel in the next boundary, we conclude

$$\partial^2(x) = 0 \quad \text{for all } x \in L_n.$$

Definition 3.10. An n-chain is called a *cycle* if its boundary is zero. The set of all n-cycles of K over $\mathbb{Z}$ is denoted Z_n, and we have

$$Z_n = \operatorname{Ker}(\partial) \subseteq L_n.$$

From our example, every boundary is a cycle — in fact, this is exactly the content of the property $\partial^2 = 0$. This principle will be central in defining homology groups, where we compare cycles to those that

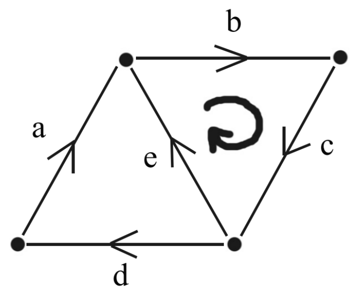

Figure 2: Boundaries?

arise as boundaries.

Up to this point, we have seen that the group Z_n consists of all n-chains whose boundary is zero—these are our *n-cycles*. Such cycles are "closed" in the sense that they have no edges left open, but they do not all carry the same topological significance. Some cycles are genuinely interesting because they encircle a *hole* in the space, while others are topologically trivial: they bound a higher-dimensional region entirely contained within the complex.

To make this idea concrete, imagine walking around the perimeter of a filled triangle in a simplicial complex. The three edges form a 1-cycle, but if we include the filled interior of the triangle (a 2-simplex), this boundary can be "shrunk" to a point without leaving the space. In this case, the cycle tells us nothing about the presence of a hole—it merely outlines something that is already filled in.

This motivates a refinement of our classification of cycles: we wish to declare cycles that arise as boundaries of higher-dimensional chains as *equivalent to zero* in homology, because they do not detect any new topological feature.

Definition 3.11. We say that an n-cycle x of a k-complex K is *homologous to zero* if it is the boundary of an $(n+1)$-chain of K, for $n = 0, 1, \dots, k-1$. A *boundary* is then any cycle that is homologous to zero. We denote this relation by $x \sim 0$, and the subgroup of Z_n consisting of boundaries is written B_n. By the definition of the boundary operator, we have

$$B_n = \mathrm{Im}(\partial).$$

Less formally, a cycle belongs to B_n if it *bounds* an $(n+1)$-dimensional portion of the complex. For example, in Figure 2, the chain $b + c + e$ encloses a collection of 2-simplices and is therefore a boundary, while $a + d + e$ does not bound any filled-in region and so is not.

The relation $x \sim 0$ immediately yields an *equivalence relation*: for two cycles x and y,

$$(x - y) \sim 0 \quad \Longrightarrow \quad x \sim y,$$

and in this case we say that x and y are *homologous*. Two homologous cycles differ only by something that "bounds" and hence represent the same topological feature.

Since B_n is a subgroup of Z_n, we can form the quotient group that measures the "essential" cycles.

Definition 3.12. The *n-dimensional homology group* of the complex K over $\mathbb{Z}$ is

$$H_n = Z_n / B_n.$$

Equivalently, using the chain complex formulation,

$$H_n \;=\; \mathrm{Ker}(\partial) \;/\; \mathrm{Im}(\partial).$$

This quotient group is the key object in homology theory.

Remark 3.1. It is important to interpret what the expression $H_n = Z_n/B_n$ is telling us. We start with all n–cycles in Z_n. Inside this group sit the boundaries B_n, which are cycles that arise as the boundary of some $(n+1)$–chain. When we form the quotient Z_n/B_n, we are declaring that two cycles should be regarded as *the same* if their difference is a boundary. In particular, every boundary is identified with 0 in H_n. From this point of view, H_n consists of those cycle classes that survive after all boundaries have been collapsed — these are exactly the "non–bounding" cycles, and so H_n measures the n–dimensional holes in the space.

There is a very concrete way to package the boundary maps we have just defined. Pick an ordering of all n–simplices $\{\sigma_1, \dots, \sigma_{N_n}\}$ (this is a basis of L_n) and an ordering of all $(n-1)$–simplices $\{\tau_1, \dots, \tau_{N_{n-1}}\}$ (a basis of L_{n-1}). For each σ_j, write its boundary in the $(n-1)$–simplex basis:

$$\partial_n(\sigma_j) = \sum_{i=1}^{N_{n-1}} \varepsilon_{ij}\, \tau_i, \qquad \varepsilon_{ij} \in \{-1, 0, +1\}.$$

By definition, $\varepsilon_{ij} = +1$ if τ_i is a face of σ_j with the induced orientation, $\varepsilon_{ij} = -1$ if it is a face with the opposite orientation, and $\varepsilon_{ij} = 0$ otherwise. Placing the coefficient lists $(\varepsilon_{1j}, \dots, \varepsilon_{N_{n-1},j})$ as the j-th column gives an integer matrix

$$[\partial_n] \in M_{N_{n-1}\times N_n}(\mathbb{Z})$$

that *is* the map ∂_n in coordinates. An n–chain $x = a_1\sigma_1 + \cdots + a_{N_n}\sigma_{N_n}$ corresponds to the column vector $x = (a_1, \ldots, a_{N_n})^\top$, and

$$\partial_n(x) = 0 \iff [\partial_n]\, x = 0, \qquad x \in \operatorname{im}\partial_{n+1} \iff x = [\partial_{n+1}]\, y \text{ for some } y.$$

So

$$Z_n = \mathrm{Ker}\partial_n = \mathrm{Ker}([\partial_n]), \qquad B_n = \operatorname{im}\partial_{n+1} = \operatorname{im}([\partial_{n+1}]).$$

This is nothing new conceptually—it is the same boundary map we defined geometrically—but it turns "find all cycles" into "solve a homogeneous integer linear system," and "find all boundaries" into "describe the image of a linear map."

Remark 3.2. Two ideas to keep in mind:

1. changing the ordering (or flipping orientations of basis simplices) multiplies $[\partial_n]$ on the left or right by an invertible integer matrix with determinant ± 1; kernels and images change by a change of basis, so the homology groups do not depend on these choices, only their coordinate descriptions do;

2. we are working over $\mathbb{Z}$ here, so solutions are integer vectors (linear *Diophantine* systems). For the examples we meet next (like S^1 and the tetrahedral S^2), this causes no extra subtlety; later, if torsion appears, one can use Smith normal form to read off the free and torsion parts directly from these matrices.

Let us *do* the check in a tiny case to see everything explicitly.

Example 3.3. Take S^1 triangulated by three vertices v_0, v_1, v_2 and three oriented edges

$$e_0 = (v_0, v_1), \qquad e_1 = (v_1, v_2), \qquad e_2 = (v_2, v_0).$$

Use the edge basis (e_0, e_1, e_2) for $L_1 \cong \mathbb{Z}^3$ and the vertex basis (v_0, v_1, v_2) for $L_0 \cong \mathbb{Z}^3$. By the definition of ∂_1 on an oriented edge (a, b), we have $\partial_1(a, b) = b - a$, so

$$\partial_1(e_0) = v_1 - v_0, \qquad \partial_1(e_1) = v_2 - v_1, \qquad \partial_1(e_2) = v_0 - v_2.$$

In coordinates relative to (v_0, v_1, v_2) these are the column vectors

$$(-1, 1, 0)^\top, \qquad (0, -1, 1)^\top, \qquad (1, 0, -1)^\top,$$

hence

$$[\partial_1] = \begin{pmatrix} -1 & 0 & 1 \\ 1 & -1 & 0 \\ 0 & 1 & -1 \end{pmatrix}.$$

Now *solve the kernel over the integers.* Write a general 1–chain as $x = a\, e_0 + b\, e_1 + c\, e_2$; the cycle condition $[\partial_1]\,(a, b, c)^\top = 0$ gives the three equations

$$-a + c = 0, \qquad a - b = 0, \qquad b - c = 0.$$

From $a - b = 0$ and $b - c = 0$ we get $a = b = c$, and then $-a + c = 0$ is automatic. Thus every integer solution is $(a, b, c) = k(1, 1, 1)$ with $k \in \mathbb{Z}$. In other words

$$Z_1 = \mathrm{Ker}\partial_1 = \mathbb{Z} \cdot (e_0 + e_1 + e_2).$$

Since this S^1 triangulation has no 2–simplices, $B_1 = \mathrm{im}\, \partial_2 = 0$, so

$$H_1(S^1) = Z_1/B_1 \cong \mathbb{Z}.$$

(Exactly the whole loop class you expect.)

It is also instructive to see H_0 in this language. The image $\operatorname{im}\partial_1 \subseteq L_0 \cong \mathbb{Z}^3$ is generated by the columns of $[\partial_1]$, i.e.

$$v_1 - v_0, \quad v_2 - v_1, \quad v_0 - v_2.$$

Every vector $(x_0, x_1, x_2) \in \operatorname{im}\partial_1$ has coordinate sum $x_0 + x_1 + x_2 = 0$ (each column sums to 0), so $\operatorname{im}\partial_1$ lies in the subgroup $\{(x_0, x_1, x_2) \in \mathbb{Z}^3 : \ x_0 + x_1 + x_2 = 0\}$. Conversely, the three displayed generators span that subgroup over $\mathbb{Z}$ (indeed, $(1, -1, 0)$, $(0, 1, -1)$, and $(-1, 0, 1)$ generate all triples with sum 0). Hence $\operatorname{im}\partial_1$ is exactly the "sum–zero" subgroup, and the quotient

$$H_0(S^1) = L_0/\operatorname{im}\partial_1 \cong \mathbb{Z}$$

is detected by the "total sum" map $(x_0, x_1, x_2) \mapsto x_0 + x_1 + x_2$. This matches the geometric intuition: H_0 counts connected components (here, one).

The same procedure scales to higher dimensions without any new ideas: choose orientations and an ordering of the n– and $(n-1)$–simplices, fill the integer matrix $[\partial_n]$ by the $\{-1, 0, +1\}$ incidence numbers, compute $\operatorname{Ker}([\partial_n])$ and $\operatorname{im}([\partial_{n+1}])$, and form the quotient $H_n = Z_n/B_n$. Soon, when we pass from homology to cohomology, this coordinate picture becomes even more valuable: the coboundary maps will be given by the *transposes* of these same matrices, so the familiar "kernel/image" computations carry over verbatim to the cochain side; and later, in de Rham cohomology, the exterior derivative plays the role of this matrix in an infinite–dimensional setting.

Before we proceed further, it will be helpful to recall a few structural notions that will appear repeatedly when studying complexes.

Definition 3.13. A *subcomplex* of a simplicial complex K is a subset S of the simplices of K such that S is itself a simplicial complex.

Example 3.4. The *n-skeleton* of a simplicial complex K is the set of all simplices in K of dimension at most n. From the definition of a subcomplex, it follows that the n-skeleton is itself a subcomplex.

We also need a way to talk about how the pieces of a complex "hold together":

Definition 3.14. A simplicial complex K is said to be *connected* if it cannot be represented as the disjoint union of two or more non-empty subcomplexes. A geometric complex is *path-connected* if there exists a path made entirely of 1-simplices from any vertex to any other vertex.

These two notions are, in fact, equivalent for simplicial complexes:

Proposition 3.1. A simplicial complex is path-connected if and only if it is connected.

Proof. ($\Longrightarrow$) Suppose K is not connected. Then K can be expressed as the disjoint union $K = L \cup M$ of two non-empty subcomplexes L and M. Assume, for contradiction, that there exists a path between some vertex $l_0 \in L$ and $m_0 \in M$. Let l_i be the last vertex in the path that lies in L. The next vertex in the path must lie in M, so the 1-simplex joining l_i to it lies in both L and M, contradicting the assumption that $L \cap M = \varnothing$.

($\Longleftarrow$) Suppose there exist vertices l_0 and m_0 in K with no path between them. Let L be the path-connected subcomplex of K containing l_0, and M the path-connected subcomplex containing m_0. If $v_0 \in L \cap M \neq \varnothing$, then there would exist a path from l_0 to v_0 within L, and a path from v_0 to m_0 within M. Concatenating these paths yields a path from l_0 to m_0, a contradiction. Therefore $L \cap M = \varnothing$, so K is disconnected. ■

So far, we have only computed H_n for a single connected space, the circle S^1. But it is not hard to imagine what would happen if our space had more than one disconnected piece. For instance, if we took two disjoint copies of S^1, it feels natural that the H_1 group should have one $\mathbb{Z}$ summand for each circle, and the H_0 group should have one $\mathbb{Z}$ summand for each connected component.

This is no coincidence. In fact, the same principle holds in *any* dimension n: the homology of a complex is just the homology of each connected component, placed side–by–side in a direct sum.

Why should this be true? Think about chains: an n–chain in the whole space is simply a sum of chains, one from each component, and the boundary map never mixes them — the boundary of a chain in one component stays inside that component. As a result, cycles and boundaries also decompose componentwise, and the quotient Z_n/B_n inherits the same splitting. Once you see this, the statement of the theorem below is not surprising.

Note that the symbol "$\oplus$" denotes a *direct sum*: putting groups side by side and keeping track of each piece separately. For example, if A and B are groups, then $A \oplus B$ consists of all ordered pairs (a, b) with

$a \in A$ and $b \in B$, with addition done componentwise:

$$(a_1, b_1) + (a_2, b_2) = (a_1 + a_2,\, b_1 + b_2).$$

So $H_n^1 \oplus \cdots \oplus H_n^p$ just means we have one copy of H_n^i for each component K_i, and an element of the whole group is a tuple of elements, one from each copy.

Theorem 3.3. Let $K_1, \ldots, K_p$ be the connected components of a simplicial complex K, and let H_n and H_n^i denote the n–th homology groups of K and K_i, respectively. Then

$$H_n(K) \;\cong\; H_n^1 \;\oplus\; \cdots \;\oplus\; H_n^p.$$

Proof. Let L_n be the group of n–chains of K, and L_n^i the group of n–chains of K_i. Each L_n^i is a subgroup of L_n, and any n–chain in K can be written uniquely as a sum of chains from the different components:

$$L_n = L_n^1 \;\oplus\; \cdots \;\oplus\; L_n^p.$$

The boundary map ∂_n respects this decomposition, because the boundary of a chain in K_i lies entirely in K_i.

For the boundaries, define

$$B_n^i \;:=\; \partial_n(L_{n+1}^i) \;\subseteq\; L_n^i.$$

If $x \in L_{n+1}$ is written as $x = x_1 + \cdots + x_p$ with $x_i \in L_{n+1}^i$, then

$$\partial_n(x) \;=\; \partial_n(x_1) + \cdots + \partial_n(x_p),$$

where each $\partial_n(x_i) \in B_n^i$. Thus

$$B_n = B_n^1 \;\oplus\; \cdots \;\oplus\; B_n^p.$$

For the cycles, define

$$Z_n^i \;:=\; \mathrm{Ker}(\partial_n) \;\cap\; L_n^i.$$

If $x \in Z_n$ is written as $x = x_1 + \cdots + x_p$ with $x_i \in L_n^i$, then $\partial_n(x) = 0$ means

$$\partial_n(x_1) + \cdots + \partial_n(x_p) = 0.$$

But since each $\partial_n(x_i)$ lies in a different summand L^i_{n-1}, this sum can be 0 only if each $\partial_n(x_i) = 0$. Thus each $x_i \in Z^i_n$, and

$$Z_n = Z^1_n \oplus \cdots \oplus Z^p_n.$$

Since both Z_n and B_n decompose componentwise, so does their quotient:

$$H_n(K) = \frac{Z_n}{B_n} \cong \frac{Z^1_n}{B^1_n} \oplus \cdots \oplus \frac{Z^p_n}{B^p_n} = H^1_n \oplus \cdots \oplus H^p_n.$$

■

The theorem above explains how H_n behaves when a space has multiple connected components. It is worth pausing to understand in detail what H_0 looks like for a *connected* space, since this is the simplest case of the theorem and it admits a very concrete description.

Recall that H_0 is the quotient Z_0/B_0: 0–cycles modulo 0–boundaries. A 0–chain is just an integer combination of vertices, and the boundary map ∂_1 sends an oriented edge (a, b) to $b - a$. This means that 0–boundaries are precisely the combinations of vertices whose coefficients sum to zero along each connected piece.

This "sum of coefficients" is important enough to name:

Definition 3.15. If $x = \sum_{i=1}^{k} g_i\, A^0_i$ is a 0–chain (where A^0_i are vertices), its *index* is defined as

$$I(x) := \sum_{i=1}^{k} g_i.$$

The index simply adds up the integer coefficients of all the vertices in the chain. Because the boundary of an edge (a, b) is $b - a$, every boundary has index zero: $I(b - a) = 1 - 1 = 0$. This already suggests that boundaries lie inside the kernel of I. The following proposition makes the relationship precise.

Proposition 3.2. If K is a connected complex, then for a 0–chain x,

$$I(x) = 0 \iff x \sim 0,$$

and therefore $H_0(K) \cong \mathbb{Z}$.

Proof. First, suppose $x \sim 0$, so $x = \partial_1(y)$ for some 1–chain y. Write $y = \sum_{j=1}^{m} g_j\, A_j^1$, where $A_j^1 = (a_j, b_j)$ are oriented edges. Then

$$\partial_1(A_j^1) = b_j - a_j,$$

so

$$I(\partial_1(A_j^1)) = I(b_j - a_j) = 1 - 1 = 0.$$

By linearity of I, we have $I(\partial_1(y)) = 0$, hence $I(x) = 0$.

Conversely, suppose $I(x) = 0$. Let v and w be any two vertices of K. Since K is connected, there is a path

$$(a_0, a_1),\ (a_1, a_2),\ \ldots,\ (a_{q-1}, a_q),$$

with $a_0 = v$ and $a_q = w$. Consider the 1–chain

$$y = g\,(a_0, a_1) + g\,(a_1, a_2) + \cdots + g\,(a_{q-1}, a_q).$$

Its boundary is

$$\partial_1(y) = g(a_1 - a_0) + g(a_2 - a_1) + \cdots + g(a_q - a_{q-1}) = g\,w - g\,v.$$

Thus $g\,w \sim g\,v$ in H_0. By moving along paths, any vertex is homologous to any other. In particular, any 0–chain x is homologous to a multiple of a single vertex, say v, and that multiple is exactly $I(x)$. So if $I(x) = 0$, we get $x \sim 0$.

We have shown that $B_0 = \mathrm{Ker}(I)$. Since $I(Z_0) = \mathbb{Z}$ (take $g\,v$ as a cycle of index g), the First Isomorphism Theorem gives

$$H_0(K) = Z_0/B_0 \ \cong\ \mathbb{Z}.$$

■

Before stating the general result, it helps to picture what happens to 0–chains when the complex splits into several disconnected pieces. On each connected component we can measure the "total mass" of a 0–chain by summing the integer coefficients of the vertices that lie in that component. These component–wise sums are the only invariants that survive after we mod out by boundaries, and that is exactly why the number of connected components shows up in H_0.

Theorem 3.4. Let $K_1, \dots, K_p$ be the connected components of a simplicial complex K. Then the zero–dimensional homology group of K with integer coefficients is

$$H_0(K;\mathbb{Z}) \;\cong\; \underbrace{\mathbb{Z}\oplus\cdots\oplus\mathbb{Z}}_{p \text{ copies}} \;=\; \mathbb{Z}^p.$$

Proof. Label the components $K_1, \dots, K_p$. Every vertex of K lies in exactly one component, so any 0–chain

$$x = \sum_{v \text{ vertex}} g_v\, v$$

can be split uniquely as a sum $x = x_1 + \cdots + x_p$ where x_i is the part of x supported on the vertices of K_i.

For each component K_i define the *component–index* map

$$I_i : Z_0(K) \longrightarrow \mathbb{Z}, \qquad I_i\Big(\sum_v g_v\, v\Big) \;:=\; \sum_{v\in K_i} g_v,$$

that is, $I_i(x)$ adds the coefficients of the vertices lying in K_i. Putting these together gives a homomorphism

$$I = (I_1, \dots, I_p) : Z_0(K) \longrightarrow \mathbb{Z}^p.$$

First note that every boundary has zero component–index. Indeed, for an oriented edge (a, b) lying in some component, $\partial_1(a, b) = b - a$ has index $1 - 1 = 0$ on that component and 0 on every other component. Hence $B_0 \subseteq \mathrm{Ker} I$.

Conversely, suppose $x \in Z_0(K)$ satisfies $I(x) = (0, \dots, 0)$. Using connectivity of each K_i exactly as in Proposition 2.12, each part x_i (the restriction of x to K_i) is homologous in K_i to an integer multiple of any fixed vertex $v_i \in K_i$, and that multiple is precisely $I_i(x) = 0$. Thus each x_i is a boundary in K_i, so $x = x_1 + \cdots + x_p$ is a boundary in K. Therefore $\mathrm{Ker} I \subseteq B_0$, and we conclude $\mathrm{Ker} I = B_0$.

Finally, I is surjective: given any tuple $(n_1, \dots, n_p) \in \mathbb{Z}^p$, pick for each i a vertex $v_i \in K_i$ and send $(n_1, \dots, n_p)$ to the 0–chain $\sum_i n_i v_i$, which is a cycle whose image under I is exactly $(n_1, \dots, n_p)$. By the First Isomorphism Theorem for groups,

$$H_0(K) = Z_0(K)/B_0 \;\cong\; Z_0(K)/\mathrm{Ker} I \;\cong\; \mathrm{im}\, I \;=\; \mathbb{Z}^p.$$

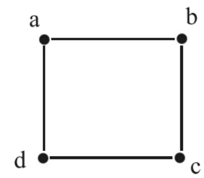

Figure 3: A simplicial structure on the circle

This completes the proof. ■

Consider another example with S^1.

Example 3.5. Take a simplicial model of the circle made of four vertices a, b, c, d and four oriented edges as in Figure 3

$$e_1 = (a, b), \quad e_2 = (b, c), \quad e_3 = (c, d), \quad e_4 = (d, a).$$

A general 0–chain is

$$x = g_1 a + g_2 b + g_3 c + g_4 d, \qquad g_i \in \mathbb{Z}.$$

Recall that the boundary of an oriented edge (u, v) is $\partial(u, v) = v - u$. We will use boundaries to move coefficients from one vertex to another, thereby showing that every 0–chain is homologous to a single integer multiple of one chosen vertex.

For example, eliminate the d-coefficient by using the boundary of the edge (d, a):

$$\partial(d, a) = a - d \quad \Longrightarrow \quad \partial\big(g_4\,(d, a)\big) = g_4 a - g_4 d.$$

Subtracting this boundary from x gives

$$x - \partial\big(g_4(d, a)\big) = (g_1 + g_4)\,a + g_2 b + g_3 c + 0 \cdot d.$$

Next, eliminate the b-coefficient by using the boundary of (a, b) (note $\partial(a, b) = b - a$):

$$\partial\big(g_2(a, b)\big) = g_2 b - g_2 a,$$

so subtracting $\partial(g_2(a,b))$ yields

$$\big(x - \partial(g_4(d,a))\big) - \partial(g_2(a,b)) = (g_1 + g_4 - g_2)\,a + 0 \cdot b + g_3 c.$$

Finally, move the c-coefficient to a along the path $c \rightsquigarrow d \rightsquigarrow a$ (or directly via (b,c) then (a,b)); after two boundary subtractions we obtain a chain supported only at a:

$$x \sim \big(g_1 + g_2 + g_3 + g_4\big)\,a.$$

(The intermediate algebra depends on the path/orientations chosen, but the final integer coefficient at a is always the sum $g_1 + g_2 + g_3 + g_4$.)

Thus every 0–cycle is homologous to a single multiple of the vertex a. The integer that appears is precisely the *index* $I(x) = \sum_i g_i$ discussed above, so different 0–cycles are distinguished in homology only by this total sum. Therefore

$$H_0(S^1) \;\cong\; \mathbb{Z},$$

with generator represented by a single vertex (or, equivalently, the class of any vertex).

3.2.1 Computations of some General Homology Groups

We now calculate some more general homology groups.

Example 3.6. Compute $H_n(S^n)$ using the boundary of an $(n+1)$-simplex. This triangulation is the smallest natural simplicial model of S^n and makes the argument purely combinatorial.

Let Δ^{n+1} be an $(n+1)$-simplex with vertex set $\{v_0, v_1, \ldots, v_{n+1}\}$. Its boundary $\partial\Delta^{n+1}$ is homeomorphic to S^n and consists of the $n+2$ oriented n-simplices obtained by omitting one vertex:

$$\sigma_i \;=\; (v_0, \ldots, \widehat{v_i}, \ldots, v_{n+1}), \qquad i = 0, \ldots, n+1.$$

Write a general n-chain as

$$x \;=\; \sum_{i=0}^{n+1} g_i\,\sigma_i, \qquad g_i \in \mathbb{Z}.$$

We ask: when is $\partial_n(x) = 0$, i.e. when is x an n-cycle?

Fix any $(n-1)$-face τ of the boundary. By construction τ is a face of exactly two of the σ's;

say τ sits in σ_k and σ_ℓ. The induced orientations of τ coming from σ_k and σ_ℓ are opposite, so the contribution of x to the coefficient of τ in $\partial_n(x)$ is (up to the global sign convention)

$$g_k - g_\ell.$$

Thus the vanishing $\partial_n(x) = 0$ forces $g_k = g_\ell$ for every pair (k, ℓ) of indices corresponding to two n-simplices sharing an $(n-1)$-face.

Now observe that the adjacency graph of the σ_i's (vertices = the σ_i's, edges = shared $(n-1)$-faces) is connected: any two n-simplices can be joined by a chain of face-adjacent simplices. The equalities $g_k = g_\ell$ therefore propagate along such chains, and we deduce

$$g_0 = g_1 = \cdots = g_{n+1} =: g.$$

Hence every n-cycle has the form

$$x \;=\; g \sum_{i=0}^{n+1} \sigma_i \;=\; g\Sigma, \qquad \Sigma := \sum_{i=0}^{n+1} \sigma_i,$$

so $Z_n(\partial\Delta^{n+1}) \cong \mathbb{Z} \cdot \Sigma$.

Finally, $\partial\Delta^{n+1}$ contains no $(n+1)$-simplices, so $\operatorname{Im} \partial_{n+1} = B_n = 0$. Therefore

$$H_n(S^n) = Z_n/B_n \cong \mathbb{Z},$$

generated by the class $[\Sigma]$, the top-dimensional (fundamental) homology class of the sphere.

Example 3.7. Compute $H_n(D^n)$. Equip the n-disk D^n with the simplest simplicial model: a single n-simplex Δ^n. This is the most economical triangulation and makes the computation immediate.

With this model the only n-simplex is Δ^n, so every n-chain has the form

$$x \;=\; g\,\Delta^n, \qquad g \in \mathbb{Z},$$

and the group $L_n \cong \mathbb{Z}$ is generated by Δ^n. There are no $(n+1)$-simplices in the complex, hence

$\operatorname{Im}\partial_{n+1} = B_n = 0$. Therefore

$$H_n(D^n) = Z_n/B_n = Z_n.$$

It remains to determine which multiples $g\Delta^n$ are cycles. Applying the boundary map gives

$$\partial_n(g\Delta^n) \ = \ g\,\partial_n(\Delta^n),$$

and $\partial_n(\Delta^n)$ is the nonzero sum of the oriented $(n-1)$-faces of Δ^n. In particular $\partial_n(\Delta^n) \neq 0$ as an $(n-1)$-chain, so $\partial_n(g\Delta^n) = 0$ forces $g = 0$. Hence $Z_n = 0$, and therefore

$$H_n(D^n) = 0.$$

Geometrically this matches the expectation: the disk is contractible and has no n-dimensional "hole," so its top-dimensional homology vanishes.

3.3 Singular Homology

Up to now, all our homology computations have lived in the clean, discrete world of simplicial complexes. We chop our space into vertices, edges, triangles, and higher-dimensional simplices, write down the boundary maps, and let the algebra tell us about the topology.

This has been working beautifully for the examples we've seen, but notice something: so far we've always had a fixed triangulation in hand. If we re–triangulate the same space in a completely different way, the chain groups change. So here's a question that ought to make you pause:

If two spaces are "really" the same topologically (say, they are homeomorphic), do they have the same homology groups? And how would we prove it?

Our geometric intuition screams "yes." A square loop and a round circle are homeomorphic, so of course their H_1 should match — after all, both have exactly one "loop hole." But proving this from the simplicial definition is not so obvious. Why? Because simplicial homology is tied to a particular triangulation: different triangulations give you different chain groups L_n, and it is not immediate how to compare them directly.

We need a definition of homology that is *intrinsic to the space itself*, not to how we chop it up. That way, a homeomorphism between two spaces will give a direct correspondence between their chains, and the rest will follow automatically.

How might we do that?

Well, in simplicial homology, an n-simplex $(v_0, \ldots, v_n)$ is just an abstract copy of the standard n-simplex Δ^n (the set of points $(t_0, \ldots, t_n)$ with $t_i \geq 0$ and $\sum t_i = 1$), together with a map telling us where in our space X the vertices go. But nothing in the definition of "simplex" requires it to be perfectly straight — what really matters is which points of Δ^n land where in X.

So why not take this to the extreme: let a simplex in X be *any* continuous map from Δ^n into X? It might be wildly bent or stretched, but that doesn't matter; topology doesn't care about straightness.

This is the key idea of *singular homology*.

Definition 3.16. A *singular n-simplex* in a topological space X is a continuous map

$$\sigma : \Delta^n \longrightarrow X.$$

Now we make chains exactly as before:

Definition 3.17. Let $C_n(X)$ be the free abelian group on the set of all singular n-simplices in X. An element of $C_n(X)$ is a *singular n-chain*, a finite sum $\sum_i g_i \, \sigma_i$ with $g_i \in \mathbb{Z}$.

And the boundary? We already know how to take the boundary of a simplex: restrict it to each of its faces and alternate the signs.

Definition 3.18. The boundary operator $\partial_n : C_n(X) \to C_{n-1}(X)$ is

$$\partial_n(\sigma) \;=\; \sum_{i=0}^{n} (-1)^i \, \sigma|[\nu_0, \ldots, \widehat{\nu_i}, \ldots, \nu_n],$$

where ν_i is the i-th vertex map $\Delta^0 \to X$ of σ, and $\widehat{\nu_i}$ means that vertex is omitted.

Because this formula is exactly the same as in the simplicial case, the familiar property $\partial_{n-1} \circ \partial_n = 0$ still holds. So we again get cycles and boundaries, and we define:

Definition 3.19. The *singular homology group* of X is

$$H_n(X) \;:=\; \operatorname{Ker}(\partial_n) \big/ \operatorname{im}(\partial_{n+1}).$$

From now on we will write H_n^Δ for the simplicial homology groups to distinguish them from the singular ones H_n.

Now here's the beauty of this new approach: If $f : X \to Y$ is a homeomorphism, then given any singular simplex $\sigma : \Delta^n \to X$, we can simply compose with f to get a singular simplex $f \circ \sigma : \Delta^n \to Y$. This extends linearly to a map $C_n(X) \to C_n(Y)$ that commutes with the boundary maps, and hence induces a map $H_n(X) \to H_n(Y)$. Doing the same with f^{-1} shows this map is an isomorphism. So *homeomorphic spaces automatically have isomorphic singular homology groups.* No triangulation-comparison needed.

One last concern: these definitions look suspiciously parallel to those for simplicial homology, so shouldn't $H_n(X)$ and $H_n^\Delta(X)$ agree when X is a nice simplicial complex? Our intuition says yes — they are both trying to measure the same "holes" — but at first glance, the chain groups are very different: $C_n(X)$ has one generator for *every* continuous map $\Delta^n \to X$, which is usually uncountable, while the simplicial chain group L_n has only finitely many generators for a finite complex. That these give the same result is not obvious at all.

And yet, for spaces that can be triangulated, it is true:

$$H_n(X) \;\cong\; H_n^\Delta(X).$$

The proof of this equivalence is beyond the scope of this paper. For now, think of singular homology as a more flexible version of our old theory: it agrees with simplicial homology when both make sense, but it applies to *any* topological space and plays nicely with homeomorphisms right out of the box.

Before we move on, let's sanity–check our new singular theory against the simplicial results we already know. We said earlier that singular homology should match simplicial homology when both are defined, so the most basic theorems ought to carry over without change. Let's test that.

Think back to the combinatorial case: if a simplicial complex split into several connected components, the homology groups simply split into one copy of each component's group. That worked because the boundary of a chain in one component could never "jump" into another. The same logic should work here, right?

Indeed, there's an even stronger restriction now: a singular simplex is a continuous map from a standard simplex Δ^n into X, and the image of Δ^n is always path–connected. So a single singular simplex can only ever live inside *one* path–connected component of X.

Proposition 3.3. Let X be a topological space with path–connected components $X_1, \dots, X_p$. Then

$$H_n(X) \;\cong\; H_n(X_1) \;\oplus\; \cdots \;\oplus\; H_n(X_p).$$

Proof. Since each singular simplex lies entirely in a single X_i, the chain group $C_n(X)$ splits as a direct sum

$$C_n(X) \;=\; C_n(X_1) \;\oplus\; \cdots \;\oplus\; C_n(X_p).$$

The boundary map ∂_n is defined simplex–by–simplex, so it preserves this splitting. Thus both $\mathrm{Ker}(\partial_n)$ (the n–cycles) and $\mathrm{im}(\partial_{n+1})$ (the n–boundaries) also split in the same way. Taking the quotient Ker/ im gives the same direct–sum decomposition for $H_n(X)$. ■

In particular, we can revisit the H_0 story from simplicial homology. There, H_0 was just one copy of $\mathbb{Z}$ for each connected component. Here the word "connected" gets replaced by "path–connected" (which for nice spaces is the same), but otherwise nothing changes.

Proposition 3.4. The zero–dimensional singular homology group of a space X is

$$H_0(X) \;\cong\; \underbrace{\mathbb{Z} \;\oplus\; \cdots \;\oplus\; \mathbb{Z}}_{\text{one for each path–component of } X}.$$

Proof. It suffices to handle the case where X is path–connected, because the general case follows from the direct–sum proposition above.

If x is a singular 0–simplex (a continuous map $\Delta^0 \to X$), its boundary vanishes automatically, so every 0–simplex is a cycle. Thus

$$H_0(X) \;=\; C_0(X) \big/ \mathrm{im}(\partial_1).$$

We recall the "index" map from the simplicial setting:

$$I : C_0(X) \longrightarrow \mathbb{Z}, \qquad I\left(\sum_i g_i \sigma_i\right) = \sum_i g_i.$$

In the simplicial case, we proved that $\mathrm{Ker}(I) = \mathrm{im}(\partial_1)$, so that $H_0(X) \cong \mathbb{Z}$. The exact same argument works here: the only change is that the 0–simplices σ_i are now singular 0–simplices rather than vertices of a triangulation. Thus, for a path–connected X, $H_0(X) \cong \mathbb{Z}$. ■

So our new theory passes these basic tests: it still breaks apart over path–components, and H_0 still counts them, just as before.

3.4 Chain Complexes, Exact Sequences, and Relative Homology Groups

Up to this point, our computations of homology have followed a concrete recipe: start with the building blocks of a space (vertices, edges, faces, etc.), form chains in each dimension, apply the boundary map, and then pass to the quotient Z_n/B_n. At each stage, we were really doing the same kind of thing — only the geometric ingredients changed. If we now take a step back, we notice that a very general algebraic pattern has been hiding in plain sight.

In every example, we have:

- an abelian group C_n consisting of all formal sums of n-dimensional pieces,
- a boundary operator $\partial_n : C_n \to C_{n-1}$ lowering dimension by one,
- and the property $\partial_n \circ \partial_{n+1} = 0$, meaning: *the boundary of a boundary is empty.*

Moreover, these groups are connected together in an infinite sequence:

$$\cdots \longrightarrow C_{n+1} \xrightarrow{\partial_{n+1}} C_n \xrightarrow{\partial_n} C_{n-1} \xrightarrow{\partial_{n-1}} \cdots \xrightarrow{\partial_1} C_0 \longrightarrow 0.$$

Each arrow tells you "how to pass from an object in one dimension to its boundary in the next lower dimension." The final 0 at the end is the trivial group — there is nothing below dimension 0.

It is worth giving this common structure its own name.

Definition 3.20. A *chain complex* is a sequence of abelian groups

$$\cdots \longrightarrow C_{n+1} \xrightarrow{\partial_{n+1}} C_n \xrightarrow{\partial_n} C_{n-1} \longrightarrow \cdots$$

such that $\partial_n \circ \partial_{n+1} = 0$ for all n. The maps ∂_n are called *boundary operators*.

From this abstract point of view, a chain complex is a machine that:

1. holds a group of objects in each dimension,
2. has a process for lowering dimension,

3. and the property that the composition of two consecutive boundary maps is the zero map:

$$\partial_n \circ \partial_{n+1} = 0 \quad \text{for all } n.$$

In other words, if you start with any $(n+1)$-chain, take its boundary to get an n-chain, and then take the boundary again, you always obtain the zero element of C_{n-1}.

Our familiar singular chain groups $(C_n(X), \partial_n)$ form such a complex:

$$\cdots \longrightarrow C_{n+1}(X) \xrightarrow{\partial_{n+1}} C_n(X) \xrightarrow{\partial_n} C_{n-1}(X) \longrightarrow \cdots \longrightarrow C_0(X) \longrightarrow 0.$$

The homology groups are then just:

$$H_n(C_*) \;=\; \operatorname{Ker}(\partial_n)\big/\operatorname{im}(\partial_{n+1}),$$

exactly as before — only now, the definition applies to *any* chain complex, not just those coming from geometry.

Why is this abstraction useful? Because it lets us compare completely different constructions of homology. For example, if one may wish to compare simplicial homology H_n^{Δ} and singular homology H_n. Even though their chain groups look nothing alike, both form chain complexes, and we can ask whether there is a *map of chain complexes* between them.

Definition 3.21. Let (A_*, ∂_A) and (B_*, ∂_B) be chain complexes. A *chain map* $f : A_* \to B_*$ is a sequence of homomorphisms

$$f_n : A_n \longrightarrow B_n \quad \text{for each } n,$$

such that for all n,

$$\partial_B \circ f_n \;=\; f_{n-1} \circ \partial_A.$$

The condition $\partial_B f_n = f_{n-1}\partial_A$ says that the two possible paths

“first apply f, then take the boundary” and “first take the boundary, then apply f”

always give the same result. In diagram form:

$$\begin{array}{cccccc} \cdots & \to & A_{n+1} & \xrightarrow{\partial_A} & A_n & \to \cdots \\ & & \downarrow f_{n+1} & & \downarrow f_n & \\ \cdots & \to & B_{n+1} & \xrightarrow{\partial_B} & B_n & \to \cdots \end{array}$$

Why insist on this commutativity? Because it guarantees that f sends cycles to cycles and boundaries to boundaries — precisely the ingredients that define homology.

Theorem 3.5. A chain map $f : A_* \to B_*$ induces well-defined homomorphisms

$$f_* : H_n(A_*) \longrightarrow H_n(B_*)$$

for all n.

Proof. If $x \in A_n$ is a cycle ($\partial_A x = 0$), then

$$\partial_B f_n(x) = f_{n-1}\partial_A(x) = f_{n-1}(0) = 0,$$

so $f_n(x)$ is a cycle in B_n. If x is a boundary, $x = \partial_A y$, then

$$f_n(x) = f_n(\partial_A y) = \partial_B f_{n+1}(y),$$

so $f_n(x)$ is a boundary in B_n. Thus f_n sends cycles to cycles and boundaries to boundaries, and therefore descends to a homomorphism

$$f_* : H_n(A_*) \to H_n(B_*).$$

■

We now return to the case of singular homology, and apply Theorem 3.5 to actual topological spaces. Let X and Y be topological spaces, and let $f : X \to Y$ be any continuous map. Our goal is to understand how f gives rise to a map between their singular chain complexes.

Recall that an n-chain in $C_n(X)$ is a finite $\mathbb{Z}$-linear combination of singular n-simplices $\sigma : \Delta^n \to X$.

Given such a simplex, there is a natural way to make it "live" in Y: simply follow σ with f to obtain

$$f_a(\sigma) \;:=\; f \circ \sigma : \Delta^n \longrightarrow Y.$$

In words: map the standard simplex into X, and then use f to send it on to Y.

We extend this rule linearly to all of $C_n(X)$:

$$f_a\left(\sum_i g_i\,\sigma_i\right) \;:=\; \sum_i g_i\,(f \circ \sigma_i).$$

This defines a homomorphism $f_a : C_n(X) \to C_n(Y)$ in each dimension n.

Why does this matter? Because the family $(f_a)_n$ forms a *chain map*. Indeed, if we take the boundary of a simplex in X and then apply f, we get the same result as applying f first and then taking the boundary in Y. This is simply the statement:

$$\partial\, f_a = f_a\,\partial$$

in every degree, which follows directly from the definition of the boundary operator.

In diagram form:

$$\begin{array}{ccccccccc} \cdots & \to & C_{n+1}(X) & \xrightarrow{\partial} & C_n(X) & \xrightarrow{\partial} & C_{n-1}(X) & \to \cdots \\ & & \downarrow f_a & & \downarrow f_a & & \downarrow f_a & \\ \cdots & \to & C_{n+1}(Y) & \xrightarrow{\partial} & C_n(Y) & \xrightarrow{\partial} & C_{n-1}(Y) & \to \cdots \end{array}$$

By Theorem 3.5, any such chain map induces homomorphisms on homology:

$$f_* : H_n(X) \longrightarrow H_n(Y) \quad \text{for all } n.$$

Intuitively, f_* is the "shadow" that f casts on the homology level — it tells us how f transforms cycles and boundaries when we view them up to homology.

A particularly simple case occurs when f is a homeomorphism. Then f has a continuous inverse, which also induces maps between the chain complexes in each direction. These two induced maps are inverses of each other on homology, so f_* is an isomorphism. Thus, homeomorphic spaces have isomorphic singular homology groups — a fact that aligns perfectly with our geometric intuition that homeomorphic spaces are "the same" from the viewpoint of topology. Before moving on, it is worth introducing a slightly richer notion than a chain complex, one that will later become indispensable for connecting different homology theories: the *exact sequence*.

Recall that in a chain complex, the key property was that the image of one map lies inside the kernel of the next. An exact sequence is simply the case where this containment is as tight as possible — in fact, the image *equals* the kernel at every step.

Definition 3.22. A sequence of abelian groups and homomorphisms

$$\cdots \xrightarrow{\alpha_{n+1}} A_n \xrightarrow{\alpha_n} A_{n-1} \xrightarrow{\alpha_{n-1}} \cdots$$

is called *exact* if

$$\mathrm{Ker}(\alpha_n) = \mathrm{Im}(\alpha_{n+1}) \quad \text{for all } n.$$

This condition means two things at once:

1. Because $\mathrm{Im}(\alpha_{n+1}) \subset \mathrm{Ker}(\alpha_n)$ always holds for any sequence of maps, exactness forces the inclusion to be an equality. Thus, exact sequences are automatically chain complexes with the *strongest possible* relation between consecutive terms: every element that dies under α_n must have come from the previous group.

2. If the sequence is exact, then the homology groups $\mathrm{Ker}(\alpha_n)/\mathrm{Im}(\alpha_{n+1})$ vanish everywhere, because the numerator and denominator are the same. In other words, exactness means "no homology survives" — nothing is a cycle without already being a boundary.

Exactness can encode familiar algebraic properties:

- $0 \to A \xrightarrow{a} B$ is exact $\iff$ a is injective (its kernel is zero).
- $A \xrightarrow{a} B \to 0$ is exact $\iff$ a is surjective (its image is all of B).
- $0 \to A \xrightarrow{a} B \to 0$ is exact $\iff$ a is an isomorphism.

A particularly important case is the *short exact sequence*:

$$0 \longrightarrow A \xrightarrow{a} B \xrightarrow{b} C \longrightarrow 0,$$

which is exact if and only if a is injective, b is surjective, and $\mathrm{Ker}(b) = \mathrm{Im}(a)$. In this case we can think of C as the quotient B/A.

Why introduce this now? Because exact sequences are the natural language for relating the homology of different spaces — especially when one space sits inside another.

3.4.1 Relative Homology

Suppose $A \subset X$. Chains in $C_n(A)$ can be thought of as chains in X that "live entirely inside" A. We define the *relative chain group*

$$C_n(X, A) := C_n(X)/C_n(A),$$

where we identify all chains in A with the zero element. Geometrically, we are treating A as if it has been "shrunk to a point" inside X when measuring chains.

The boundary map $\partial : C_n(X) \to C_{n-1}(X)$ passes naturally to the quotient, giving us

$$\partial : C_n(X, A) \to C_{n-1}(X, A),$$

with the same property $\partial \circ \partial = 0$ as before. Thus, $(C_n(X, A), \partial)$ is itself a chain complex, and we can define its homology groups:

$$H_n(X, A) := \mathrm{Ker}(\partial_n)/\mathrm{Im}(\partial_{n+1}),$$

called the *relative homology groups* of the pair (X, A).

These groups capture cycles in X *up to* chains in A: two cycles in X represent the same relative homology class if their difference bounds a chain in X whose boundary lies entirely in A.

We now want to understand how the relative homology groups $H_n(X, A)$ are related to the absolute homology groups $H_n(X)$ and $H_n(A)$. Intuitively, $H_n(X, A)$ measures n–dimensional cycles in X that are allowed to intersect A, but where anything lying entirely inside A is regarded as trivial. If a cycle lives completely in A, it should be detected by $H_n(A)$ and disappear in the relative group. If it does not, it should represent something genuinely new in $H_n(X, A)$. From this perspective, it is natural to expect that there is some systematic way to pass between:

$$H_n(A), \quad H_n(X), \quad \text{and} \quad H_n(X, A)$$

in all dimensions. The remarkable fact is that this can be done in one elegant construction: the *long exact sequence of a pair*. It packages these three kinds of homology groups into a single chain of maps where the image of one map is exactly the kernel of the next.

To see where this comes from, we begin at the chain level. There is a natural short exact sequence

$$0 \longrightarrow C_n(A) \xrightarrow{i} C_n(X) \xrightarrow{j} C_n(X,A) \longrightarrow 0,$$

where i is the inclusion map (chains in A are chains in X) and j is the quotient map (we collapse all chains in A to zero). Because i and j commute with the boundary operator ∂, each of these groups fits into a diagram where the rows are chain complexes and the columns are exact:

$$\begin{array}{ccccccccc}
 & & \vdots & & \vdots & & \vdots & & \\
\cdots & \xrightarrow{\partial} & C_{n+1}(A) & \xrightarrow{i} & C_{n+1}(X) & \xrightarrow{j} & C_{n+1}(X,A) & \xrightarrow{\partial} & \cdots \\
 & & \downarrow\partial & & \downarrow\partial & & \downarrow\partial & & \\
\cdots & \xrightarrow{\partial} & C_{n}(A) & \xrightarrow{i} & C_{n}(X) & \xrightarrow{j} & C_{n}(X,A) & \xrightarrow{\partial} & \cdots \\
 & & \downarrow\partial & & \downarrow\partial & & \downarrow\partial & & \\
\cdots & \xrightarrow{\partial} & C_{n-1}(A) & \xrightarrow{i} & C_{n-1}(X) & \xrightarrow{j} & C_{n-1}(X,A) & \xrightarrow{\partial} & \cdots \\
 & & \vdots & & \vdots & & \vdots & &
\end{array}$$

The vertical sequences are short exact, the rows are chain complexes, and everything commutes.

From this purely algebraic setup, there is a standard way to extract a *connecting homomorphism* on homology, and the result is the long exact sequence:

$$\cdots \longrightarrow H_n(A) \xrightarrow{i_*} H_n(X) \xrightarrow{j_*} H_n(X,A) \xrightarrow{\delta} H_{n-1}(A) \longrightarrow \cdots \longrightarrow H_0(X,A) \longrightarrow 0.$$

The meaning of each map is straightforward when you think in terms of cycles and boundaries. The map i_* simply regards a cycle in A as a cycle in X. The map j_* takes a cycle in X and views it in the relative group, killing any part of it that lies entirely in A. The connecting homomorphism δ takes a relative cycle in $H_n(X,A)$—which by definition has its boundary inside A—and sends it to the class of that boundary in $H_{n-1}(A)$.

Exactness means that at each stage, the cycles that become trivial in the next group are exactly the ones that come from the previous group. For example, a cycle in X maps to zero in $H_n(X,A)$ precisely when it can be represented by a cycle lying entirely in A, i.e. when it comes from $H_n(A)$. Similarly, a relative cycle maps to zero in $H_{n-1}(A)$ exactly when it is the image under j_* of some absolute class in $H_n(X)$.

In this way, the long exact sequence of the pair (X, A) acts as a bridge, allowing us to move between the homology of A, the homology of X, and the relative homology that measures what is in X but not in A.

Up to this point, we have introduced the long exact sequence in homology for a pair (X, A), and we have defined the connecting homomorphism

$$\delta : H_n(C) \longrightarrow H_{n-1}(A)$$

by tracing how a cycle in C can be "lifted" back into B and then measuring its boundary inside A. Before we proceed to the long exact sequence itself, it is worth checking that this map really behaves well algebraically.

Intuitively, we expect δ to be a group homomorphism: adding two homology classes in $H_n(C)$ should produce a class whose connecting image in $H_{n-1}(A)$ is just the sum of the images of the two original classes. The construction of δ is quite geometric — it takes a cycle in C, picks a preimage in B, then takes its boundary, which (by exactness of the short sequence) lands in A. Since all these steps are linear at the chain level, the map should be additive. Let us record this formally.

Proposition 3.5. The map $\delta : H_n(C) \to H_{n-1}(A)$ is a homomorphism.

Proof. Take two classes $[c_1], [c_2] \in H_n(C)$. By definition of the connecting homomorphism δ, we choose $b_1, b_2 \in B_n$ with $j(b_i) = c_i$, and then we set

$$\delta[c_i] = [a_i] \quad \text{where} \quad i(a_i) = \partial b_i.$$

Now $[c_1] + [c_2]$ is represented by $j(b_1 + b_2)$. Applying δ gives $[a_1 + a_2]$ because

$$i(a_1 + a_2) = \partial b_1 + \partial b_2 = \partial(b_1 + b_2).$$

Thus $\delta([c_1] + [c_2]) = \delta[c_1] + \delta[c_2]$, and hence δ is a homomorphism.

■

Having established that δ is a homomorphism, we are now ready to bring together the maps i_*, j_*, and δ into a single structure — the long exact sequence in homology. Our goal is to verify the exactness at every stage. This will require showing, for example, that $\mathrm{Im}(i_*) = \ker(j_*)$, that $\mathrm{Im}(j_*) = \ker(\delta)$, and so on. The idea is that each piece of information about cycles and boundaries in A, B, and C fits perfectly with the short exact sequence of chain complexes we started with.

Proposition 3.6. The sequence

$$\cdots \xrightarrow{i_*} H_n(B) \xrightarrow{j_*} H_n(C) \xrightarrow{\delta} H_{n-1}(A) \xrightarrow{i_*} H_{n-1}(B) \xrightarrow{j_*} \cdots$$

is exact.

Proof. We proceed one step at a time, checking the definition of exactness at each position.

1. $\mathrm{Im}(i_*) \subset \ker(j_*)$*:* If $[a] \in H_n(A)$, then $i_*[a]$ is represented by $i(a) \in B_n$. Since $i(a)$ comes entirely from A, it vanishes in C when we apply j, so $j_* i_*[a] = 0$.

2. $\ker(j_*) \subset \mathrm{Im}(i_*)$*:* If $[b] \in H_n(B)$ with $j_*[b] = 0$, then $j(b)$ is a boundary in C, say $j(b) = \partial c$ for some $c \in C_{n+1}$. Surjectivity of j at the chain level gives $c = j(b')$ for some $b' \in B_{n+1}$. Then $b - \partial b'$ lies in $\mathrm{Im}(i)$, say $b - \partial b' = i(a)$ with a a cycle. Thus $[b] = i_*[a]$.

3. $\mathrm{Im}(j_*) \subset \ker(\delta)$*:* If $[c] = j_*[b]$ with b a cycle, then $\delta[c]$ is defined by $i(a) = \partial b$. But $\partial b = 0$ because b is a cycle, so a represents the zero class in $H_{n-1}(A)$.

4. $\ker(\delta) \subset \mathrm{Im}(j_*)$*:* If $\delta[c] = 0$, then c has a preimage b in B whose boundary ∂b lies in A and is a boundary there. We can adjust b by subtracting something from $i(A)$ to make it a cycle, showing $[c]$ comes from $j_*[b]$.

5. $\mathrm{Im}(\delta) \subset \ker(i_*)$*:* If $[a] = \delta[c]$, then $i_*[a]$ is represented by $i(a) = \partial b$ for some $b \in B_n$, hence is trivial in $H_{n-1}(B)$.

6. $\ker(i_*) \subset \mathrm{Im}(\delta)$*:* If $i_*[a] = 0$, then $i(a) = \partial b$ for some $b \in B_n$. Then $j(b)$ is a cycle in C whose connecting image under δ is exactly $[a]$.

All inclusions match up, and the sequence is exact.

■

Finally, we apply this to the case we care about: the pair (X, A). Here $B = X$, C becomes the relative complex X/A, and the same reasoning gives the standard long exact sequence in relative homology.

Proposition 3.7. For any pair (X, A) of topological spaces, there is a long exact sequence in singular homology

$$\cdots \xrightarrow{i_*} H_n(X) \xrightarrow{j_*} H_n(X, A) \xrightarrow{\delta} H_{n-1}(A) \xrightarrow{i_*} H_{n-1}(X) \to \cdots \to H_0(X, A) \to 0.$$

Proof. This is the previous proposition with the identifications $B = X$ and $C = X/A$, together with the definition of relative homology as the homology of the quotient chain complex $C_*(X, A) = C_*(X)/C_*(A)$. All the exactness arguments carry over unchanged. ■

At this point, we have assembled most of the major tools for working with homology: the simplicial and singular theories, the way homology behaves on disconnected spaces, the relationship between absolute and relative homology, and the long exact sequence that ties these ideas together. Before we leave the subject, it is worth recording one final, extremely useful result. It describes a situation in which we can "cut out" a portion of our space without affecting the homology of the pair we are studying.

Theorem 3.6 (Excision)**.** Let $Y \subset A \subset X$ be spaces such that the closure of Y is contained in the interior of A. Then the inclusion of pairs

$$(X - Y,\, A - Y) \hookrightarrow (X,\, A)$$

induces isomorphisms

$$H_n(X - Y,\, A - Y) \xrightarrow{\cong} H_n(X,\, A)$$

for all n.

The statement itself is appealingly simple: if a subset Y sits "deep inside" A, we can remove Y from both X and A without altering the relative homology. Intuitively, this makes sense — from the perspective of relative cycles and boundaries, the interior portion Y is invisible once A is being quotiented out. However, turning this intuition into a rigorous proof is far from trivial; it requires delicate manipulations of chain complexes and subdivision arguments. Because of its technical nature, we will not attempt the proof here, but will instead take it as a powerful tool that can be applied in later contexts.

Note that the homologies discussed are equivalent.

Theorem 3.7. For all n, the homomorphisms $H_n^{\Delta}(X) \longrightarrow H_n(X)$ are isomorphisms. Thus the singular and simplicial homology groups are equivalent.

The proof is out of scope for this paper.

3.5 Relation to Homotopy

Up to this point, our study of homology has been purely in terms of cycles and boundaries: chains of simplices, their boundaries, and the resulting groups $H_n(X)$. However, there is another classical invariant of a topological space, introduced long before homology: the *fundamental group*.

Let us begin informally. Fix a point x_0 in a space X. A *path* from x_0 to another point x_1 is just a continuous map

$$\gamma : [0,1] \longrightarrow X$$

such that $\gamma(0) = x_0$ and $\gamma(1) = x_1$. If $x_1 = x_0$, we call γ a *loop based at* x_0. Two loops γ and γ' are said to be *homotopic* (as based loops) if there is a continuous family of loops deforming γ into γ', while keeping the basepoint fixed at all times. We write $[\gamma]$ for the homotopy class of γ.

The set of all based loop classes at x_0 can be given a natural group structure: we can *concatenate* two loops γ_1 and γ_2 by first traversing γ_1 and then γ_2. The inverse of a loop is the same path traversed in the opposite direction. With these operations, the set of based loop classes becomes a group, denoted $\pi_1(X, x_0)$ and called the *fundamental group* of X at x_0.

It is important to note that $\pi_1(X, x_0)$ need not be abelian: the order in which we traverse loops can matter. This is already visible in spaces where loops can "link" in a nontrivial way, such as in a figure-eight space.

Now recall what we have learned about $H_1(X)$: it measures 1–dimensional holes by taking 1–cycles (formal sums of edges whose boundaries vanish) and identifying those that differ by a boundary. A loop in X is certainly a 1–cycle when we regard it as a singular 1–simplex (or a sum of such simplices) with no boundary. It is then natural to ask:

Can every element of $H_1(X)$ be represented by a loop? And how does the group structure in $\pi_1(X, x_0)$ compare with the abelian group $H_1(X)$?

Let us make the connection precise. Given a based loop $\gamma : S^1 \to X$, we can triangulate S^1 with a finite number of edges and regard γ as a singular 1–cycle in X. If γ and γ' are based–homotopic, the homotopy between them can be seen as a singular 2–chain whose boundary is $\gamma - \gamma'$. Thus, homotopic loops define the same homology class in $H_1(X)$. In other words, we have a well–defined map

$$h : \pi_1(X, x_0) \longrightarrow H_1(X; \mathbb{Z}),$$

called the *Hurewicz map* in dimension 1.

One immediate observation is that any *commutator* loop

$$\gamma_1 \cdot \gamma_2 \cdot \gamma_1^{-1} \cdot \gamma_2^{-1}$$

bounds a 2–chain: we can picture a square whose horizontal edges trace γ_1 and vertical edges trace γ_2. The boundary of this square is exactly the commutator loop. Since boundaries vanish in homology, h sends every commutator to 0. This means h factors through the *abelianization* of $\pi_1(X, x_0)$:

$$\pi_1(X, x_0)^{\mathrm{ab}} \;:=\; \pi_1(X, x_0) \;/\; [\pi_1(X, x_0), \pi_1(X, x_0)],$$

where $[\pi_1, \pi_1]$ denotes the subgroup generated by commutators.

We are thus led to the following important fact:

Proposition 3.8. If X is path–connected, the first homology group is the abelianization of the fundamental group:

$$H_1(X; \mathbb{Z}) \;\cong\; \pi_1(X, x_0)^{\mathrm{ab}}.$$

Proof. We have already argued that h factors through $\pi_1(X, x_0)^{\mathrm{ab}}$, so it suffices to show that the induced map

$$\bar{h} : \pi_1(X, x_0)^{\mathrm{ab}} \longrightarrow H_1(X)$$

is an isomorphism.

Surjectivity: Let $[z] \in H_1(X)$ be represented by a singular 1–cycle $z = \sum_i g_i \sigma_i$, where $\sigma_i : \Delta^1 \to X$. Since X is path–connected, we can choose paths from x_0 to each $\sigma_i(0)$ and from each $\sigma_i(1)$ back to x_0. By attaching these "whiskers" to σ_i, we obtain based loops γ_i whose homology classes agree with those of the σ_i. The sum $\sum_i g_i[\gamma_i]$ in $\pi_1(X)^{\mathrm{ab}}$ maps to $[z]$ under $\bar{h}$.

Injectivity: Suppose a loop γ maps to 0 in $H_1(X)$. Then the 1–cycle corresponding to γ is a boundary, say ∂c for some 2–chain c. We can realize c as a formal sum of 2–simplices in X, and by tracking their edges, one sees that γ is homotopic (as a based loop) to a product of commutators. Thus its class in the abelianization is trivial. ■

This tells us that H_1 is, in a very literal sense, "the abelian part" of the fundamental group: it records how loops interact up to deformation, but ignores any non–abelian ordering information.

The idea of relating homotopy groups to homology groups extends beyond π_1. In higher dimensions, we can consider the set $\pi_n(X, x_0)$ of based maps from the n–sphere S^n into X, up to based homotopy. Just as before, such a map can be interpreted as an n–cycle, giving a *Hurewicz homomorphism*

$$h : \pi_n(X, x_0) \longrightarrow H_n(X).$$

In general, h need not be an isomorphism, but if X is k–connected (i.e., all $\pi_i(X) = 0$ for $1 \leq i \leq k$), then h is an isomorphism in degree $k+1$ and a surjection in degree $k+2$. Thus for simply connected spaces, the first nontrivial homotopy group is detected exactly by homology.

Theorem 3.8 (Whitehead, simple form)**.** Let $f : X \to Y$ be a map between simply connected simplicial complexes. If $f_* : H_n(X) \to H_n(Y)$ is an isomorphism for all n, then f is a homotopy equivalence.

Heuristic: When π_1 vanishes, the higher homotopy groups behave "linearly" enough that they can be reconstructed from homology. If f induces isomorphisms on all homology groups, one can build a homotopy inverse step–by–step over the skeleta of X and Y, with no obstructions arising.

Remark 3.3. Two simply connected spaces with isomorphic homology groups need not be homotopy equivalent — the theorem requires an actual map inducing those isomorphisms.

With this, we conclude our exploration of homology. We began with the concrete picture of simplicial chains, generalized to singular homology to handle arbitrary spaces, and then examined how these groups behave under maps, decompositions, and relative constructions and finally related it to homotopy classes. Along the way, we have developed enough machinery to compute homology in many settings and to understand its structural properties. The stage is now set for us to turn to cohomology and its differential form incarnation, de Rham cohomology.

4 Cohomology Theory

In our study of homology, we have always worked with *chains*: formal sums of simplices which we could "walk along" in our space, measure their boundaries, and detect holes. This viewpoint is geometric and concrete — a chain is something you can draw. But there is another, equally important perspective: instead of building objects in the space, we can assign *numbers* to them.

More precisely, imagine we want to attach an integer (or some other coefficient) to each chain in a consistent, linear way. If two chains are added together, the numbers we assign should add as well. Such an assignment is called a *cochain*, and it is in many ways the "dual" of a chain. Where a chain is an actual geometric object in the space, a cochain is a *measurement* of it.

This dual viewpoint turns out to be more than just a curiosity. Cohomology carries additional algebraic structure — we will later see a multiplication (the cup product) which allows us to combine classes in ways homology cannot. Cohomology also interacts more naturally with maps between spaces: while homology pushes chains forward along maps (covariance), cohomology *pulls back* measurements along maps (contravariance). In practice, this means cohomology often fits better with constructions from analysis and geometry.

Perhaps most intriguingly, cohomology will provide the natural setting for *differential forms* and the *generalised Stokes theorem*. The abstract coboundary operator we will define shortly will, in that context, become the familiar exterior derivative d from calculus, and evaluating a cohomology class on a homology class will become the act of integrating a differential form over a cycle.

Thus, cohomology can be viewed as the algebraic language which prepares the ground for the de Rham theory to come. In the remainder of this section, we will first define cochains, then construct the coboundary operator, and finally build the cohomology groups themselves, paralleling the constructions for homology.

4.1 Cochains, Cocycles and Coboundaries

We now make the dual picture precise, working in the setting of *singular* cohomology. As before, let X be a topological space, and let $C_n(X)$ denote the free abelian group generated by all singular n–simplices $\sigma : \Delta^n \to X$. Recall that an element of $C_n(X)$ is a finite formal sum

$$c = \sum_{i=1}^{k} g_i \, \sigma_i,$$

with coefficients $g_i \in \mathbb{Z}$.

From the cohomological point of view, instead of *building* such chains, we wish to *measure* them. That is, we want to assign an integer to each n–chain, in a way that is compatible with addition.

Definition 4.1. A *singular n–cochain* on X is a group homomorphism

$$\varphi : C_n(X) \longrightarrow \mathbb{Z}.$$

The set of all singular n–cochains is denoted

$$C^n(X) := \mathrm{Hom}_{\mathbb{Z}}\big(C_n(X), \mathbb{Z}\big),$$

and is called the *n–th cochain group* of X (with integer coefficients).

This definition may look abstract at first, but the idea is simple: since $C_n(X)$ is generated by singular simplices $\sigma : \Delta^n \to X$, a cochain is completely determined by its values on these simplices. For example, a 0–cochain assigns an integer to each vertex of a singular 0–simplex, and a 1–cochain assigns an integer to each singular 1–simplex (edge), in a way that extends linearly to all 1–chains.

Example 4.1. If X is a graph, then:

- A 0–cochain assigns an integer to each vertex.
- A 1–cochain assigns an integer to each oriented edge (with sign reversed when the edge is reversed).

In both cases, the value of the cochain on a sum of simplices is the sum of its values.

We now need a way to *differentiate* cochains, just as we took boundaries of chains. Given that a cochain is a function on chains, the natural way to do this is to *pre-compose* with the boundary operator on chains.

Definition 4.2. The *coboundary operator* is the homomorphism

$$\delta : C^n(X) \longrightarrow C^{n+1}(X)$$

defined by

$$(\delta\varphi)(c) := \varphi\big(\partial c\big),$$

for every $\varphi \in C^n(X)$ and $c \in C_{n+1}(X)$.

In words: to evaluate $\delta\varphi$ on an $(n+1)$–chain c, we first take the boundary ∂c (which is an n–chain), and then apply φ to it. This is exactly the dual construction to the boundary in homology.

Remark 4.1. The notation δ here is traditional for the cohomological coboundary. It plays the same algebraic role as ∂ in homology, but "points" in the opposite direction:

$$\cdots \xrightarrow{\delta} C^{n-1}(X) \xrightarrow{\delta} C^n(X) \xrightarrow{\delta} C^{n+1}(X) \xrightarrow{\delta} \cdots$$

We immediately have the cohomological analogue of the fundamental fact $\partial^2 = 0$:

Proposition 4.1. For every n, we have $\delta \circ \delta = 0$.

Proof. Let $\varphi \in C^n(X)$ and $c \in C_{n+2}(X)$. Then

$$(\delta\delta\varphi)(c) = (\delta\varphi)(\partial c) = \varphi(\partial\partial c) = \varphi(0) = 0,$$

since $\partial \circ \partial = 0$ in the chain complex $C_\bullet(X)$. Thus $\delta^2 = 0$. ■

This allows us to define the cohomological analogues of cycles and boundaries.

Definition 4.3. An n–cochain φ is called:

- a *cocycle* if $\delta\varphi = 0$ (it lies in $\ker\delta$),
- a *coboundary* if there exists $\psi \in C^{n-1}(X)$ such that $\varphi = \delta\psi$ (it lies in $\operatorname{im}\delta$).

The group of n–cocycles is denoted $Z^n(X) := \ker\big(\delta : C^n \to C^{n+1}\big)$, and the group of n–coboundaries is denoted $B^n(X) := \operatorname{im}\big(\delta : C^{n-1} \to C^n\big)$.

Just as in homology, $B^n(X) \subset Z^n(X)$, thanks to $\delta^2 = 0$. This inclusion lets us form the quotient:

Definition 4.4. The *n–th singular cohomology group* of X (with integer coefficients) is

$$H^n(X) := \frac{Z^n(X)}{B^n(X)} = \frac{\ker\big(\delta : C^n \to C^{n+1}\big)}{\operatorname{im}\big(\delta : C^{n-1} \to C^n\big)}.$$

Intuitively:

- Cocycles are "measurements" on n–chains that vanish on all boundaries.
- Coboundaries are measurements that *come from* measuring $(n-1)$–chains and extending via δ.
- Cohomology measures the cocycles *modulo* those that are "trivial" in the sense of being coboundaries.

Example 4.2. Let us examine $H^1(S^1)$. Recall that $C_1(S^1)$ is the free abelian group generated by all singular 1–simplices $\sigma : \Delta^1 \to S^1$. A 1–cochain $\varphi \in C^1(S^1)$ assigns an integer to each such σ, and is determined entirely by these values (linearly extended to sums).

The condition for φ to be a *cocycle* is that $\delta\varphi = 0$. Unwinding the definition, this means:

$$(\delta\varphi)(\tau) = \varphi(\partial\tau) = 0$$

for every 2–simplex $\tau : \Delta^2 \to S^1$. Geometrically, this says: *the sum of the values of φ around the oriented edges of any singular triangle in S^1 must be zero.* This is the cohomological analogue of "the sum of oriented edge lengths around a triangle is zero".

But here is the key observation: S^1 has no non-degenerate singular 2–simplices that "wrap" in a nontrivial way. Any continuous map from a 2–simplex into the circle must collapse its interior to a 1-dimensional subset. Thus $\partial\tau$ is always degenerate, and the cocycle condition becomes *vacuous*. Therefore:

$$Z^1(S^1) = C^1(S^1),$$

meaning every 1–cochain is automatically a cocycle.

Next, let us identify the coboundaries. By definition, a 1–coboundary is something of the form δf for some 0–cochain $f \in C^0(S^1)$. A 0–cochain assigns an integer to each singular 0–simplex $\sigma : \Delta^0 \to S^1$, which is the same as assigning an integer to each point of S^1. Given such an f, the 1–cochain δf is defined by:

$$(\delta f)(\sigma) = f(\partial\sigma),$$

where σ is now a singular 1–simplex (an oriented edge). The boundary $\partial\sigma$ is the formal difference of its two endpoints:

$$\partial\sigma = \sigma(v_1) - \sigma(v_0),$$

so

$$(\delta f)(\sigma) = f\big(\sigma(v_1)\big) - f\big(\sigma(v_0)\big).$$

Thus coboundaries are exactly those 1–cochains that measure the *difference* of a function f along each edge.

From this perspective, $B^1(S^1)$ consists of "gradient-like" 1–cochains: they record how a vertex–value function changes from start to end of an edge. If you imagine walking around the circle and summing such differences, the total will always be zero: there is no net "accumulation" over a closed loop.

The quotient $H^1(S^1) = Z^1(S^1)/B^1(S^1)$ therefore measures which 1–cochains are *not* gradients of 0–cochains. On S^1, there is exactly one independent way to fail to be a gradient: assign a constant nonzero value to every positively oriented edge around the loop. Such a 1–cochain cannot be written as differences of vertex–values without creating a "jump" somewhere, which is impossible for a continuous assignment on S^1.

Algebraically, this "loop–measuring" cochain represents the generator of $H^1(S^1) \cong \mathbb{Z}$. Its integer value counts how many times you wind around the circle, much like the winding number in complex analysis.

Remark 4.2. This example foreshadows the deep connection between $H^1(S^1)$ and the fundamental group $\pi_1(S^1)$. The generator of H^1 detects the essential 1–dimensional hole in the circle, in perfect agreement with the fact that $\pi_1(S^1) \cong \mathbb{Z}$.

4.2 Functoriality

Up to now, we have thought of cochains as "measurement rules" for our space: given a singular simplex, a cochain assigns an integer (or other coefficient) to it, and this assignment is linear in the same way homology chains were. But so far, all our cochains have lived on a *fixed* space.

What if we have two spaces X and Y and a continuous map $f : X \to Y$? In homology, we already know the answer: the geometry of a simplex in X can be pushed forward along f to produce a simplex in Y, and from there we can compute its boundary or its homology class. This was the "covariant" nature of homology: maps between spaces induce maps in the *same* direction between homology groups.

Cohomology, however, lives on the "dual" side of the picture. Our objects are not geometric simplices, but rules for measuring them. If we have a cochain ψ on Y, it already knows how to evaluate any singular simplex in Y. But suppose we want to use ψ to measure a simplex σ in X. The only way to do this sensibly is to *first* transport σ to Y via f, and *then* let ψ measure it. In other words, the measurement rule stays fixed while the object being measured is moved. This is the opposite direction from homology: cohomology is *contravariant*.

This simple observation is enough to guess the definition:

Definition 4.5. Let $f : X \to Y$ be continuous. For each n, the *pullback on cochains*

$$f^* : C^n(Y) \longrightarrow C^n(X)$$

is the group homomorphism defined on a singular n–simplex $\sigma : \Delta^n \to X$ by

$$\big(f^*\psi\big)(\sigma) \;:=\; \psi\big(f \circ \sigma\big), \qquad \psi \in C^n(Y),$$

and extended linearly to all of $C_n(X)$.

Two immediate features are worth noticing.

(1) Linearity. Because everything is defined simplexwise and then extended linearly, f^* is a group homomorphism: $f^*(\psi + \phi) = f^*\psi + f^*\phi$ and $f^*(a\psi) = a\, f^*\psi$.

(2) Reversal of arrows. Chains move *with* f (covariantly), but cochains move *against* f (contravariantly):

$$f_\# : C_n(X) \to C_n(Y) \quad \text{whereas} \quad f^* : C^n(Y) \to C^n(X).$$

This simply encodes the idea that a measurement must follow the object it measures.

The pullback interacts perfectly with the coboundary. The key identity is nothing more than "boundary commutes with composition" at the level of simplices.

Lemma 4.0.1 (Naturality of the coboundary)**.** For any continuous $f : X \to Y$ and any $\psi \in C^n(Y)$,

$$\delta\big(f^*\psi\big) \;=\; f^*\big(\delta\psi\big) \in C^{n+1}(X).$$

Proof. Evaluate both sides on an $(n+1)$–simplex $\sigma : \Delta^{n+1} \to X$:

$$\big(\delta f^*\psi\big)(\sigma) = f^*\psi(\partial\sigma) = \psi\big(f \circ \partial\sigma\big) = \psi\big(\partial(f \circ \sigma)\big) = (\delta\psi)(f \circ \sigma) = \big(f^*\delta\psi\big)(\sigma).$$

We used $\partial(f \circ \sigma) = f \circ (\partial\sigma)$ and the definition of δ on cochains. ∎

This one-line computation has two important consequences.

Proposition 4.2 (Induced maps on cohomology)**.** If $f : X \to Y$ is continuous, then for each n the map $f^* : C^n(Y) \to C^n(X)$ sends cocycles to cocycles and coboundaries to coboundaries. Hence it descends to a well-defined homomorphism on cohomology

$$f^* : H^n(Y) \longrightarrow H^n(X), \qquad f^*[\psi] := [f^*\psi].$$

Proof. If $\psi \in Z^n(Y) = \ker\delta$, then $\delta(f^*\psi) = f^*(\delta\psi) = f^*(0) = 0$, so $f^*\psi \in Z^n(X)$. If $\psi = \delta\theta \in B^n(Y)$, then $f^*\psi = f^*(\delta\theta) = \delta(f^*\theta) \in B^n(X)$. Thus f^* preserves the subgroups Z^n and B^n and induces a homomorphism on the quotient $H^n = Z^n/B^n$. ∎

Just as importantly, pullback behaves functorially with respect to composition and identities.

Proposition 4.3 (Functoriality)**.** For continuous maps $X \xrightarrow{f} Y \xrightarrow{g} Z$ and any n,

$$(\mathrm{id}_X)^* = \mathrm{id}_{H^n(X)} \quad \text{and} \quad (g \circ f)^* = f^* \circ g^* : H^n(Z) \longrightarrow H^n(X).$$

Proof. On cochains, $(\mathrm{id}_X)^*\psi(\sigma) = \psi(\mathrm{id}_X \circ \sigma) = \psi(\sigma)$. For composition, $(g \circ f)^*\psi(\sigma) = \psi(g \circ f \circ \sigma) = g^*\psi(f \circ \sigma) = (f^*g^*\psi)(\sigma)$. Passing to cohomology uses the previous proposition. ∎

It helps to see all this happen in concrete cases.

Example 4.3 (Restriction along an inclusion)**.** Let $i : U \hookrightarrow X$ be the inclusion of a subspace. For a cochain $\psi \in C^n(X)$, the pullback $i^*\psi \in C^n(U)$ is just ψ evaluated on simplices that land in U: $(i^*\psi)(\sigma) = \psi(i \circ \sigma) = \psi(\sigma)$ for $\sigma : \Delta^n \to U$. Thus $i^* : H^n(X) \to H^n(U)$ is the familiar *restriction* map in cohomology.

Example 4.4 (Degree k map on the circle)**.** Let $f_k : S^1 \to S^1$ be the map $f_k(e^{2\pi it}) = e^{2\pi ikt}$ with $k \in \mathbb{Z}$. We know $H^1(S^1) \cong \mathbb{Z}$. A convenient generator $[\omega] \in H^1(S^1)$ can be described intuitively as the class that measures the *winding number* of loops in S^1: when a loop goes once counterclockwise, $[\omega]$ evaluates to 1. Now pull back this measurement along f_k. Given any loop $\gamma : \Delta^1 \to S^1$ in the domain,

$$(f_k^*\omega)(\gamma) = \omega(f_k \circ \gamma).$$

But $f_k \circ \gamma$ winds around the target circle exactly k times as often as γ does. Hence the pulled-back class detects k-times the winding: $f_k^*[\omega] = k\,[\omega]$ in $H^1(S^1)$. In particular, on $H^1(S^1) \cong \mathbb{Z}$, the induced map is multiplication by k.

Example 4.5 (A quick vanishing)**.** If $c : X \to Y$ is a constant map into a path component of Y, then for every $n \geq 1$, $c^* : H^n(Y) \to H^n(X)$ is the zero map. Indeed, any singular n–simplex $\sigma : \Delta^n \to X$ composes with c to the degenerate simplex $c \circ \sigma : \Delta^n \to Y$ with image a point, on which $(n \geq 1)$–cochains evaluate trivially in cohomology. (One may view this as the cohomological counterpart of the fact that higher homology vanishes on a point.)

All of the above can be remembered as a single principle: *cohomology is a contravariant measurement theory.* A map $f : X \to Y$ sends objects in X forward to Y; to keep measuring those objects by the same rule, we must pull the rule back along f. The algebra (definitions, naturality, and functoriality) is nothing more than a precise encoding of this idea.

4.3 Relative Cohomology

Up to now, our cochains have been defined on *all* singular simplices in a space X. But what if part of the space is "uninteresting" to us — say, a subspace $A \subset X$ on which we do not want our cochains to register anything at all? Perhaps A is already well understood, and we wish to study only the new behaviour that arises in X beyond A. This leads naturally to the notion of *relative cochains.*

The idea is simple: a relative cochain is just a cochain on X that *vanishes* on every simplex lying entirely inside A. In other words, it measures chains in X while ignoring those supported in A.

Definition 4.6. Let $A \subset X$ be a subspace. The group of *n–dimensional relative cochains* on (X, A) is

$$C^n(X, A) := \{\varphi \in C^n(X) \mid \varphi(\sigma) = 0 \text{ for all } \sigma \in C_n(A)\}.$$

Equivalently, $C^n(X, A)$ is the kernel of the restriction map $j : C^n(X) \to C^n(A)$.

From this point of view, $C^n(X, A)$ consists of all cochains that are "blind" to A. Since the coboundary of a cochain that vanishes on A still vanishes on A, the coboundary map $\delta : C^n(X) \to C^{n+1}(X)$ restricts to a well-defined map $\delta : C^n(X, A) \to C^{n+1}(X, A)$. Thus the relative cochains form a cochain complex in their own right:

$$\cdots \xrightarrow{\delta} C^{n-1}(X, A) \xrightarrow{\delta} C^n(X, A) \xrightarrow{\delta} C^{n+1}(X, A) \xrightarrow{\delta} \cdots$$

Definition 4.7. The *n–th relative cohomology group* of the pair (X, A) is

$$H^n(X, A) := \frac{\ker\big(\delta : C^n(X, A) \to C^{n+1}(X, A)\big)}{\operatorname{im}\big(\delta : C^{n-1}(X, A) \to C^n(X, A)\big)}.$$

This group measures precisely the obstruction to making a relative cocycle (i.e. a measurement on X that ignores A and is consistent on boundaries) into a relative coboundary (i.e. the coboundary of a cochain on X that vanishes on A).

To relate $H^n(X, A)$ to the absolute cohomology of X and A, we note that there is a natural restriction map

$$j : C^n(X) \longrightarrow C^n(A), \quad j(\varphi) := \varphi|_{C_n(A)}.$$

Its kernel is exactly $C^n(X, A)$. The inclusion $i : C^n(X, A) \hookrightarrow C^n(X)$ is also natural, and together they form a short exact sequence of cochain groups:

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

Exactness here simply says: every cochain on X that vanishes on A is a relative cochain; every cochain on X restricts to a cochain on A; and every cochain on A is the restriction of some cochain on X.

The important point is that this is a short exact sequence *of cochain complexes*. Applying the general homological algebra machinery (which we have already seen in the chain complex setting), we obtain a *long exact sequence in cohomology*:

Theorem 4.1 (Long exact sequence of a pair). For any subspace $A \subset X$, there is a natural long exact sequence

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

where i^* is induced by the inclusion $C^\bullet(X,A) \hookrightarrow C^\bullet(X)$, j^* is induced by restriction to A, and δ is the connecting homomorphism.

A striking feature here is the reversal of direction compared to the homology long exact sequence: this is a direct consequence of cohomology being *contravariant*.

Example 4.6. Consider the pair (D^2, S^1), where D^2 is the closed disk and S^1 its boundary. One can check that

$$H^n(D^2) \cong \begin{cases} \mathbb{Z}, & n = 0, \\ 0, & n > 0, \end{cases} \quad \text{and} \quad H^n(S^1) \cong \begin{cases} \mathbb{Z}, & n = 0, \\ \mathbb{Z}, & n = 1, \\ 0, & n > 1. \end{cases}$$

The long exact sequence for the pair in low degrees reads

$$0 \longrightarrow H^1(D^2,S^1) \xrightarrow{i^*} H^1(D^2) \xrightarrow{j^*} H^1(S^1) \xrightarrow{\delta} H^2(D^2,S^1) \xrightarrow{i^*} H^2(D^2) = 0.$$

Here $H^1(D^2) = 0$, so exactness forces $i^* : H^1(D^2,S^1) \to 0$ to be trivial and $j^* : 0 \to H^1(S^1) \cong \mathbb{Z}$ to be the zero map. Thus $\delta : H^1(S^1) \to H^2(D^2,S^1)$ is injective. But $H^1(S^1) \cong \mathbb{Z}$, so we find

$$H^2(D^2,S^1) \cong \mathbb{Z}.$$

Geometrically, this generator corresponds to the "fundamental class" of the disk, viewed relative to its boundary.

In practice, the long exact sequence of a pair is one of the most powerful tools for computing cohomology. It links together the absolute cohomology of X and A with the relative cohomology of the pair, and it does so in a way that reflects the precise algebraic relationship between measurements on X, measurements on A, and measurements that ignore A entirely.

4.4 Mayer–Vietoris Sequence

Up to this point, our definitions and tools for cohomology have been quite general, but when it comes to *actual computations*, we need techniques that allow us to break a space into simpler pieces and then reassemble the results. If X is complicated, it is often much easier to cover it with open sets U and V whose individual cohomology groups are well-understood, together with their overlap $U \cap V$. The natural question is:

Given $H^*(U)$, $H^*(V)$, and $H^*(U \cap V)$, can we recover $H^*(X)$?

It turns out that the answer is yes — at least in a way that is both precise and computationally effective. The key tool is the **Mayer–Vietoris sequence**, which relates the cohomology of X to that of U, V, and $U \cap V$ through a long exact sequence.

To understand where such a sequence might come from, we start from the level of *cochains*. Given a cochain φ on X, we can restrict it to U and to V. This gives a natural map

$$r \colon C^n(X) \longrightarrow C^n(U) \oplus C^n(V), \qquad r(\varphi) = (\varphi|_U,\, \varphi|_V).$$

Here $C^n(X)$ denotes the group of singular n–cochains on X, and the restriction $\varphi|_U$ means "evaluate only on singular simplices lying in U". This restriction is clearly compatible with addition of cochains and with scalar multiplication, so r is a homomorphism of abelian groups.

Now suppose we have a *pair* $(\alpha, \beta) \in C^n(U) \oplus C^n(V)$ coming from the restriction of a single φ on X. On $U \cap V$, the two restrictions α and β must agree: for every singular n–simplex σ with image in $U \cap V$, we have $\alpha(\sigma) = \beta(\sigma)$. This suggests defining the "difference on the overlap" map

$$d \colon C^n(U) \oplus C^n(V) \longrightarrow C^n(U \cap V), \qquad d(\alpha, \beta) = \alpha|_{U \cap V} - \beta|_{U \cap V}.$$

If (α, β) comes from a global φ on X, then $d(\alpha, \beta) = 0$. Conversely, if $d(\alpha, \beta) = 0$, then the two agree on the overlap and can be glued together to form a cochain on all of X.

Note that a singular simplex $\sigma \colon \Delta^n \to X$ need not lie entirely inside U or V; it may cross between them. Consequently, the injectivity of r in $(*)$ *fails on the nose*: a cochain on X can vanish on all simplices contained in U and in V, yet be nonzero on simplices that straddle $U \cup V$. The standard remedy is to *subdivide* simplices until each little simplex lies entirely in U or entirely in V.

Concretely, there is a natural chain map

$$\mathrm{sd}\colon C_\bullet(X) \longrightarrow C_\bullet(X)$$

(barycentric subdivision) with the properties: (i) sd is chain homotopic to the identity; hence it induces the identity on cohomology, and (ii) for any open cover $\mathcal{U}$ of X, some iterate sd^N subdivides every singular simplex as a chain supported in the cover, i.e. each summand simplex is entirely contained in some $U \in \mathcal{U}$.

Precomposing cochains with sd^N does not change cohomology, but it ensures that values of cochains on X are determined by their values on simplices lying in U or in V.

With this understood, the following becomes true and (importantly) sufficient for Mayer–Vietoris.

Proposition 4.4 (Short exact sequence after subdivision). Let $X = U \cup V$ with U, V open. For $N \gg 0$, the maps

$$r_N(\varphi) := (\varphi \circ \mathrm{sd}^N)|_U, \quad (\varphi \circ \mathrm{sd}^N)|_V, \qquad d(\alpha, \beta) := \alpha|_{U\cap V} - \beta|_{U\cap V}$$

assemble (degreewise) into a short exact sequence of cochain complexes

$$0 \longrightarrow C^\bullet(X) \xrightarrow{r_N} C^\bullet(U) \oplus C^\bullet(V) \xrightarrow{d} C^\bullet(U \cap V) \longrightarrow 0.$$

Moreover, r_N is a cochain homotopy equivalence to r_0 (the naive restriction), so passing to cohomology does not depend on N.

Proof. Since sd^N is a chain map, precomposition preserves coboundaries, so r_N is a cochain map. Exactness at the middle term means $\ker(d) = \mathrm{im}(r_N)$. If $(\alpha, \beta) \in \ker(d)$, then α and β agree on $U \cap V$. Because every n–simplex of $\mathrm{sd}^N(\sigma)$ lies entirely in U or entirely in V, we can define φ on a generator σ by summing the contributions of those little simplices: use α on the ones in U and β on the ones in V. Linearity then gives a well-defined $\varphi \in C^n(X)$ with $r_N(\varphi) = (\alpha, \beta)$. Thus $\ker(d) \subset \mathrm{im}(r_N)$; the reverse inclusion $d \circ r_N = 0$ is immediate.

Surjectivity of d is easy: given $\gamma \in C^n(U \cap V)$, extend it to U by defining a cochain that agrees with γ on simplices in $U \cap V$ and is 0 on simplices in $U \setminus V$; take $\beta := 0$ on V. Then $d(\alpha, 0) = \gamma$.

Injectivity of r_N follows from the support property: if $r_N(\varphi) = (0, 0)$, then φ evaluates to 0 on every subdivided simplex lying in U or in V, hence on every summand of $\mathrm{sd}^N(\sigma)$ for every σ, so $\varphi \circ \mathrm{sd}^N = 0$. Since sd^N is chain homotopic to the identity, this forces $\varphi = 0$ on cohomology (and, degreewise, is enough for exactness of complexes after choosing N uniformly). ∎

Remark 4.3. If one prefers to avoid subdivision in the proof, an equivalent and classical route is to derive Mayer–Vietoris from the long exact sequence of a pair together with *excision*: apply excision to the pair (X, V) with the excised set $X \setminus U$; combine the long exact sequences for (X, V) and $(U, U \cap V)$ and chase the diagram. We keep the concrete subdivision picture here because it matches the computational intuition: after enough subdivision, each simplex "sits entirely inside" one piece.

With Proposition 4.4 in hand, we can invoke the general homological algebra fact that a short exact sequence of cochain complexes induces a long exact sequence in cohomology.

Theorem 4.2 (Mayer–Vietoris sequence for singular cohomology)**.** If $X = U \cup V$ with U, V open, there is a natural long exact sequence

$$\cdots \longrightarrow H^{n-1}(U\cap V) \xrightarrow{\delta} H^n(X) \xrightarrow{(i_U^*, i_V^*)} H^n(U)\oplus H^n(V) \xrightarrow{j_U^* - j_V^*} H^n(U\cap V) \xrightarrow{\delta} H^{n+1}(X) \longrightarrow \cdots$$

where the maps are induced by the inclusions of the corresponding spaces, and δ is the connecting homomorphism.

Proof. We sketch the standard zig–zag construction, which mirrors the one for relative cohomology.

Connecting map. Let $[\gamma] \in H^{n-1}(U \cap V)$ be represented by a cocycle $\gamma \in C^{n-1}(U \cap V)$. By surjectivity of d in Proposition 4.4, choose $(\alpha, \beta) \in C^{n-1}(U) \oplus C^{n-1}(V)$ with $d(\alpha, \beta) = \gamma$; that is, $\alpha|_{U\cap V} - \beta|_{U\cap V} = \gamma$. Apply the coboundary: $\delta\alpha$ and $\delta\beta$ agree on $U \cap V$ (since $\delta\gamma = 0$), so by exactness at the middle term, they glue to a global cocycle $\Phi \in C^n(X)$ with $r_N(\Phi) = (\delta\alpha, \delta\beta)$. Define $\delta[\gamma] := [\Phi] \in H^n(X)$.

One checks that different choices of (α, β) change Φ by a coboundary, so δ is well defined on cohomology classes.

Exactness. We verify kernels equal images at the three relevant spots.

(i) $\mathrm{im}(\delta) = \ker(i_U^*, i_V^*)$: If $[\Phi] = \delta[\gamma]$ as above, then by construction $r_N(\Phi) = (\delta\alpha, \delta\beta)$, hence both restrictions are coboundaries and $(i_U^*, i_V^*)([\Phi]) = (0, 0)$. Conversely, if $[\Phi] \in H^n(X)$ restricts trivially to both $H^n(U)$ and $H^n(V)$, pick cochains α, β with $\delta\alpha = \Phi|_U$ and $\delta\beta = \Phi|_V$; then $\alpha|_{U\cap V} - \beta|_{U\cap V}$ is a cocycle in $C^{n-1}(U \cap V)$, and the construction above shows $\delta[\alpha| - \beta|] = [\Phi]$.

(ii) $\operatorname{im}(i_U^*, i_V^*) = \ker(j_U^* - j_V^*)$: If $[\Phi] \in H^n(X)$, then $(i_U^*, i_V^*)([\Phi]) = ([\Phi|_U], [\Phi|_V])$, whose difference restricts to zero on $U \cap V$ by functoriality of restriction. Conversely, if $([\alpha], [\beta])$ have equal restrictions on $H^n(U \cap V)$, choose cocycles α, β representing them with $\alpha|_{U\cap V} = \beta|_{U\cap V}$. Exactness of Proposition 4.4 glues them to a global cocycle Φ with $r_N(\Phi) = (\alpha, \beta)$, so $([\alpha], [\beta])$ is in the image of $H^n(X)$.

(iii) $\operatorname{im}(j_U^* - j_V^*) = \ker(\delta)$: If $[\gamma] = (j_U^* - j_V^*)([\alpha], [\beta])$ with α, β cocycles, the construction of δ uses (α, β) to produce Φ cohomologous to 0 (since $\delta\alpha$ and $\delta\beta$ glue to a coboundary), so $\delta[\gamma] = 0$. Conversely, if $\delta[\gamma] = 0$, unraveling the definition shows γ is the difference of restrictions of classes from U and V. ■

Example 4.7 (A first computation: $H^*(S^1)$)**.** Cover the circle by two open arcs U and V whose intersection is the disjoint union of two open arcs. Each arc is contractible, so

$$H^0(U) \cong \mathbb{Z}, \quad H^k(U) = 0 \ (k > 0), \qquad H^0(V) \cong \mathbb{Z}, \quad H^k(V) = 0 \ (k > 0).$$

The intersection has two components, so $H^0(U \cap V) \cong \mathbb{Z} \oplus \mathbb{Z}$ and $H^k(U \cap V) = 0$ for $k > 0$. The relevant segment of the Mayer–Vietoris sequence is

$$0 \longrightarrow H^0(S^1) \xrightarrow{(i_U^*, i_V^*)} \mathbb{Z} \oplus \mathbb{Z} \xrightarrow{j_U^* - j_V^*} \mathbb{Z} \oplus \mathbb{Z} \xrightarrow{\delta} H^1(S^1) \longrightarrow 0.$$

The first map is the diagonal $z \mapsto (z, z)$. One checks that $j_U^* - j_V^*$ has image a rank-1 subgroup (differences of values on the two components), so its cokernel is $\mathbb{Z}$. Exactness then gives

$$H^0(S^1) \cong \mathbb{Z}, \qquad H^1(S^1) \cong \mathbb{Z}.$$

Remark 4.4 (Homology version)**.** There is a completely analogous Mayer–Vietoris sequence for *homology*. Because homology is covariant (chains are pushed forward), the arrows run in the opposite direction:

$$\cdots \longrightarrow H_n(U \cap V) \longrightarrow H_n(U) \oplus H_n(V) \longrightarrow H_n(X) \xrightarrow{\partial} H_{n-1}(U \cap V) \longrightarrow \cdots$$

The same subdivision/excision ideas underlie its proof; only the variance changes.

4.5 The Ring Structure

Up to now, our cohomology groups $H^n(X)$ have been sitting side by side, one for each n, like a shelf of separate books. Each group measures something about X, but they have not yet interacted with each other.

There is, however, a remarkable and entirely *natural* way to combine cohomology classes of different degrees into new ones. This is not some extra decoration we decide to add; it is *forced* upon us once we take seriously the idea that a cochain is a *measurement* on simplices.

Think about it this way: suppose φ is a p–cochain, and ψ is a q–cochain. Given a $(p+q)$–simplex σ, φ knows how to measure p–dimensional simplices, and ψ knows how to measure q–dimensional simplices. The $(p+q)$–simplex naturally splits into two pieces: a *front* p–face and a *back* q–face. We can simply let φ measure the front, let ψ measure the back, and multiply the two readings. That multiplication gives us a perfectly good $(p+q)$–cochain.

If you draw a triangle and imagine $p = 1$ and $q = 1$, the front face is the first edge, the back face is the second edge. Measuring each edge separately and multiplying the results produces a number for the whole triangle. This is the geometric seed of the construction.

Definition 4.8 (Cup product on singular cochains). Let X be a topological space. Take $\varphi \in C^p(X)$ and $\psi \in C^q(X)$. For a singular $(p+q)$–simplex $\sigma : \Delta^{p+q} \to X$, write

$$\sigma|[v_0, \dots, v_p] \quad \text{and} \quad \sigma|[v_p, \dots, v_{p+q}]$$

for its front p–face and back q–face. We define a $(p+q)$–cochain $\varphi \smile \psi \in C^{p+q}(X)$ by

$$(\varphi \smile \psi)(\sigma) := \varphi\big(\sigma|[v_0, \dots, v_p]\big) \cdot \psi\big(\sigma|[v_p, \dots, v_{p+q}]\big),$$

and extend this definition bilinearly to all $(p+q)$–chains.

So the picture is:

$$\text{front face} \xrightarrow{\varphi} \text{number}, \quad \text{back face} \xrightarrow{\psi} \text{number}, \quad \text{multiply to get the measurement of the whole.}$$

At this point you might ask: does this multiplication behave well with respect to the coboundary δ? This is an essential question, because if it does not, then multiplying representatives of cohomology classes might give something that *does not* represent the product of the classes.

The answer turns out to be yes, in the following precise way.

Proposition 4.5 (Leibniz rule for the cup product)**.** For $\varphi \in C^p(X)$ and $\psi \in C^q(X)$,

$$\delta(\varphi \smile \psi) = \delta\varphi \smile \psi \; + \; (-1)^p \, \varphi \smile \delta\psi$$

as $(p+q+1)$–cochains.

The significance of this identity is immediate: if either φ or ψ is a cocycle, then so is $\varphi \smile \psi$. Thus the product descends to cohomology classes.

Corollary 4.1. If $[\varphi] \in H^p(X)$ and $[\psi] \in H^q(X)$ are represented by cocycles, then $\varphi \smile \psi$ is also a cocycle, and the cohomology class $[\varphi \smile \psi]$ depends only on $[\varphi]$ and $[\psi]$. Hence there is a well-defined product

$$\smile : H^p(X) \times H^q(X) \longrightarrow H^{p+q}(X), \quad [\varphi] \smile [\psi] := [\varphi \smile \psi].$$

From the definition, several algebraic properties follow almost at once.

Proposition 4.6 (Bilinearity, unit, and naturality)**.** The following are true:

1. **Bilinearity:** $\smile$ is linear in each variable.
2. **Unit:** The constant 0–cochain 1 (value 1 on each vertex) represents a class $1 \in H^0(X)$ which acts as a multiplicative unit: $1 \smile \alpha = \alpha = \alpha \smile 1$.
3. **Naturality:** If $f : X \to Y$ is continuous, then for $\alpha \in H^p(Y)$ and $\beta \in H^q(Y)$,

$$f^*(\alpha \smile \beta) = f^*\alpha \smile f^*\beta.$$

One of the most elegant properties is *graded commutativity*: swapping two factors introduces a predictable sign.

Theorem 4.3 (Graded commutativity)**.** If $\alpha \in H^p(X)$ and $\beta \in H^q(X)$, then

$$\alpha \smile \beta = (-1)^{pq}\beta \smile \alpha.$$

Geometric heuristic. On a $(p+q)$–simplex, swapping the roles of a p–cochain and a q–cochain amounts to moving p front vertices past q back vertices. Each swap of two vertices introduces a sign -1, and there are pq swaps, giving the factor $(-1)^{pq}$. The full combinatorial proof shows the difference between the two orders is actually a coboundary, so they agree in cohomology. ■

Example 4.8 ($H^*(S^n)$ as a ring)**.** For $n \geq 1$, $H^0(S^n) \cong \mathbb{Z}$, $H^n(S^n) \cong \mathbb{Z}$, and all other groups vanish. The only nontrivial products involve the unit in degree 0; if u generates $H^n(S^n)$, then $u \smile u$ lies in $H^{2n}(S^n) = 0$. Thus

$$H^*(S^n) \cong \mathbb{Z} \oplus \mathbb{Z}\langle u\rangle, \quad u^2 = 0, \quad |u| = n.$$

Example 4.9 (The torus T^2)**.** Let $T^2 = S^1 \times S^1$ with projections π_1, π_2. If $a \in H^1(S^1)$ is the generator, define $x := \pi_1^* a$ and $y := \pi_2^* a$ in $H^1(T^2)$. These detect winding in each circle factor. Graded commutativity forces $x \smile x = 0 = y \smile y$ and $x \smile y = -y \smile x$. One finds $x \smile y$ generates $H^2(T^2) \cong \mathbb{Z}$, so

$$H^*(T^2) \cong \Lambda_{\mathbb{Z}}(x, y),$$

the exterior algebra on two degree–1 generators.

Remark 4.5. Cohomology groups alone tell you how many holes of each dimension exist; the cup product tells you how those holes *interact*. In de Rham cohomology, this product will be realised concretely as the wedge product of differential forms.

4.6 Relation to Homotopy

Up to this point we have related *homology* to *homotopy* (via the Hurewicz map), and we have learned to think of *cohomology* as linear "measurements" on homology classes. It is natural, then, to ask how cohomology reflects homotopy. A good first guess is: if homology in degree n records the n–dimensional "shapes" in X, then cohomology in degree n should record all *linear functionals* on those shapes. This intuition leads directly to the universal coefficient theorem.

To make the guess precise, remember what a class $[\varphi] \in H^n(X)$ *does*: pick a singular n–cycle z (a closed "shape"); evaluate any cocycle representative φ on z to get an integer; if you change z by a boundary or change φ by a coboundary, that integer does not change. In other words, $[\varphi]$ is a well-defined, additive assignment of integers to $H_n(X)$—at least when no hidden torsion effects interfere. The theorem below says that this picture is essentially complete.

Theorem 4.4 (Universal Coefficient Theorem for Cohomology)**.** For any topological space X, abelian group G, and any $n \geq 0$, there is a natural short exact sequence

$$0 \longrightarrow \mathrm{Ext}\big(H_{n-1}(X),\, G\big) \longrightarrow H^n(X;G) \longrightarrow \mathrm{Hom}\big(H_n(X),\, G\big) \longrightarrow 0,$$

which always splits (not canonically). In particular, if $H_n(X)$ is free abelian, then

$$H^n(X;G) \;\cong\; \mathrm{Hom}\big(H_n(X),\, G\big).$$

Pause and unpack this. The group $\mathrm{Hom}(H_n(X), G)$ is exactly the space of all additive measurements of n–dimensional homology classes with values in G—this matches our intuition for cocycles. The extra term $\mathrm{Ext}(H_{n-1}(X), G)$ is the correction that appears when there is torsion in $H_{n-1}(X)$: it contributes "hidden" cohomology classes that cannot be seen just by evaluating on $H_n(X)$. Over $G = \mathbb{Z}$ this term picks up torsion in $H_{n-1}(X)$ (for example, $\mathrm{Ext}(\mathbb{Z}/m, \mathbb{Z}) \cong \mathbb{Z}/m$).

A first, extremely useful consequence is the clean identification in the torsion–free case.

Corollary 4.2. If $H_n(X)$ is free abelian (e.g. $H_n(X) \cong \mathbb{Z}^r$), then

$$H^n(X) \;\cong\; \mathrm{Hom}\big(H_n(X), \mathbb{Z}\big).$$

Concretely, a cohomology class is determined by the integers it assigns to a basis of $H_n(X)$, and any such assignment extends uniquely and linearly.

Now bring homotopy back into the picture. You have already met the Hurewicz map

$$h : \pi_n(X) \longrightarrow H_n(X),$$

which sends a based map $S^n \to X$ to the n–dimensional homology class it represents. When X is

$(n-1)$–connected (no homotopy in lower degrees), the Hurewicz theorem tells us that h is an isomorphism. Combining this with the universal coefficient theorem gives the promised bridge.

Corollary 4.3 (Cohomology reads homotopy in the highly connected range). If X is $(n-1)$–connected and $H_n(X)$ is free abelian (as happens, for instance, for $X = S^n$), then

$$H^n(X) \;\cong\; \mathrm{Hom}\big(H_n(X), \mathbb{Z}\big) \;\cong\; \mathrm{Hom}\big(\pi_n(X), \mathbb{Z}\big).$$

Thus degree-n cohomology classes are exactly the integer-valued homomorphisms on $\pi_n(X)$.

This is a satisfying stopping point for the big picture: in good connectivity, cohomology literally "linearises" homotopy. Let us test the statement on familiar spaces.

Example 4.10. For the n–sphere S^n, we know $\pi_n(S^n) \cong \mathbb{Z}$ by basic homotopy theory. Since S^n is $(n-1)$–connected, the Hurewicz theorem gives $H_n(S^n) \cong \mathbb{Z}$, and the universal coefficient theorem identifies $H^n(S^n) \cong \mathrm{Hom}(\mathbb{Z}, \mathbb{Z}) \cong \mathbb{Z}$. The generator of $H^n(S^n)$ can be thought of as the "degree" map measuring how many times a map wraps S^n around itself.

Remark 4.6. This point of view also makes it clear why cohomology, in general, cannot distinguish all homotopy types. It captures only the "abelianised" part of the homotopy information: once $\pi_n(X)$ is replaced by its abelianisation via the Hurewicz map, cohomology records linear maps out of it. Nonetheless, in many important cases — especially for simply connected spaces — this is enough to recover significant geometric information.

In summary, one should keep in mind:

$$\textit{Cohomology} = \textit{linear functionals on homology} + \textit{a torsion correction,}$$

and, where Hurewicz applies, these linear functionals can be read directly from the corresponding homotopy group. This is exactly the bridge we will cross again in de Rham theory, where the "functionals" are realised as integration of differential forms over cycles.

With this, we bring our exploration of cohomology to a close. We began by reversing our homological viewpoint: instead of building geometric chains, we learned to measure them via cochains, built the coboundary operator, and constructed cohomology groups as a dual theory to homology. Along the way, we uncovered the structural richness of cohomology — its functorial behaviour, its role in long exact sequences and Mayer–Vietoris decompositions, and most strikingly, its graded ring structure through the cup product. We have even seen how cohomology reflects certain homotopy-theoretic properties of spaces, and how the universal coefficient theorem links it back to homology.

Yet, cohomology is more than an abstract algebraic formalism. In many cases of interest — especially in geometry and analysis — cohomology classes can be represented concretely by *differential forms*. Here the abstract coboundary δ will reveal itself as the familiar exterior derivative d, and the pairing of cohomology with homology will appear as the act of integrating a form over a cycle. In this guise, cohomology becomes a bridge between topology and calculus, culminating in the generalised Stokes theorem.

It is to this concrete, geometric incarnation — the world of differential forms and de Rham cohomology — that we now turn.

5 Differential Forms and the Generalised Stokes' Theorem

We have now completed our study of cohomology in its algebraic form. Starting from the dual picture of cochains as "measurements" on chains, we developed the coboundary operator, constructed cohomology groups, and explored their algebraic structures, functoriality, and long exact sequences. In many ways, this theory mirrors the homology we began with, but it also carries richer structure, such as the cup product, which reveals how different pieces of a space interact.

However, our discussion so far has been entirely *abstract*: our cochains were arbitrary functions from chains to coefficients, subject only to linearity and the cocycle condition. Nothing in the definition required our space to be smooth, or even to have a geometry at all. As a result, we have not yet made use of the powerful tools of calculus and differential geometry.

When the space *is* a smooth manifold, we can do better. In this setting, there is a particularly elegant and geometric incarnation of cohomology, built not from arbitrary cochains but from *differential forms*: smooth, multilinear, alternating objects that can be integrated over chains. These forms interact beautifully with the manifold's smooth structure: they can be differentiated (via the *exterior derivative*), multiplied (via the *wedge product*), and pulled back along smooth maps in a way that respects these operations.

Most strikingly, there is a single, unifying statement — the *generalised Stokes theorem* — which contains as special cases the fundamental theorems of calculus, Green's theorem, the divergence theorem, and classical Stokes' theorem from vector calculus. This theorem is, in the smooth setting, the exact analogue of the algebraic relation that $\delta^2 = 0$ in cohomology, and it leads naturally to the definition of *de Rham cohomology.*

Our goal in this section is to first introduce differential forms and their basic operations, then explore how to integrate them over smooth chains, and finally state and interpret the generalised Stokes theorem. This will set the stage for the de Rham cohomology groups, which bridge the gap between the algebraic topology we have developed and the analysis and geometry of smooth manifolds.

5.1 Differential Forms

In our study of singular cochains, we treated them as abstract assignments of numbers to chains, with the only requirements being linearity and the compatibility enforced by the coboundary operator. This generality was powerful, but also completely unconstrained: a cochain could behave in any strange way so long as it respected linearity. If our space X had extra structure — say, if X were a smooth manifold —

then we would expect our cochains to reflect this smoothness.

To see what this means, recall how things work in ordinary calculus. A smooth function $f : \mathbb{R}^n \to \mathbb{R}$ assigns to each point p a real number $f(p)$, varying smoothly from point to point. From f we can extract its *differential* df, which at each point is a linear functional on the tangent space $T_p\mathbb{R}^n$. Given a vector v based at p, $df_p(v)$ tells us the instantaneous rate of change of f in the direction of v. In coordinates $(x^1, \ldots, x^n)$ we write

$$df = \frac{\partial f}{\partial x^1}\, dx^1 + \frac{\partial f}{\partial x^2}\, dx^2 + \cdots + \frac{\partial f}{\partial x^n}\, dx^n,$$

where each dx^i is the basic "coordinate differential" which extracts the i–th component of a vector.

This familiar object df is our first clue: it is not itself a function, but rather something that eats a vector and produces a number. Such objects are called *1–forms*, and they are the simplest examples of differential forms beyond functions.

Definition 5.1. A 0*–form* on a smooth manifold M is a smooth function

$$f : M \longrightarrow \mathbb{R}.$$

A 0–form assigns a single real number to each point of M, with no directional component. It is "dimensionless" in the sense that it measures nothing along any curve, surface, or higher–dimensional piece of M — it just evaluates the point itself. Nevertheless, 0–forms are the starting point of our hierarchy: as soon as we take their differential, they give rise to 1–forms.

Definition 5.2. A 1*–form* ω on M assigns to each point $p \in M$ a linear functional

$$\omega_p : T_pM \longrightarrow \mathbb{R}$$

that varies smoothly with p.

That is, for each p, ω_p takes in a tangent vector $v \in T_pM$ and returns a real number, in a way that is linear in v. The smoothness condition says that if p moves slightly, the coefficients of ω in any coordinate chart also change smoothly.

In coordinates $(x^1, \ldots, x^n)$ on $\mathbb{R}^n$, a 1–form can always be written as

$$\omega = a_1(x)\, dx^1 + a_2(x)\, dx^2 + \cdots + a_n(x)\, dx^n,$$

where each $a_i : \mathbb{R}^n \to \mathbb{R}$ is a smooth function. The dx^i here are the basic coordinate 1–forms: at a point p, dx^i_p is the functional which sends a tangent vector $v = (v^1, \ldots, v^n)$ to its i–th component v^i.

Example 5.1. On $\mathbb{R}^2$, define

$$\omega = y\,dx - x\,dy.$$

At a point (x, y), a tangent vector (v_x, v_y) is measured by

$$\omega(v_x, v_y) = y \cdot v_x - x \cdot v_y.$$

If (v_x, v_y) is tangent to the circle $x^2 + y^2 = r^2$, this expression is proportional to the angular speed around the origin. Thus ω measures rotation rather than translation.

From 1–forms it is a natural step to consider objects that take *several* tangent vectors at once. Just as a 1–form can be thought of as a tool for measuring oriented lengths, a 2–form will measure oriented areas, a 3–form will measure oriented volumes, and so on.

The correct algebraic way to capture this is to require *multilinearity* (linear in each vector separately) and *antisymmetry* (swapping two vectors reverses the sign).

Definition 5.3. A *k–form* ω on a smooth manifold M assigns to each point $p \in M$ an alternating k–linear map

$$\omega_p : T_pM \times \cdots \times T_pM \longrightarrow \mathbb{R}$$

that is linear in each argument and changes sign when any two arguments are exchanged. The assignment $p \mapsto \omega_p$ is smooth in the sense that, in local coordinates, the coefficient functions vary smoothly.

The antisymmetry ensures that a k–form vanishes whenever two of its input vectors are the same or linearly dependent. Geometrically, this matches our intuition: a k–dimensional volume should be zero if the k vectors fail to span a k–dimensional parallelepiped.

In coordinates $(x^1, \ldots, x^n)$, every k–form can be expressed as

$$\omega = \sum_{1 \le i_1 < \cdots < i_k \le n} a_{i_1 \cdots i_k}(x)\, dx^{i_1} \wedge \cdots \wedge dx^{i_k},$$

where the $a_{i_1 \cdots i_k}$ are smooth functions. The symbols $dx^{i_1} \wedge \cdots \wedge dx^{i_k}$, whose properties we will soon explore, represent the basic infinitesimal k–dimensional volume elements in the indicated coordinate directions.

Example 5.2. On $\mathbb{R}^3$, the 2–form

$$\eta = x\,dy \wedge dz$$

evaluated at a point (x, y, z) takes two tangent vectors u, v and returns x times the signed area of their projection onto the yz–plane. If $x = 0$ this measurement vanishes, showing that η is sensitive to position as well as direction.

We have now laid out the basic cast of characters: 0–forms are smooth functions, 1–forms measure lengths in given directions, and k–forms in general measure oriented k–dimensional volumes in the tangent space. In the next step, we will see how to combine these forms using the *wedge product*, which allows us to build higher–degree measurements from lower–degree ones.

5.2 The wedge product

Up to now we have learned to *measure* directions with 1–forms and to measure oriented k–dimensional parallelepipeds with k–forms. A natural question is: can we *combine* simple measurements to make more complicated ones? If one 1–form reads the "amount of motion" in the x–direction and another does so in the y–direction, there ought to be a canonical way to produce from them a 2–form that reads off oriented *area* in the xy–plane. The operation that does exactly this is the *wedge product*.

Definition 5.4 (Wedge of 1–forms at a point)**.** Let V be a real vector space. Given linear functionals $\alpha_1, \ldots, \alpha_k \in V^*$, their *wedge* (or *exterior*) product

$$\alpha_1 \wedge \cdots \wedge \alpha_k \;\in\; \mathrm{Alt}^k(V^*)$$

is the alternating k–linear form defined on vectors $v_1, \ldots, v_k \in V$ by

$$(\alpha_1 \wedge \cdots \wedge \alpha_k)(v_1, \ldots, v_k) \;:=\; \det\bigl(\alpha_i(v_j)\bigr)_{1 \le i,j \le k}.$$

This formula encapsulates two key ideas you already know from linear algebra:

- *Multilinearity:* the determinant is linear in each column, so the wedge is linear in each vector argument.

- *Alternation:* swapping two columns changes the sign of the determinant, hence the wedge vanishes as soon as two input vectors are equal or linearly dependent.

Because each α_i is itself linear, the definition above is also *bilinear* in each α_i. On a smooth manifold M, we apply this construction pointwise in $V = T_pM$ and let the coefficients vary smoothly in p.

Definition 5.5 (Wedge product of differential forms). If $\omega \in \Omega^p(M)$ and $\eta \in \Omega^q(M)$ are differential forms, their *wedge product* $\omega \wedge \eta \in \Omega^{p+q}(M)$ is defined pointwise by

$$(\omega \wedge \eta)_p \;:=\; \mathrm{Alt}\big((\omega_p \otimes \eta_p)\big),$$

that is, for $v_1, \dots, v_{p+q} \in T_pM$,

$$(\omega \wedge \eta)_p(v_1, \dots, v_{p+q}) = \frac{1}{p!\,q!} \sum_{\sigma \in S_{p+q}} \mathrm{sgn}(\sigma)\, \omega_p(v_{\sigma(1)}, \dots, v_{\sigma(p)})\, \eta_p(v_{\sigma(p+1)}, \dots, v_{\sigma(p+q)}).$$

Equivalently, when $\omega = \alpha_1 \wedge \cdots \wedge \alpha_p$ and $\eta = \beta_1 \wedge \cdots \wedge \beta_q$ are simple forms (wedges of 1–forms),

$$\omega \wedge \eta \;=\; \alpha_1 \wedge \cdots \wedge \alpha_p \wedge \beta_1 \wedge \cdots \wedge \beta_q,$$

and we extend by bilinearity.

In local coordinates $(x^1, \dots, x^n)$ this becomes completely concrete. Every form can be written as a linear combination of the elementary wedges $dx^{i_1} \wedge \cdots \wedge dx^{i_k}$ with $i_1 < \cdots < i_k$, and the product is determined by

$$dx^i \wedge dx^j = -\, dx^j \wedge dx^i, \qquad dx^i \wedge dx^i = 0,$$

extended bilinearly and associatively. Thus, if

$$\omega = \sum_I a_I(x)\, dx^{i_1} \wedge \cdots \wedge dx^{i_p}, \qquad \eta = \sum_J b_J(x)\, dx^{j_1} \wedge \cdots \wedge dx^{j_q},$$

then

$$\omega \wedge \eta = \sum_{I,J} a_I(x)\, b_J(x)\, dx^{i_1} \wedge \cdots \wedge dx^{i_p} \wedge dx^{j_1} \wedge \cdots \wedge dx^{j_q},$$

and you reorder the indices into increasing order, inserting the corresponding sign; any repeated index forces the term to vanish.

Example 5.3 (Oriented area in $\mathbb{R}^2$). On $\mathbb{R}^2$ with coordinates (x, y),

$$(dx \wedge dy)\big((v_x, v_y), (w_x, w_y)\big) = \det \begin{pmatrix} v_x & w_x \\ v_y & w_y \end{pmatrix} = v_x w_y - v_y w_x.$$

Thus $dx \wedge dy$ measures the signed area of the parallelogram spanned by two vectors. If you reverse the order, $dy \wedge dx = -\, dx \wedge dy$, the sign (orientation) flips.

Example 5.4 (A quick coordinate computation). On $\mathbb{R}^2$, let $\alpha = x^2\, dx + y\, dy$ and $\beta = x\, dy$. Then

$$\alpha \wedge \beta = (x^2\, dx + y\, dy) \wedge (x\, dy) = x^3\, dx \wedge dy \;+\; yx\, dy \wedge dy = x^3\, dx \wedge dy,$$

since $dy \wedge dy = 0$. The result is a 2–form (an area density) weighted by x^3.

The wedge product satisfies the basic algebraic laws you would expect of a multiplication, with a twist in the commutativity that reflects orientation. However, we must introduce a key lemma.

Lemma 5.1 (Alternation identities on a vector space). Let V be a real vector space and $T^k(V^*)$ the space of k–tensors on V. Write $\mathrm{Alt}_k : T^k(V^*) \to \mathrm{Alt}^k(V^*)$ for the alternation (skew–symmetrization)

$$\mathrm{Alt}_k(\tau)(v_1, \ldots, v_k) := \frac{1}{k!} \sum_{\sigma \in S_k} \mathrm{sgn}(\sigma)\, \tau(v_{\sigma(1)}, \ldots, v_{\sigma(k)}).$$

Then:

1. **Projection/idempotence:** Alt_k is a projection onto the alternating tensors: $\mathrm{Alt}_k \circ \mathrm{Alt}_k = \mathrm{Alt}_k$, and $\mathrm{Alt}_k(\tau) = \tau$ if τ is alternating.

2. **Block alternation:** For $p, q \geq 0$ and $\alpha \in T^p(V^*)$, $\beta \in T^q(V^*)$,

$$\mathrm{Alt}_{p+q}\big(\mathrm{Alt}_p(\alpha) \otimes \beta\big) = \mathrm{Alt}_{p+q}(\alpha \otimes \beta) = \mathrm{Alt}_{p+q}\big(\alpha \otimes \mathrm{Alt}_q(\beta)\big).$$

3. **Block swap sign:** Let $\tau \in S_{p+q}$ be the permutation that moves the first p slots behind the

next q slots (i.e. $1, \ldots, p, p{+}1, \ldots, p{+}q \mapsto p{+}1, \ldots, p{+}q, 1, \ldots, p$). Then $\operatorname{sgn}(\tau) = (-1)^{pq}$ and

$$\operatorname{Alt}_{p+q}(\alpha \otimes \beta) = (-1)^{pq} \operatorname{Alt}_{p+q}(\beta \otimes \alpha).$$

Proof. (1) By definition Alt_k averages over all signed permutations, so applying it twice just reproduces the same average: $\operatorname{Alt}_k \circ \operatorname{Alt}_k = \operatorname{Alt}_k$. If τ is already alternating, permuting arguments only adds signs, hence $\operatorname{Alt}_k(\tau) = \tau$.

(2) In the sum defining Alt_{p+q}, the subgroup $S_p \times S_q \subset S_{p+q}$ permutes the first p and last q positions independently. Alternating in the first p slots (or the last q) before alternating in all $p{+}q$ slots does not change the final average, hence the equalities.

(3) The permutation τ is the product of pq transpositions (move each of the p entries past q entries), so $\operatorname{sgn}(\tau) = (-1)^{pq}$. Since Alt_{p+q} averages over all permutations, precomposing with τ pulls out this sign and swaps the two blocks, giving the identity. ■

Definition 5.6 (Wedge product via alternation)**.** For $\omega \in \operatorname{Alt}^p(V^*)$ and $\eta \in \operatorname{Alt}^q(V^*)$ set

$$\omega \wedge \eta \ := \ \operatorname{Alt}_{p+q}(\omega \otimes \eta) \in \operatorname{Alt}^{p+q}(V^*).$$

For a smooth manifold M, define $(\omega \wedge \eta)_p$ at each $p \in M$ by the above construction on T_pM, and note smoothness from smooth local coefficients.

Proposition 5.1 (Fundamental properties of the wedge)**.** Let $\omega \in \Omega^p(M)$, $\eta \in \Omega^q(M)$, $\theta \in \Omega^r(M)$. Then

1. **Bilinearity:** $(a\,\omega_1 + b\,\omega_2) \wedge \eta = a\,\omega_1 \wedge \eta + b\,\omega_2 \wedge \eta$, and similarly in the second slot.
2. **Degree additivity:** $\omega \wedge \eta \in \Omega^{p+q}(M)$.
3. **Associativity:** $(\omega \wedge \eta) \wedge \theta = \omega \wedge (\eta \wedge \theta)$.
4. **Graded commutativity:** $\omega \wedge \eta = (-1)^{pq}\, \eta \wedge \omega$.
5. **Naturality under pullback:** for a smooth $f : M \to N$, $f^*(\omega \wedge \eta) = f^*\omega \wedge f^*\eta$.

Proof. All identities are pointwise, so fix $p \in M$ and work in $V = T_pM$.

(1) Bilinearity follows because both the tensor product and Alt are linear in each slot:

$$(a\omega_1+b\omega_2)\wedge\eta = \mathrm{Alt}\big((a\omega_1+b\omega_2)\otimes\eta\big) = a\ \mathrm{Alt}(\omega_1\otimes\eta) + b\ \mathrm{Alt}(\omega_2\otimes\eta).$$

(2) By definition $\omega\wedge\eta = \mathrm{Alt}_{p+q}(\omega\otimes\eta)$ is alternating of degree $p+q$, hence lies in $\mathrm{Alt}^{p+q}(V^*)$.

(3) Associativity:

$$(\omega\wedge\eta)\wedge\theta = \mathrm{Alt}_{p+q+r}\big((\omega\wedge\eta)\otimes\theta\big) = \mathrm{Alt}_{p+q+r}\big(\mathrm{Alt}_{p+q}(\omega\otimes\eta)\otimes\theta\big).$$

By Lemma (2) with the block sizes $(p+q,r)$, $\mathrm{Alt}_{p+q+r}\big(\mathrm{Alt}_{p+q}(\omega\otimes\eta)\otimes\theta\big) = \mathrm{Alt}_{p+q+r}(\omega\otimes\eta\otimes\theta)$. Similarly,

$$\omega\wedge(\eta\wedge\theta) = \mathrm{Alt}_{p+q+r}\big(\omega\otimes\mathrm{Alt}_{q+r}(\eta\otimes\theta)\big) = \mathrm{Alt}_{p+q+r}(\omega\otimes\eta\otimes\theta).$$

Thus the two sides agree.

(4) Graded commutativity: Using Lemma (3) for the block swap permutation,

$$\omega\wedge\eta = \mathrm{Alt}_{p+q}(\omega\otimes\eta) = (-1)^{pq}\ \mathrm{Alt}_{p+q}(\eta\otimes\omega) = (-1)^{pq}\,\eta\wedge\omega.$$

(5) Naturality: At $p\in M$ and $v_1,\dots,v_{p+q}\in T_pM$,

$$\big(f^*(\omega\wedge\eta)\big)_p(v_1,\dots,v_{p+q}) = (\omega\wedge\eta)_{f(p)}\big(df_pv_1,\dots,df_pv_{p+q}\big).$$

By definition and Lemma (2),

$$(\omega\wedge\eta)_{f(p)} = \mathrm{Alt}_{p+q}\big(\omega_{f(p)}\otimes\eta_{f(p)}\big),$$

so the RHS equals

$$\mathrm{Alt}_{p+q}\Big(\omega_{f(p)}\circ(df_p)^{\otimes p}\ \otimes\ \eta_{f(p)}\circ(df_p)^{\otimes q}\Big)(v_1,\dots,v_{p+q}) = \big((f^*\omega)\wedge(f^*\eta)\big)_p(v_1,\dots,v_{p+q}),$$

which gives the claimed identity. ■

Two practical coordinate rules are worth recording.

Proposition 5.2 (Coordinate rules)**.** Fix local coordinates $(x^1, \dots, x^n)$.

1. $dx^{i_1} \wedge \cdots \wedge dx^{i_p} = 0$ if two indices repeat.
2. If $I = (i_1 < \cdots < i_p)$ and $J = (j_1 < \cdots < j_q)$ have disjoint images, then

$$dx^{i_1} \wedge \cdots \wedge dx^{i_p} \wedge dx^{j_1} \wedge \cdots \wedge dx^{j_q} = \epsilon(I, J)\, dx^{k_1} \wedge \cdots \wedge dx^{k_{p+q}},$$

where $(k_1 < \cdots < k_{p+q})$ is the ordered merge of I and J and $\epsilon(I, J) \in \{\pm 1\}$ is the sign of the permutation that reorders the concatenation (I, J) into increasing order.

Example 5.5 (From lengths to areas to volumes)**.** On $\mathbb{R}^3$ with coordinates (x, y, z), the 2–form $dx \wedge dy$ measures oriented area in the xy–plane, $dy \wedge dz$ does so in the yz–plane, and $dz \wedge dx$ in the zx–plane. The 3–form $dx \wedge dy \wedge dz$ measures oriented volume; if you reorder its factors, the sign keeps track of orientation.

Remark 5.1 (Analogy with the cup product)**.** Earlier we built a product on singular cohomology, the cup product $H^p(X) \smile H^q(X) \to H^{p+q}(X)$, which was bilinear, natural under pullback, and *graded commutative*. The wedge product on differential forms has exactly the same formal features. Later, de Rham's theorem will identify cohomology classes of closed forms with singular cohomology classes, and under that identification the wedge corresponds to the cup:

$$[\omega] \smile [\eta] \quad \longleftrightarrow \quad [\omega \wedge \eta].$$

You can already feel the parallel in the "front part/back part" picture for the cup product and the "concatenate and alternate" rule for the wedge.

5.3 The Exterior Derivative

At this point, we have built up a new language for talking about smooth geometry: *differential forms*. These are objects that can be evaluated on tangent vectors at each point, and which can be combined using the wedge product to produce higher–degree forms. We have seen how 0–forms (smooth functions) measure a single vector at a time through their derivative, and how 1–forms can take in a direction and output a number describing the rate of change of functions along that direction. We have also seen that k–forms measure oriented k–dimensional "parallelepipeds" in the tangent space, and that the wedge product provides a natural way to combine these measurements into higher–dimensional ones.

But up to now, our operations have been purely algebraic: we can multiply forms with $\wedge$, we can add them, and we can scale them. There is still no way to *differentiate* a k–form to produce a $(k+1)$–form, generalising the way df turns a function into a 1–form. If such a process exists, it would let us measure the infinitesimal *change* of a form across the boundary of a region, just as the derivative of a function measures its infinitesimal change along a curve.

This idea should feel familiar. In cohomology, we had the coboundary operator δ, which took an n–cochain (a measurement on n–chains) and produced an $(n+1)$–cochain that recorded the net measurement along the boundary. A natural guess is that for smooth forms we can do the same: there should be an operator

$$d : \Omega^k(M) \to \Omega^{k+1}(M)$$

that plays the role of δ, increasing degree by one and somehow encoding "boundary data".

If this is to work, we must first decide what properties such an operator *must* satisfy. On 0–forms (functions), we already know what it must do: df should be the usual differential from multivariable calculus. Next, since we can multiply forms via $\wedge$, we want d to interact with this product in the same way δ interacted with the cup product: a graded Leibniz rule, with a sign depending on the degree. Finally, if d is really a smooth counterpart to δ, then just as $\delta^2 = 0$, we must have $d^2 = 0$: the "boundary of a boundary is zero" principle.

Surprisingly, these simple demands force d to be unique if it exists, and they also tell us exactly how to compute it in coordinates. Let us write them down clearly.

Definition 5.7 (Exterior derivative)**.** Let M be a smooth manifold. The *exterior derivative* is a family of $\mathbb{R}$–linear maps

$$d : \Omega^k(M) \longrightarrow \Omega^{k+1}(M), \qquad k \geq 0,$$

satisfying:

(i) (**Agreement on functions**) If $f \in \Omega^0(M) = C^\infty(M)$, then df is the usual differential:

$$(df)_p(v) = v(f), \quad \text{for } p \in M,\ v \in T_pM.$$

(ii) (**Graded Leibniz rule**) If $\omega \in \Omega^p(M)$ and $\eta \in \Omega^q(M)$, then

$$d(\omega \wedge \eta) = d\omega \wedge \eta + (-1)^p \, \omega \wedge d\eta.$$

(iii) (**Nilpotence**) $d \circ d = 0$ on all forms.

The point of these axioms is that they do not merely describe d — they *characterise* it completely. We now prove that such an operator exists, is unique, and has a simple local formula.

Theorem 5.1 (Existence and uniqueness of the exterior derivative)**.** Let M be a smooth manifold. There exists a unique family of $\mathbb{R}$–linear maps

$$d : \Omega^k(M) \to \Omega^{k+1}(M) \quad (k \geq 0)$$

satisfying *(i)*, *(ii)*, and *(iii)*. In any local coordinate chart $(U; x^1, \dots, x^n)$ the action of d is given by the explicit formula

$$d\left(\sum_{I=(i_1<\cdots<i_k)} \omega_I(x) \, dx^{i_1} \wedge \cdots \wedge dx^{i_k} \right) = \sum_I \sum_{j=1}^n \frac{\partial \omega_I}{\partial x^j}(x) \, dx^j \wedge dx^{i_1} \wedge \cdots \wedge dx^{i_k}.$$

Proof. We proceed in two parts.

Uniqueness on a coordinate patch. Let $D : \Omega^\bullet(U) \to \Omega^{\bullet+1}(U)$ satisfy *(i)–(iii)* on an open set U with coordinates $(x^1, \dots, x^n)$. We claim that for every multiindex $I = (i_1 < \cdots < i_k)$ and smooth function f,

$$D\big(f \, dx^{i_1} \wedge \cdots \wedge dx^{i_k}\big) \;=\; (Df) \wedge dx^{i_1} \wedge \cdots \wedge dx^{i_k} \;=\; df \wedge dx^{i_1} \wedge \cdots \wedge dx^{i_k}. \tag{1}$$

Indeed, by *(i)* we have $Df = df$. Next, observe that $dx^j = d(x^j)$ by definition of the standard differential on functions, and therefore

$$D(dx^j) = D(d(x^j)) = D^2(x^j) = 0 \qquad \text{by } \textit{(iii)}.$$

Using the graded Leibniz rule *(ii)* repeatedly, we obtain

$$D\big(f \, dx^{i_1} \wedge \cdots \wedge dx^{i_k}\big) = (Df) \wedge dx^{i_1} \wedge \cdots \wedge dx^{i_k} + (-1)^0 f \, D(dx^{i_1} \wedge \cdots \wedge dx^{i_k}).$$

Expanding the second term again by *(ii)* and using $D(dx^{i_s}) = 0$ for each s shows it vanishes. Thus (1) holds. Since every k–form on U has a unique expression $\sum_I f_I \, dx^{i_1} \wedge \cdots \wedge dx^{i_k}$, linearity yields the pointwise formula

$$D\left(\sum_I f_I \, dx^{i_1} \wedge \cdots \wedge dx^{i_k}\right) = \sum_I df_I \wedge dx^{i_1} \wedge \cdots \wedge dx^{i_k},$$

so D is uniquely determined on U.

Existence on $\mathbb{R}^n$. On an open set $U \subset \mathbb{R}^n$ with standard coordinates, *define* an operator $d : \Omega^\bullet(U) \to \Omega^{\bullet+1}(U)$ by the recipe just forced by uniqueness: for a simple monomial $f \, dx^{i_1} \wedge \cdots \wedge dx^{i_k}$ set

$$d\big(f \, dx^{i_1} \wedge \cdots \wedge dx^{i_k}\big) := df \wedge dx^{i_1} \wedge \cdots \wedge dx^{i_k},$$

and extend $\mathbb{R}$–linearly. We verify *(i)–(iii)*.

(i) Agreement on functions. If $k = 0$ the rule reads $d(f) = df$, i.e. the usual differential

$$df = \sum_{j=1}^{n} \frac{\partial f}{\partial x^j} \, dx^j.$$

(ii) Graded Leibniz rule. It suffices to check the rule on simple forms and extend by bilinearity. Write $\omega = f \, \alpha$, $\eta = g \, \beta$ with $f, g \in C^\infty(U)$ and α, β wedge products of coordinate 1–forms. Then

$$\omega \wedge \eta = (fg) \, (\alpha \wedge \beta),$$

so by definition

$$d(\omega \wedge \eta) = d(fg) \wedge \alpha \wedge \beta = (f \, dg + g \, df) \wedge \alpha \wedge \beta = d\omega \wedge \eta + (-1)^{\deg \omega} \, \omega \wedge d\eta,$$

because $d\omega = df \wedge \alpha$ and $d\eta = dg \wedge \beta$, and moving df past the $\deg \alpha = \deg \omega$ one–forms in α produces the sign $(-1)^{\deg \omega}$. (The general case follows by linearity.)

(iii) Nilpotence. Again check on monomials. For $\omega = f \, dx^{i_1} \wedge \cdots \wedge dx^{i_k}$,

$$d\omega = df \wedge dx^{i_1} \wedge \cdots \wedge dx^{i_k}.$$

Then

$$d^2\omega = d(df) \wedge dx^{i_1} \wedge \cdots \wedge dx^{i_k} - df \wedge d\big(dx^{i_1} \wedge \cdots \wedge dx^{i_k}\big),$$

where the minus sign is $(-1)^{\deg df} = (-1)^1$. By our definition $d(dx^{i_s}) = 0$, hence $d(dx^{i_1} \wedge \cdots \wedge dx^{i_k}) = 0$. Meanwhile

$$d(df) = d\left(\sum_j \frac{\partial f}{\partial x^j}\, dx^j\right) = \sum_{j,\ell} \frac{\partial^2 f}{\partial x^\ell\, \partial x^j}\, dx^\ell \wedge dx^j = \sum_{\ell<j} \left(\frac{\partial^2 f}{\partial x^\ell\, \partial x^j} - \frac{\partial^2 f}{\partial x^j\, \partial x^\ell}\right) dx^\ell \wedge dx^j = 0,$$

by equality of mixed partials. Thus $d^2\omega = 0$ for monomials, and hence for all forms.

Therefore d satisfies *(i)–(iii)* on $U \subset \mathbb{R}^n$.

Gluing and global uniqueness. Let $\{(U_\alpha; x^1_\alpha, \dots, x^n_\alpha)\}$ be a smooth atlas on M. On each U_α we have just defined an operator d_α with *(i)–(iii)*. We must check that on overlaps $U_\alpha \cap U_\beta$ these definitions agree, so they glue to a global d.

Fix an overlap $W = U_\alpha \cap U_\beta$. Both d_α and d_β satisfy *(i)–(iii)* on W, and they agree on functions: for $f \in C^\infty(W)$, both give the usual differential df independent of coordinates. By the *uniqueness on a coordinate patch* argument (applied now to the open set W), an operator satisfying *(i)–(iii)* is uniquely determined by its action on functions. Hence $d_\alpha = d_\beta$ on W. The family $\{d_\alpha\}$ therefore glues to a globally defined $d : \Omega^\bullet(M) \to \Omega^{\bullet+1}(M)$ with the stated local formula.

Finally, if D is any other operator on M satisfying *(i)–(iii)*, then for each chart U_α the restrictions $D|_{U_\alpha}$ and $d|_{U_\alpha}$ both satisfy the axioms and agree on functions; by the uniqueness-on-a-patch argument they coincide on U_α, hence $D = d$ globally.

This proves existence and uniqueness. ■

Remark 5.2 (What the formula is saying)**.** In coordinates, a k–form is a linear combination of basic wedges $dx^{i_1} \wedge \cdots \wedge dx^{i_k}$ with smooth coefficients $\omega_I(x)$. The exterior derivative simply differentiates the *coefficients* and wedges in one more dx^j in front. The antisymmetry of $\wedge$ ensures signs are handled correctly, and the symmetry of second derivatives forces $d^2 = 0$.

Example 5.6. On $\mathbb{R}^2$ with coordinates (x, y),

$$d\big(f(x,y)\,dx\big) = \big(f_x\,dx + f_y\,dy\big) \wedge dx = f_y\,dy \wedge dx = -\,f_y\,dx \wedge dy,$$

and

$$d(x\,dy) = dx \wedge dy, \qquad d(y\,dx) = -\,dx \wedge dy.$$

We now establish the basic calculus of the exterior derivative $d : \Omega^k(M) \to \Omega^{k+1}(M)$. Keep the following mental picture throughout: d should (i) agree with the ordinary differential on functions, (ii) distribute over wedges with the graded sign, and (iii) square to zero. The proofs below make these expectations precise.

Proposition 5.3 (Linearity and degree shift). For each $k \geq 0$, the map $d : \Omega^k(M) \to \Omega^{k+1}(M)$ is $\mathbb{R}$-linear and raises degree by one.

Proof. Linearity: for $\omega, \eta \in \Omega^k(M)$ and $a, b \in \mathbb{R}$, the defining coordinate formula for d is linear in the coefficient functions and in partial derivatives, hence $d(a\omega + b\eta) = a\,d\omega + b\,d\eta$.

Degree: if locally $\omega = \sum_I \omega_I\,dx^{i_1} \wedge \cdots \wedge dx^{i_k}$, then

$$d\omega = \sum_{I,j} \frac{\partial \omega_I}{\partial x^j}\,dx^j \wedge dx^{i_1} \wedge \cdots \wedge dx^{i_k} \in \Omega^{k+1}(M).$$

■

Proposition 5.4 (Behaviour on functions and on 1-forms). Let $(U; x^1, \ldots, x^n)$ be a coordinate chart.

1. If $f \in \Omega^0(M) = C^\infty(M)$, then on U, $df = \sum_{j=1}^n \frac{\partial f}{\partial x^j}\,dx^j$.
2. If $\alpha = \sum_{j=1}^n a_j\,dx^j \in \Omega^1(M)$, then on U,

$$d\alpha = \sum_{1 \leq i < j \leq n} \left(\frac{\partial a_j}{\partial x^i} - \frac{\partial a_i}{\partial x^j} \right) dx^i \wedge dx^j.$$

   In particular, $d(dx^j) = 0$ for each j.

Proof. (1) is the very definition of df.

(2) Using the graded Leibniz rule (with $p = 0$ for functions) and bilinearity,

$$d\alpha = \sum_j d(a_j\,dx^j) = \sum_j da_j \wedge dx^j = \sum_{j,i} \frac{\partial a_j}{\partial x^i}\,dx^i \wedge dx^j.$$

Split the sum into $i < j$ and $i > j$, and use $dx^i \wedge dx^j = -dx^j \wedge dx^i$ to obtain

$$\sum_{i<j} \left(\frac{\partial a_j}{\partial x^i} - \frac{\partial a_i}{\partial x^j} \right) dx^i \wedge dx^j.$$

Finally, $d(dx^j) = d(d(x^j)) = 0$ by nilpotence (proved below) applied to the function x^j. ■

Proposition 5.5 (Coordinate expansion for general forms). On a chart $(U; x^1, \dots, x^n)$, every $\omega \in \Omega^k(U)$ can be written uniquely as $\omega = \sum_I \omega_I \, dx^{i_1} \wedge \cdots \wedge dx^{i_k}$, where $I = (i_1 < \cdots < i_k)$ ranges over increasing k-multiindices and $\omega_I \in C^\infty(U)$. Then

$$d\omega = \sum_I \sum_{j=1}^n \frac{\partial \omega_I}{\partial x^j} \, dx^j \wedge dx^{i_1} \wedge \cdots \wedge dx^{i_k}.$$

Proof. Write ω as stated. Using linearity and that $d(dx^{i_m}) = 0$ for each m,

$$d\omega = \sum_I d\omega_I \wedge dx^{i_1} \wedge \cdots \wedge dx^{i_k} = \sum_{I,j} \frac{\partial \omega_I}{\partial x^j} \, dx^j \wedge dx^{i_1} \wedge \cdots \wedge dx^{i_k}.$$

This is already in increasing-order form in the first wedge factor; if desired one can reorder to the canonical increasing multiindex with the appropriate sign, but the displayed expression is valid as it stands and transforms correctly on overlaps (see the next proposition). ■

Proposition 5.6 (Compatibility with restriction and coordinate changes). If $V \subset U \subset M$ are open, then for all $\omega \in \Omega^\bullet(U)$,

$$(d\omega)\big|_V = d\big(\omega\big|_V\big).$$

Equivalently, for an inclusion $\iota : V \hookrightarrow U$, $\iota^*(d\omega) = d(\iota^*\omega)$. Consequently the local coordinate formula for d above defines a global form, independent of chart.

Proof. Both $(d\omega)|_V$ and $d(\omega|_V)$ are computed in any chart on V from the same coefficient functions $\omega_I|_V$ by the same expression; hence they agree on V. For independence of chart, note that on chart overlaps both coordinate computations produce the same smooth $(k+1)$-form (they agree after restriction), so they glue to a global form. ■

Proposition 5.7 (Naturality under pullback). Let $f : M \to N$ be smooth. Then for every $\omega \in \Omega^\bullet(N)$,

$$f^*(d\omega) = d(f^*\omega).$$

Proof. Work locally on a chart $(V; y^1, \dots, y^m)$ of N. Any $\omega \in \Omega^k(V)$ can be written as $\omega = \sum_J h_J\, dy^{j_1} \wedge \cdots \wedge dy^{j_k}$, with $h_J \in C^\infty(V)$. We first check the identity on the algebra generators $h \in C^\infty(V)$ and $dy^i = d(y^i)$.

For a function h, $f^*(dh) = d(h \circ f) = d(f^*h)$ by the chain rule (the definition of df). Next, for a coordinate 1-form,

$$f^*(dy^i) = d(y^i \circ f) = d(f^*y^i),$$

again by the chain rule. Since pullback *commutes with wedge* and is $\mathbb{R}$-linear, for a simple monomial $h\, dy^{j_1} \wedge \cdots \wedge dy^{j_k}$ we have

$$\begin{aligned} f^*\big(d(h\, dy^J)\big) &= f^*\big(dh \wedge dy^J + h\, d(dy^J)\big) \qquad \text{(Leibniz and } d(dy^i) = 0\text{)} \\ &= d(f^*h) \wedge f^*(dy^J) \\ &= d\big(f^*(h\, dy^J)\big), \end{aligned}$$

where $dy^J = dy^{j_1} \wedge \cdots \wedge dy^{j_k}$. By linearity the equality holds for ω. Since the statement is local on N, the result globalizes. ∎

Proposition 5.8 (Nilpotence). For every $\omega \in \Omega^k(M)$, one has $d^2\omega = 0$.

Proof. The claim is local. In coordinates $(x^1, \dots, x^n)$, write $\omega = \sum_I \omega_I\, dx^{i_1} \wedge \cdots \wedge dx^{i_k}$. Then, by the coordinate formula,

$$d\omega = \sum_{I,j} \frac{\partial \omega_I}{\partial x^j}\, dx^j \wedge dx^{i_1} \wedge \cdots \wedge dx^{i_k}.$$

Applying d again and using bilinearity,

$$d^2\omega = \sum_{I,j,\ell} \frac{\partial^2 \omega_I}{\partial x^\ell\, \partial x^j}\, dx^\ell \wedge dx^j \wedge dx^{i_1} \wedge \cdots \wedge dx^{i_k}.$$

Now split the double sum in (j, ℓ) into ordered pairs with $j < \ell$ and $j > \ell$, and use antisymmetry

$dx^\ell \wedge dx^j = -dx^j \wedge dx^\ell$. Grouping the two contributions for each unordered pair $\{j, \ell\}$ gives

$$\frac{1}{2}\sum_{I,j,\ell}\left(\frac{\partial^2 \omega_I}{\partial x^\ell\,\partial x^j} - \frac{\partial^2 \omega_I}{\partial x^j\,\partial x^\ell}\right) dx^\ell \wedge dx^j \wedge dx^{i_1} \wedge \cdots \wedge dx^{i_k} = 0,$$

by equality of mixed partials. Hence $d^2\omega = 0$ on the chart, and therefore globally. ■

Proposition 5.9 (Product rule with functions)**.** For $f \in C^\infty(M)$ and $\omega \in \Omega^k(M)$,

$$d(f\,\omega) = df \wedge \omega + f\,d\omega.$$

Proof. This is the graded Leibniz rule with $p = 0$, so there is no sign. ■

Example 5.7. Take the following two computations for example

1. On $\mathbb{R}^2$, $\alpha = (x^2 + y)\,dx + (xy)\,dy$. Then

$$d\alpha = \left(\frac{\partial}{\partial x}(xy) - \frac{\partial}{\partial y}(x^2 + y)\right) dx \wedge dy = (y-1)\,dx \wedge dy.$$

2. On $\mathbb{R}^3$, for $\omega = P\,dy \wedge dz + Q\,dz \wedge dx + R\,dx \wedge dy$,

$$d\omega = \left(\frac{\partial P}{\partial x} + \frac{\partial Q}{\partial y} + \frac{\partial R}{\partial z}\right) dx \wedge dy \wedge dz.$$

These mirror familiar vector–calculus identities under the standard dictionary.

Remark 5.3 (Parallel with cohomology)**.** In singular cohomology, the coboundary δ took an n–cochain and produced an $(n+1)$–cochain by evaluating it on the boundary of a chain. Here, d takes a k–form and produces a $(k+1)$–form by differentiating its coefficients and wedging in an extra dx^j — a smooth analogue of the same idea. The axiom $d^2 = 0$ is the geometric incarnation of $\delta^2 = 0$.

5.3.1 Closed and Exact Forms

Up to this point we have studied the exterior derivative d as a geometric differentiation operator, built its formula, and proved its essential properties: linearity, degree shift, Leibniz rule, nilpotence, and naturality. But $d^2 = 0$ has a striking algebraic consequence that deserves its own language.

Observe: if $\omega = d\eta$ for some η, then $d\omega = d(d\eta) = 0$ automatically. Thus, *being the derivative of something* forces a form to have vanishing derivative. Forms with zero derivative are often called "curl-free" or "divergence-free" in vector calculus, depending on their degree. We now promote this into formal terminology.

Definition 5.8 (Closed and exact forms)**.** Let M be a smooth manifold and $k \geq 0$.

- A k–form $\omega \in \Omega^k(M)$ is called *closed* if $d\omega = 0$.
- A k–form $\omega \in \Omega^k(M)$ is called *exact* if there exists $\eta \in \Omega^{k-1}(M)$ such that $\omega = d\eta$.

By convention, every 0–form is closed iff it is locally constant, and there are no nonzero exact 0–forms except the zero function.

The motivation for these names is that closed forms have "no boundary" in the sense that their d–derivative vanishes, while exact forms are literally the d–derivative of some other form.

The following proposition is the basic relationship between the two notions.

Proposition 5.10 (Exact $\implies$ closed)**.** Every exact form is closed: if $\omega = d\eta$ for some η, then $d\omega = 0$.

Proof. Immediate from the nilpotence property $d^2 = 0$:

$$d\omega = d(d\eta) = 0.$$

■

Remark 5.4. The converse (closed $\implies$ exact) is *not* true on general manifolds — the obstruction is precisely what de Rham cohomology measures. However, in special domains, such as star-shaped open sets in $\mathbb{R}^n$, the converse *does* hold; this is Poincaré's Lemma, which we will prove later.

Before moving further, let us check that the space of closed forms is an algebraic subspace of $\Omega^k(M)$, and that exact forms form a subspace of closed forms.

Proposition 5.11 (Algebra of closed and exact forms). Fix $k \geq 0$. Then:

1. The set $Z^k(M) := \{\omega \in \Omega^k(M) \mid d\omega = 0\}$ of closed k–forms is a vector subspace of $\Omega^k(M)$.
2. The set $B^k(M) := \{\omega \in \Omega^k(M) \mid \omega = d\eta \text{ for some } \eta \in \Omega^{k-1}(M)\}$ of exact k–forms is a vector subspace of $\Omega^k(M)$.
3. $B^k(M) \subseteq Z^k(M)$.

Proof. (1) If ω_1, ω_2 are closed, $d(\omega_1 + \omega_2) = d\omega_1 + d\omega_2 = 0 + 0 = 0$, and for $a \in \mathbb{R}$, $d(a\omega_1) = a\, d\omega_1 = 0$.

(2) If $\omega_1 = d\eta_1$ and $\omega_2 = d\eta_2$ are exact, then for $a, b \in \mathbb{R}$,

$$a\omega_1 + b\omega_2 = a\, d\eta_1 + b\, d\eta_2 = d(a\eta_1 + b\eta_2),$$

so the combination is exact.

(3) Already proved: $d(d\eta) = 0$ for any η. ■

It is also important to know how closedness and exactness behave with respect to wedge products.

Proposition 5.12 (Wedge products of closed/exact forms). Let $\omega \in \Omega^p(M)$, $\eta \in \Omega^q(M)$.

1. If ω and η are closed, then $\omega \wedge \eta$ is closed.
2. If ω is exact and η is closed, then $\omega \wedge \eta$ is exact.

Proof. (1) If $d\omega = 0$ and $d\eta = 0$, then

$$d(\omega \wedge \eta) = d\omega \wedge \eta + (-1)^p\, \omega \wedge d\eta = 0 \wedge \eta + (-1)^p\, \omega \wedge 0 = 0,$$

so $\omega \wedge \eta$ is closed.

(2) If $\omega = d\alpha$ and $d\eta = 0$, then

$$\omega \wedge \eta = (d\alpha) \wedge \eta = d(\alpha \wedge \eta) - (-1)^{p-1}\, \alpha \wedge d\eta = d(\alpha \wedge \eta) - (-1)^{p-1}\, \alpha \wedge 0 = d(\alpha \wedge \eta),$$

which is exact. ■

Example 5.8 (Vector calculus analogy in $\mathbb{R}^3$)**.** If we identify 1–forms with vector fields via the Euclidean metric, closed 1–forms correspond to *curl-free* vector fields, and exact 1–forms correspond to *gradient* vector fields. The identity grad ∘ grad $= 0$ is the same as $d^2 = 0$, and $\operatorname{curl}(\operatorname{grad} f) = 0$ is just "exact $\implies$ closed" in disguise.

5.4 Poincaré's Lemma

Up to now we have seen that $d^2 = 0$ makes "being a derivative" (exact) imply "having zero derivative" (closed). The natural question is when the converse holds: *if a form has zero derivative, must it be a derivative of something*? In vector calculus you have met this as "curl–free ⇒ gradient" on nice domains. Poincaré's Lemma is the precise higher–degree statement in the smooth setting. The key geometric idea is that on a region which can be *linearly contracted* to a point, you can build a primitive by "integrating along the contraction." We now make that idea precise and carry out the full calculation.

Definition 5.9 (Star-shaped set)**.** An open set $U \subset \mathbb{R}^n$ is *star-shaped (with respect to* 0*)* if for every $x \in U$ and every $t \in [0,1]$ the point $t\,x$ lies in U. Equivalently, the straight line segment from 0 to x is contained in U.

The straight-line contraction $H : [0,1] \times U \to U$, $H(t,x) = t\,x$, is the tool that lets us "accumulate" a primitive. We will build an explicit linear operator $K : \Omega^k(U) \to \Omega^{k-1}(U)$ (for $k \geq 1$) such that

$$dK + Kd = \mathrm{id} \quad \text{on } \Omega^k(U)\ (k \geq 1).$$

When $d\omega = 0$, this identity reduces to $\omega = d(K\omega)$, producing a global primitive. The operator K is the rigorous version of "integrate along the radial paths from 0 to x."

Theorem 5.2 (Poincaré's Lemma on star-shaped domains)**.** Let $U \subset \mathbb{R}^n$ be star-shaped (with respect to 0). If $\omega \in \Omega^k(U)$ is closed and $k \geq 1$, then ω is exact. Equivalently, $H^k_{\mathrm{dR}}(U) = 0$ for all $k \geq 1$.

Proof. We give a completely explicit construction. Write the *radial vector field* $R(x) = \sum_{i=1}^n x_i \frac{\partial}{\partial x_i}$. For $k \geq 1$ define a linear operator $K : \Omega^k(U) \to \Omega^{k-1}(U)$ pointwise by

$$\boxed{(K\omega)_x(v_1, \ldots, v_{k-1}) \ :=\ \int_0^1 t^{\,k-1}\, \omega_{tx}\big(R_{tx},\, (dH_t)_x v_1,\, \ldots,\, (dH_t)_x v_{k-1}\big)\, dt,} \qquad (*)$$

where $H_t(x) = t\,x$ and $(dH_t)_x = t\,\mathrm{Id}$. Intuitively: at the point x we evaluate ω along the contracted point tx, feed it the radial vector R_{tx} (the velocity of $t \mapsto tx$), and transport the remaining arguments by the linear map $dH_t = t\,\mathrm{Id}$. The weight t^{k-1} is precisely the Jacobian factor ensuring the algebra works out.

Step 1: K is well-defined and smooth. For fixed x and vectors $v_1, \dots, v_{k-1}$, the integrand is smooth in t because ω is smooth and $t \mapsto tx$ is smooth; compactness of $[0,1]$ yields a smooth output in $(x; v_1, \dots, v_{k-1})$. Multilinearity and alternation are inherited from ω, hence $K\omega$ is a $(k-1)$–form.

Step 2: The homotopy identity $dK + Kd = \mathrm{id}$ on $\Omega^k(U)$ for $k \geq 1$. Fix $\omega \in \Omega^k(U)$ with $k \geq 1$ and $x \in U$. Consider the map

$$F(t) \;:=\; (H_t)^*\omega \;\in\; \Omega^k(U).$$

We compute its derivative using the chain rule for pullbacks along the flow H_t:

$$\frac{d}{dt}F(t) \;=\; \frac{d}{dt}(H_t)^*\omega \;=\; (H_t)^*\big(\mathcal{L}_R\omega\big),$$

where $\mathcal{L}_R$ denotes differentiation along the vector field R. (Indeed, $\frac{d}{dt}H_t(x) = R_{H_t(x)}$.) Now invoke the *Cartan identity* $\mathcal{L}_R\omega = d(\iota_R\omega) + \iota_R(d\omega)$, valid for all smooth forms and vector fields (it follows directly from the graded Leibniz rule for d and the definitions of pullback and contraction ι_R). We obtain

$$\frac{d}{dt}(H_t)^*\omega \;=\; (H_t)^*\big(d(\iota_R\omega)\big) \;+\; (H_t)^*\big(\iota_R(d\omega)\big) \;=\; d\big((H_t)^*\iota_R\omega\big) \;+\; (H_t)^*\iota_R(d\omega),$$

where we used that d commutes with pullback.

Integrate this identity from $t = 0$ to $t = 1$:

$$(H_1)^*\omega \;-\; (H_0)^*\omega \;=\; \int_0^1 d\big((H_t)^*\iota_R\omega\big)\,dt \;+\; \int_0^1 (H_t)^*\iota_R(d\omega)\,dt.$$

Note that $H_1 = \mathrm{id}_U$, so $(H_1)^*\omega = \omega$, while H_0 collapses U to 0, hence $(H_0)^*\omega = 0$ for $k \geq 1$ (a k–form pulled back along a constant map is zero). Thus

$$\omega \;=\; d\bigg(\int_0^1 (H_t)^*\iota_R\omega\,dt\bigg) \;+\; \bigg(\int_0^1 (H_t)^*\iota_R(d\omega)\,dt\bigg). \tag{†}$$

Finally, a straightforward comparison of $(*)$ with the definition of pullback shows that

$$\int_0^1 (H_t)^*\iota_R\omega\,dt \;=\; K\omega \qquad \text{and} \qquad \int_0^1 (H_t)^*\iota_R(d\omega)\,dt \;=\; K(d\omega).$$

Plugging these into (†) yields the homotopy identity

$$\boxed{d(K\omega) \;+\; K(d\omega) \;=\; \omega \qquad (k \geq 1).}$$

Step 3: Closed ⇒ exact. If $d\omega = 0$, the identity becomes $\omega = d(K\omega)$, so ω is exact with primitive $K\omega$.

This completes the proof. ■

Remark 5.5 (Coordinate formula for K)**.** If $\omega = \sum_{|I|=k} f_I(x)\,dx_I$ with $dx_I = dx_{i_1} \wedge \cdots \wedge dx_{i_k}$, then

$$(K\omega)(x) \;=\; \sum_{|I|=k} \sum_{j=1}^{k} (-1)^{j-1} \left(\int_0^1 t^{k-1}\, x_{i_j}\, f_I(tx)\, dt \right) dx_{i_1} \wedge \cdots \widehat{dx_{i_j}} \cdots \wedge dx_{i_k}.$$

This is exactly the "integrate coefficients along radial segments" recipe, and one can verify $dK + Kd = \mathrm{id}$ directly from this formula by differentiating under the integral sign and using the product rule.

Example 5.9 (The case $k = 1$ revisited: curl–free ⇒ gradient)**.** Let $\alpha = \sum_{i=1}^{n} a_i(x)\, dx_i$ be a 1–form on a star-shaped U with $d\alpha = 0$, which is the classical symmetry condition $\partial a_i/\partial x_j = \partial a_j/\partial x_i$. The operator K produces a potential

$$f(x) \;=\; \int_0^1 \alpha_{tx}(x)\, dt \;=\; \int_0^1 \sum_{i=1}^{n} a_i(tx)\, x_i\, dt,$$

and one checks directly that $df = \alpha$. Thus a curl–free vector field on a star-shaped domain is a gradient field, the familiar fact from multivariable calculus.

Corollary 5.1 (Contractible manifolds)**.** If M is a smooth manifold that is *contractible* (there exists a smooth homotopy $H : [0,1] \times M \to M$ from id_M to a constant map), then every closed k–form on M is exact for $k \geq 1$; in particular $H^k_{\mathrm{dR}}(M) = 0$ for $k \geq 1$.

Heuristic. Cover M by star-shaped coordinate neighborhoods; by the lemma, closed forms are locally exact. A standard partition-of-unity/gluing argument then produces a global primitive on a contractible M. Alternatively, the proof of Theorem 5.2 adapts verbatim to any smooth contraction H, replacing $H_t(x) = tx$ and R by the homotopy H and its velocity vector field $\partial_t H$; the same calculation yields $dK + Kd = \mathrm{id}$. ■

Corollary 5.2 (Closed $\Rightarrow$ exact in $\mathbb{R}^n$)**.** Every smooth closed k–form on $\mathbb{R}^n$ with $k \geq 1$ is exact. In other words,

$$H^k_{\mathrm{dR}}(\mathbb{R}^n) \;=\; 0 \quad \text{for all } k \geq 1.$$

Proof. $\mathbb{R}^n$ is star–shaped with respect to the origin. Applying Poincaré's Lemma to $U = \mathbb{R}^n$ gives the conclusion directly. ∎

Example 5.10. A constant 1–form $\alpha = a_1\,dx_1 + \cdots + a_n\,dx_n$ is closed, since all its coefficients have zero derivatives. By the corollary, $\alpha = df$ for some f. Indeed, one can take

$$f(x) = a_1x_1 + \cdots + a_nx_n.$$

Similarly, a 2–form with constant coefficients is d of a 1–form whose coefficients are linear functions of x.

Remark 5.6 (Exact forms are closed: $d^2 = 0$ revisited)**.** From our earlier work we already know that $d \circ d = 0$ on all forms. We also know that *every exact form is automatically closed.*

Poincaré's Lemma provides a satisfying converse on contractible sets: there, closed $\Rightarrow$ exact, so we have a perfect equivalence between "closed" and "exact" forms. This is precisely the d–analogue of the fact in singular cohomology that, on a contractible space, every cocycle is a coboundary.

5.5 Generalised Stokes' Theorem

Up to this point, our story has had two parallel threads.

On the *algebraic topology* side, we began with chains, boundaries, and homology groups. We then turned the picture around and looked at cochains, coboundaries, and cohomology. There, the coboundary operator δ had a simple but profound meaning: it measured how a cochain on a region was determined by its values on the *boundary* of that region.

On the *smooth* side, we have just built the language of differential forms: smoothly varying multilinear functionals that can be wedged together, pulled back along smooth maps, and differentiated using the exterior derivative d. Here too, d raises degree by one and satisfies $d^2 = 0$. Already, the parallel is impossible to miss:

$$\delta \quad \longleftrightarrow \quad d.$$

The generalised Stokes theorem is the moment when these two threads meet: it says, in effect, that integrating $d\omega$ over a region is exactly the same as integrating ω over the boundary of that region. In other words, d *is* the geometric incarnation of the cohomological coboundary.

Seen from this angle, Stokes' theorem is not a bolt from the blue, but the inevitable geometric realisation of the algebraic principles we already know. If ω is a $(n-1)$–form on an n–dimensional space, it is the smooth analogue of an $(n-1)$–cochain. Integrating ω over the boundary ∂M is the smooth analogue of evaluating a cochain on a boundary chain. And integrating $d\omega$ over M itself is the analogue of evaluating the coboundary $\delta\omega$ on the chain M.

To make this vivid, let us recall some familiar special cases, each of which is a different face of the same principle.

- The **Fundamental Theorem of Calculus** states that for a smooth function F on an interval $[a, b]$,

$$\int_a^b F'(x)\,dx = F(b) - F(a).$$

  This is nothing more than Stokes' theorem for 0–forms on a 1–dimensional manifold: $dF = F'\,dx$ is a 1–form, and the boundary of $[a, b]$ is $\{b\} - \{a\}$.

- In $\mathbb{R}^2$, **Green's theorem** says:

$$\oint_{\partial R} P\,dx + Q\,dy = \iint_R \left(\frac{\partial Q}{\partial x} - \frac{\partial P}{\partial y} \right) dx\,dy.$$

  Here $P\,dx + Q\,dy$ is a 1–form, its exterior derivative is $d(P\,dx + Q\,dy) = \left(\frac{\partial Q}{\partial x} - \frac{\partial P}{\partial y} \right) dx \wedge dy$, and the theorem asserts that integrating $d\omega$ over the region is the same as integrating ω over its boundary.

- In $\mathbb{R}^3$, the **Divergence theorem** says:

$$\iint_{\partial V} \mathbf{F} \cdot \mathbf{n}\,dS = \iiint_V (\nabla \cdot \mathbf{F})\,dV.$$

  If we regard $\mathbf{F}$ as a 2–form via the Hodge star, this is again the same pattern.

Each time, we see the same structure:

$$\int_{\partial(\text{region})} \omega \quad = \quad \int_{\text{region}} d\omega.$$

The power of the generalised Stokes theorem is that it works on *any* oriented smooth manifold, not just subsets of Euclidean space, and for *any* degree of form, not just those corresponding to vector fields or functions.

Remark 5.7. From the cohomological perspective, Stokes' theorem is the statement that integration provides a pairing between de Rham cohomology and singular homology that is perfectly compatible with the coboundary–boundary duality. From the geometric perspective, it says that d computes the infinitesimal change of a quantity, and integrating over a region "accumulates" this change into a net flux across the boundary.

Before we can state the theorem in its final form, we must carefully set up:

1. what it means to integrate a differential form over a manifold (possibly with boundary),
2. how to orient boundaries consistently with the interior,
3. and how to interpret the boundary ∂M of a general smooth manifold M.

We turn to these points now.

Definition 5.10 (Manifold with boundary)**.** An n–dimensional smooth manifold with boundary is a space in which every point has a neighbourhood diffeomorphic to an open set in the closed half-space

$$\mathbb{H}^n = \{(x_1, \dots, x_n) \in \mathbb{R}^n \mid x_n \geq 0\}.$$

The *boundary* ∂M consists of those points mapping to $\{x_n = 0\}$ under some (and hence any) such chart.

Example 5.11. The closed unit disk $D^2 \subset \mathbb{R}^2$ is a 2–manifold with boundary $\partial D^2 = S^1$. The cylinder $S^1 \times [0, 1]$ is a 2–manifold with boundary consisting of two disjoint circles $S^1 \times \{0\}$ and $S^1 \times \{1\}$.

Definition 5.11 (Oriented manifold with boundary and boundary orientation). Let M be a smooth n–manifold with (possibly empty) boundary ∂M. An *orientation* on M is a choice of orientation class on each tangent space that varies smoothly. For $p \in \partial M$ there is a canonical orientation on ∂M, called the *boundary orientation*: a basis $(v_1, \ldots, v_{n-1})$ of $T_p(\partial M)$ is positively oriented if and only if $(\nu_p, v_1, \ldots, v_{n-1})$ is a positively oriented basis of T_pM, where ν_p is a *chosen outward pointing* normal to ∂M in M at p.

Example 5.12. In the interval $[a, b] \subset \mathbb{R}$ with the standard orientation, the boundary $\{a, b\}$ inherits the orientation in which b is counted positively and a negatively. In the unit disk D^2 with counterclockwise orientation, the induced orientation on $\partial D^2 = S^1$ is the counterclockwise one you expect.

Definition 5.12 (Integration of top forms and compatibility with pullback). If M is an oriented n–manifold (with or without boundary) and $\omega \in \Omega^n(M)$ has compact support, define $\int_M \omega$ by choosing an oriented atlas $\{(U_\alpha, \varphi_\alpha)\}$ with $\varphi_\alpha : U_\alpha \to V_\alpha \subset \mathbb{R}^n$ and a partition of unity $\{\rho_\alpha\}$ subordinate to $\{U_\alpha\}$:

$$\int_M \omega \;:=\; \sum_\alpha \int_{V_\alpha} (\varphi_\alpha^{-1})^*(\rho_\alpha\, \omega).$$

This is independent of choices. If $f : N \to M$ is an orientation-preserving diffeomorphism and $\eta \in \Omega^n(M)$ has compact support, then $\displaystyle\int_N f^*\eta = \int_M \eta$.

Remark 5.8. If ω is not of top degree ($\deg \omega < n$), it cannot be integrated over all of M, but it *can* be integrated over a k–dimensional oriented submanifold $S \subset M$ with $k = \deg \omega$. In that case we use the inclusion map $i : S \hookrightarrow M$ and integrate the pullback $i^*\omega$ over S.

Example 5.13. Let $M = \mathbb{R}^3$ with its standard orientation and let S be the unit circle in the xy–plane oriented counterclockwise. If $\omega = x\,dy$ is a 1–form on $\mathbb{R}^3$, then $i^*\omega$ on S is the 1–form $\cos t\, d(\sin t) = \cos^2 t\, dt$ in the usual parametrisation $t \mapsto (\cos t, \sin t, 0)$. Integrating over S means integrating this pullback over $t \in [0, 2\pi]$.

Finally, note the definition of compact support.

Definition 5.13 (Compact support). Let X be a topological space and $f : X \to \mathbb{R}$ a continuous function. We say that f has *compact support* if the closure of the set

$$\mathrm{supp}(f) \ := \ \overline{\{x \in X \mid f(x) \neq 0\}}$$

is a compact subset of X. Similarly, a differential form ω on a smooth manifold M is said to have compact support if $\mathrm{supp}(\omega)$ is compact in M.

Compact support intuitively means that the object "vanishes outside a bounded region" in M and does so in a way that leaves no stray limit points. This is crucial in integration theory: it ensures the integral is finite even on non-compact domains and allows localisation arguments to work cleanly without convergence issues.

We now have all the geometric ingredients to state Stokes' theorem in full generality. The only input is a differential form ω of degree $n-1$ on an oriented n–manifold M with boundary. On the left-hand side of the formula will be the integral of $d\omega$ over the interior, and on the right-hand side the integral of ω over the boundary with its induced orientation.

Theorem 5.3 (Generalized Stokes' Theorem). Let M be a compact, oriented smooth n–manifold with boundary ∂M, and let $\omega \in \Omega^{n-1}(M)$. Then

$$\int_M d\omega \ = \ \int_{\partial M} \iota^*\omega,$$

where $\iota : \partial M \hookrightarrow M$ is the inclusion and the integral on ∂M uses the boundary orientation.

We prove the theorem by reducing to the elementary case of regions in $\mathbb{R}^n$ via a sequence of local statements. For clarity, we isolate these as lemmas.

Lemma 5.3.1 (Local normal form near the boundary). For every $p \in \partial M$ there exists a coordinate chart $\varphi : U \to W \subset \mathbb{H}^n := \{x \in \mathbb{R}^n : \ x_n \geq 0\}$ such that $\varphi(U \cap \partial M) = W \cap \{x_n = 0\}$, φ is an orientation-preserving diffeomorphism onto its image, and the boundary orientation on $U \cap \partial M$ corresponds to the standard orientation on $\{x_n = 0\} \cong \mathbb{R}^{n-1}$ (i.e. the one for which $(-\partial_{x_n}, e_1, \ldots, e_{n-1})$ is positively oriented in $\mathbb{R}^n$).

Proof. This is the standard *collar chart* construction. Choose any boundary chart sending p to 0 with image in $\mathbb{H}^n$ and such that ∂M goes to $\{x_n = 0\}$. Precompose with an orientation-preserving linear change of variables (if necessary) so that the differential at p carries an outward normal to $-\partial_{x_n}$. Because the outward normal at the boundary points into the region $x_n < 0$, the rule from the definition ensures the induced boundary orientation is the standard one on $\{x_n = 0\}$. Shrinking the chart if needed yields the claim. ■

Lemma 5.3.2 (Stokes on a rectangular box in $\mathbb{R}^n$)**.** Let $Q = \prod_{i=1}^n [a_i, b_i] \subset \mathbb{R}^n$ with the standard orientation, and let $\alpha \in \Omega^{n-1}(Q)$ be smooth. Then

$$\int_Q d\alpha \;=\; \int_{\partial Q} \alpha,$$

where ∂Q is oriented by the outward normal first convention.

Proof. Write any $(n-1)$–form on Q uniquely as

$$\alpha = \sum_{i=1}^{n} (-1)^{i-1} A_i(x)\, dx^1 \wedge \cdots \wedge \widehat{dx^i} \wedge \cdots \wedge dx^n,$$

with smooth coefficient functions $A_i : Q \to \mathbb{R}$. A direct computation gives

$$d\alpha \;=\; \sum_{i=1}^{n} \frac{\partial A_i}{\partial x_i}(x)\, dx^1 \wedge \cdots \wedge dx^n.$$

Thus

$$\int_Q d\alpha \;=\; \sum_{i=1}^{n} \int_Q \frac{\partial A_i}{\partial x_i}(x)\, dx \;=\; \sum_{i=1}^{n} \left(\int_{\prod_{j\neq i}[a_j,b_j]} \int_{a_i}^{b_i} \frac{\partial A_i}{\partial x_i}\, dx_i\, dx_{\hat{i}} \right),$$

where dx abbreviates $dx^1 \cdots dx^n$ and $dx_{\hat{i}}$ is the product over $j \neq i$. By the Fundamental Theorem of Calculus,

$$\int_{a_i}^{b_i} \frac{\partial A_i}{\partial x_i}(x)\, dx_i = A_i(x_1, \ldots, x_{i-1}, b_i, x_{i+1}, \ldots, x_n) - A_i(x_1, \ldots, x_{i-1}, a_i, x_{i+1}, \ldots, x_n).$$

Hence

$$\int_Q d\alpha = \sum_{i=1}^{n} \left(\int_{\prod_{j\neq i}[a_j,b_j]} A_i(\cdots, b_i, \cdots)\, dx_{\hat{i}} - \int_{\prod_{j\neq i}[a_j,b_j]} A_i(\cdots, a_i, \cdots)\, dx_{\hat{i}} \right).$$

Now compare with the boundary integral. The face F_i^+ given by $x_i = b_i$ inherits orientation so that (ν, basis of TF_i^+) matches the orientation of Q, with outward normal $\nu = +\partial_{x_i}$. The induced oriented volume element on F_i^+ is $dx^1 \wedge \cdots \wedge \widehat{dx^i} \wedge \cdots \wedge dx^n$, and $\alpha|_{F_i^+}$ restricts to $(-1)^{i-1} A_i(\cdots, b_i, \cdots)\, dx^1 \wedge \cdots \wedge \widehat{dx^i} \wedge \cdots \wedge dx^n$. Thus

$$\int_{F_i^+} \alpha \;=\; \int_{\prod_{j\neq i}[a_j,b_j]} A_i(\cdots, b_i, \cdots)\, dx_{\hat{i}}.$$

Similarly, for the face F_i^- given by $x_i = a_i$, the outward normal is $-\partial_{x_i}$, and the induced orientation contributes a minus sign:

$$\int_{F_i^-} \alpha \;=\; -\int_{\prod_{j\neq i}[a_j,b_j]} A_i(\cdots, a_i, \cdots)\, dx_{\hat{i}}.$$

Summing over all faces,

$$\int_{\partial Q} \alpha = \sum_{i=1}^{n} \left(\int_{F_i^+} \alpha + \int_{F_i^-} \alpha \right) = \sum_{i=1}^{n} \Big(\int_{\prod_{j\neq i}[a_j,b_j]} A_i(\cdots, b_i, \cdots)\, dx_{\hat{i}} - \int_{\prod_{j\neq i}[a_j,b_j]} A_i(\cdots, a_i, \cdots)\, dx_{\hat{i}} \Big),$$

which matches $\int_Q d\alpha$ obtained above. This proves the lemma. ■

Lemma 5.3.3 (Change of variables for top forms)**.** Let $U, V \subset \mathbb{R}^n$ be open, $\Phi : U \to V$ a diffeomorphism, and $\eta \in \Omega^n(V)$ compactly supported. If Φ is orientation-preserving, then

$$\int_U \Phi^*\eta \;=\; \int_V \eta.$$

If Φ is orientation-reversing, the integral acquires a minus sign.

Proof. Write $\eta = f(y)\, dy^1 \wedge \cdots \wedge dy^n$ in local coordinates on V with f compactly supported. Then $\Phi^*\eta = (f \circ \Phi)\, \det(D\Phi)\, dx^1 \wedge \cdots \wedge dx^n$, and the claim is exactly the classical multidimensional change-of-variables formula from advanced calculus. The sign is the sign of $\det(D\Phi)$. ■

Lemma 5.3.4 (Local-to-global via partition of unity)**.** Let M be as in the theorem. Suppose there exists an open cover $\{U_\alpha\}$ of M such that for every α and every $\omega \in \Omega^{n-1}(M)$ with $\mathrm{supp}(\omega) \subset U_\alpha$ we have

$$\int_M d\omega = \int_{\partial M} \iota^*\omega.$$

Then the equality holds for all $\omega \in \Omega^{n-1}(M)$.

Proof. Choose a smooth partition of unity $\{\rho_\alpha\}$ subordinate to $\{U_\alpha\}$. Because d is linear and obeys the Leibniz rule,

$$d(\rho_\alpha\,\omega) = d\rho_\alpha \wedge \omega + \rho_\alpha\, d\omega.$$

Summing over α and using $\sum_\alpha \rho_\alpha \equiv 1$,

$$\sum_\alpha d(\rho_\alpha\,\omega) = \Big(\sum_\alpha d\rho_\alpha\Big) \wedge \omega + \Big(\sum_\alpha \rho_\alpha\Big) d\omega = 0 \wedge \omega + d\omega = d\omega.$$

Integrate over M and use linearity:

$$\int_M d\omega \;=\; \sum_\alpha \int_M d(\rho_\alpha\,\omega).$$

By hypothesis each $\rho_\alpha\,\omega$ is supported in U_α, hence $\displaystyle\int_M d(\rho_\alpha\,\omega) = \int_{\partial M} \iota^*(\rho_\alpha\,\omega)$. Summing and using ι^* linear,

$$\int_M d\omega = \sum_\alpha \int_{\partial M} \iota^*(\rho_\alpha\,\omega) = \int_{\partial M} \iota^*\Big(\sum_\alpha \rho_\alpha\,\omega\Big) = \int_{\partial M} \iota^*\omega.$$

∎

Lemma 5.3.5 (Straightening a single chart to a half–box)**.** Let (U,φ) be a boundary chart as in Lemma 5.3.1, with $\varphi : U \to W \subset \mathbb{H}^n$. Let $K \subset U$ be compact and $\omega \in \Omega^{n-1}(M)$ with $\mathrm{supp}(\omega) \subset K \subset U$. Then there exists a bounded Lipschitz domain $D \subset \mathbb{H}^n$ with piecewise smooth boundary, and an orientation-preserving diffeomorphism $\Psi : W \to D$ such that

$$\int_U d\omega \;=\; \int_D d\big(\Psi^*\varphi^*\omega\big), \qquad \int_{\partial M\cap U} \iota^*\omega \;=\; \int_{\partial D\cap\{x_n=0\}} \iota_0^*\big(\Psi^*\varphi^*\omega\big),$$

where $\iota_0 : \partial D \cap \{x_n = 0\} \hookrightarrow D$ is the inclusion with the induced boundary orientation.

Proof. Since $\mathrm{supp}(\omega)$ is compact in U, we may choose a compact $K \Subset U$ such that $\mathrm{supp}(\omega) \subset K$

and $\varphi(K) \subset W$ is compact in W. By standard smoothing/straightening (Whitney's theorem or direct polygonal approximation), there exists a bounded domain $D \subset \mathbb{H}^n$ with piecewise smooth boundary and a diffeomorphism $\Psi : W \to D$ that is orientation-preserving and equals the identity on a neighbourhood of $\varphi(K)$ (so that supports are unaffected). Then, using Lemma 5.3.3 twice (for φ and Ψ) and the naturality of pullback,

$$\int_U d\omega = \int_W d(\varphi^*\omega) = \int_D d(\Psi^*\varphi^*\omega).$$

Likewise, ∂U is sent to $\{x_n = 0\}$ by φ with boundary orientation matching the standard one (Lemma 5.3.1), and Ψ preserves orientation, so change of variables on the $(n-1)$–dimensional boundary gives

$$\int_{\partial M \cap U} \iota^*\omega = \int_{\{x_n=0\}\cap W} \iota_0^*(\varphi^*\omega) = \int_{\{x_n=0\}\cap D} \iota_0^*(\Psi^*\varphi^*\omega).$$

■

Proof of the Theorem. By Lemma 5.3.4 it suffices to prove the identity for $(n-1)$–forms whose support is contained in a single boundary chart or a single interior chart.

Interior support. If $\mathrm{supp}(\omega) \subset U$ where U is a chart mapped diffeomorphically to an open set $W \subset \mathbb{R}^n$, then by change of variables (Lemma 5.3.3)

$$\int_M d\omega = \int_U d\omega = \int_W d(\varphi^*\omega).$$

But ∂M does not meet U, so $\iota^*\omega \equiv 0$ on ∂M and the right-hand side is 0. On the other hand, the integral of $d(\varphi^*\omega)$ over W vanishes by a standard cutoff argument (expand support slightly to a box and use Stokes on the box together with cancellation on the artificial boundary); alternatively, this is the special case of Lemma 5.3.2 where the boundary integral is zero because ω vanishes near the (artificial) boundary. Hence $\int_M d\omega = 0 = \int_{\partial M} \iota^*\omega$.

Boundary-supported case. Assume $\mathrm{supp}(\omega) \subset U$ for a boundary chart (U, φ) as in Lemma 5.3.1. By Lemma 5.3.5 we may replace (U, ω) by a domain $D \subset \mathbb{H}^n$ and a form $\widetilde{\omega} := \Psi^*\varphi^*\omega$ supported in D, so that

$$\int_U d\omega = \int_D d\widetilde{\omega}, \qquad \int_{\partial M \cap U} \iota^*\omega = \int_{\partial D \cap \{x_n=0\}} \iota_0^*\widetilde{\omega}.$$

Now cover the compact set $\mathrm{supp}(\widetilde{\omega}) \subset D$ by finitely many closed rectangular boxes lying inside D and meeting the boundary only along the plane $\{x_n = 0\}$, with pairwise overlaps of measure zero (a finite grid suffices). Apply Lemma 5.3.2 on each box and sum; interior faces cancel in pairs (their orientations are opposite), leaving only the contributions from the portions of the boundary lying in $\{x_n = 0\}$. Precisely,

$$\int_D d\widetilde{\omega} \;=\; \int_{\partial D\cap\{x_n>0\}} \widetilde{\omega} \;+\; \int_{\partial D\cap\{x_n=0\}} \widetilde{\omega}.$$

But our boxes are chosen so that $\widetilde{\omega}$ has support disjoint from the artificial faces in $\{x_n > 0\}$, hence their contribution vanishes and we obtain

$$\int_D d\widetilde{\omega} \;=\; \int_{\partial D\cap\{x_n=0\}} \iota_0^*\widetilde{\omega}.$$

Finally, undo the straightening via Lemma 5.3.5:

$$\int_U d\omega \;=\; \int_{\partial M\cap U} \iota^*\omega.$$

Combining the interior and boundary-supported cases and invoking Lemma 5.3.4 with a partition of unity $\{\rho_\alpha\}$ subordinate to a finite atlas adapted to ∂M, we conclude

$$\int_M d\omega = \int_{\partial M} \iota^*\omega$$

for every $\omega \in \Omega^{n-1}(M)$, as claimed. ■

Remark 5.9 (On the notation for the boundary term). In the formal statement of the theorem, the term on the right is written as

$$\int_{\partial M} \iota^*\omega,$$

where $\iota : \partial M \hookrightarrow M$ is the inclusion map. This is because ω is defined on M, while the integral is over ∂M, so one must first *pull back* ω to the boundary before integrating. In many contexts this pullback is not written explicitly, and the theorem is presented as

$$\int_M d\omega \;=\; \int_{\partial M} \omega,$$

with the understanding that the form on ∂M is obtained by restricting ω to tangent vectors lying in the boundary. Both notations express the same content; the formal one simply makes the restriction explicit.

The generalized Stokes theorem now stands before us in its full generality. It is worth pausing to take stock of what we have actually achieved.

At the outset of this section, we motivated the theorem by recalling three classical results: the Fundamental Theorem of Calculus, Green's theorem, and the divergence theorem. We now see clearly that these are not separate facts, but rather special cases of a single, elegant principle: the integral of a derivative over a region equals the integral of the original object over the region's boundary. In each case, the "derivative" and the "object" live in the appropriate space of differential forms, and the Stokes formula is the bridge between them.

One striking feature is how little the proof depended on the global geometry of M. The argument was local: we worked in coordinate charts and pieced the result together using a partition of unity. Yet the conclusion is global, relating an n–dimensional integral over all of M to an $(n-1)$–dimensional integral over its entire boundary. This interplay of local and global is a hallmark of modern geometry and topology.

We should also stress the role of orientation. The formula

$$\int_M d\omega \;=\; \int_{\partial M} \iota^*\omega$$

depends crucially on the convention that the boundary inherits the *outward normal first* orientation. Reversing the orientation of M reverses the sign of both integrals, while reversing the boundary orientation alone negates the right-hand side. Many sign errors in vector calculus stem from neglecting this point.

Another important corollary concerns manifolds with empty boundary. If $\partial M = \varnothing$, Stokes' theorem immediately gives

$$\int_M d\omega = 0$$

for every $(n-1)$–form ω on M. In other words, exact n–forms always have zero total integral over a closed manifold. This fact will reappear later in a cohomological guise: integration over a closed manifold depends only on the cohomology class of the top-degree form.

While we have stated and proved the theorem for compact oriented manifolds with smooth boundary, the statement extends in various directions. One can allow manifolds with corners, non-compact manifolds with forms of compact support, and more general domains in Euclidean space. The proof adapts with minor modifications in each case.

Finally, it is worth foreshadowing where we go from here. The operator d we have been using is, from the algebraic point of view, a coboundary map on the complex of differential forms:

$$0 \longrightarrow \Omega^0(M) \xrightarrow{d} \Omega^1(M) \xrightarrow{d} \Omega^2(M) \xrightarrow{d} \cdots$$

The identity $d^2 = 0$ means we can speak of closed forms and exact forms, and take their quotient to form cohomology groups. In the next section, we will see that these *de Rham cohomology groups* are topological invariants of M, and that Stokes' theorem is precisely the statement that integration of forms pairs naturally with the homology of M. This will complete the bridge from our earlier study of singular cohomology to the analytic language of differential forms.

6 De Rham Cohomology

Up to now, our story has run along two intertwining threads.

On one side, we have built *cohomology* from a purely algebraic recipe: start with a graded collection of objects, apply a coboundary operator δ satisfying $\delta^2 = 0$, and take the quotient of things killed by δ (the cocycles) by those coming from something in a lower degree (the coboundaries). This construction was entirely combinatorial — our cochains were abstract functions on simplices, their meaning coming only from the algebraic structure.

On the other side, our detour into smooth manifolds brought us into the realm of *differential forms*, objects that live directly on the manifold and know how to measure infinitesimal pieces of geometry. Here too we encountered a graded collection $(\Omega^\bullet(M), \wedge)$ and a distinguished operator, the exterior derivative d, which — intriguingly — satisfies $d^2 = 0$.

The resemblance is not an accident. In fact, d behaves exactly like a coboundary operator: it takes an n–form and produces an $(n+1)$–form, and applying it twice always gives zero. The generalized Stokes theorem already hinted that exact forms (those of the form $d\eta$) vanish when integrated over cycles, and that closed forms (those with $d\omega = 0$) are the natural ones to integrate. If you felt a déjà vu when hearing those terms, you should: this is precisely the closed/coboundary dichotomy from cohomology theory, but now in the smooth, analytic world.

At this point the mathematician's instinct is clear: if the algebraic pattern is the same, let us build the same quotient as before, but with *forms* instead of cochains. The result will be a new cohomology theory, one that speaks the language of calculus on manifolds. We call it *de Rham cohomology.*

Why bother? Because the punchline, proved by de Rham in the 1930s, is astonishing: for smooth manifolds, this analytic cohomology is *isomorphic* to singular cohomology with real coefficients. Two completely different worlds — one discrete and combinatorial, the other smooth and analytic — turn out to encode exactly the same topological information. This bridge allows us to use theorems from one side to solve problems on the other, and it will eventually unify the entire theory we've been developing.

To set the stage, here is the analogy to keep in mind:

Singular cohomology	de Rham cohomology
Cochains $C^n(X)$	n-forms $\Omega^n(M)$
$\delta : C^n \to C^{n+1}$	$d : \Omega^n \to \Omega^{n+1}$
$\delta^2 = 0$	$d^2 = 0$
$\ker \delta / \operatorname{im} \delta$	$\ker d / \operatorname{im} d$

The plan is simple: we will define de Rham cohomology in direct parallel to the singular case, then explore its properties, compute examples, and finally prove its equivalence to singular cohomology — thus closing one of the most elegant circles in mathematics.

6.1 Definitions

We already have the blueprint from singular cohomology: take our graded collection of objects, apply the special differential d that satisfies $d^2 = 0$, and look at the kernel modulo the image. The only change is that now our objects are not abstract cochains but smooth differential forms.

Definition 6.1 (de Rham cohomology)**.** Let M be a smooth manifold. The space of *closed* n–forms is

$$Z^n_{\mathrm{dR}}(M) \;:=\; \{\, \omega \in \Omega^n(M) \mid d\omega = 0 \,\}.$$

The space of *exact* n–forms is

$$B^n_{\mathrm{dR}}(M) \;:=\; \{\, \omega \in \Omega^n(M) \mid \omega = d\eta \text{ for some } \eta \in \Omega^{n-1}(M) \,\}.$$

The *nth de Rham cohomology group* of M is the quotient

$$H^n_{\mathrm{dR}}(M) \;:=\; Z^n_{\mathrm{dR}}(M) \,/\, B^n_{\mathrm{dR}}(M).$$

Exactly as before, two closed n–forms represent the same class if their difference is exact. The group $H^n_{\mathrm{dR}}(M)$ measures the obstruction to solving $d\eta = \omega$ for a given closed ω.

Remark 6.1 (Parallel with singular cohomology)**.** If you mentally replace $\Omega^n(M)$ with $C^n(X)$ and d with δ, this definition is identical to that of singular cohomology with real coefficients. The difference is that the de Rham version lives entirely in the smooth world: forms can be integrated, multiplied by the wedge product, and manipulated with calculus. Later, de Rham's theorem will tell us that $H^\bullet_{\mathrm{dR}}(M) \cong H^\bullet(M;\mathbb{R})$.

Before we launch into theorems, let us build intuition through examples.

Example 6.1 ($\mathbb{R}^n$)**.** By Poincaré's lemma, every closed form on $\mathbb{R}^n$ is exact. Therefore

$$H^0_{\mathrm{dR}}(\mathbb{R}^n) \cong \mathbb{R} \quad \text{and} \quad H^k_{\mathrm{dR}}(\mathbb{R}^n) = 0 \text{ for } k \geq 1.$$

Geometrically: $\mathbb{R}^n$ has one connected component and no holes of any dimension.

Example 6.2 (S^1)**.** A closed 0–form is a locally constant function, hence constant on S^1. So $H^0_{\mathrm{dR}}(S^1) \cong \mathbb{R}$.

For degree 1, consider $\omega = d\theta$ on the unit circle (here θ is the angular coordinate). It is closed: $d\omega = d^2\theta = 0$. If $\omega = df$ were exact, f would have to be single-valued, but θ itself is not well-defined globally on S^1. Thus ω is *not* exact, and in fact $H^1_{\mathrm{dR}}(S^1) \cong \mathbb{R}$, generated by the class of $d\theta$. All higher-degree cohomology vanishes. This matches the familiar picture: S^1 has one 1–dimensional hole and nothing else.

Example 6.3 (The torus $T^2 = S^1 \times S^1$)**.** Let θ_1, θ_2 be the angular coordinates on the two S^1 factors. Then $d\theta_1$ and $d\theta_2$ are closed 1–forms which are not exact. They form a basis for $H^1_{\mathrm{dR}}(T^2) \cong \mathbb{R}^2$.

The wedge $d\theta_1 \wedge d\theta_2$ is a closed 2–form, and it generates $H^2_{\mathrm{dR}}(T^2) \cong \mathbb{R}$. Again the cohomology perfectly reflects the geometry: two independent 1–cycles and one 2–dimensional fundamental class.

Already in these basic cases, the pattern is clear: de Rham cohomology counts the "independent closed forms modulo those that are d of something smaller," and the result mirrors the hole-counting intuition from singular cohomology, but now in a smooth, analytic setting.

The definition of $H^n_{\mathrm{dR}}(M)$ hides a familiar pattern: just as singular cohomology came from the coboundary operator δ on cochains, de Rham cohomology comes from the exterior derivative d acting on differential forms. Let us make that structure explicit.

Definition 6.2 (The de Rham complex)**.** For a smooth manifold M, consider the graded vector space of differential forms

$$\Omega^\bullet(M) := \bigoplus_{k\geq 0} \Omega^k(M),$$

equipped with the exterior derivative $d : \Omega^k(M) \to \Omega^{k+1}(M)$. The *de Rham complex* of M is the cochain complex

$$0 \longrightarrow \Omega^0(M) \xrightarrow{d} \Omega^1(M) \xrightarrow{d} \cdots \xrightarrow{d} \Omega^k(M) \xrightarrow{d} \Omega^{k+1}(M) \xrightarrow{d} \cdots \longrightarrow 0.$$

Proposition 6.1 ($d^2 = 0$ and the cochain condition)**.** For every $k \geq 0$ we have $d \circ d = 0 : \Omega^k(M) \to \Omega^{k+2}(M)$. Hence $(\Omega^\bullet(M), d)$ is a cochain complex.

Proof. This was established when we constructed d (Leibniz rule, locality, and $d^2 = 0$). Briefly: in a chart, write a k–form as a linear combination of wedge products of coordinate 1–forms with smooth coefficient functions. Applying d twice produces only second mixed partials, which cancel by symmetry of mixed derivatives and antisymmetry of the wedge, giving $d^2 = 0$. ■

Remark 6.2 (Cycles, boundaries, and the cohomology of the complex)**.** With $d^2 = 0$, each inclusion $\mathrm{Im}(d : \Omega^{k-1} \to \Omega^k) \subseteq \ker(d : \Omega^k \to \Omega^{k+1})$ holds. Thus the closed and exact forms we defined above are exactly the k*–cocycles* and k*–coboundaries* of this cochain complex, and

$$H^k_{\mathrm{dR}}(M) = \frac{\ker\big(d : \Omega^k(M) \to \Omega^{k+1}(M)\big)}{\mathrm{Im}\big(d : \Omega^{k-1}(M) \to \Omega^k(M)\big)}$$

is the cohomology of the de Rham complex in degree k. This perfectly parallels the singular cohomology quotient Z^k/B^k with δ in place of d.

Example 6.4 (Seeing the arrows in low degrees)**.** On $\mathbb{R}^2$ with coordinates (x, y), the beginning of the de Rham complex reads

$$0 \to \Omega^0(\mathbb{R}^2) \xrightarrow{d} \Omega^1(\mathbb{R}^2) \xrightarrow{d} \Omega^2(\mathbb{R}^2) \to 0,$$

that is,

$$f \overset{d}{\longmapsto} f_x\,dx + f_y\,dy \overset{d}{\longmapsto} (f_{xy} - f_{yx})\,dx \wedge dy \;=\; 0.$$

So every 1–form of the form df is automatically closed, and by Poincaré's lemma on $\mathbb{R}^2$ every closed 1–form is of this form locally (indeed globally on $\mathbb{R}^2$).

Example 6.5 (A circle snapshot)**.** For S^1 with angular coordinate θ,

$$0 \to \Omega^0(S^1) \xrightarrow{d} \Omega^1(S^1) \xrightarrow{d} 0, \quad \text{so} \quad f(\theta) \overset{d}{\longmapsto} f'(\theta)\,d\theta \overset{d}{\longmapsto} 0.$$

Here $d\theta$ is closed but not exact, and it represents a nonzero class in $H^1_{\mathrm{dR}}(S^1)$.

6.2 Functoriality

In singular cohomology, functoriality meant that a continuous map $f\colon X \to Y$ gave a pullback on cochains, $f^*\colon C^\bullet(Y) \to C^\bullet(X)$, which commuted with the coboundary operator δ. The slogan was: *"chains push forward, cochains pull back"*. For de Rham cohomology, the slogan is exactly the same — only now the "cochains" are differential forms, and the coboundary is the exterior derivative d.

Let us recall the basic geometric picture. A differential form on N is a rule that assigns a multilinear, alternating "measurement" to each point of N, acting on tangent vectors at that point. If $f\colon M \to N$ is smooth, and ω is a form on N, the only natural way to make a form on M from it is to *pull back* those measurements along f: at $p \in M$, apply ω not to tangent vectors of M directly, but to their images under the differential df_p in $T_{f(p)}N$.

This leads to the familiar pullback of forms:

Definition 6.3 (Pullback of a differential form)**.** Let $f\colon M \to N$ be smooth and let $\omega \in \Omega^k(N)$. The *pullback* $f^*\omega \in \Omega^k(M)$ is defined pointwise by

$$(f^*\omega)_p(v_1, \ldots, v_k) := \omega_{f(p)}\big(df_p(v_1), \ldots, df_p(v_k)\big),$$

for all $p \in M$ and $v_1, \ldots, v_k \in T_pM$.

Proposition 6.2 (Basic properties of the pullback). Let $f \colon M \to N$ be smooth, and let $\omega, \eta \in \Omega^\bullet(N)$. Then:

1. f^* is degree-preserving: if ω is a k–form, so is $f^*\omega$.
2. f^* is $\mathbb{R}$–linear: $f^*(a\omega + b\eta) = a\, f^*\omega + b\, f^*\eta$.
3. f^* respects the wedge product: $f^*(\omega \wedge \eta) = f^*\omega \wedge f^*\eta$.
4. f^* commutes with the exterior derivative: $f^*(d\omega) = d(f^*\omega)$.

Proof. (1) and (2) follow immediately from the definition: f^* does not change the number of arguments the form takes, and linearity is preserved pointwise.

(3) For $p \in M$ and tangent vectors $v_1, \dots, v_{k+\ell} \in T_pM$, by definition of the wedge product we have

$$\begin{aligned}
(f^*(\omega \wedge \eta))_p(v_1, \dots, v_{k+\ell}) &= (\omega \wedge \eta)_{f(p)}\big(df_p v_1, \dots, df_p v_{k+\ell}\big) \\
&= \frac{1}{k!\,\ell!} \sum_{\sigma \in S_{k+\ell}} \operatorname{sgn}(\sigma)\; \omega_{f(p)}(df_p v_{\sigma(1)}, \dots, df_p v_{\sigma(k)}) \\
&\quad \cdot \eta_{f(p)}(df_p v_{\sigma(k+1)}, \dots, df_p v_{\sigma(k+\ell)}) \\
&= (f^*\omega \wedge f^*\eta)_p(v_1, \dots, v_{k+\ell}),
\end{aligned}$$

since df_p is linear and the same permutation sum appears in both definitions.

(4) For $p \in M$ and $v_0, \dots, v_k \in T_pM$,

$$\begin{aligned}
(f^*(d\omega))_p(v_0, \dots, v_k) &= (d\omega)_{f(p)}(df_p v_0, \dots, df_p v_k) \\
&= \sum_{i=0}^{k} (-1)^i\, df_p v_i \big[\omega_{f(p)}(df_p v_0, \dots, \widehat{df_p v_i}, \dots, df_p v_k)\big] \\
&\quad + \sum_{0 \le i < j \le k} (-1)^{i+j}\, \omega_{f(p)}\big([df_p v_i, df_p v_j], df_p v_0, \dots, \widehat{df_p v_i}, \dots, \widehat{df_p v_j}, \dots, df_p v_k\big).
\end{aligned}$$

Since df_p is a linear map of tangent spaces and respects the Lie bracket of vector fields ($df_p\,[X, Y] = [df_p X, df_p Y]$), each term above matches the corresponding term in

$$(d(f^*\omega))_p(v_0, \dots, v_k),$$

showing $f^*d\omega = df^*\omega$. ■

The last property is the crucial one: it says the pullback f^* is a *cochain map* from the de Rham complex of N to that of M:

$$\begin{array}{ccccccc} \cdots & \xrightarrow{d} & \Omega^k(N) & \xrightarrow{d} & \Omega^{k+1}(N) & \xrightarrow{d} & \cdots \\ & & \downarrow f^* & & \downarrow f^* & & \\ \cdots & \xrightarrow{d} & \Omega^k(M) & \xrightarrow{d} & \Omega^{k+1}(M) & \xrightarrow{d} & \cdots \end{array}$$

Since f^* sends closed forms to closed forms and exact forms to exact forms, it descends to cohomology:

Corollary 6.1 (Functoriality of de Rham cohomology)**.** A smooth map $f\colon M \to N$ induces a graded ring homomorphism

$$f^*\colon H^\bullet_{\mathrm{dR}}(N) \longrightarrow H^\bullet_{\mathrm{dR}}(M)$$

by $[\,\omega\,] \mapsto [\,f^*\omega\,]$.

Remark 6.3 (Parallel with singular cohomology)**.** In singular cohomology, a continuous map $f\colon X \to Y$ also induces a pullback on cochains $f^*\colon C^\bullet(Y) \to C^\bullet(X)$ commuting with δ. The exterior derivative d here plays the role of δ, and $\Omega^\bullet(M)$ plays the role of $C^\bullet(M)$. Thus functoriality of $H^\bullet_{\mathrm{dR}}$ is not an extra miracle but a built-in feature of the cochain-complex viewpoint.

6.3 Mayer-Vietoris Sequence

Up to this point, we have defined de Rham cohomology, explored its functoriality, and hinted at its potential for computation. But as in singular cohomology, the raw definitions alone are not always computationally friendly: given a complicated manifold, how can we hope to find its cohomology groups directly from the definition?

One of the most powerful tools in this situation is the *Mayer–Vietoris sequence*. In singular cohomology, it arose from gluing information about two overlapping subspaces. The same philosophy applies here: if we can cover a manifold by two simpler open sets, then we can compute the cohomology of the whole by relating it to the cohomologies of the pieces and their intersection.

Let us begin by setting up the scene. Suppose M is a smooth manifold, and let $U, V \subset M$ be open subsets such that $M = U \cup V$. We know that $\Omega^*(M)$, $\Omega^*(U)$, $\Omega^*(V)$, and $\Omega^*(U \cap V)$ are all differential graded algebras under the wedge product and exterior derivative. Moreover, the inclusions

$$\iota_U : U \hookrightarrow M, \quad \iota_V : V \hookrightarrow M, \quad \iota_{UV} : U \cap V \hookrightarrow U, \quad \iota_{VU} : U \cap V \hookrightarrow V$$

induce pullbacks of forms

$$\iota_U^* : \Omega^*(M) \to \Omega^*(U), \quad \iota_V^* : \Omega^*(M) \to \Omega^*(V), \quad \iota_{UV}^* : \Omega^*(U) \to \Omega^*(U \cap V), \quad \iota_{VU}^* : \Omega^*(V) \to \Omega^*(U \cap V).$$

Definition 6.4 (Mayer–Vietoris cochain complex)**.** Define the map

$$\Phi : \Omega^*(M) \longrightarrow \Omega^*(U) \oplus \Omega^*(V), \quad \Phi(\omega) = (\omega|_U, \omega|_V),$$

and the map

$$\Psi : \Omega^*(U) \oplus \Omega^*(V) \longrightarrow \Omega^*(U \cap V), \quad \Psi(\alpha, \beta) = \alpha|_{U\cap V} - \beta|_{U\cap V}.$$

Intuitively, Φ restricts a global form to the two pieces U and V, while Ψ compares these local forms on their overlap. The minus sign is not an arbitrary choice: it reflects the alternating nature of Čech-type constructions and is crucial for exactness.

A key fact — proved exactly as in the singular case, but now in the category of differential forms — is:

Lemma 6.1. The sequence of cochain complexes

$$0 \longrightarrow \Omega^*(M) \xrightarrow{\Phi} \Omega^*(U) \oplus \Omega^*(V) \xrightarrow{\Psi} \Omega^*(U \cap V) \longrightarrow 0$$

is exact in each degree, and each map is a morphism of cochain complexes, i.e. it commutes with d.

Proof. Exactness is straightforward to check:

- Injectivity of Φ: If $\omega|_U = 0$ and $\omega|_V = 0$, then ω vanishes on $M = U \cup V$, so $\omega = 0$.
- $\mathrm{Im}(\Phi) = \ker(\Psi)$: If $\omega \in \Omega^*(M)$, then its restrictions agree on $U \cap V$, so $\Psi(\Phi(\omega)) = 0$. Conversely, if (α, β) satisfies $\alpha|_{U\cap V} = \beta|_{U\cap V}$, then we can glue α and β to obtain a global $\omega \in \Omega^*(M)$ agreeing with each on its domain, by smooth partition of unity.

- Surjectivity of Ψ: Given $\gamma \in \Omega^*(U \cap V)$, choose a smooth partition of unity $\{\rho_U, \rho_V\}$ subordinate to $\{U, V\}$ and define

$$\alpha := \rho_V \gamma \quad \text{on } U, \qquad \beta := -\rho_U \gamma \quad \text{on } V.$$

  Then $\alpha|_{U\cap V} - \beta|_{U\cap V} = \gamma$.

Commutation with d follows from the fact that restriction of forms and d are compatible: $d(\omega|_U) = (d\omega)|_U$. ■

With this short exact sequence of cochain complexes in hand, the machinery from algebraic topology tells us that we obtain a long exact sequence in cohomology.

Theorem 6.1 (Mayer–Vietoris sequence for de Rham cohomology)**.** Let M be a smooth manifold, $U, V \subset M$ open with $M = U \cup V$. Then there is a natural long exact sequence

$$\cdots \longrightarrow H^{k-1}_{\mathrm{dR}}(U \cap V) \xrightarrow{\delta} H^k_{\mathrm{dR}}(M) \xrightarrow{(\mathrm{res}_U, \mathrm{res}_V)} H^k_{\mathrm{dR}}(U) \oplus H^k_{\mathrm{dR}}(V) \xrightarrow{\mathrm{diff}} H^k_{\mathrm{dR}}(U \cap V) \longrightarrow \cdots$$

where the maps are induced from Φ and Ψ above, and δ is the connecting homomorphism from the short exact sequence of complexes.

The proof follows the standard snake-lemma construction: since the short exact sequence of cochain complexes is natural and commutes with d, passing to cohomology yields the long exact sequence. The connecting homomorphism δ is constructed exactly as in the singular case, by taking a cohomology class in $H^{k-1}(U \cap V)$, lifting it to a cochain in $\Omega^{k-1}(U) \oplus \Omega^{k-1}(V)$, applying d to get a cocycle in $\Omega^k(M)$, and checking well-definedness.

Example 6.6 (A warm-up computation)**.** Let $M = S^1$ and choose U, V to be open arcs covering the circle, with $U \cap V$ consisting of two disjoint arcs (homotopy equivalent to two points). We have

$$H^0_{\mathrm{dR}}(U) \cong \mathbb{R}, \quad H^0_{\mathrm{dR}}(V) \cong \mathbb{R}, \quad H^0_{\mathrm{dR}}(U \cap V) \cong \mathbb{R} \oplus \mathbb{R},$$

and all higher groups vanish on these contractible opens. The Mayer–Vietoris sequence then reduces to

$$0 \to H^0_{\mathrm{dR}}(S^1) \to \mathbb{R} \oplus \mathbb{R} \xrightarrow{\mathrm{diff}} \mathbb{R} \oplus \mathbb{R} \xrightarrow{\delta} H^1_{\mathrm{dR}}(S^1) \to 0,$$

from which one reads off $H^1_{\mathrm{dR}}(S^1) \cong \mathbb{R}$, recovering our earlier computation by a different method.

Remark 6.4 (Comparison with singular cohomology). The Mayer–Vietoris sequence in de Rham cohomology is formally identical to that in singular cohomology. The only differences are the underlying cochain complexes and the fact that we now glue and restrict *smooth forms* instead of singular cochains. In particular, the exactness proof at the cochain level is analytic, using partitions of unity, while the singular case is purely combinatorial.

6.4 Integration Map

So far, de Rham cohomology has been built entirely from the calculus of smooth differential forms: forms live on manifolds, the exterior derivative relates them, and their cohomology groups record global obstructions to exactness. But in the back of our minds, we know there is another, older theory — singular cohomology — which we developed earlier. The two theories look very different: singular cohomology starts with singular simplices, chains, and cochains, while de Rham cohomology begins from a geometric–analytic perspective.

The natural question is: *can we compare them directly?* Even more, can we construct a canonical map

$$H^*_{\mathrm{dR}}(M) \longrightarrow H^*_{\mathrm{sing}}(M;\mathbb{R})$$

and hope for it to be an isomorphism? This is the content of the celebrated *de Rham theorem* — but before we can prove it, we must define such a map. The most obvious bridge is integration.

Recall that a singular k–simplex in M is a continuous map

$$\sigma : \Delta^k \longrightarrow M$$

from the standard k–simplex to M. Given a smooth k–form $\omega \in \Omega^k(M)$, it is tempting to measure it on σ by pulling it back and integrating over Δ^k:

$$\omega \quad \mapsto \quad \int_{\Delta^k} \sigma^*\omega.$$

This expression is well–defined because $\sigma^*\omega$ is a k–form on Δ^k, which is a subset of $\mathbb{R}^k$ with its standard smooth structure, and we know how to integrate such forms from multivariable calculus.

Definition 6.5 (Integration map on cochains)**.** Let M be a smooth manifold. For each $k \geq 0$, define a map

$$I_k : \Omega^k(M) \longrightarrow C^k_{\text{sing}}(M;\mathbb{R})$$

by

$$(I_k\omega)(\sigma) := \int_{\Delta^k} \sigma^*\omega,$$

for every singular k–simplex $\sigma : \Delta^k \to M$.

In words: a smooth k–form becomes a singular k–cochain by "evaluating" on each simplex via integration. This fits our earlier picture: cochains in singular cohomology assign numbers to simplices; here we assign the *integral* of the pulled–back form.

The next step is to check compatibility with the cochain differentials. On the de Rham side, the differential is d; on the singular side, it is the coboundary δ. If these are to correspond, we must have

$$I_{k+1}(d\omega) = \delta(I_k\omega).$$

But this is nothing other than the *generalised Stokes theorem*! Indeed, for any singular $(k+1)$–simplex Σ,

$$\begin{aligned}
\big(\delta(I_k\omega)\big)(\Sigma) &= (I_k\omega)(\partial\Sigma) \\
&= \sum_{i=0}^{k+1} (-1)^i (I_k\omega)(\Sigma_i) \\
&= \sum_{i=0}^{k+1} (-1)^i \int_{\Delta^k} \Sigma_i^*\omega \quad \text{(definition of } I_k) \\
&= \int_{\partial\Delta^{k+1}} (\Sigma|_{\partial\Delta^{k+1}})^*\omega \\
&= \int_{\Delta^{k+1}} \Sigma^*(d\omega) \quad \text{(generalised Stokes theorem)} \\
&= (I_{k+1}(d\omega))(\Sigma).
\end{aligned}$$

Thus $I_\bullet$ is a *cochain map.*

Proposition 6.3 (Naturality and cochain map property)**.** The maps I_k satisfy:

1. *Cochain map:* $I_{k+1} \circ d = \delta \circ I_k$ for all $k \geq 0$.

2. *Naturality:* If $f : M \to N$ is smooth, then $I_k(f^*\omega) = f^*_{\text{sing}}(I_k\omega)$ for all $\omega \in \Omega^k(N)$, where f^*_{sing} is the pullback of singular cochains along f.

Proof. (1) The calculation above is exactly the verification of $I_{k+1} \circ d = \delta \circ I_k$, using the generalised Stokes theorem and the identification of the coboundary as integration over the boundary.

(2) For naturality, let $\sigma : \Delta^k \to M$ be a singular simplex. Then

$$(I_k(f^*\omega))(\sigma) = \int_{\Delta^k} \sigma^*(f^*\omega) = \int_{\Delta^k} (f \circ \sigma)^*\omega = (I_k\omega)(f \circ \sigma) = (f^*_{\text{sing}}(I_k\omega))(\sigma).$$

■

Because $I_\bullet$ is a cochain map, it induces a homomorphism on cohomology.

Definition 6.6 (Integration map in cohomology)**.** The *integration map* in de Rham cohomology is

$$I^* : H^*_{\text{dR}}(M) \longrightarrow H^*_{\text{sing}}(M;\mathbb{R}), \qquad I^*([\omega]) := [I_\bullet\omega],$$

where $[I_\bullet\omega]$ denotes the singular cohomology class of the cochain $I_k\omega$.

The definition is well–posed: if ω differs from ω' by an exact form $d\eta$, then $I_\bullet(\omega - \omega') = I_\bullet(d\eta) = \delta(I_\bullet\eta)$ is a coboundary in the singular complex, so the cohomology classes agree.

Example 6.7 (A 1–simplex on S^1: integration detects winding)**.** View S^1 as $\mathbb{R}/\mathbb{Z}$ via the covering map

$$p : \mathbb{R} \longrightarrow S^1, \qquad t \longmapsto e^{2\pi i t}.$$

There is a unique smooth 1–form $\alpha \in \Omega^1(S^1)$ such that $p^*\alpha = dt$ on $\mathbb{R}$. (Equivalently, α is the invariant form that "measures angle" along the circle; it is closed and not exact.)

Fix an integer $m \neq 0$, and define a singular 1–simplex $\sigma : \Delta^1 \to S^1$ by

$$\sigma(s) := e^{2\pi i m s}, \qquad s \in \Delta^1 = [0,1].$$

This path winds m times around the circle. Compute the integration cochain $I_1(\alpha)$ on σ:

$$(I_1\alpha)(\sigma) = \int_{\Delta^1} \sigma^*\alpha = \int_0^1 (\alpha)_{\sigma(s)}\Big(\tfrac{d}{ds}\sigma(s)\Big)\, ds.$$

Pull back along the lift $s \mapsto ms$ on $\mathbb{R}$: since $p^*\alpha = dt$,

$$\sigma^*\alpha \;=\; (p \circ (ms))^*\alpha \;=\; (ms)^*dt \;=\; m\,ds.$$

Therefore

$$(I_1\alpha)(\sigma) \;=\; \int_0^1 m\,ds \;=\; m.$$

This is the key phenomenon behind the de Rham map in degree 1: integrating α over a loop returns its *winding number*. In particular, the induced map $I^1 : H^1_{\mathrm{dR}}(S^1) \to H^1_{\mathrm{sing}}(S^1;\mathbb{R})$ sends $[\alpha]$ to the standard generator that evaluates a loop by its degree.

Example 6.8 (A 2–simplex on the torus T^2: a determinant appears)**.** Identify the torus as $T^2 = \mathbb{R}^2/\mathbb{Z}^2$ with covering map $\pi : \mathbb{R}^2 \to T^2$. The 1–forms dx, dy on $\mathbb{R}^2$ are translation–invariant, hence descend to global forms (still denoted) $dx, dy \in \Omega^1(T^2)$. Consider the standard area form $\omega := dx \wedge dy \in \Omega^2(T^2)$.

Write the standard 2–simplex as

$$\Delta^2 = \{(u,v) \in \mathbb{R}^2 : \; u \geq 0, \; v \geq 0, \; u + v \leq 1\}$$

(with barycentric coordinate $w = 1 - u - v$ understood). Fix real constants a, b, c, d, and define a singular 2–simplex $\sigma : \Delta^2 \to T^2$ by the affine rule

$$\sigma(u,v) \;:=\; \big(au + bv\,,\, cu + dv\big) \mod \mathbb{Z}^2.$$

(If $a, b, c, d \in \mathbb{Z}$, this is the projection of a linear map $\mathbb{R}^2 \to \mathbb{R}^2$ to T^2; geometrically, it wraps the triangle through lattice directions.)

Compute the pullback of ω:

$$\sigma^*dx \;=\; a\,du + b\,dv, \qquad \sigma^*dy \;=\; c\,du + d\,dv,$$

so

$$\sigma^*\omega \;=\; \sigma^*(dx \wedge dy) \;=\; (a\,du + b\,dv) \wedge (c\,du + d\,dv) \;=\; (ad - bc)\,du \wedge dv.$$

Therefore the integration cochain gives

$$(I_2\omega)(\sigma) \;=\; \int_{\Delta^2} \sigma^*\omega \;=\; (ad-bc)\int_{\Delta^2} du\wedge dv \;=\; (ad-bc)\cdot \mathrm{Area}(\Delta^2).$$

Since $\mathrm{Area}(\Delta^2) = \frac{1}{2}$, we get

$$(I_2\omega)(\sigma) \;=\; \frac{ad-bc}{2}.$$

If you triangulate the unit square by two standard 2–simplices and map each by the same linear rule, the values add and you recover exactly $\det\left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)$. This matches the intuition that integrating $dx\wedge dy$ over a fundamental domain counts oriented area (and, for integer data, winding multiplicity) on the torus.

Remark 6.5 (First glimpse at the de Rham theorem). We now have a concrete, natural, and functorial map

$$I^* : H^*_{\mathrm{dR}}(M) \longrightarrow H^*_{\mathrm{sing}}(M;\mathbb{R}),$$

sending a smooth form to the singular cohomology class of its integration functional. The de Rham theorem will assert that I^* is an *isomorphism* for every smooth manifold M. That is, every real singular cohomology class can be represented by a smooth closed form, and two such forms represent the same class precisely when they differ by an exact form. Our next steps will be to prove this remarkable fact, but even now we can see how the worlds of geometry and topology are meeting: the generalised Stokes theorem ensures compatibility of differentials, and partitions of unity will let us glue local primitives into global ones.

6.5 The de Rham Theorem

At this point we possess two parallel cohomology theories on a smooth manifold M: on the one hand, *singular cohomology* $H^*_{\mathrm{sing}}(M;\mathbb{R})$, built from cochains that assign numbers to singular simplices; on the other, *de Rham cohomology* $H^*_{\mathrm{dR}}(M)$, built from smooth differential forms and the exterior derivative. We also constructed the *integration map* (a natural cochain map)

$$I^* : H^*_{\mathrm{dR}}(M) \;\longrightarrow\; H^*_{\mathrm{sing}}(M;\mathbb{R}), \qquad [\omega] \longmapsto \Big[\sigma \mapsto \int_{\Delta^k} \sigma^*\omega\Big],$$

whose compatibility with differentials is nothing but the generalised Stokes theorem. The de Rham theorem asserts that this bridge is, in fact, a perfect match.

Theorem 6.2 (de Rham)**.** For every smooth manifold M, the integration cochain map

$$I^k : \Omega^k(M) \longrightarrow C^k_{\text{sing}}(M;\mathbb{R}), \qquad I^k(\omega)(\sigma) := \int_{\Delta^k} \sigma^*\omega,$$

induces an isomorphism of graded abelian groups

$$I^* : H^*_{\text{dR}}(M) \xrightarrow{\cong} H^*_{\text{sing}}(M;\mathbb{R}).$$

We begin by isolating local statements and the glueing principle we shall use.

Lemma 6.2.1 (Homotopy invariance of singular cohomology via a cochain homotopy)**.** Let X be a space and $f_0, f_1 : X \to Y$ be homotopic via a continuous map $H : X \times [0,1] \to Y$ with $H(-,0) = f_0$, $H(-,1) = f_1$. There exists a degree (-1) map $S : C^k(Y;\mathbb{R}) \to C^{k-1}(X;\mathbb{R})$ such that

$$\delta S + S\delta \;=\; f_1^* - f_0^*.$$

Consequently $f_0^* = f_1^*$ on singular cohomology.

Proof. We construct S from a *prism operator* on chains. For $\sigma : \Delta^m \to X$ define

$$\widehat{\sigma} : \Delta^m \times [0,1] \longrightarrow Y, \qquad \widehat{\sigma}(u,t) := H(\sigma(u),t).$$

Triangulate the prism $\Delta^m \times [0,1]$ into $(m+1)$ oriented $(m{+}1)$–simplices $Q_i : \Delta^{m+1} \to \Delta^m \times [0,1]$, $i = 0,\dots,m$, by the standard "front/back" decomposition: in barycentric coordinates $\{e_0,\dots,e_m\}$ on Δ^m and $\{0,1\}$ on $[0,1]$, the vertices of Q_i are

$$(e_0,0),\,(e_1,0),\dots,(e_i,0),(e_i,1),(e_{i+1},1),\dots,(e_m,1),$$

with the induced orientation. Define the *prism chain operator* $P : C_m(X;\mathbb{R}) \to C_{m+1}(Y;\mathbb{R})$ by

$$P(\sigma) \;:=\; \sum_{i=0}^{m} (-1)^i\, \widehat{\sigma} \circ Q_i.$$

We claim the *chain homotopy identity*

$$\partial P + P\partial \;=\; H_{1\#} - H_{0\#} \quad \text{on } C_\bullet(X;\mathbb{R}), \tag{2}$$

where $H_t := H(-,t)$ and $H_{t\#}$ denotes pushforward of chains by $H_t \circ (\cdot)$. To prove (2), compute $\partial(\widehat{\sigma} \circ Q_i)$ by the alternating sum of its $(m{+}2)$ faces; these are of three types: the bottom face ($t = 0$), the top face ($t = 1$), and the (m) "lateral" faces. The signed sum of bottom faces over i yields $H_{0\#}(\sigma)$; the signed sum of top faces yields $-\,H_{1\#}(\sigma)$ (orientation comparison produces the minus sign); and the signed sum of lateral faces is precisely $-\,P(\partial\sigma)$ (this is the standard cancellation of the "prism of a boundary" with the "boundary of a prism"; one checks that the lateral face of Q_i glued along the jth face of Δ^m appears with the opposite sign to the corresponding face in Q_j glued along the ith face, so they assemble to $-P(\partial\sigma)$). Summing over i gives (2).

Now pass to cochains by precomposition with P. Define $S : C^k(Y;\mathbb{R}) \to C^{k-1}(X;\mathbb{R})$ via

$$S(\varphi)(\tau) \;:=\; \varphi\big(P(\tau)\big), \qquad \tau \in C_{k-1}(X).$$

For $\varphi \in C^k(Y)$ and $\tau \in C_k(X)$,

$$\big(\delta S(\varphi) + S(\delta\varphi)\big)(\tau) = S(\varphi)(\partial\tau) + \delta\varphi\big(P(\tau)\big) = \varphi\big(P(\partial\tau)\big) + \varphi\big(\partial P(\tau)\big) = \varphi\big((\partial P + P\partial)(\tau)\big),$$

and by (2) this equals $\varphi\big(H_{1\#}(\tau) - H_{0\#}(\tau)\big) = (H_1^*\varphi - H_0^*\varphi)(\tau)$. Thus $\delta S + S\delta = H_1^* - H_0^*$. Exactness of cohomology then implies $H_1^* = H_0^*$ on $H^*_{\text{sing}}(X;\mathbb{R})$. ■

Lemma 6.2.2 (Local identification on contractible open sets). Let $U \subset \mathbb{R}^n$ be a nonempty contractible open set. Then

$$I^* : H^*_{\text{dR}}(U) \xrightarrow{\cong} H^*_{\text{sing}}(U;\mathbb{R}).$$

Proof. On the de Rham side, Poincaré's Lemma shows $H^k_{\text{dR}}(U) = 0$ for $k \geq 1$, and $H^0_{\text{dR}}(U) \cong \mathbb{R}$ (connected case; otherwise a product over components), since closed 0–forms are locally constant functions. On the singular side, by Lemma 6.2.1 applied to the homotopy from id_U to a constant map, we have the same result: $H^k_{\text{sing}}(U;\mathbb{R}) = 0$ for $k \geq 1$ and $H^0_{\text{sing}}(U;\mathbb{R}) \cong \mathbb{R}$. It remains to check that I^* identifies the generators in degree 0. A degree-0 de Rham class is represented by a constant function $c \in \mathbb{R}$ on each path component; then for a 0–simplex $\sigma : \Delta^0 \to U$ we have $I^0([c])([\sigma]) = \int_{\Delta^0} \sigma^* c = c$. Hence I^* is an isomorphism in all degrees. ∎

Lemma 6.2.3 (Good covers exist)**.** Every smooth manifold M admits a *good cover*: an open cover $\{U_i\}_{i \in I}$ such that every nonempty finite intersection $U_{i_0 \cdots i_r}$ is diffeomorphic to an open ball in $\mathbb{R}^n$ (in particular, contractible).

Proof. Equip M with a Riemannian metric (standard existence theorem). For each $p \in M$, let $\operatorname{inj}(p) > 0$ be its injectivity radius. Choose $0 < r(p) < \frac{1}{2}\operatorname{inj}(p)$ and let U_p be the geodesic ball of radius $r(p)$ about p; then the exponential map $\exp_p$ restricts to a diffeomorphism from the Euclidean ball $B_{r(p)}(0) \subset T_p M$ onto U_p, and U_p is *strongly convex*: any two points in U_p are joined by a unique minimizing geodesic that lies entirely in U_p. The family $\{U_p\}_{p \in M}$ is an open cover. By paracompactness of M we may extract a locally finite refinement by such balls (e.g. choose a countable subcover by standard arguments). The intersection of finitely many strongly convex balls is strongly convex: if x, y lie in $\bigcap_{j=1}^r U_{p_j}$ then the unique minimizing geodesic γ joining x to y lies in each U_{p_j}, hence in the intersection. Strongly convex sets are diffeomorphic to Euclidean balls via normal coordinates, so each nonempty finite intersection is diffeomorphic to a ball. This yields a good cover. ∎

Lemma 6.2.4 (Mayer–Vietoris and naturality of integration)**.** Let $M = U \cup V$ with U, V open. There are long exact sequences in cohomology

$$\cdots \to H^{k-1}_{\mathrm{dR}}(U \cap V) \xrightarrow{\delta_{\mathrm{dR}}} H^k_{\mathrm{dR}}(M) \to H^k_{\mathrm{dR}}(U) \oplus H^k_{\mathrm{dR}}(V) \xrightarrow{r_U^* - r_V^*} H^k_{\mathrm{dR}}(U \cap V) \to \cdots$$

and

$$\cdots \to H^{k-1}_{\mathrm{sing}}(U \cap V; \mathbb{R}) \xrightarrow{\delta_{\mathrm{sing}}} H^k_{\mathrm{sing}}(M; \mathbb{R}) \to H^k_{\mathrm{sing}}(U; \mathbb{R}) \oplus H^k_{\mathrm{sing}}(V; \mathbb{R}) \xrightarrow{j_U^* - j_V^*} H^k_{\mathrm{sing}}(U \cap V; \mathbb{R}) \to \cdots$$

and the integration maps provide a commutative diagram of these exact sequences. In particular,

$$I^k \circ \delta_{\mathrm{dR}} \;=\; \delta_{\mathrm{sing}} \circ I^{k-1}.$$

Proof. We first recall cochain-level constructions.

De Rham side: Consider the short exact sequence of complexes

$$0 \to \Omega^\bullet(M) \xrightarrow{(r_U,\; r_V)} \Omega^\bullet(U) \oplus \Omega^\bullet(V) \xrightarrow{r^U_{U\cap V} - r^V_{U\cap V}} \Omega^\bullet(U \cap V) \to 0,$$

where r_U denotes restriction. Exactness is clear: a pair (α, β) lies in the kernel of the right map iff $\alpha|_{U\cap V} = \beta|_{U\cap V}$, i.e. iff it comes from a unique $\omega \in \Omega^\bullet(M)$ with $\omega|_U = \alpha$, $\omega|_V = \beta$. The associated long exact sequence in cohomology is the de Rham Mayer–Vietoris sequence. The connecting morphism δ_{dR} can be described explicitly: given a class $[\theta] \in H^{k-1}_{\mathrm{dR}}(U \cap V)$, choose a partition of unity $\{\rho_U, \rho_V\}$ subordinate to $\{U, V\}$ and set

$$\alpha := \rho_U\,\theta \in \Omega^{k-1}(U \cap V), \qquad \beta := -\,\rho_V\,\theta \in \Omega^{k-1}(U \cap V).$$

Extend α by zero to a form on U and β by zero to a form on V (still denoted α, β); then on $U \cap V$ we have $\alpha - \beta = \theta$. By exactness there exists $\omega \in \Omega^k(M)$ with $d\omega$ representing $\delta_{\mathrm{dR}}[\theta]$ and satisfying $r_U\omega = d\alpha$, $r_V\omega = d\beta$. Concretely, one may take

$$\delta_{\mathrm{dR}}[\theta] \;=\; \left[\, d(\widetilde{\rho_U\theta}) \,\right] \tag{3}$$

where $\widetilde{\rho_U\theta}$ denotes any extension of $\rho_U\theta$ by zero to M (different choices differ by a form supported in U with exact differential on U, hence yield the same cohomology class on M).

Singular side: Consider the short exact sequence of cochain complexes

$$0 \to C^\bullet_{\mathrm{sing}}(M) \xrightarrow{(j_U^*,\; j_V^*)} C^\bullet_{\mathrm{sing}}(U) \oplus C^\bullet_{\mathrm{sing}}(V) \xrightarrow{r} C^\bullet_{\mathrm{sing}}(U \cap V) \to 0,$$

where $j_U : U \hookrightarrow M$ and $j_V : V \hookrightarrow M$ are inclusions and $r(\Phi, \Psi) := \Phi|_{U\cap V} - \Psi|_{U\cap V}$. Exactness is immediate. The connecting morphism δ_{sing} is defined as follows: for $[\varphi] \in H^{k-1}_{\mathrm{sing}}(U \cap V)$, choose cochain lifts $\Phi \in C^{k-1}(U)$, $\Psi \in C^{k-1}(V)$ with $\Phi|_{U\cap V} - \Psi|_{U\cap V} = \varphi$; then $(\delta\Phi, \delta\Psi)$ lies in the image of $C^k(M)$ under (j_U^*, j_V^*), hence there is a unique $\Xi \in C^k(M)$ with $j_U^*\Xi = \delta\Phi$ and $j_V^*\Xi = \delta\Psi$. Define $\delta_{\mathrm{sing}}[\varphi] := [\Xi]$.

Commutativity with integration: We must show $I^k \circ \delta_{\mathrm{dR}} = \delta_{\mathrm{sing}} \circ I^{k-1}$. Let $[\theta] \in H^{k-1}_{\mathrm{dR}}(U \cap V)$ with θ closed; define α, β from θ and a subordinate partition of unity as above, extend by zero, and let $\omega := d\alpha$ on U and $\omega := d\beta$ on V, which glue to a global closed form representing $\delta_{\mathrm{dR}}[\theta]$ (this is (3)). On the singular side, define cochains

$$\Phi(\sigma) := \int_{\Delta^{k-1}} \sigma^*\alpha, \qquad \Psi(\sigma) := \int_{\Delta^{k-1}} \sigma^*\beta.$$

By construction $\Phi|_{U\cap V} - \Psi|_{U\cap V} = I^{k-1}(\theta)$. Compute $\delta\Phi$ on a k–simplex $\tau : \Delta^k \to U$:

$$(\delta\Phi)(\tau) = \Phi(\partial\tau) = \sum_{i=0}^{k}(-1)^i \int_{\Delta^{k-1}} (\tau \circ \iota_i)^*\alpha = \int_{\partial\Delta^k} \tau^*\alpha = \int_{\Delta^k} \tau^* d\alpha = \int_{\Delta^k} \tau^*\omega = I^k(\omega)(\tau),$$

where ι_i is the ith face inclusion and we used Stokes. The same holds on V with Ψ and $d\beta$, hence the unique $\Xi \in C^k(M)$ with $j_U^*\Xi = \delta\Phi$, $j_V^*\Xi = \delta\Psi$ is *exactly* $I^k(\omega)$, because $I^k(\omega)$ restricts to those values on U and V by the same computation. Therefore $\delta_{\text{sing}}(I^{k-1}[\theta]) = [\Xi] = [I^k(\omega)] = I^k(\delta_{\text{dR}}[\theta])$, as required. The commutativity of the other squares (restrictions and difference) follows directly from functoriality of pullback and the definition of I^k. ■

Lemma 6.2.5 (Five Lemma)**.** Consider a commutative diagram of abelian groups with exact rows

$$\begin{array}{ccccccccc} A_1 & \to & A_2 & \to & A_3 & \to & A_4 & \to & A_5 \\ \downarrow^{\cong} & & \downarrow^{\cong} & & \downarrow & & \downarrow^{\cong} & & \downarrow^{\cong} \\ B_1 & \to & B_2 & \to & B_3 & \to & B_4 & \to & B_5 \end{array}$$

If the four outer vertical maps are isomorphisms and the diagram commutes, then the middle vertical map $A_3 \to B_3$ is an isomorphism.

Proof. Standard diagram chase: to prove injectivity, take $x \in A_3$ mapping to 0 in B_3; exactness and isomorphisms show x lies in the image of $A_2 \to A_3$ and its preimage must be 0, hence $x = 0$. For surjectivity, given $y \in B_3$, lift along $B_2 \to B_3$ and push across isomorphisms to obtain a preimage in A_3. The commuting squares ensure well-definedness. (A complete elementwise chase is routine and omitted here for brevity.) ■

Proof of Theorem 6.2. Fix a good cover $\{U_i\}_{i\in I}$ of M (Lemma 6.2.3). We prove that I^* is an isomorphism by iterated use of Mayer–Vietoris and Lemma 6.2.5.

First, by Lemma 6.2.2, for each i and, more generally, for each nonempty finite intersection $U_{i_0\cdots i_r}$, the map

$$I^* : H^*_{\text{dR}}(U_{i_0\cdots i_r}) \to H^*_{\text{sing}}(U_{i_0\cdots i_r}; \mathbb{R})$$

is an isomorphism. Consider two opens $A, B \subset M$ for which I^* is known to be an isomorphism on A, B, and $A \cap B$. Lemma 6.2.4 furnishes a commutative diagram of Mayer–Vietoris long exact sequences for

$A \cup B$, with vertical maps I^*. The outer vertical maps in the relevant five-term extract are isomorphisms by hypothesis; the Five Lemma then forces $I^* : H^k_{\mathrm{dR}}(A \cup B) \to H^k_{\mathrm{sing}}(A \cup B; \mathbb{R})$ to be an isomorphism for all k.

Now order the index set so that $M_r := U_1 \cup \cdots \cup U_r$ increases to M. The base case $r = 1$ holds by Lemma 6.2.2. Assume the claim holds for M_{r-1}. Since $M_r = M_{r-1} \cup U_r$ and $M_{r-1} \cap U_r$ is a union of finite intersections of the good cover (hence a disjoint union of sets on which I^* is an isomorphism by Lemma 6.2.2), the previous paragraph applies to yield that I^* is an isomorphism on M_r. By induction over finitely many r (using a finite subcover for each connected component, as Mayer–Vietoris is local and our argument works one component at a time), we conclude I^* is an isomorphism on M in all degrees. ■

Remark 6.6 (Role of Stokes and of locality). The cochain map identity $\delta \circ I^k = I^{k+1} \circ d$ is the analytic heart: it is exactly the generalised Stokes theorem on singular simplices. Poincaré's Lemma (local exactness of the de Rham complex) and homotopy invariance (local triviality of singular cohomology on contractibles) identify the two theories on small pieces. Mayer–Vietoris is the mechanism that propagates this local agreement to a global isomorphism; the Five Lemma provides the formal passage from commutative exact diagrams to isomorphisms. No multiplicative structure was used anywhere in the argument.

Theorem 6.2 is far more than a technical identification of two cohomology theories. It is a striking unification of ideas: the left-hand side, $H^*_{\mathrm{dR}}(M)$, is born entirely from calculus — differential forms, exterior derivatives, and the geometry of smooth manifolds; the right-hand side, $H^*_{\mathrm{sing}}(M; \mathbb{R})$, arises from the combinatorial world of simplices, chains, and purely topological invariants. The integration map I^* weaves these two worlds together, translating analytic data into topological information, and vice versa, with perfect fidelity.

The theorem serves as a culmination of all the machinery we have built: wedge products, exterior derivatives, Stokes' theorem, Mayer–Vietoris — each plays a precise role. It is also a doorway: by recognising that smooth and singular cohomology coincide over $\mathbb{R}$, we gain access to every computational and conceptual tool from both domains. In this sense, de Rham's Theorem is not merely an endpoint; it is a bridge, opening the path toward deeper interactions between geometry, topology, and analysis.

6.6 The Ring Structure

One of the most remarkable features of de Rham cohomology is that it is not just a graded family of vector spaces: it naturally carries the structure of a *graded-commutative* $\mathbb{R}$*-algebra*. This structure is inherited directly from the wedge product of differential forms, which we studied earlier.

The motivation is straightforward: If ω measures some geometric quantity on p-dimensional pieces of M and η measures some other quantity on q-dimensional pieces, then their wedge product $\omega \wedge \eta$ measures the combination of the two, on $(p+q)$-dimensional pieces. Since de Rham cohomology is built from closed forms modulo exact forms, we ask whether this operation respects cohomology classes — and indeed, it does.

Definition 6.7 (Cup product in de Rham cohomology)**.** Let $[\omega] \in H^p_{\mathrm{dR}}(M)$ and $[\eta] \in H^q_{\mathrm{dR}}(M)$ be represented by closed forms $\omega \in \Omega^p(M)$ and $\eta \in \Omega^q(M)$. We define their *product* to be the class

$$[\omega] \smile_{\mathrm{dR}} [\eta] \;:=\; [\omega \wedge \eta] \in H^{p+q}_{\mathrm{dR}}(M).$$

The first question is: *is this well-defined?* That is, does the cohomology class of $\omega \wedge \eta$ depend only on the cohomology classes of ω and η, not on the particular representatives we choose?

Proposition 6.4 (Well-definedness)**.** If $\omega' = \omega + d\alpha$ and $\eta' = \eta + d\beta$ are cohomologous to ω and η respectively, then $\omega' \wedge \eta'$ is cohomologous to $\omega \wedge \eta$. Hence the above definition does not depend on the choice of representatives.

Proof. We expand:

$$\omega' \wedge \eta' = (\omega + d\alpha) \wedge (\eta + d\beta) = \omega \wedge \eta \;+\; \omega \wedge d\beta \;+\; d\alpha \wedge \eta \;+\; d\alpha \wedge d\beta.$$

We now use the graded Leibniz rule for d:

$$d(\alpha \wedge \eta) = d\alpha \wedge \eta + (-1)^p \alpha \wedge d\eta,$$

where $\deg(\alpha) = p - 1$ and $d\eta = 0$ since η is closed. Thus $d\alpha \wedge \eta = d(\alpha \wedge \eta)$ is exact.

Similarly,

$$d(\omega \wedge \beta) = d\omega \wedge \beta + (-1)^p \omega \wedge d\beta,$$

and since $d\omega = 0$, we have $\omega \wedge d\beta = (-1)^p d(\omega \wedge \beta)$, also exact.

Finally, $d\alpha \wedge d\beta$ is exact because

$$d(\alpha \wedge d\beta) = d\alpha \wedge d\beta - (-1)^{p-1} \alpha \wedge d^2\beta = d\alpha \wedge d\beta,$$

using $d^2 = 0$.

Therefore,

$$\omega' \wedge \eta' - \omega \wedge \eta = (\text{exact form}),$$

so $\omega' \wedge \eta'$ and $\omega \wedge \eta$ define the same cohomology class. ■

Having established well-definedness, the rest of the structure follows from the algebraic properties of the wedge product.

Proposition 6.5 (Algebraic structure)**.** The operation $\smile_{\mathrm{dR}}$ makes $H^*_{\mathrm{dR}}(M)$ into a graded-commutative $\mathbb{R}$-algebra:

1. **Bilinearity:** $([\omega_1] + [\omega_2]) \smile_{\mathrm{dR}} [\eta] = [\omega_1] \smile_{\mathrm{dR}} [\eta] + [\omega_2] \smile_{\mathrm{dR}} [\eta]$ and similarly in the second slot.
2. **Associativity:** $([\omega] \smile_{\mathrm{dR}} [\eta]) \smile_{\mathrm{dR}} [\theta] = [\omega] \smile_{\mathrm{dR}} ([\eta] \smile_{\mathrm{dR}} [\theta])$.
3. **Graded commutativity:** $[\omega] \smile_{\mathrm{dR}} [\eta] = (-1)^{pq} [\eta] \smile_{\mathrm{dR}} [\omega]$.
4. **Unit:** The constant function 1 in $\Omega^0(M)$ is closed and represents the identity element in $H^0_{\mathrm{dR}}(M)$.

Proof. Each property follows directly from the corresponding property of the wedge product on forms, which we have already proved earlier. For example, graded commutativity is inherited because for forms $\omega \in \Omega^p$ and $\eta \in \Omega^q$,

$$\omega \wedge \eta = (-1)^{pq} \eta \wedge \omega.$$

Since d respects the wedge product via the graded Leibniz rule, this property survives the passage to cohomology classes. ■

Remark 6.7 (Connection with singular cohomology)**.** Under the isomorphism of de Rham's Theorem, this wedge product corresponds exactly to the cup product on singular cohomology with real coefficients:

$$\smile_{\mathrm{dR}} \quad \longleftrightarrow \quad \smile.$$

Thus $H^*_{\mathrm{dR}}(M)$ is not just isomorphic to $H^*_{\mathrm{sing}}(M;\mathbb{R})$ as graded vector spaces, but as graded rings. This reinforces the deep compatibility between the analytic and topological viewpoints.

Example 6.9 (Ring structure of $H^*_{\mathrm{dR}}(S^n)$)**.** From earlier computations, $H^0_{\mathrm{dR}}(S^n) \cong \mathbb{R}$, $H^n_{\mathrm{dR}}(S^n) \cong \mathbb{R}$, and all other groups vanish. Let u be a generator of $H^n_{\mathrm{dR}}(S^n)$. Since $u \wedge u$ would have degree $2n > n$, it must vanish. Thus

$$H^*_{\mathrm{dR}}(S^n) \cong \mathbb{R} \oplus \mathbb{R}\langle u\rangle, \qquad u^2 = 0,$$

as a graded ring.

Example 6.10 (Ring structure of $H^*_{\mathrm{dR}}(T^2)$)**.** Let $T^2 = S^1 \times S^1$ with projection maps π_1, π_2. Let $\alpha, \beta \in H^1_{\mathrm{dR}}(S^1)$ be the standard generators. Set $x := \pi_1^*\alpha$ and $y := \pi_2^*\beta$ in $H^1_{\mathrm{dR}}(T^2)$. Then $x \wedge x = 0$, $y \wedge y = 0$, and $x \wedge y = -y \wedge x$ generates $H^2_{\mathrm{dR}}(T^2) \cong \mathbb{R}$. Thus

$$H^*_{\mathrm{dR}}(T^2) \cong \Lambda_{\mathbb{R}}(x, y),$$

the exterior algebra on two degree-one generators.

The graded ring structure is a powerful refinement: where the vector space structure of cohomology detects *how many* holes there are in each dimension, the ring structure captures *how those holes interact*. It is in this interaction that much of the deeper geometry and topology of a manifold is encoded.

6.7 Relative de Rham Cohomology

So far, our study of de Rham cohomology has measured global properties of a manifold M as a whole. But in many situations, we are interested in information about M *relative* to a subspace $A \subseteq M$. For instance:

- In integration problems, one may want to fix a portion of the boundary to have a prescribed value.

- In topology, relative invariants appear when studying spaces obtained by gluing along a subspace.
- In algebraic topology, we have already seen the concept of *relative singular cohomology* $H^*(X, A)$, which measures cocycles on X that vanish on A.

The de Rham version follows exactly the same spirit: we look at differential forms on M that vanish when restricted to A.

Definition 6.8 (Relative differential forms)**.** Let M be a smooth manifold and $A \subseteq M$ a smooth submanifold (possibly with boundary). We define the space of *relative k-forms* by

$$\Omega^k(M, A) \ := \ \{\, \omega \in \Omega^k(M) \mid \iota^*\omega = 0 \,\},$$

where $\iota : A \hookrightarrow M$ is the inclusion map. That is, ω vanishes on all k-tuples of tangent vectors based at points of A.

It is immediate that d preserves this condition: if ω vanishes on A, then $d\omega$ vanishes on A as well (because pullback commutes with d and $\iota^*0 = 0$). Hence the spaces $\Omega^\bullet(M, A)$ form a subcomplex of the usual de Rham complex.

Definition 6.9 (Relative de Rham cohomology)**.** The *relative de Rham cohomology* of the pair (M, A) is

$$H^k_{\mathrm{dR}}(M, A) \ := \ \frac{\ker\big(d : \Omega^k(M, A) \to \Omega^{k+1}(M, A)\big)}{\operatorname{im}\big(d : \Omega^{k-1}(M, A) \to \Omega^k(M, A)\big)}.$$

In other words, a relative cohomology class is represented by a closed form ω on M that vanishes on A, with two such forms considered equivalent if they differ by an exact form $d\eta$ whose η also vanishes on A.

Remark 6.8 (Parallel with singular cohomology)**.** This matches the singular cohomology definition:

$$H^k(X, A) = \frac{\{k\text{-cocycles on } X \text{ vanishing on } A\}}{\{k\text{-coboundaries with this property}\}}.$$

The inclusion $\iota : A \hookrightarrow M$ plays the same role in both theories, and the map ι^* is just the restriction of forms.

We can now establish the long exact sequence of the pair in the de Rham setting, just as in the singular case.

Theorem 6.3 (Long exact sequence of a pair in de Rham cohomology)**.** For a smooth manifold M and a smooth submanifold $A \subseteq M$, there is a natural long exact sequence

$$\cdots \longrightarrow H^{k-1}_{\mathrm{dR}}(A) \xrightarrow{\delta} H^k_{\mathrm{dR}}(M, A) \longrightarrow H^k_{\mathrm{dR}}(M) \xrightarrow{\iota^*} H^k_{\mathrm{dR}}(A) \longrightarrow \cdots$$

Proof. We follow the standard short exact sequence of cochain complexes argument.

Consider the inclusion of complexes:

$$0 \longrightarrow \Omega^\bullet(M, A) \xrightarrow{\text{inc}} \Omega^\bullet(M) \xrightarrow{\iota^*} \Omega^\bullet(A) \longrightarrow 0,$$

where ι^* is restriction to A. We first check exactness:

- At $\Omega^\bullet(M, A)$: the inclusion is injective by definition.
- At $\Omega^\bullet(M)$: $\ker(\iota^*)$ consists precisely of forms vanishing on A, i.e. $\Omega^\bullet(M, A)$.
- At $\Omega^\bullet(A)$: surjectivity of ι^* follows because every form on A can be extended to a form on M (using a tubular neighbourhood of A and a partition of unity).

Since d commutes with both inclusion and pullback, this is a short exact sequence of cochain complexes. By the general snake lemma for cohomology (or the long exact sequence of cohomology for a short exact sequence of cochain complexes), we obtain a long exact sequence in cohomology:

$$\cdots \to H^{k-1}_{\mathrm{dR}}(A) \xrightarrow{\delta} H^k_{\mathrm{dR}}(M, A) \to H^k_{\mathrm{dR}}(M) \xrightarrow{\iota^*} H^k_{\mathrm{dR}}(A) \to \cdots$$

Here δ is the usual connecting homomorphism induced by the short exact sequence of complexes. ■

Example 6.11. Let $M = D^n$ be the n-dimensional closed ball and $A = S^{n-1} = \partial D^n$. We know that $H^k_{\mathrm{dR}}(D^n) \cong 0$ for $k > 0$, and $H^0_{\mathrm{dR}}(D^n) \cong \mathbb{R}$. The long exact sequence of the pair (D^n, S^{n-1}) gives

$$0 \to H^{n-1}_{\mathrm{dR}}(S^{n-1}) \xrightarrow{\delta} H^n_{\mathrm{dR}}(D^n, S^{n-1}) \to 0,$$

and $H^{n-1}_{\mathrm{dR}}(S^{n-1}) \cong \mathbb{R}$, so $H^n_{\mathrm{dR}}(D^n, S^{n-1}) \cong \mathbb{R}$. This matches the intuition: a relative n-form is a volume form on D^n whose restriction to the boundary vanishes, and its cohomology class records the total oriented volume.

Remark 6.9 (Relation to integration and Stokes' theorem). Relative cohomology is the natural setting for Stokes' theorem with prescribed boundary conditions: if $\omega \in \Omega^n(M, A)$ and $d\omega = 0$, then $\int_M \omega$ depends only on the class $[\omega] \in H^n_{\mathrm{dR}}(M, A)$. This viewpoint becomes particularly useful in manifold theory and gauge theory, and later, in the proof of Poincaré duality.

6.8 Relation to Homotopy

In our earlier study of singular cohomology, we proved that cohomology is a *homotopy invariant*: if two maps $f, g : X \to Y$ are homotopic, they induce the same map on cohomology. This property lies at the heart of why cohomology is a topological invariant rather than merely a geometric one.

It is natural to ask: does the same hold for de Rham cohomology, which is defined in terms of *smooth* differential forms? At first glance, one might worry: smooth forms seem sensitive to the precise geometric shape of a manifold, while homotopy is a purely topological notion. The beauty of de Rham's theorem (which we have already proven) tells us these worlds agree — and homotopy invariance can be proven directly in the smooth setting, without appealing to the singular theory.

Theorem 6.4 (Homotopy invariance of de Rham cohomology). Let M and N be smooth manifolds, and let $f, g : M \to N$ be smooth maps that are smoothly homotopic. Then the induced pullbacks on de Rham cohomology are equal:

$$f^* = g^* : H^\bullet_{\mathrm{dR}}(N) \longrightarrow H^\bullet_{\mathrm{dR}}(M).$$

Proof. Suppose $H : M \times [0, 1] \to N$ is a smooth homotopy from f to g, so $H(x, 0) = f(x)$ and $H(x, 1) = g(x)$.

Our goal is to show that for any closed k-form ω on N, the pullbacks $f^*\omega$ and $g^*\omega$ differ by an *exact* form on M. This will prove they define the same cohomology class.

The key idea is to use a *chain homotopy operator* K between the pullbacks, defined by integrating along the homotopy direction.

Let $\pi_M : M \times [0, 1] \to M$ be projection onto M, and $\pi_I : M \times [0, 1] \to [0, 1]$ be projection onto the interval. Given $\omega \in \Omega^k(N)$, pull it back to $M \times [0, 1]$ via $H^*\omega$. Now decompose any form $\eta \in \Omega^k(M \times [0, 1])$

uniquely as

$$\eta = \alpha(t) + dt \wedge \beta(t),$$

where $\alpha(t)$ and $\beta(t)$ are families of forms on M depending smoothly on t.

We now define

$$K : \Omega^k(N) \longrightarrow \Omega^{k-1}(M)$$

by the formula

$$(K\omega)_x := \int_0^1 \beta_x(t)\, dt,$$

where $\beta(t)$ is the component of $H^*\omega$ in the dt-direction. The integral is taken pointwise in $x \in M$, producing a smooth $(k-1)$-form on M.

Claim: This K satisfies the chain homotopy relation

$$g^* - f^* = d \circ K + K \circ d.$$

To see this, note that $H^*\omega = \alpha(t) + dt \wedge \beta(t)$ by decomposition. The pullback at $t = 0$ is $f^*\omega = \alpha(0)$ and at $t = 1$ is $g^*\omega = \alpha(1)$. Thus

$$g^*\omega - f^*\omega = \alpha(1) - \alpha(0) = \int_0^1 \frac{d}{dt}\alpha(t)\, dt.$$

Cartan's magic formula for the exterior derivative on a product with an interval tells us

$$d_{M\times I}\big(\alpha(t) + dt \wedge \beta(t)\big) = d_M\alpha(t) + dt \wedge \frac{d\alpha}{dt} \;-\; dt \wedge d_M\beta(t) \;+\; (\text{no } dt \text{ term here}).$$

Comparing dt-components gives $\frac{d\alpha}{dt} = d_M\beta(t) \pm \beta'(t)$ appropriately, and one checks that the integral over t matches exactly $d(K\omega) + K(d\omega)$.

If ω is closed, $d\omega = 0$, so

$$g^*\omega - f^*\omega = d(K\omega),$$

showing they differ by an exact form. Passing to cohomology classes gives $f^* = g^*$ on $H^k_{\mathrm{dR}}(N)$ for all k. ■

Example 6.12. If M is contractible, then the identity map id_M is homotopic to the constant map $c : M \to \{p\} \subset M$. Homotopy invariance shows that $H^\bullet_{\mathrm{dR}}(M) \cong H^\bullet_{\mathrm{dR}}(\mathrm{pt})$, so $H^0_{\mathrm{dR}}(M) \cong \mathbb{R}$ and all higher groups vanish. This recovers the familiar fact that contractible manifolds have trivial de Rham cohomology in positive degrees.

6.9 Key de Rham Cohomology Computations

We now consolidate the theory by computing $H^\bullet_{\mathrm{dR}}$ for several fundamental spaces. The plan is always the same: identify good covers by simple (preferably contractible) open sets, use Poincaré's Lemma (d–closed $\Rightarrow$ exact on contractible sets), apply Mayer–Vietoris when helpful, and certify nontrivial classes by pairing with cycles via integration (Generalized Stokes ensures exact forms integrate to 0 on cycles).

6.9.1 Contractible open sets $U \subset \mathbb{R}^n$

Proposition 6.6. If U is nonempty and contractible, then

$$H^0_{\mathrm{dR}}(U) \cong \mathbb{R}, \qquad H^k_{\mathrm{dR}}(U) = 0 \quad (k \geq 1).$$

Proof. A 0–form is a smooth function. If $df = 0$ on a connected set, f is locally constant, hence constant; thus $H^0 \cong \mathbb{R}$. For $k \geq 1$, Poincaré's Lemma on contractible U says every closed k–form is exact; hence $Z^k = B^k$ and $H^k = 0$. ■

6.9.2 The circle S^1

Cover S^1 by two overlapping open arcs U, V whose closures are proper arcs (no full loop). Each of U, V is contractible, so

$$H^0(U) \cong H^0(V) \cong \mathbb{R}, \quad H^k(U) = H^k(V) = 0 \ (k \geq 1).$$

The intersection $U \cap V$ is a disjoint union of two arcs, hence has two path components, so

$$H^0(U \cap V) \cong \mathbb{R} \oplus \mathbb{R}, \qquad H^k(U \cap V) = 0 \ (k \geq 1).$$

The Mayer–Vietoris long exact sequence in cohomology for the cover $S^1 = U \cup V$ includes

$$0 \longrightarrow H^0(S^1) \xrightarrow{\rho} H^0(U) \oplus H^0(V) \xrightarrow{\delta_0} H^0(U \cap V) \xrightarrow{\partial} H^1(S^1) \longrightarrow 0,$$

because the H^1–terms of $U, V, U \cap V$ vanish. Here ρ is restriction and $\delta_0(a, b) = a|_{U\cap V} - b|_{U\cap V}$. Concretely, for constants $a, b \in \mathbb{R}$, $\delta_0(a, b) = (a - b, a - b) \in \mathbb{R} \oplus \mathbb{R}$. Thus

$$\mathrm{Im}(\delta_0) = \{(t, t) : t \in \mathbb{R}\} \subset \mathbb{R} \oplus \mathbb{R}.$$

Exactness at $H^0(U) \oplus H^0(V)$ implies $\mathrm{Im}(\rho) = \ker(\delta_0) = \{(a, a)\}$, so ρ identifies $H^0(S^1) \cong \mathbb{R}$ with the diagonal. Exactness at $H^0(U \cap V)$ gives

$$H^1(S^1) \cong \frac{H^0(U \cap V)}{\mathrm{Im}(\delta_0)} \cong \frac{\mathbb{R} \oplus \mathbb{R}}{\{(t, t)\}} \cong \mathbb{R}.$$

Higher degrees vanish by dimension reasons, so

$$H^0(S^1) \cong \mathbb{R}, \qquad H^1(S^1) \cong \mathbb{R}, \qquad H^k(S^1) = 0 \ (k \geq 2).$$

Identifying a generator. On the universal cover $\pi : \mathbb{R} \to S^1$, $t \mapsto e^{2\pi i t}$, the 1–form dt is closed. Although dt does not descend to the circle as a global 1–form, its *period* on the deck transformation $t \mapsto t + 1$ is 1, and the corresponding class on S^1 is represented in coordinates by $d\theta$ (in any angular chart). Its integral over the positively oriented loop is 1, so it generates $H^1(S^1) \cong \mathbb{R}$.

6.9.3 The n–sphere S^n $(n \geq 1)$

Cover S^n by two open hemispheres U, V (remove the north/south poles); each is diffeomorphic to $\mathbb{R}^n$, so $H^0(U) \cong H^0(V) \cong \mathbb{R}$ and $H^k(U) = H^k(V) = 0$ for $k \geq 1$. Their intersection $U \cap V$ deformation retracts onto the equator S^{n-1}, hence $H^\bullet(U \cap V) \cong H^\bullet(S^{n-1})$.

Mayer–Vietoris gives, for each k,

$$\cdots \to H^k(S^n) \xrightarrow{\rho} H^k(U) \oplus H^k(V) \xrightarrow{\delta_k} H^k(U \cap V) \xrightarrow{\partial} H^{k+1}(S^n) \to \cdots$$

Using $H^k(U) = H^k(V) = 0$ for $k \geq 1$, we get for $1 \leq k \leq n-2$ an isomorphism

$$H^k(U \cap V) \xrightarrow{\sim} H^{k+1}(S^n).$$

Inducting on n starting from S^1 (computed above), we conclude

$$H^k(S^n) = \begin{cases} \mathbb{R}, & k = 0, n, \\ 0, & \text{otherwise.} \end{cases}$$

Top class. The orientation (volume) form vol_{S^n} is closed and integrates to the volume of S^n; after scaling to have total integral 1, its class generates $H^n(S^n) \cong \mathbb{R}$.

6.9.4 Punctured Euclidean space $\mathbb{R}^n \setminus \{0\}$

The radial deformation retraction $\mathbb{R}^n \setminus \{0\} \to S^{n-1}$ is smooth and a homotopy equivalence, so homotopy invariance yields

$$H^\bullet_{\mathrm{dR}}(\mathbb{R}^n \setminus \{0\}) \cong H^\bullet_{\mathrm{dR}}(S^{n-1}).$$

Hence $H^0 \cong \mathbb{R}$, $H^{n-1} \cong \mathbb{R}$, and all other degrees vanish.

Explicit generator in degree $n-1$. Let $r = \sqrt{x_1^2 + \cdots + x_n^2}$ and set

$$\omega_{n-1} \;=\; \frac{1}{\mathrm{Vol}(S^{n-1})} \sum_{i=1}^{n} (-1)^{i-1} \frac{x_i}{r^n}\, dx_1 \wedge \cdots \wedge \widehat{dx_i} \wedge \cdots \wedge dx_n.$$

One checks directly $d\omega_{n-1} = 0$ on $\mathbb{R}^n \setminus \{0\}$, and by pulling back the standard volume form on S^{n-1} along the radial projection, $\displaystyle\int_{S^{n-1}} \omega_{n-1} = 1$. If $\omega_{n-1} = d\eta$, Stokes would give $\int_{S^{n-1}} \omega_{n-1} = \int_{S^{n-1}} d\eta = 0$, a contradiction. Thus $[\omega_{n-1}]$ is the generator of H^{n-1}.

6.9.5 The 2–torus $T^2 = S^1 \times S^1$

View T^2 as $\mathbb{R}^2/\mathbb{Z}^2$ with coordinates (x, y) (mod 1). The 1–forms dx and dy are translation–invariant on $\mathbb{R}^2$ and descend to smooth forms on T^2. They satisfy $d(dx) = d(dy) = 0$.

They are not exact. Let $\gamma_x(t) = (t, 0)$ and $\gamma_y(t) = (0, t)$ be the fundamental loops (mod 1). Then

$$\int_{\gamma_x} dx = 1, \int_{\gamma_x} dy = 0, \qquad \int_{\gamma_y} dx = 0, \int_{\gamma_y} dy = 1.$$

If $dx = dF$ for a global function F on T^2, Stokes would give $\int_{\gamma_x} dx = \int_{\gamma_x} dF = 0$, contradiction. Similarly for dy. Hence $[dx], [dy]$ define two linearly independent classes in $H^1(T^2)$.

By de Rham's theorem, $H^1(T^2) \cong \mathrm{Hom}(H_1(T^2;\mathbb{R}), \mathbb{R})$ and $H_1(T^2;\mathbb{R}) \cong \mathbb{R}^2$, so $\dim H^1(T^2) = 2$. Therefore $\{[dx], [dy]\}$ is a basis:

$$H^1(T^2) \cong \mathbb{R}\langle [dx], [dy]\rangle \;\cong\; \mathbb{R}^2.$$

For degree 2, note $dx \wedge dy$ is closed, $\displaystyle\int_{T^2} dx \wedge dy = 1$ (after suitable normalization of the fundamental domain), so $[dx \wedge dy] \neq 0$. Since $\dim H^2(T^2) = 1$ (again by de Rham's theorem / Künneth for $\mathbb{R}$–coefficients), this class generates:

$$H^0(T^2) \cong \mathbb{R}, \quad H^1(T^2) \cong \mathbb{R}^2, \quad H^2(T^2) \cong \mathbb{R}.$$

(As a preview of the ring structure already discussed, $[dx] \smile [dy] = [dx \wedge dy]$ and $[dx] \smile [dx] = 0 = [dy] \smile [dy]$, so $H^*(T^2) \cong \Lambda_{\mathbb{R}}([dx], [dy])$.)

6.9.6 Real projective space $\mathbb{RP}^n$ (with $\mathbb{R}$–coefficients)

Cover $\mathbb{RP}^n$ by the two standard affine charts

$$U_0 = \{[x_0 : \cdots : x_n] \mid x_0 \neq 0\} \cong \mathbb{R}^n, \qquad U_1 = \{x_1 \neq 0\} \cong \mathbb{R}^n,$$

so $H^k(U_i) = 0$ for $k \geq 1$ and $H^0(U_i) \cong \mathbb{R}$. Their intersection $U_0 \cap U_1$ deformation retracts onto $\mathbb{RP}^{n-1}$: indeed, $U_0 \cap U_1$ is diffeomorphic to $(\mathbb{R}^n \setminus \{0\})/\{\pm 1\}$, whose radial retraction descends to the quotient and lands on $\mathbb{RP}^{n-1}$.

Mayer–Vietoris gives, for $k \geq 1$,

$$0 \to H^k(\mathbb{RP}^n) \to H^k(U_0) \oplus H^k(U_1) \to H^k(U_0 \cap U_1) \to H^{k+1}(\mathbb{RP}^n) \to 0,$$

i.e.

$$0 \to H^k(\mathbb{RP}^n) \to 0 \oplus 0 \to H^k(\mathbb{RP}^{n-1}) \to H^{k+1}(\mathbb{RP}^n) \to 0.$$

Thus for $1 \le k \le n-2$ we have isomorphisms

$$H^k(\mathbb{RP}^{n-1}) \cong H^{k+1}(\mathbb{RP}^n).$$

Together with $H^0(\mathbb{RP}^n) \cong \mathbb{R}$, this recursion shows (by induction on n and using $\mathbb{RP}^1 \cong S^1$) that

$$H^k(\mathbb{RP}^n;\mathbb{R}) = \begin{cases} \mathbb{R}, & k=0, \\ 0, & 0<k<n, \\ \mathbb{R}, & k=n \text{ and } n \text{ odd}, \\ 0, & k=n \text{ and } n \text{ even}. \end{cases}$$

(The parity in top degree comes from orientability: $\mathbb{RP}^n$ is orientable iff n is odd; an orientation yields a nonzero top-degree class given by a global volume form.)

6.9.7 Closed orientable surfaces Σ_g of genus g (outline with precise invariants)

Let Σ_g be a compact, connected, oriented surface of genus g. Cutting along $2g$ disjoint essential loops gives a $4g$–gon with paired edges. Cover Σ_g by two open sets U, V obtained by slightly thickening two complementary subskeleta of this cell structure: U deformation retracts to a wedge of g loops, as does V; the intersection $U \cap V$ retracts to a wedge of $2g$ loops. Using Mayer–Vietoris and Poincaré duality in dimension 2 (or by a careful count with the LES), one finds

$$H^0(\Sigma_g) \cong \mathbb{R}, \qquad H^1(\Sigma_g) \cong \mathbb{R}^{2g}, \qquad H^2(\Sigma_g) \cong \mathbb{R}.$$

Concretely, you can build $2g$ closed 1–forms whose periods produce a basis dual to the $2g$ fundamental 1–cycles, and the oriented area form (normalized to integrate to 1) generates H^2. (We omit the full polygonal Mayer–Vietoris arithmetic here; the LES is completely analogous to the T^2 computation, just with $2g$ generators.)

How to certify nontriviality (general recipe). Whenever you have a candidate closed form ω on a manifold M, pick a cycle C in the corresponding degree and compute $\int_C \omega$. If this integral is nonzero, then $[\omega] \neq 0$ in cohomology, because exact forms integrate to 0 on cycles by Stokes. Conversely, on many basic spaces the *only* obstruction to exactness is a nonzero period, so these integrals both detect and normalize generators (as we did with S^1, S^n, $\mathbb{R}^n\setminus\{0\}$, and T^2).

6.9.8 Summary of computations.

Manifold M	$H^\bullet_{\mathrm{dR}}(M)$
Contractible U	$H^0 \cong \mathbb{R}, \quad H^{k\geq 1} = 0$
S^1	$H^0 \cong \mathbb{R},\ H^1 \cong \mathbb{R},\ H^{k\geq 2} = 0$
S^n $(n \geq 1)$	$H^0 \cong \mathbb{R},\ H^n \cong \mathbb{R},$ else 0
$\mathbb{R}^n \setminus \{0\}$	$H^0 \cong \mathbb{R},\ H^{n-1} \cong \mathbb{R},$ else 0
T^2	$H^0 \cong \mathbb{R},\ H^1 \cong \mathbb{R}^2,\ H^2 \cong \mathbb{R}$
$\mathbb{RP}^n$	$H^0 \cong \mathbb{R},\ H^k = 0\ (0 < k < n),\ H^n \cong \mathbb{R}$ iff n odd

Closing remark. These examples show the full toolkit at work: Poincaré's Lemma for contractible pieces, Mayer–Vietoris to glue information, homotopy invariance to simplify spaces, and integration pairings to certify generators. They also show the ring structure (e.g. T^2) and the tight link with homology provided by the de Rham isomorphism: on each space above, the ranks of H^k_{dR} coincide with those of H_k and the natural pairings detect bases on both sides.

With this, we bring our exploration of de Rham cohomology to a close. We began by grounding ourselves in the language of differential forms, built the cochain complex $(\Omega^\bullet(M), d)$, and studied its functorial behaviour, computational tools such as Mayer–Vietoris, and the integration map. The de Rham theorem then revealed the remarkable bridge between smooth and singular worlds, identifying our smooth cohomology with the singular cohomology of the underlying topological space. From there we enriched the theory with a graded ring structure, extended it to relative settings, and examined explicit computations on familiar manifolds.

7 Conclusion

We have travelled a long path, beginning with the concrete geometry of simplicial complexes and ending with the smooth elegance of de Rham cohomology. Our first steps were firmly anchored in the world of *chains* and *boundaries*, learning how to detect and classify holes in a space through homology. From there we turned the perspective inside out, passing to *cochains* and *coboundaries* to build cohomology, a dual theory that not only detects the same topological features but also carries a rich algebraic structure via the cup product.

This algebraic insight paved the way for geometry to return to the stage. By introducing differential forms, wedge products, and the exterior derivative, we gained a calculus on manifolds capable of encoding topological information. The generalised Stokes theorem unified familiar results from vector calculus into a single, far-reaching principle. In the smooth category, these constructions assembled into the de Rham complex, whose cohomology groups measure the global geometry of a manifold.

The de Rham theorem provided the key revelation: that the smooth world and the topological world are intimately linked, and that differential forms, despite their analytic nature, capture precisely the same information as singular cohomology with real coefficients. Through functoriality, Mayer–Vietoris, relative cohomology, and the graded ring structure, we saw how de Rham cohomology is both a powerful computational tool and a deep conceptual bridge. Concrete computations on familiar manifolds illustrated that these ideas are not merely abstract, but tangible and computable.

At the end of this journey, we are left with a picture of mathematics where geometry, topology, and algebra are inseparably intertwined. Homology gave us the language of shape; cohomology revealed a new dual perspective enriched with algebraic structure; and de Rham theory united smooth analysis with purely topological invariants. It is in this interplay that much of the beauty of modern mathematics lies: a single phenomenon viewed from multiple angles, each viewpoint illuminating the others.

Though our exploration ends here, the tools developed open the door to further landscapes — characteristic classes, Hodge theory, and beyond — where the themes of duality, algebraic structure, and geometric meaning continue to play in harmony. The bridge built here between analysis and topology is not the end of the story, but a gateway to many others.

References

[1] Prerna Nadathur. *An Introduction to Homology*. 2007. URL: https://www.math.uchicago.edu/~may/VIGRE/VIGRE2007/REUPapers/INCOMING/Project.pdf.

[2] Pekka Pankka. *Introduction to de Rham cohomology*. 2013. URL: https://www.mv.helsinki.fi/home/pankka/deRham2013.

[3] Ulrike Tillmann. *Cohomology theories*. Oxford University. URL: https://people.maths.ox.ac.uk/tillmann/coho.pdf.