Algebraic topology

Travis Dillon

Contents

0 A very brief introduction 1

1 The definition of homology 1

2 The basics of category theory 3
- 2.1 A lot of definitions . . . 4
- 2.2 Categorifying homology . . . 6

3 Some explicit homology groups 8
- 3.1 This section has a point . . . 8
- 3.2 A bit of topology . . . 9
- 3.3 A little bit of magic . . . 11
- 3.4 Deformation retracts . . . 13

4 Homology gets bigger and better 13
- 4.1 Relative homology and the long exact sequence . . . 13
- 4.2 Excision . . . 15
- 4.3 Fundamental properties of homology: A list . . . 16
- 4.4 Our first nontrivial homologies . . . 17
- 4.5 Locality and another theorem . . . 18
- 4.6 More homology calculations . . . 19

5 CW homology 20
- 5.1 Pushouts . . . 20
- 5.2 CW-complexes . . . 22
- 5.3 Homology groups of CW-complexes . . . 24
- 5.4 Using the cellular chain complex . . . 26
- 5.5 Homology of real projective space . . . 28
- 5.6 Geometric realization . . . 29
- 5.7 Euler characteristic . . . 29

6 Homology with coefficients 30
- 6.1 What are coefficients? . . . 30
- 6.2 The categorical algebra of R-modules . . . 32
  - 6.2.1 The category of R-modules . . . 32
  - 6.2.2 Tensor products . . . 33
  - 6.2.3 Exactness properties of the Hom functor and tensor product . . . 35
- 6.3 Quasi-isomorphisms and free resolutions . . . 35
- 6.4 Tensor product of chain complexes . . . 37
- 6.5 Universal coefficient theorem . . . 39

7 Cohomology 42
- 7.1 Homology & Co. . . . 42
- 7.2 The (n)Ext functor . . . 43
- 7.3 The cohomology ring . . . 44
- 7.4 Cohomology of composite spaces . . . 48
- 7.5 Cohomology ring as an invariant . . . 49
- 7.6 Poincaré duality . . . 50

0. A VERY BRIEF INTRODUCTION

There is a subject called algebraic topology. Its goal is to ~~overload notation as much as possible~~ distinguish topological spaces through algebraic invariants. You may be familiar with the fundamental group; this is one such invariant. The goal of (most) of this course is to develop a different invariant: homology.

1. THE DEFINITION OF HOMOLOGY

DEFINITION 1.1. A *semisimplicial set* $\mathcal{X}$ is a sequence of sets $X_0, X_1, \dots$ and, for each $n \in \mathbb{N}_0$, a collection of $n+1$ maps $d_i^n : X_n \to X_{n-1}$ for $0 \le i \le n$ that satisfy the *simplicial identities*

$$d_i^{n-1} \circ d_j^n = d_{j-1}^{n-1} \circ d_i^n$$

whenever $i < j$.

The idea here is that a semisimplicial set will represent the triangulation of a topological space, where X_0 is the set of points, X_1 is the set of edges, X_2 the faces, and so on. In this triangulation, we can label the vertices with natural numbers, and the map d_i^n consists of "deleting the ith smallest vertex on an n-dimensional face." For example,

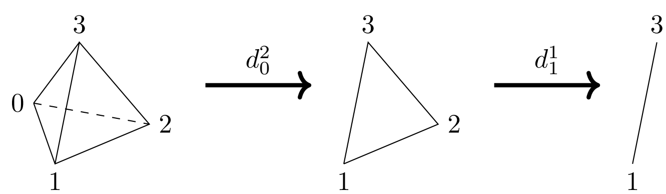

EXERCISE 1.2. Confirm the simplicial identity $d_0^1 \circ d_0^2 = d_1^1 \circ d_0^2$ in this picture.

EXERCISE 1.3. Prove that a semisimplicial set satisfies the identities

$$d_i^{n-1} \circ d_j^n = d_j^{n-1} \circ d_{i+1}^n$$

whenever $i \ge j$.

Because the superscripts are a bit cumbersome, and because algebraic topology has a long and storied history of overloading notation, the functions d_i^n are usually denoted simply d_i, with the domain and range understood from context.

Triangulations are nice, but they're pretty hard to construct for an arbitrary topological space. If someone says "I have here a topological space X," you certainly can't triangulate the space. That's where the generality of semisimplicial sets comes in useful. We can *always* define a canonical semisimplicial set on a topological space.

DEFINITION 1.4. For each $n \in \mathbb{N}_0$, the *standard n-simplex*, denoted Δ^n, is the convex hull of the $n+1$ standard basis vectors in $\mathbb{R}^n$.

DEFINITION 1.5. Let X be a topological space. We denote by $\mathrm{Sing}_n(X)$ the set of all functions from Δ^n to X.

Each map in $\mathrm{Sing}_n(X)$ is called a *singular n-simplex*; the adjective simply indicates that it may have cusps or any number of singularities. Given a map $\Delta^n \to X$, we can easily construct a different map $\Delta^{n-1} \to x$ by restricting to a single facet[1] of Δ^n. (Of course, we need to relabel the vertices, but this is done by simply maintaining their order.) If $f\colon \Delta^n \to X$, we write $d_i(f)$ for the map obtained by restricting to the facet that does not contain the ith vertex. Perhaps you can already see where this is going ...

DEFINITION 1.6. If X is a topological space, $\mathrm{Sing}(X)$ is the sequence of sets $\mathrm{Sing}_0(X), \mathrm{Sing}_1(X), \ldots$ together with the maps $d_i\colon \mathrm{Sing}_n(X) \to \mathrm{Sing}_{n-1}(X)$ for each $0 \leq i \leq n$.

It's not too hard to verify that $\mathrm{Sing}(X)$ is a semisimplicial set.

EXERCISE 1.7. Verify it.

Of course, it's only very loosely like a triangulation, in that it's not really one at all—it's some wildly infinite beast. But it's what we have to work with. You'll learn to love it.

And now we introduce the algebraic machinery—can't do algebraic topology without it.

DEFINITION 1.8. The abelian group $S_n(X)$ of *singular n-chains* is the free abelian group generated by $\mathrm{Sing}_n(X)$.

An *n-chain* is simply an element of $S_n(X)$: A finite linear combination of n-simplices:

$$\sum_{i=1}^{k} a_i\sigma_i \qquad a_i \in \mathbb{Z} \text{ and } \sigma_i \in \mathrm{Sing}_n(X).$$

If $n < 0$, then $\mathrm{Sing}_n(X) = \emptyset$ (by convention), so $S_n(X) = \{0\}$.

DEFINITION 1.9. The *boundary operator* $\partial_n\colon S_n(X) \to S_{n-1}(X)$ is defined on $\mathrm{Sing}_n(X)$ by

$$\partial_n(\sigma) = \sum_{k=0}^{n} (-1)^k d_k(\sigma)$$

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

This is a weird definition, but the idea is this: A union of n-simplices $\sigma_1 \cup \cdots \cup \sigma_m$ forms a "boundary" exactly when $\partial(\sigma_1 + \cdots + \sigma_m) = 0$. This is not quite right (we can't really take the union of functions!), but it essentially is. The next exercises go into more detail, but you can skip them if you want.

EXERCISE 1.10. Suppose $\sigma_1, \ldots, \sigma_m \in \pm\,\mathrm{Sing}_n(X)$ (each σ_i is either an n-simplex or its additive inverse in $S_n(X)$) and let ∇_i denote the image of σ_i. Some of the facets of the ∇_i might overlap or coincide. Prove that if $\partial_n(\sigma_1 + \cdots \sigma_m) = 0$, the $m(n+1)$ facets of the ∇_i can be paired so that the facets in each pair are exactly the same. (In other words, if $\partial_n \sum_i \sigma_i = 0$, then $\bigcup_i \nabla_i$ has no "hanging" $(n-1)$-dimensional boundaries.

EXERCISE 1.11. Suppose that $\nabla_1, \ldots, \nabla_m$ are the images of n-simplices in X that have a matching as described in the previous exercise. Prove that there are n-simplices $\sigma_1, \ldots, \sigma_m \in \pm\,\mathrm{Sing}_n(X)$ such that $\sigma_i = \nabla_i$ for every i and $\partial_n(\sigma_1 + \cdots + \sigma_m) = 0$. [HINT: This is hard! Here's one method: Finding these maps σ_i essentially corresponds to labelling each facet with ± 1 so that there is at most one more $+1$ than -1 or vice versa, and when you add up the labels on the facets that coincide, you get 0. (Why?) To find one of these, transform it to a graph problem. Form an undirected graph G with the vertex set $\{1, 2, \ldots, m\}$ and draw an edge between i and j for each facet of ∇_i that's paired with a facet of ∇_j. (So if ∇_1 and ∇_2 share 2 facets, then there are 2 edges between the vertices 1 and 2 in G.) Now you want to show that it's possible to direct the edges of G so that $|\mathrm{indeg}(v) - \mathrm{outdeg}(v)| \leq 1$ for every vertex v. (Of course, the first step is figuring out *why* this is what you want to do.)]

[1] A *facet* of a simplex is what you get by taking the convex hull of all but one of its vertices.

In short, if $c \in S_n(X)$ and $\partial_n(c) = 0$, there's good reason to think of c as "having no poky bits." Here's some associated definitions.

DEFINITION 1.12. An *n-cycle* in X is an n-chain c such that $\partial_n(c) = 0$. An *n-boundary* the image of an $(n+1)$-chain. The corresponding sets are:

$$\begin{array}{rl} n\text{-cycles} & Z_n(X) = \ker(\partial_n) \\ n\text{-boundaries} & B_n(X) = \operatorname{im}(\partial_{n+1}). \end{array}$$

EXERCISE 1.13. Prove that $Z_n(X)$ and $B_n(X)$ are subgroups of $S_n(X)$.

If everything is as it should be, then $\partial_n(\sigma)$ should be an $(n-1)$-cycle for every $\sigma \in \mathrm{Sing}_n(X)$. Fortunately, this is the case.

PROPOSITION 1.14. *$\partial_{n-1} \circ \partial_n = 0$ for every $n \geq 1$.*

Proof. This is just a matter of record-keeping with the simplicial identities. To start:

$$(\partial_{n-1} \circ \partial_n)(\sigma) = \sum_{k=0}^{n} (-1)^k \partial_{n-1} d_k(\sigma) = \sum_{k=0}^{n} (-1)^k \sum_{r=0}^{n-1} (-1)^r d_r d_k(\sigma).$$

Splitting by whether $r < k$ or $r \geq k$ and applying the simplicial identities to the first case, this becomes

$$\sum_{k=1}^{n} \sum_{0 \leq r < k} (-1)^{r+k} d_{k-1} d_r(\sigma) + \sum_{k=1}^{n-1} \sum_{k \leq r \leq n-1} (-1)^{r+k} d_r d_k(\sigma).$$

Substituting $k = \ell + 1$, the left sum becomes

$$\sum_{\ell=0}^{n-1} \sum_{0 \leq r \leq \ell} (-1)^{r+\ell+1} d_\ell d_r(\sigma);$$

switching the order of summation in the second sum, we get

$$\sum_{r=0}^{n-1} \sum_{0 \leq k \leq r} (-1)^{r+k} d_r d_k(\sigma).$$

These expressions differ only by a factor of -1, so the entire sum vanishes. □

The boundaries are in some sense the trivial cycles—we can quotient by them to get the interesting behavior.

DEFINITION 1.15. The *nth singular homology group* of X is

$$H_n(X) = {Z_n(X)}/{B_n(X)}.$$

Intuitively, two n-cycles in $Z_n(X)$ become equivalent in $H_n(X)$ if the space between them can be filled in by an $(n+1)$-chain. Since you can't fill a hole, this is (very roughly) how $H_n(X)$ detects them. Exactly *how* it detects them is a question best postponed.

Remark. Homology groups can be defined for any semisimplicial set $\mathcal{X}$ by mimicking the progression here: Define $S_n(\mathcal{X})$ to be the free group on X_n and then copy the definitions of ∂_n, $Z_n(\mathcal{X})$, $B_n(\mathcal{X})$, and $H_n(\mathcal{X})$ directly.

2. THE BASICS OF CATEGORY THEORY

Otherwise known as "abstract nonsense." It codifies a lot of assertions we have about functions. This means there will be a lot of vocabulary.

2.1. A LOT OF DEFINITIONS

DEFINITION 2.1. A *category* $\mathcal{C}$ consists of

1. a class[2] of *objects*, denoted $\mathrm{ob}(\mathcal{C})$,
2. a set $\mathrm{Hom}_{\mathcal{C}}(X, Y)$ of *morphisms* for every pair of objects $X, Y \in \mathrm{ob}(\mathcal{C})$,
3. an *identity morphism* $1_X \in \mathrm{Hom}_{\mathcal{C}}(X, X)$ for every $X \in \mathrm{ob}(\mathcal{C})$, and
4. an associative *composition operation* $\circ\colon \mathrm{Hom}_{\mathcal{C}}(X, Y) \times \mathrm{Hom}_{\mathcal{C}}(Y, Z)$, usually written $(f, g) \mapsto g \circ f$.

There are lots of categories. Here are some common ones.

EXAMPLE 2.2.

- The objects of the category Set are, naturally enough, sets, and a morphism from X to Y is a function.
- The objects of Ab are the abelian groups, and the morphisms are group homomorphisms.
- The objects of Top are topological spaces, and the morphisms are continuous maps. ◊

EXERCISE 2.3. Show that every object has exactly one identity morphism.

EXERCISE 2.4. A morphism $f\colon X \to Y$ is called an *isomorphism* if there is a morphism $g\colon Y \to X$ such that $f \circ g = 1_Y$ and $g \circ f = 1_X$. In this case, g is called an *inverse* of f. Prove that every morphism has at most one inverse.

Now we want to look at relationships between categories.

DEFINITION 2.5. Suppose $\mathcal{C}$ and $\mathcal{D}$ are two categories. A *functor* F from $\mathcal{C}$ to $\mathcal{D}$, denoted $F\colon \mathcal{C} \to \mathcal{D}$, assigns

- an object $F(X) \in \mathrm{ob}\,\mathcal{D}$ for every object $X \in \mathrm{ob}\,\mathcal{C}$ and
- a morphism $F(f) \in \mathrm{Hom}_{\mathcal{D}}(F(X), F(Y))$ for every morphism $f \in \mathrm{Hom}_{\mathcal{C}}(X, Y)$

such that

- $F(1_X) = 1_{F(X)}$ for every $X \in \mathrm{ob}\,\mathcal{C}$ and
- $F(g \circ f) = F(g) \circ F(f)$ for every pair of morphisms $f, g \in \mathcal{C}$ which can be composed.

EXAMPLE 2.6.

- The map which sends every element and map to itself is the *identity functor*.
- The map $\mathsf{Top} \to \mathsf{Set}$ that sends topological spaces to the set of path-connected components and the continuous map $f\colon X \to Y$ to the map $f^*\colon \pi_0(X) \to \pi_0(Y)$ where $f^*(C)$ is the component that contains $f(C)$—this is a functor.
- If $f\colon X \to Y$ is a continuous map and $\sigma \in \mathrm{Sing}_n(X)$, then $f \circ \sigma \in \mathrm{Sing}_n(Y)$. If we define $\mathrm{Sing}_n(f) = f \circ \sigma$, then the map Sing_n becomes a functor from Top to Set.
- Extending this, S_n is a functor from Top to Ab. ◊

There is, in fact, a "category" CAT of all categories, whose morphisms are functors. This raises some set-theoretic issues—how could CAT be an element of itself? In short, it's not. For us, suffice it to say that CAT is a "larger" type of category.[4]

But anyway, that's completely irrelevant. Category theory is an exercise in going down the rabbit hole, so let's pursue that instead. You have a category of categories? I raise you a category of functors.

[2] This is a "collection" that's bigger than a set. (We need this because we want to consider, for example, the category of sets, which is not, itself, a set.[3]) For all practical purposes, just think "collection" and you'll be fine.
[3] Why? If X were the set of all sets, then $X = 2^X$, but Cantor told us that there's no bijection between X and its power set.
[4] To be precise, Definition 2.1 is a definition of *locally small* categories. CAT is a type of category where the collection of morphisms can be bigger than a set. You could show this, for example, by proving that there are at least as many functors from Set to Set as there are objects.

Definition 2.7. Suppose $F, G\colon \mathcal{C} \to \mathcal{D}$. A *natural transformation* $\Theta\colon F \to G$ consists of a morphism $\Theta_X\colon F(X) \to G(X)$ for every $X \in \operatorname{ob}\mathcal{C}$ such that the diagram below commutes (which means that all the ways of getting from one point to another are equal; in this case, that $G(f) \circ \Theta_X = \Theta_Y \circ F(f)$).

$$\begin{array}{ccc} F(X) & \xrightarrow{\Theta_X} & G(X) \\ {\scriptstyle F(f)}\downarrow & & \downarrow{\scriptstyle G(f)} \\ F(Y) & \xrightarrow{\Theta_Y} & G(Y) \end{array}$$

One way to think of this is that a natural transformation $\Theta\colon F \to G$ provides a way of transitioning from F to G. We start out with this picture:

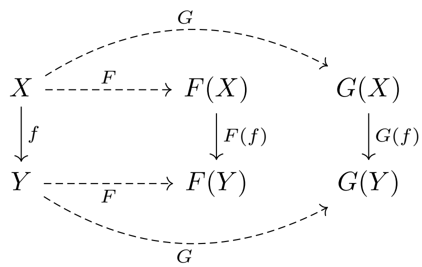

and we fill in the missing bits with a map so that the non-dashed portion commutes:

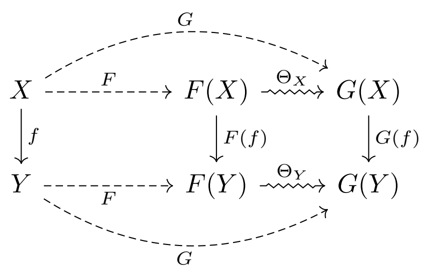

(It doesn't make sense to say that the whole diagram commutes, for example $\Theta_X \circ F = G$, because F and G are functors while Θ_X is a morphism; you can't compose them.) It's nice to think of a natural transformation schematically, like this:

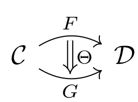

Exercise 2.8. If Θ_X is an isomorphism for every $X \in \operatorname{ob}\mathcal{C}$, then Θ is called a *natural isomorphism*. Suppose this is the case and define a new collection of maps $\bar{\Theta}$ with $\bar{\Theta}'_X = \Theta_X^{-1}$. Prove that $\bar{\Theta}$ is also a natural isomorphism.

Example 2.9. Remember that Sing_n is a functor from Top to Set. For each $0 \le i \le n$, the map d_i is a natural transformation from Sing_n to Sing_{n-1}, which you can verify by checking that this square commutes for every $X, Y \in \operatorname{ob}\mathsf{Top}$ and continuous map $f\colon X \to Y$:

$$\begin{array}{ccc} \mathrm{Sing}_n(X) & \xrightarrow{d_i} & \mathrm{Sing}_{n-1}(X) \\ {\scriptstyle \mathrm{Sing}_n(f)}\downarrow & & \downarrow{\scriptstyle \mathrm{Sing}_{n-1}(f)} \\ \mathrm{Sing}_n(Y) & \xrightarrow[d_i]{} & \mathrm{Sing}_{n-1}(Y) \end{array}$$

◊

Anyway, here's your crazy new category:

DEFINITION 2.10. Given two categories $\mathcal{C}$ and $\mathcal{D}$, their *functor category* is denoted $\mathrm{Fun}(\mathcal{C}, \mathcal{D})$; its objects are the functors from $\mathcal{C}$ to $\mathcal{D}$, and the morphisms are the natural transformations.

As a *digestif*, let's introduce a final, easier construction.

DEFINITION 2.11. Given a category $\mathcal{C}$, its *opposite category* $\mathcal{C}^{\mathrm{op}}$ has the same objects as $\mathcal{C}$ and a morphism $f^{\mathrm{op}}\colon Y \to X$ for each morphism $f\colon X \to Y$ in $\mathcal{C}$ such that $g^{\mathrm{op}} \circ f^{\mathrm{op}} = (f \circ g)^{\mathrm{op}}$.

In short, $\mathcal{C}^{\mathrm{op}}$ is formed by reversing the arrows in $\mathcal{C}$; it's a sort of dual to $\mathcal{C}$. The basic idea is that sometimes we have a map $\mathcal{C} \to \mathcal{D}$ that's almost a functor, but it switches the directions of the arrows. In this case, we can write it as an *actual* functor $\mathcal{C}^{\mathrm{op}} \to \mathcal{D}$.

EXAMPLE 2.12. Consider the category $\mathsf{Vect}_{\mathbb{R}}$ of real vector spaces with linear transformations as homomorphisms. Recall that the *dual* of a vector space V is defined $V^* = \mathrm{Hom}_{\mathsf{Vect}_{\mathbb{R}}}(V, \mathbb{R})$. Each morphism $\varphi\colon W \to V$ in $\mathsf{Vect}_{\mathbb{R}}$ induces a map $V^* \to W^*$ by sending $\psi \in V^*$ to $\psi \circ \varphi \in W^*$. So $(-)^*$ is a functor from $\mathsf{Vect}_{\mathbb{R}}^{\mathrm{op}}$ to $\mathsf{Vect}_{\mathbb{R}}$. ♢

2.2. CATEGORIFYING HOMOLOGY

Now we apply all these definitions to the homology from Section 1. First, of course, we introduce a new category.

DEFINITION 2.13. The category Δ_{inj} has the sets $[i] = \{0, 1, \ldots, i\}$, for each $i \in \mathbb{N}_0$, as objects, and the morphisms between objects are all order-preserving injections.

EXERCISE 2.14. Prove that Δ_{inj} has $\binom{b+1}{a+1}$ morphisms from $[a]$ to $[b]$.

It turns out that we can conceptualize semisimplicial sets equally well as functors $\Delta_{\mathrm{inj}}^{\mathrm{op}} \to \mathsf{Set}$. This is because the simplicial identities are baked into the definition of $\Delta_{\mathrm{inj}}^{\mathrm{op}}$.

EXERCISE 2.15. Let φ_i^a denote the injective order-preserving function $[a-1] \to [a]$ whose image does not contain i. Prove that

$$(\varphi_i^{a-1})^{\mathrm{op}} \circ (\varphi_j^a)^{\mathrm{op}} = (\varphi_{j-1}^{a-1})^{\mathrm{op}} \circ (\varphi_i^a)^{\mathrm{op}}$$

whenever $i < j$.

Starting with a semisimplicial set $\mathcal{X}$, define $F([i]) = X_i$ for each $i \in \mathbb{N}_0$, and for a map $\varphi\colon [a] \to [a+1]$ whose image does not contain j, define $F(\varphi^{\mathrm{op}}) = d_j^{a+1}$. Then extend F to all morphisms in Δ_{inj} by composition.

EXERCISE 2.16. Check that this is well-defined: it's not clear that if $\varphi_1 \circ \varphi_2 = \psi_1 \circ \psi_2$, then $F(\varphi_1) \circ F(\varphi_2) = F(\psi_1) \circ F(\psi_2)$. [HINT: Show that each morphism $[b] \to [a]$ in $\Delta_{\mathrm{inj}}^{\mathrm{op}}$ can be written uniquely in the form $(\varphi_{i_1}^b)^{\mathrm{op}} \circ \cdots \circ (\varphi_{i_{b-a}}^{a+1})^{\mathrm{op}}$ such that $i_1 \geq i_2 \geq \cdots \geq i_{b-a}$.]

Conversely, if F is a functor from $\Delta_{\mathrm{inj}}^{\mathrm{op}}$, Exercise 2.15 shows that the maps induced by $(\varphi_i^a)^{\mathrm{op}}$ satisfy the simplicial identities. All told, then, we've transformed semisimplicial sets from a combinatorial idea to category-ey one:

PROPOSITION 2.17. *A semisimplicial set is a functor* $\Delta_{\mathrm{inj}}^{\mathrm{op}} \to \mathsf{Set}$.

Of course, we can now look at its category: $\mathrm{Fun}(\Delta_{\mathrm{inj}}^{\mathrm{op}}, \mathsf{Set})$, which we'll call, unsurprisingly, the category of semisimplicial sets. A morphism in this category is a natural transformation; in other words, it's a collection of maps from the objects of one semisimplicial complex to another that commute with the maps d_i^n.

EXAMPLE 2.18. Remember that $\mathrm{Sing}(X)$ is the semisimplicial set derived from X (see Definition 1.6). We can extend this to continuous maps in this way: If $f\colon X \to Y$, then $\mathrm{Sing}(f)$ is the natural transformation with constituent maps $\mathrm{Sing}(f)_{[n]} = \mathrm{Sing}_n(f)$. In other words, it's exactly what you'd expect (here, X_n is short for $\mathrm{Sing}_n(X)$):

$$\begin{array}{ccccccccc}
\cdots & \xrightarrow{d_i} & X_2 & \xrightarrow{d_i} & X_1 & \xrightarrow{d_i} & X_0 & \xrightarrow{d_i} & 0 \\
 & & \downarrow{\scriptstyle \mathrm{Sing}_2(f)} & & \downarrow{\scriptstyle \mathrm{Sing}_1(f)} & & \downarrow{\scriptstyle \mathrm{Sing}_0(f)} & & \downarrow{\scriptstyle 0} \\
\cdots & \xrightarrow[d_i]{} & Y_2 & \xrightarrow[d_i]{} & Y_1 & \xrightarrow[d_i]{} & Y_0 & \xrightarrow[d_i]{} & 0
\end{array}$$

This relies crucially on the fact that d_i is a natural transformation from Sing_n to Sing_{n-1}!

We claim that Sing, as it's defined here, is actually a functor (from Top to $\mathrm{Fun}(\Delta^{\mathrm{op}}_{\mathrm{inj}}, \mathsf{Set})$). The proof of this is essentially recursion to previous results. Is $\mathrm{Sing}(1_X) = 1_{\mathrm{Sing}(X)}$? Well, since Sing_n is a functor for every n, then $\mathrm{Sing}_n(1_X) = 1_{\mathrm{Sing}_n(X)}$, which means that, yes, $\mathrm{Sing}(1_X)$ is the identity on $\mathrm{Sing}(X)$. And Sing distributes across composition because each of the Sing_n do. So there you go: A functor from Top to the set of semisimplicial sets. $\Diamond$

EXAMPLE 2.19. The function S_n forms an abelian group from a topological space. We can extend it to a functor as follows. We can define a group homomorphism from $S_n(X)$ to $S_n(Y)$ by designating its values on the basis $\mathrm{Sing}_n(X)$. So, given a continuous map $f\colon X \to Y$, define $S_n(f)$ as the group homomorphism that sends $\sigma\colon \Delta^n \to X$ to $f \circ \sigma\colon \Delta^n \to Y$. Clearly $S_n(1_X) = 1_{S_n(X)}$, and it's easy to check that S_n distributes over composition. So $S_n\colon \mathsf{Top} \to \mathsf{Ab}$. $\Diamond$

EXERCISE 2.20. Show that Z_n, B_n, and H_n also extend to functors $\mathsf{Top} \to \mathsf{Ab}$. (Make sure to check that the map $\sigma \mapsto f \circ \sigma$ sends each element $Z_n(X)$ to an element of $Z_n(Y)$, each element of $B_n(X)$ to an element of $B_n(Y)$, and each element of $H_n(X)$ to an element of $H_n(Y)$.)

Our ultimate goal is the homology groups. To get one step closer, let's introduce one more specialty category.

DEFINITION 2.21. The objects of the category Fil are the integers, and it has exactly one morphism from m to n if $m \geq n$ and no morphism from m to n if $m < n$.[5]

Quick check: Does this category actually exist? I've asserted that it has, but I haven't given any explicit morphisms so that you can check the axioms of a category. But actually building it is not too hard: If φ_a is the morphism from m to $m-1$, then the unique morphism from m to n (when $m < n$) is $\varphi_{n+1} \circ \varphi_{n+2} \circ \cdots \circ \varphi_m$. (When $m = n$, the unique morphism is, of course, the identity 1_m.)

We can think of a functor $A\colon \mathsf{Fil} \to \mathsf{Ab}$ as a diagram that looks like this:

$$\cdots \xrightarrow{\partial_3} A_2 \xrightarrow{\partial_2} A_1 \xrightarrow{\partial_1} A_0 \xrightarrow{\partial_0} A_{-1} \xrightarrow{\partial_{-1}} \cdots$$

where each of the A_i is an abelian group and each of the ∂_i is a group homomorphism.

DEFINITION 2.22. A *chain complex* of abelian groups is a functor $\mathsf{Fil} \to \mathsf{Ab}$ such that $\partial_i \circ \partial_{i+1} = 0$ for every $i \in \mathbb{Z}$.

EXERCISE 2.23. Verify that $\partial_i \circ \partial_{i+1} = 0$ if and only if $\mathrm{im}(\partial_{i+1}) \subseteq \ker(\partial_i)$.

Now, of course, we group all of these into a category.

DEFINITION 2.24. We denote by chAb the category of chain complexes, which is a subcategory of $\mathrm{Fun}(\mathsf{Fil}, \mathsf{Ab})$.

[5] "Fil" is shorthand for *filtration*, though I've no idea why. I imagine the answer is a quick web search away.

In particular, the morphisms in chAb are natural transformations. A morphism $f\colon A_\bullet \to B_\bullet$ looks like this:

$$\begin{array}{ccccccccccc}
A_\bullet & & \cdots & \longrightarrow & A_2 & \longrightarrow & A_1 & \longrightarrow & A_0 & \longrightarrow & A_{-1} & \longrightarrow & \cdots \\
{\scriptstyle f}\big\downarrow & & & & \big\downarrow & & \big\downarrow & & \big\downarrow & & \big\downarrow & & \\
B_\bullet & & \cdots & \longrightarrow & B_2 & \longrightarrow & B_1 & \longrightarrow & B_0 & \longrightarrow & B_{-1} & \longrightarrow & \cdots
\end{array}$$

In other words, it's a collection of group homomorphisms $A_i \to B_i$ so that every square commutes.

EXAMPLE 2.25. Given any semisimplicial set $\mathcal{X}$, we can form the chain complex

$$\cdots \xrightarrow{\partial_3} \mathbb{Z}X_2 \xrightarrow{\partial_2} \mathbb{Z}X_1 \xrightarrow{\partial_1} \mathbb{Z}X_0 \xrightarrow{\partial_0} 0 \xrightarrow{\partial_{-1}} 0 \xrightarrow{\partial_{-2}} \cdots$$

where $\partial_n = \sum_{k=0}^{n}(-1)^k d_k$ if $n \geq 0$ and $\partial_n = 0$ if $n < 0$. (Here $\mathbb{Z}X_i$ denotes the free abelian group generated by the elements of X_i.) We'll call this map $S_\bullet$. We can extend this to a functor as follows.

A morphism $f\colon \mathcal{X} \to \mathcal{Y}$ of semisimplicial sets is a collection of morphisms $f_i\colon X_i \to Y_i$ such that $d_k \circ f_i = f_{i-1} \circ d_k$ for every k and i. (This is what it means for every square to commute.) Each map f_i extends uniquely to a group homomorphism $\bar{f}_i\colon \mathbb{Z}X_i \to \mathbb{Z}Y_i$, and it's straightforward to check that $\partial_k \circ \bar{f}_i = \bar{f}_{i-1} \circ \partial_k$. So these maps comprise a morphism between chain complexes—in short, $S_\bullet\colon \operatorname{Fun}(\Delta_{\text{inj}}^{\text{op}}, \mathsf{Set}) \to \mathsf{Ab}$ is a functor. ♢

EXAMPLE 2.26. We can extend the map Z_n to arbitrary chain complexes by analogy: It takes in a chain complex $A_\bullet$ and outputs the kernel of ∂_n. In other words, it's a map from the objects of chAb to the objects of Ab. Does it extend to a functor? Well, suppose $f\colon A_\bullet \to B_\bullet$, so we get a diagram like this:

$$\begin{array}{ccccccc}
\cdots & \xrightarrow{\partial_{n+1}^A} & A_n & \xrightarrow{\partial_n^A} & A_{n-1} & \xrightarrow{\partial_{n-1}^A} & \cdots \\
 & & \big\downarrow{\scriptstyle f_n} & & \big\downarrow{\scriptstyle f_{n-1}} & & \\
\cdots & \xrightarrow[\partial_{n+1}^B]{} & B_n & \xrightarrow[\partial_n^B]{} & B_{n-1} & \xrightarrow[\partial_{n-1}^B]{} & \cdots
\end{array}$$

We'd like to define $Z_n(f)$ to be the restriction $f_n|_{\ker(\partial_n^A)}$, but we need to check that this is well-defined: Is $f_n()$ actually a subset of $\ker(\partial_n^B)$?

This is relatively straightforward to check: Choose any $x \in \ker(\partial_n^A)$. Using the diagram's commutativity, we have

$$\partial_n^B\big(f_n(x)\big) = f_{n-1}\big(\partial_n^A(x)\big) = f_{n-1}(0) = 0,$$

so $x \in \ker(\partial_n^B)$. So indeed, we can restrict f_n to be a homomorphism $Z_n(A_\bullet) \to Z_n(B_\bullet)$, which makes Z_n a functor. ♢

EXERCISE 2.27. Show that B_n and H_n also extend to functors $\mathsf{chAb} \to \mathsf{Ab}$.

3. SOME EXPLICIT HOMOLOGY GROUPS

Let's take a brief break from the extreme abstraction and get our hands dirty. Our eventual goal is to be able to (fairly) easily calculate $H_i(X)$ for many different topological spaces, since we want to be able to use them as an effective invariant.

3.1. THIS SECTION HAS A POINT

Let's begin with the easiest possible topological space: A single point $X = \{0\}$. There's only a single function $\Delta^n \to X$; we'll call this function a_n. This means that $\operatorname{Sing}_n(X) = \{a_n\}$, and there's really no choice for the simplicial maps: d_i^n is defined by $a_n \mapsto a_{n-1}$ for every n and i.

So $S_n(X)$ is also quite simple: it's the free group on the single element a_n. Our chain complex looks like this:

$$\cdots \xrightarrow{\partial_3} \mathbb{Z}a_2 \xrightarrow{\partial_2} \mathbb{Z}a_1 \xrightarrow{\partial_1} \mathbb{Z}a_0 \xrightarrow{\partial_0} 0.$$

The boundary maps, too, are easy to calculate:

$$\begin{aligned}
\partial_0(a_0) &= 0 \\
\partial_1(a_1) &= a_0 - a_0 = 0 \\
\partial_2(a_2) &= a_1 - a_1 + a_1 = a_1 \\
\partial_3(a_3) &= a_2 - a_2 + a_2 - a_2 = 0 \\
&\cdots
\end{aligned}$$

so $\partial_n(a_n) = 0$ if n is odd or $n = 0$ and $\partial_n(a_n) = a_{n-1}$ if $n \geq 2$ is even. So

$$Z_n(X) = \begin{cases} \mathbb{Z}a_n & \text{if } n \text{ is odd or } n = 0 \\ 0 & \text{if } n \geq 2 \text{ is even} \end{cases} \quad \text{and} \quad B_n(X) = \begin{cases} 0 & \text{if } n \text{ is odd or } n = 0 \\ \mathbb{Z}a_{n-1} & \text{if } n \geq 2 \text{ is even} \end{cases}$$

Putting it all together, we have

$$\begin{aligned}
H_0(X) &= \mathbb{Z}a_0/0 &&= \mathbb{Z}a_0 \\
H_{2n}(X) &= 0/0 &&= 0 \\
H_{2n+1}(X) &= \mathbb{Z}a_{2n+1}/\mathbb{Z}a_{2n+1} &&= 0.
\end{aligned}$$

In other words, $H_n(X) = 0$ if $n > 0$, which makes sense—we certainly don't think of a point as having any holes. But what's up with $H_0(X)$? Does it mean a point has a 1-dimensional hole (whatever that means)?[6] Nothing so bizarre as that.

Exercise 3.1. Let X be a topological space and $\pi_0(X)$ be the set of path-connected components of X. Prove that $H_0(X)$ is isomorphic to the free abelian group generated by the elements of $\pi_0(X)$. (More concisely: $\mathbb{Z}\pi_0(X) \cong H_0(X)$.)

Exercise 3.2. Let $\bigcup_{i\in I} X_i$ denote the disjoint union of the topological spaces $\{X_i\}$. Prove that $H_n\Big(\bigcup_{i\in I} X_i\Big) \cong \bigoplus_{i\in I} H_n(X_i)$ for every n.

3.2. A bit of topology

When are two maps "the same"? Well, when they're identical. But we can adopt a looser perspective once we put on our topology hats.

Definition 3.3. Let X and Y be topological spaces and $f, g\colon X \to Y$ be two continuous maps. A *homotopy* from f to g is a continuous function $h\colon X \times [0,1] \to Y$ such that $h(x,0) = f(x)$ and $h(x,1) = g(x)$. If a homotopy from f to g exists, then we say that f and g are *homotopic* and write $f \simeq g$.

One nice way to think of a homotopy is as a continuous deformation of one continuous function to another, where the interval $[0,1]$ indicates a time parameter. "Deformation" is a good way to think of homotopy: We can stretch and contract things, but we can't tear them.

Example 3.4. Let $f\colon [0,1] \to [0,1]$ be defined by $f(x) = 0$ for all $x \in [0,1]$ and let ι denote the identity function on $[0,1]$. Then actually $f \simeq \iota$, since

$$h(x,t) = tx$$

is a homotopy from f to ι. ◊

[6] Remember that we think of $H_n(X)$ as an account of the $(n+1)$-dimensional holes.

In fact, this example can be generalized: Let X be any topological space and $Y \subseteq \mathbb{R}^n$ be a convex set equipped with the subspace topology. Then *any* two maps $f, g\colon X \to Y$ are homotopic, since

$$h(x,t) = tf(x) + (1-t)g(x)$$

is a homotopy between them.

Of course, not all spaces are convex, and not all functions are homotopic.

EXAMPLE 3.5. Equip $Y = \{0,1\}$ with the discrete topology and let X be any topological space. We define $f, g\colon X \to Y$ to be the constant functions on 0 and 1 respectively. There is no homotopy from f to g. To show this, fix a point $x \in X$ and take any function $h\colon X \times [0,1] \to Y$. Consider the function $\varphi(t) = h(x,t)$. Since $\varphi^{-1}(0) \sqcup \varphi^{-1}(1) = [0,1]$, at least one of these sets is not open. But that means that φ is not continuous, which means that h is not continuous—and therefore not a homotopy. $\lozenge$

EXERCISE 3.6. Show that $\simeq$ is an equivalence relation on the set of continuous maps between two fixed topological spaces.

Okay, so what's the deal with homotopy? One reason is that it provides a looser, but still reasonable, definition of similarity.

DEFINITION 3.7. Two topological spaces X and Y are called *homotopy equivalent* if there are two continuous maps $f\colon X \to Y$ and $g\colon Y \to X$ such that $g \circ f \simeq 1_X$ and $f \circ g \simeq 1_Y$.

Certainly any homeomorphic spaces are homotopy equivalent, but many more pairs are homotopy equivalent than are homeomorphic. For example, the you can use Example 3.4 to show that a single point is homotopy equivalent to a closed segment, but these two spaces are definitely not homeomorphic—there's not even a bijection from one to the other!

EXERCISE 3.8. Verify that $[0,1]$ is homotopy equivalent to a single point.

But here's another reason, more firmly settled in algebraic topology.

THEOREM 3.9. *Suppose f and g are two morphisms $X \to Y$ in* Top, *and let H_n denote the canonical functor* $\mathsf{Top} \to \mathsf{chAb}$. *If $f \simeq g$, then $H_n(f) = H_n(g)$ for every $n \in \mathbb{N}_0$.*

In other words, when f and g are "similar" in this particular way, then the morphisms of chain complexes that they induce are actually equal. Theorem 3.9 is usually proved by showing that $S_n(gd)$ and $S_n(g)$ are something called "chain homotopic" and then applying Proposition 3.13.

From this, we can form a new category. (And there I was telling you that this was going to be concrete. Sorry about that.)

DEFINITION 3.10. The *homotopy category* HoTop has the same objects as Top, but the morphisms from X to Y are the homotopy equivalence classes of continuous functions from X to Y; symbolically:

$$\mathrm{Hom}_{\mathsf{HoTop}}(X,Y) = \mathrm{Hom}_{\mathsf{Top}}(X,Y)/\simeq.$$

There's a canonical functor $\Phi\colon \mathsf{Top} \to \mathsf{HoTop}$ that sends each object to itself and each morphism to its equivalence class. We can restate Theorem 3.9 in terms of this category by saying: There exists a functor $\mathsf{HoTop} \to \mathsf{Ab}$ such that the following diagram commutes for each $n \in \mathbb{N}_0$.

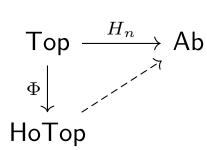

Because we bond by overloading notation in this class, we'll also call this functor H_n.

We can rephrase homotopy equivalence in this category: A continuous map $f\colon X \to Y$ in Top is a homotopy equivalence if $\Phi(f)$ is an isomorphism, and two spaces X and Y are homotopy equivalent if they're isomorphic in HoTop.

COROLLARY 3.11. *If X and Y are homotopy equivalent, then $H_n(X) = H_n(Y)$ for every $n \in \mathbb{N}_0$.*

This leads us to a Guiding Principle of algebraic topology:

> Homology does not distinguish homotopic maps or homotopy equivalent spaces.

Let's next look at how we can use this to actually compute homologies.

3.3. A LITTLE BIT OF MAGIC

Here's a definition.

DEFINITION 3.12. Suppose that $C_\bullet$ and $D_\bullet$ are chain complexes and $f_0, f_1\colon C_\bullet \to D_\bullet$. A *chain homotopy* $h\colon f_0 \simeq f_1$ is a collection of homomorphisms $C_n \to D_{n+1}$ such that $\partial h + h\partial = f_1 - f_0$.

You could now reasonably ask the question: Wait, what? Here's what a chain homotopy looks like:

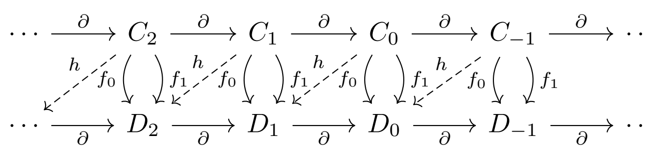

where the diagram doesn't commute, but instead satisfies this strange relationship that $\partial h + h\partial = f_1 - f_0$.

Of course, the real problem is that it's not at all clear why anyone should care about such a thing. In our case, the best motivation is simply that it works.

PROPOSITION 3.13. *If $f_0, f_1\colon C_\bullet \to D_\bullet$ are chain homotopic, then $H_n(f_0)$ and $H_n(f_1)$ are identical group homomorphisms $H_n(C_\bullet) \to H_n(D_\bullet)$ for every $n \in \mathbb{N}_0$.*

Proof. We need to show that $\big(H_n(f_0)\big)(c) = \big(H_n(f_1)\big)(c)$ for every $c \in Z_n(C_\bullet)$. This is equivalent to saying that $f_0(c) - f_1(c) \in B_n(C_\bullet)$ for every $c \in Z_n(C_\bullet)$. We just unpack the definitions:

$$\begin{aligned} f_1(c) - f_0(c) &= (\partial h + h\partial)(c) \\ &= \partial h(c) + h\partial(c) \\ &= \partial\big(h(c)\big) + 0, \end{aligned}$$

which is definitely an element of $B_n(C_\bullet)$. □

Okay, so now we have a not-very-enlightening proof of a proposition that, frankly, it's not entirely clear why we care about. Here's why we care: Chain homotopies allow us to calculate the homology groups of many more spaces. For example, consider this.

DEFINITION 3.14. A set $X \subseteq \mathbb{R}^n$ is called *star-shaped* with respect to a point $b \in X$ if the interval $\{(tb + (1-t)x : t \in [0,1]\}$ is contained in X for every point $x \in X$. (It's sort of "convex from a base point.")

EXERCISE 3.15. Let X be a star-shaped region with respect to the point b (equipped with the subset topology) and $Y = \{b\}$ be a one-point topological space. Prove that X is homotopy equivalent to Y.

If you believe Theorem 3.9, then Exercise 3.15 implies that the homology groups of any star-shaped region are the same as the homology groups of a point. The reason this is true is that homotopic maps induce chain homotopic chain complexes. We'll delay a proof of this, mostly because we need this exact result in the proof. So for now, we'll prove that star-shaped regions and a single point have the same homology groups by directly using chain homotopies.

Theorem 3.16. *If X is star-shaped, then $H_n(X) \cong H_n(\bullet)$ for every $n \in \mathbb{N}_0$.*

Proof. Let b be a point that X is star-shaped with respect to. Let $A_\bullet$ denote the chain complex where $A_n = \{0\}$ when $n \neq 0$ and $A_0 = \mathbb{Z}$ (there's only one possible set of boundary maps for this complex). We define two maps $\epsilon \colon S_\bullet(X) \to A$ and $\eta \colon A_\bullet \to S_\bullet(X)$.

$$\begin{array}{ccccccccccc} \cdots & \longrightarrow & S_2 & \longrightarrow & S_1 & \longrightarrow & S_0 & \longrightarrow & 0 & \longrightarrow & \cdots \\ & & \eta \uparrow \downarrow \epsilon & & \eta \uparrow \downarrow \epsilon & & \eta_0 \uparrow \downarrow \epsilon_0 & & \eta \uparrow \downarrow \epsilon & & \\ \cdots & \longrightarrow & 0 & \longrightarrow & 0 & \longrightarrow & \mathbb{Z} & \longrightarrow & 0 & \longrightarrow & \cdots \end{array}$$

There's only one way to define these maps except for η_0 and ϵ_0, which we define by $\epsilon_0(x) = 1$ for every $x \in X$; and $\eta_0(1) = b$.

We claim that $H_n(\eta)$ and $H_n(\epsilon)$ are function inverses. If this is the case, then $H_n(X) \cong H_n(A_\bullet)$ for every $n \in \mathbb{N}_0$, so $H_n(X) \cong H_n(\bullet)$. Let's get to it.

First, $\epsilon \circ \eta$ is the identity $1_{A_\bullet}$; since H_n is a functor, this means that $H_n(\epsilon) \circ H_n(\eta) = 1_{H_n(A_\bullet)}$. Now the reverse: We prove that $\eta \circ \epsilon \simeq 1_{S_\bullet(X)}$. We do this by defining a chain homotopy $h \colon \eta \circ \epsilon \simeq 1_{S_\bullet(X)}$. Such a map looks like this:

$$\begin{array}{ccccccccc} \cdots & \longrightarrow & S_{n+1} & \longrightarrow & S_n & \longrightarrow & S_{n-1} & \longrightarrow & \cdots \\ & \swarrow_{h_{n+1}} & & \swarrow_{h_n} & & \swarrow_{h_{n-1}} & & \swarrow_{h_{n-2}} & \\ \cdots & \longrightarrow & S_{n+1} & \longrightarrow & S_n & \longrightarrow & S_{n-1} & \longrightarrow & \cdots \end{array}$$

where $\delta h + h\delta = 1_{S_\bullet(X)} - \eta \circ \epsilon$. Each map should take in a linear combination of functions $\sigma \colon \Delta^n \to X$ and output a linear combination of functions $\Delta^{n+1} \to X$. Here's how we define it:

$$h_n(\sigma) \colon (t_0, t_1, \ldots, t_{n+1}) \mapsto t_0 b + (1 - t_0)\sigma\left(\frac{1}{1 - t_0}(t_1, \ldots, t_{n+1})\right).$$

The denominator makes the entries sum to 1 so that we can apply σ. Essentially, h_n turns an n-simplex into an $(n+1)$-simplex by adding b as new vertex. We can calculate the effect of d_0 on h by substituting $t_0 = 0$: We get $d_0 h(\sigma) = \sigma$. Likewise, by substituting $t_i = 0$ for some $1 \leq i \leq n$, we find that $d_i h(\sigma) = h\big(d_{i-1}(\sigma)\big)$.

Plugging this into the boundary operator formula, we find that, when $n \geq 1$,

$$\partial_n\big(h(\sigma)\big) = \sigma - h\big(\partial_{n-1}(\sigma)\big);$$

when $n = 0$, the formula is

$$\partial_0\big(h(\sigma)\big) = \sigma - b.$$

These two formulas combine in the single formula

$$(\partial_n h + h\partial_n)(\sigma) = \sigma - (\eta\epsilon)(\sigma).$$

Or, in other words, $\partial h + h\partial = 1_{S_\bullet(X)} - \eta\epsilon$, so h is a homotopy. □

Don't worry if that proof doesn't make much sense. Basically you do a some things, say some magic words, and the theorem is proved. That is, unfortunately, just the way this proof is.

3.4. DEFORMATION RETRACTS

Let's talk about something more concrete.

DEFINITION 3.17. A subset $A \subseteq X$ (equipped with the subset topology) is a *deformation retract* if there is a continuous map $h\colon X \times [0,1] \to X$ such that

- $h(x,0) = x$ for every $x \in X$,
- $h(x,1) \in A$ for every $x \in X$, and
- $h(a,t) = a$ for every $a \in A$ and $t \in [0,1]$.

The idea is that a deformation retract is a kind of "nice compression" of the bigger space into the smaller one. For example:

EXERCISE 3.18. Show that the point $\{0\}$ is a deformation retract of $[0,1]$.

EXAMPLE 3.19. The n-sphere S^n is a deformation retract of the space $\mathbb{R}^{n+1} \setminus \{0\}$ via the map $h(x,t) = t\frac{x}{|x|} + (1-t)x$. ♢

Deformation retracts are nice because they preserve homology groups, and they're usually one of the easier maps to find.

EXERCISE 3.20. Prove that any deformation retract of X is homotopy equivalent to X.

For example, the central point is a deformation retract of any star-shaped region, so their homologies are equal.

4. HOMOLOGY GETS BIGGER AND BETTER

⚠ Warning! This section was made in haste while I was tired. Expect a higher proportion of errors and oversights and a lower proportion of motivation and explanation.

4.1. RELATIVE HOMOLOGY AND THE LONG EXACT SEQUENCE

DEFINITION 4.1. Suppose that $C_\bullet$ is a chain complex and $S_i \subseteq C_i$ for every $i \in \mathbb{Z}$. If $S_\bullet$ forms a chain complex with the maps induced from $C_\bullet$, then $S_\bullet$ is called a *subcomplex* of $C_\bullet$.

EXERCISE 4.2. Suppose that $D_\bullet$ is a subcomplex of $C_\bullet$. Show that the induced maps $\partial_n\colon C_n/D_n \to C_{n-1}/D_{n-1}$ are well-defined and form a chain complex.

The resulting chain complex is called the *quotient* of $C_\bullet$ by $D_\bullet$, denoted $C_\bullet/D_\bullet$. Here's why we introduce it:[7]

EXERCISE 4.3. Suppose that $D_\bullet$ is a subcomplex of $C_\bullet$. Show that

$$0 \to D_n \to C_n \to C_n/D_n \to 0$$

is a short exact sequence for every n.

DEFINITION 4.4. If $B_\bullet$, $C_\bullet$, and $D_\bullet$ are three chain complexes such that

$$0 \to B_n \to C_n \to D_n \to 0$$

is exact for every n, then we say that the sequence

$$0 \to B_\bullet \to C_\bullet \to D_\bullet \to 0$$

is *exact*.

[7] At this point you should look up what an exact sequence is if you don't know.

So $C_\bullet/D_\bullet$ somehow "completes" the short exact sequence. In fact, this extends to homology:

THEOREM 4.5. *If* $0 \to B_\bullet \xrightarrow{f} C_\bullet \xrightarrow{g} D_\bullet \to 0$ *is exact, then there is an infinite exact sequence*

$$\cdots \xrightarrow{H_{n+1}(g)} H_{n+1}(D) \to H_n(B) \xrightarrow{H_n(f)} H_n(C) \xrightarrow{H_n(g)} H_n(D) \to H_{n-1}(B) \xrightarrow{H_{n-1}(f)} \cdots$$

This sequence is called the long exact sequence*.*

COROLLARY 4.6. *If* $A \subseteq X$*, then there is an exact sequence*

$$\cdots \to H_{n+1}(X, A) \to H_n(A) \to H_n(X) \to H_n(X, A) \to H_{n-1}(A) \to \cdots$$

with the maps $H_n(A) \to H_n(X)$ *and* $H_n(X) \to H_n(X, A)$ *induced by* $S_\bullet(A) \hookrightarrow S_\bullet(X)$ *and* $S_\bullet(X) \to S_\bullet(X)/S_\bullet(A)$*.*

I won't prove this theorem; you can instead see a proof here. In essence, it relies on a double application of the Snake Lemma, which is itself proved via diagram chase.

In practice, it turns out that you don't really need to know much about the degree-lowering maps $H_n(X, A) \to H_{n-1}(A)$ except the fact that they make the sequence exact. We'll see examples of this in the next section. The maps themselves are often colloquially called the "snake maps," which is why I'll denote them by 🐍$_n$.

For now, let's compute a very simple relative homology: $H_n(X, \{p\})$, where p is a point in X. To do this, we need to get a handle on the sequence

$$\cdots \xrightarrow{\partial_{n+2}} S_{n+1}(X)\big/S_{n+1}(p) \xrightarrow{\partial_{n+1}} S_n(X)\big/S_n(p) \xrightarrow{\partial_n} S_{n-1}(X)\big/S_{n-1}(p) \xrightarrow{\partial_{n-1}} \cdots$$

We'll use ∂_n^S to denote the underlying boundary maps $S_n(X) \to S_{n-1}(X)$ and σ_n to denote the unique map $\Delta^n \to \{p\}$. Since $d_i(\sigma_n) = \sigma_{n-1}$ for every $0 \le i \le n$, we have

$$\partial_n^S(\sigma_n) = \begin{cases} 0 & \text{if } n \text{ is odd} \\ \sigma_{n-1} & \text{if } n \text{ is even.} \end{cases}$$

To evaluate the kernel of ∂_n, we consider instead the map ∂_n^S. Let c denote an arbitrary element of $S_n(X)$ and $a \in \mathbb{Z}$. For every $n > 0$:

- If n is odd, then $\partial_n^S(c + a\sigma_n) = \partial_n^S(c)$, so the equivalence class $c + S_n(p) \in \ker(\partial_n)$ if and only if every element of $c + S_n(p)$ is. This means that $\ker(\partial_n) = \ker(\partial_n^S)/S_n(p)$.
- If n is even, then $\partial_n^S(c + a\sigma_n) = \partial_n^S(c) + a\sigma_{n-1}$. So for each class $c + S_n(p) \in \ker(\partial_n)$, there is exactly one element $c + b\sigma_n$ such that $\partial_n^S(c + b\sigma_n) = 0$. This means that $\ker(\partial_n) = \big(\ker(\partial_n^S) \oplus S_n(p)\big)/S_n(p)$.

You can use the same sort of reasoning to figure out that

- If n is odd, then $\operatorname{im}(\partial_n) = \big(\operatorname{im}(\partial_n) \oplus S_{n-1}(p)\big)/S_{n-1}(p)$.
- If n is even, then $\operatorname{im}(\partial_n) = \operatorname{im}(\partial_n^S)/S_{n-1}(p)$.

We can conclude that, if $n > 0$, we have

$$H_n(X, \{p\}) = \frac{\ker(\partial_n^S)/S_n(p)}{\operatorname{im}(\partial_{n+1}^S)/S_n(p)} \cong H_n(X)$$

if n is odd and

$$H_n(X, \{p\}) = \frac{\big(\ker(\partial_n^S) \oplus S_n(p)\big)/S_n(p)}{\big(\operatorname{im}(\partial_n) \oplus S_{n-1}(p)\big)/S_{n-1}(p)} \cong H_n(X)$$

if n is even. In short, if $n > 0$, then $H_n(X, \{p\}) = H_n(X)$. What about $n = 0$? Well, $\ker(\partial_0^S) = S_0(X)$, so $\ker(\partial_0) = \ker(\partial_0^S)/S_0(p)$; and we already know that $\operatorname{im}(\partial_1^S) = \big(\operatorname{im}(\partial_1) \oplus S_0(p)\big)/S_0(p)$. So

$$H_0(X, \{p\}) = \frac{\ker(\partial_0^S)/S_0(p)}{\big(\operatorname{im}(\partial_1) \oplus S_0(p)\big)/S_0(p)} \cong \frac{\ker(\partial_0^S)}{\operatorname{im}(\partial_1) \oplus S_0(p)}.$$

Now, $S_0(p) \cong \mathbb{Z}$, which means that

$$H_0(X) \cong H_0(X, \{p\}) \oplus \mathbb{Z}.$$

We sum this up in a nice proposition:

PROPOSITION 4.7. *The relative homology of a point $p \in X$ is given by*

$$H_n(X) \cong \begin{cases} H_n(X, \{p\}) & \text{if } n > 0 \\ H_n(X, \{p\}) \oplus \mathbb{Z} & \text{if } n = 0. \end{cases}$$

EXERCISE 4.8. Use the long exact sequence to show that $H_n(X, \{p\}) = H_n(X)$ for every $n > 1$. [HINT: You already know the homology of a point.]

As a last note, we can, essentially as a reflex at this point, assemble this structure into a category:

DEFINITION 4.9. The category Top_2 consists of the objects (X, A) where $X \in \mathsf{Top}$ and $A \subseteq X$, and the morphisms from (X, A) to (Y, B) consist of the continuous functions $f : X \to Y$ such that $f(A) \subseteq B$.

This category contains Top as a subcategory: Just restrict to the elements of the form $(X, \emptyset)$. Moreover, H_n is a functor $\mathsf{Top}_2 \to \mathsf{Ab}$ for every n.

4.2. EXCISION

Let's give some thought to the geometric pictures that correspond to elements of $H_n(X, A)$. Elements of $Z_n\big(S_n(X)/S_n(A)\big)$ are n-chains whose boundary lies in the set A. And elements of $B_n\big(S_n(X)/S_n(A)\big)$ are the usual boundaries, modulo the parts in $S_{n-1}(A)$. So it seems that $H_n(X, A)$ behaves very much like we would expect $H_n(X/A)$ to. Indeed, that is the content of the Excision Theorem.

THEOREM 4.10 (Excision). *Suppose that X is a topological space and $A \subseteq B \subseteq X$ such that $\bar{A} \subseteq \operatorname{int}(B)$ and A is a deformation retraction of B. In this case,*

$$H_n(X, A) = H_n(X/A, \{p\})$$

for every $n \in \mathbb{N}_0$.

The conditions on A and B are relatively mild; for most spaces, we'll have no trouble finding such a B. It essentially guarantees that A isn't "too nasty."

How is the excision theorem proved? Well, actually, it uses a different version of the theorem; we might call Theorem 4.10 "Excision in practice" and Theorem 4.11 "Excision in theory."

THEOREM 4.11. *If X is a topological space and $U \subseteq A \subseteq X$ such that $\bar{A} \subseteq \operatorname{int}(B)$, then the inclusion $(X \setminus U, A \setminus U) \hookrightarrow (X, A)$ induces a homology isomorphism*

$$H_n(X, A) \cong H_n(X \setminus U, A \setminus U).$$

We won't prove this theorem. If you want, you can find a proof in Haynes Miller's algebraic topology notes or a high-level overview on the Wikipedia page.

What we will do is show how Theorem 4.11 implies Theorem 4.10.

To do so, we'll need the an algebraic result:

LEMMA 4.12 (Five lemma). *If the following diagram of abelian groups is commutative, both rows are exact, and f_1, f_2, f_4, and f_5 are isomorphisms, then f_3 is, too.*

$$\begin{array}{ccccccccc} A_1 & \longrightarrow & A_2 & \longrightarrow & A_3 & \longrightarrow & A_4 & \longrightarrow & A_5 \\ \downarrow{\scriptstyle f_1} & & \downarrow{\scriptstyle f_2} & & \downarrow{\scriptstyle f_3} & & \downarrow{\scriptstyle f_4} & & \downarrow{\scriptstyle f_5} \\ B_1 & \longrightarrow & B_2 & \longrightarrow & B_3 & \longrightarrow & B_4 & \longrightarrow & B_5 \end{array}$$

The argument is just a bit of diagram chasing and is better relegated to a footnote.[8] Now we're ready.

Theorem 4.11 implies Theorem 4.10. Suppose that A and B satisfy the conditions of Theorem 4.10. We use the following diagram in Top_2:

$$\begin{array}{ccccc}
(X, A) & \xhookrightarrow{\;i\;} & (X, B) & \xhookleftarrow{\;j\;} & (X - A, B - A) \\
\downarrow{\scriptstyle f} & & \downarrow{\scriptstyle g} & & \downarrow{\scriptstyle h} \\
(X/A, \bullet) & \xrightarrow[\bar{\imath}]{} & (X/A, B/A) & \xleftarrow[\bar{\jmath}]{} & (X/A - \bullet, X/B - \bullet)
\end{array}$$

Here, we temporarily write $-$ to denote set subtraction for clarity, and $\bullet$ represents the point that A is contracted to in the quotient. The maps i and j are the inclusion maps and $\bar{\imath}$ and $\bar{\jmath}$ are the induced maps on the quotients. The maps f, g, and h are the projections onto the quotient spaces.

We want to show that f is a homology isomorphism. We know that j and $\bar{\jmath}$ are isomorphisms by Theorem 4.11, and h is an isomorphism in Top_2 since it's a homeomorphism in both coordinates.

Now we show that i is a homology isomorphism. We use the long exact sequence:

$$\begin{array}{ccccccccccc}
\cdots \longrightarrow & H_n(A) & \longrightarrow & H_n(X) & \longrightarrow & H_n(X, A) & \longrightarrow & H_{n-1}(A) & \longrightarrow & H_{n-1}(X) & \longrightarrow \cdots \\
& \downarrow{\scriptstyle i_1} & & \downarrow{\scriptstyle i_2} & & \downarrow{\scriptstyle i} & & \downarrow{\scriptstyle i_3} & & \downarrow{\scriptstyle i_4} & \\
\cdots \longrightarrow & H_n(B) & \longrightarrow & H_n(X) & \longrightarrow & H_n(X, B) & \longrightarrow & H_{n-1}(B) & \longrightarrow & H_{n-1}(X) & \longrightarrow \cdots
\end{array}$$

The maps i_1 and i_3 are homology isomorphisms because they are deformation retractions, and the maps i_2 and i_4 are identity maps. By the five lemma, i is a homology isomorphism. You can argue similarly that $\bar{\imath}$ is a homology isomorphism, since A is a deformation retract of B if and only if $\bullet$ is a deformation retract of B/A.

Since our initial diagram commutes, we have $f = \bar{\imath}^{-1} \circ \bar{\jmath} \circ h \circ \jmath^{-1} \circ i$; since each of these is a homology isomorphism, so is f. □

4.3. FUNDAMENTAL PROPERTIES OF HOMOLOGY: A LIST

Let's summarize what we have so far: There are functors $H_n \colon \mathsf{Top}_2 \to \mathsf{Ab}$ and natural transformations $\partial_n \colon H_n(X, A) \to H_{n-1}(A, \emptyset) =: H_{n-1}(A)$ for every $n \in \mathbb{N}_0$ such that

1. The sequence

$$\cdots \xrightarrow{\partial_{n+1}} H_n(A) \longrightarrow H_n(X) \longrightarrow H_n(X, A) \xrightarrow{\partial_n} H_n(A) \longrightarrow \cdots$$

   is exact for every $(X, A) \in \mathsf{Top}_2$. (The unmarked maps are induced by the homology functor.)
2. If $f_1, f_2 \colon (X, A) \to (Y, B)$ are homotopic, then $H_n(f) = H_n(g)$ for every n.
3. If $U \subseteq A \subseteq X$ and $\bar{U} \subseteq \operatorname{int}(A)$, then the inclusion $(X \setminus U, A \setminus U)$ induces a homology isomorphism $H_n(X, A) \cong H_n(X \setminus U, A \setminus U)$ for every n.
4. $H_n\left(\bigsqcup_{i \in I} X_i\right) = \bigoplus_{i \in I} H_n(X_i)$.
5. $H_n(\bullet) \cong \{0\}$ if $n > 0$ and $H_0(\bullet) \cong \mathbb{Z}$.

[8] For example, here's how you show that f_3 is injective, suppose that $f_3(x) = 0$. We let $\alpha_i \colon A_i \to A_{i+1}$ and $\beta_i \colon B_i \to B_{i+1}$. Then $\beta_3(f_3(x)) = 0$, so $f_4(\alpha_3(x)) = 0$; this means that there is an element $y \in a_2$ such that $\alpha_2(y) = x$ (by exactness). But then $\beta_2(f_2(y)) = f_3(\alpha_2(y)) = 0$, so there is an element $z \in B_1$ such that $\beta_1(z) = f_2(y)$. It has a unique preimage w in A_1. Now $f_2(\alpha_1(w)) = \beta_2(f_1(w)) = \beta_2(z) = f_2(y)$; since f_2 is an isomorphism, $\alpha_1(w) = y$. But then $x = \alpha_2(\alpha_1(w)) = 0$ (the composition $\alpha_i \circ \alpha_{i+1} = 0$ by exactness). So f_3 is injective.

(Actually we haven't defined homotopy in Top_2, but it's exactly what you would expect: f_1 and f_2 are homotopic if there is a map $h\colon X \times [0,1] \to Y$ such that $h(x,0) = f_1(x)$, $h(x,1) = f_2(x)$, and $h(a,t) \in B$ for every $a \in A$ and $t \in [0,1]$.)

These are the five basic properties of the homology operators, and a theorem by Eilenberg and Steenrod states that these five properties actually *characterize* the homology operators in sufficiently nice spaces (think, for example "spaces that can be triangulated"). For this reason, these properties are called the *Eilenberg–Steenrod axioms*. Other theories have been proposed that satisfy every axiom but the fifth—the so-called *dimension axiom*. These are called *extraordinary homology theories* and include K-theory, bordism, topological modular forms, and Morava E-theory.

But that's incidental. The takeaway here is that to calculate homologies of spaces, you really only need to know these five properties. So study them, take them to heart, and let's get our calculate on.

4.4. OUR FIRST NONTRIVIAL HOMOLOGIES

First up, the circle S^1. To take advantage of our results so far, we want to write S^1 as a quotient of two spaces we know the homologies of. That's not so hard: S^1 is homeomorphic to $[0,1]/\{0,1\}$ (taking the interval and gluing the endpoints together). So, set $X = [0,1]$ and $A = \{0,1\}$. We have a long exact sequence

$$\cdots \to H_n(\{0,1\}) \to H_n([0,1]) \to H_n([0,1],\{0,1\}) \to H_{n-1}(\{0,1\}) \to H_{n-1}([0,1]) \to \cdots$$

Since $\{0,1\}$ is the disjoint union of two points, its homology group is

$$H_n(\{0,1\}) = H_n(\bullet) \oplus H_n(\bullet) = \begin{cases} 0 & \text{if } n > 0 \\ \mathbb{Z}^2 & \text{if } n = 0. \end{cases}$$

And $[0,1]$ deformation retracts onto a point, so

$$H_n([0,1]) = H_n(\bullet) = \begin{cases} 0 & \text{if } n > 0 \\ \mathbb{Z} & \text{if } n = 0. \end{cases}$$

And if $n > 0$, we have $H_n([0,1],\{0,1\}) = H_n(S^1)$.

That is to say, for $n \geq 2$, our diagram looks like this:

$$\cdots \to 0 \to 0 \to H_n([0,1],\{0,1\}) \to 0 \to 0 \to \cdots$$

so $H_n(S^1) \cong 0$ if $n \geq 2$.

Now let's unpack the map $\varphi\colon H_0(\{0,1\}) \to H_0([0,1])$. The domain is the free abelian group generated by the maps $\Delta^0 \to 0$ and $\Delta^0 \to 1$. The latter is the free abelian group generated by the single equivalence class $[\Delta^0 \to p]$. (Recall Exercise 3.1.) We can fix the an isomorphism $H_0(\{0,1\}) \to \mathbb{Z} \oplus \mathbb{Z}$ by sending $(\Delta^0 \to 0) \mapsto (1,0)$ and $(\Delta^0 \to 1) \mapsto (0,1)$; we can also define the isomorphism $H_0([0,1]) \to \mathbb{Z}$ by sending the equivalence class $[\Delta^0 \to p] \mapsto 1$. In this parlance, then, the map φ is given by $(a,b) \mapsto a+b$.

The kernel of φ is therefore the free abelian group generated by $(1,-1)$. Focusing on $n = 1$ in the long exact sequence, we have

$$\cdots 0 \to 0 \to H_1(S^1) \to \mathbb{Z} \oplus \mathbb{Z} \xrightarrow{\varphi} \mathbb{Z} \to \cdots$$

The map $H_n(S^1) \to \mathbb{Z} \oplus \mathbb{Z}$ is injective (by exactness), and its image is the kernel of φ, which is isomorphic to $\mathbb{Z}$. So $H_1(S^1) = \mathbb{Z}$. And $H_0(S^1) \cong \mathbb{Z}$, since S^1 has exactly one path component. (Be careful, though! $H_0([0,1],\{0,1\}) \cong 0$.) Altogether, we have

PROPOSITION 4.13. *The homology of the circle is*

$$H_n(S^1) \cong \begin{cases} 0 & \text{if } n \geq 2 \\ \mathbb{Z} & \text{if } n = 0, 1 \end{cases}$$

EXERCISE 4.14. Show that

$$H_n(S^q) \cong \begin{cases} \mathbb{Z} & \text{if } n = 0, q \\ 0 & \text{otherwise.} \end{cases}$$

COROLLARY 4.15. *S^p and S^q are not homeomorphic if $p \neq q$.*

Proof. Any homeomorphism is a homotopy equivalence, so homeomorphic spaces have the same homology groups, which spheres of different dimensions do not. □

COROLLARY 4.16. *If $p \neq q$, then $\mathbb{R}^p$ and $\mathbb{R}^q$ are not homeomorphic.*

Proof. If $\mathbb{R}^p$ and $\mathbb{R}^q$ were homeomorphic, then $\mathbb{R}^p \setminus \{0\}$ and $\mathbb{R}^q \setminus \{0\}$ would be homeomorphic. But these spaces deformation retract onto S^p and S^q, which do not share their homology groups. □

Moreover, we already get a very famous theorem:

THEOREM 4.17 (Brouwer's fixed point theorem). *If $f\colon D^n \to D^n$ is continuous, then there is a point $x \in D^n$ such that $f(x) = x$.*

Inevitably, the person introducing this theorem mentions stirring coffee, but I'll refrain, because honestly, who stirs n-dimensional coffee?

Proof of Theorem 4.17. Suppose that no point is fixed, so that $f(x) \neq x$ for every $x \in D^n$. We define a new function $g\colon D^n \to S^{n-1}$ setting $g(x)$ to be the point where the ray from $f(x)$ to x hits the boundary of D^n. Note that g is continuous and $g(x) = x$ for every $x \in S^{n-1}$. So the composition

$$S^{n-1} \hookrightarrow D^n \xrightarrow{g} S^{n-1}$$

is the identity. But applying the functor H_{n-1} yields that $\mathbb{Z} \to 0 \to \mathbb{Z}$ is the identity, which is clearly false: Contradiction. □

4.5. LOCALITY AND ANOTHER THEOREM

The main part of proving Theorem 4.11 consists of proving the Locality Principle.

DEFINITION 4.18. A collection $\mathcal{A}$ of subsets of X is called a *cover* if $X = \bigcap_{A \in \mathcal{A}} \operatorname{int}(A)$. Given a cover, an n-simplex $\sigma\colon \Delta^n \to X$ is called $\mathcal{A}$-*small* if $\operatorname{im}(\sigma)$ lies entirely inside one set of $\mathcal{A}$. The set $\operatorname{Sing}_n^{\mathcal{A}}(X)$ is the set of all $\mathcal{A}$-small n-simplices, and $S_n^{\mathcal{A}}(X)$ is the free abelian group it generates.

THEOREM 4.19 (Locality principle). *If $\mathcal{A}$ is a cover of X, then the inclusion of chain complexes $S_\bullet^{\mathcal{A}}(X) \subseteq S_\bullet(X)$ is a homology isomorphism.*

In other words, it doesn't change the homology of a space if we only consider "small" simplices (as measured by $\mathcal{A}$). The proof of this is not so complicated in its overall shape: The basic idea is to replace each "big" n-simplex with a sum of smaller n-simplices whose boundaries sum to the original n-simplex. The difficulty is in doing this in an arbitrary topological space, and the details are somewhat involved.

But with it, we can prove a new theorem about sequence homologies that is at times easier to use than the excision theorem.

THEOREM 4.20 (Mayer-Vietoris). *If X is a topological space and $\mathcal{A} = \{A, B\}$ is a cover of X, then there is an exact sequence*

$$\cdots \xrightarrow{\partial_{n+1}} H_n(A \cap B) \xrightarrow{f_n} H_n(A) \oplus H_n(B) \xrightarrow{g_n} H_n(X) \xrightarrow{\partial_n} \cdots$$

whose maps are defined as follows: Given the inclusions

$$\begin{aligned} i &\colon A \cap B \hookrightarrow A \qquad & k &\colon A \hookrightarrow X \\ j &\colon A \cap B \hookrightarrow B \qquad & \ell &\colon B \hookrightarrow X, \end{aligned}$$

we have

$$\begin{aligned} f_n &= H_n(i) \oplus H_n(j) \\ g_n &= H_n(k) - H_n(\ell). \end{aligned}$$

The condition that $\{A, B\}$ is a cover of X is exactly the same as requiring that $\bar{A} \subseteq \operatorname{int}(B)$. Looks, familiar, eh?

Proof of Theorem 4.20. The maps are chosen exactly so that the sequence

$$0 \to S_\bullet(A \cap B) \xrightarrow{f} S_\bullet(A) \oplus S_\bullet(B) \xrightarrow{g} S_\bullet^{\mathcal{A}}(X) \to 0$$

is exact. Just expand this out into a long exact sequence of homology groups and use the locality principle to note that $H_n^{\mathcal{A}}(X) \cong H_n(X)$. □

4.6. MORE HOMOLOGY CALCULATIONS

We have a new tool; let's use it.

THE CIRCLE

We'll first calculate the homology of S^1 again using this new method. We need to choose A and B so that their interiors cover S^1 and so that we know their homologies and that of $A \cap B$. That's not so hard:

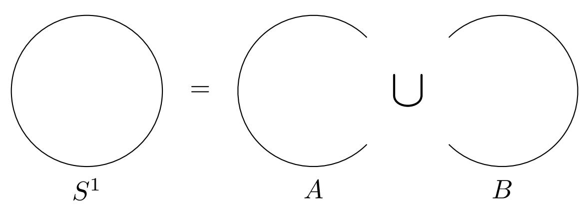

Importantly, we can't take two half-circles, because their interiors don't cover S^1. Now, $A \cap B$ is homeomorphic to two line segments, which deformation retracts to two points. So $H_n(A \cap B) \cong H_n(\bullet) \oplus H_n(\bullet)$. And $H_n(A) \oplus H_n(B) \cong H_n(\bullet) \oplus H_n(\bullet)$. At this point, the mechanics are largely the same. The long exact sequence expands out to $0 \to H_n(S^1) \to 0$ if $n \geq 2$, and we know that $H_0(S^1) \cong \mathbb{Z}$, so the we really only need to pay special attention to the case $n = 1$. In this case, we need to look at g_0 in the following diagram:

$$H_1(A) \oplus H_1(B) \to H_1(S^1) \to H_0(A \cap B) \xrightarrow{f_0} H_0(A) \oplus H_0(B)$$

The group $H_0(A \cap B)$ is the free abelian group generated by a two equivalence classes: $[\Delta^0 \to p_u]$ (the maps to a point in the upper part of $A \cap B$) and $[\Delta^0 \to p_\ell]$ (the maps to a point in the lower part of $A \cap B$). The group $H_0(A) \oplus H_0(B)$ is likewise generated by the two equivalence classes $[\Delta^0 \to p_A]$ and $[\Delta^0 \to p_B]$. If we call these equivalence classes u, ℓ, a, and b, respectively, then f_0 sends $nu + m\ell$ to $(n+m)a + (n+m)b$. So its kernel is isomorphic to $\mathbb{Z}$. Noting that $H_1(A) \oplus H_1(B) \cong \{0\}$, this means that $H_1(S^1) \cong \mathbb{Z}$.

The bagel

I won't Ti*k*Z any pictures here, so hopefully my vivid prose is vivid enough. We set $X = S^1 \times S^1$, the torus, and choose A and B to be two noodle-shaped tubes that together cover the torus. (It's like slicing a bagel in half the "wrong way," but then elongating both sides so that their interiors cover the whole ~~bagel~~ torus.) Then A and B are both homeomorphic to cylinders, and $A \cap B$ is homeomorphic to the disjoint union of two cylinders. A cylinder deformation retracts onto a circle, and we know a circle's homology. So for $n \geq 3$, both $H_n(A) \cap H_n(B)$ and $H_{n-1}(A \cap B)$ are the zero group; in the long exact sequence, we get

$0 \to H_n($$) \to 0,$

so $H_n($$) \cong \{0\}$ for every $n \geq 3$.[9] For $n = 2$, we get the exact sequence

$0 \to H_2($$) \to H_1(A \cap B) \xrightarrow{f_1} H_1(A) \oplus H_1(B).$

The analysis of f_1 here is very similar to the analysis of f_0 for the sphere; the upshot is that $H_2($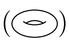$) \cong \mathbb{Z}$. Now focus on the diagram

$H_1(A \cap B) \xrightarrow{f_1} H_1(a) \oplus H_1(B) \xrightarrow{g_1} H_1($$) \xrightarrow{\partial_1} H_0(A \cap B) \xrightarrow{f_0} H_0(A) \oplus H_0(B).$

We know that $\mathrm{im}(f_1) \cong \mathbb{Z}$, and it's just as simple to check that $\ker(f_0) \cong \mathbb{Z}$. Since the sequence is exact, this means that $\mathrm{im}(\partial_1) \cong \mathbb{Z}$ and $\ker(g_1) \cong \mathbb{Z}$. Now, the sequence

$0 \to \ker(g_1) \xrightarrow{g_1} H_1($$) \xrightarrow{\partial_1} \mathrm{coker}(\partial_1) \to 0$

is always exact. Since $H_1(A \cap B) \cong \mathbb{Z} \oplus \mathbb{Z}$, the cokernel of ∂_1 is isomorphic to $\mathbb{Z}$. Therefore $H_1($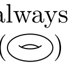$) \cong \mathbb{Z} \oplus \mathbb{Z}$, because of this here exercise:

Exercise 4.21. Prove that if $0 \to \mathbb{Z} \to A \to \mathbb{Z} \to 0$ is an exact sequence of abelian groups, then $A \cong \mathbb{Z} \oplus \mathbb{Z}$.

Finally, $H_0($$) \cong \mathbb{Z}$ because the surface of a bagel is path-connected. To summarize,

Proposition 4.22.

$H_n($$) \cong \begin{cases} \mathbb{Z} \oplus \mathbb{Z} & \text{if } n = 2 \\ \mathbb{Z} & \text{if } n = 0, 1 \\ 0 & \text{if } n \geq 3. \end{cases}$

I'll conclude with a warning: It's not always possible to recover a group from its surrounding exact sequence.

Exercise 4.23. Find two nonisomorphic groups A for which there is an exact sequence $0 \to \mathbb{Z}/2\mathbb{Z} \to A \to \mathbb{Z}/2\mathbb{Z} \to 0$.

5. CW HOMOLOGY

5.1. PUSHOUTS

We start, as is not so unusual, with a definition.

Definition 5.1. In a given category, a *pushout* of the diagram on the left is a commuting diagram on the right

[9] Gotcha! I'm gonna draw a Ti*k*Zpicture anyway!

$$\begin{array}{ccc} A & \xrightarrow{f} & B \\ {\scriptstyle g}\big\downarrow & & \\ C & & \end{array} \qquad\qquad \begin{array}{ccc} A & \xrightarrow{f} & B \\ {\scriptstyle g}\big\downarrow & & \big\downarrow{\scriptstyle p_B} \\ C & \xrightarrow[p_C]{} & P \end{array}$$

with the universal property that, for any commuting diagram

$$\begin{array}{ccc} A & \xrightarrow{f} & B \\ {\scriptstyle g}\big\downarrow & & \big\downarrow{\scriptstyle q_B} \\ C & \xrightarrow[q_C]{} & Q \end{array}$$

there is a unique map $P \to Q$ such that the following extended diagram commutes:

$$\begin{array}{cccc} A & \xrightarrow{f} & B & \\ \big\downarrow{\scriptstyle g} & & \big\downarrow{\scriptstyle p_B} & \searrow{\scriptstyle q_B} \\ C & \xrightarrow[p_C]{} & P & \\ & \searrow{\scriptstyle q_C} & & \dashrightarrow Q \end{array}$$

You can show that if (p_B, p_C, P) and (q_B, q_C, Q) are both pushouts, then the map $P \to Q$ is in fact an isomorphism, so a pushout is unique (up to, well, isomorphism).

In Top, the pushout of the first diagram in Definition 5.1 is the topological space $(B \sqcup C)/\{f(a) \sim g(a) : a \in A\}$.

EXAMPLE 5.2. The pushout of the diagram

$$\begin{array}{ccc} \emptyset & \longrightarrow & B \\ \big\downarrow & & \\ C & & \end{array}$$

is the disjoint union $B \sqcup C$. ◊

EXAMPLE 5.3. The pushout of

$$\begin{array}{ccc} A & \xrightarrow{f} & B \\ \big\downarrow & & \\ \bullet & & \end{array}$$

is $B/\operatorname{im}(f)$. ◊

In short, a pushout is (in Top, at least) a way to combine disjoint union and quotient in a single framework.

DEFINITION 5.4. If P is a pushout of the form

$$\begin{array}{ccc} \coprod_{i \in I} S^{n-1} & \xrightarrow{f} & B \\ \big\downarrow & & \big\downarrow \\ \coprod_{i \in I} D^n & \longrightarrow & P \end{array}$$

then we say that P is obtained from B by *attaching n-cells*.

EXAMPLE 5.5. The "figure eight" is obtained from a point by attaching 1-cells:

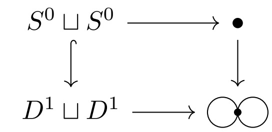

◊

EXAMPLE 5.6. If $f(-1) = 1/3$ and $f(1) = 2/3$, then this shape is also obtained by attaching a 1-cell:

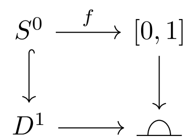

◊

EXAMPLE 5.7. The torus can be obtained by attaching a 2-cell to the figure 8:

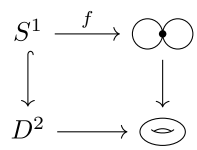

Here, the function f consists of beginning at the central point, going around the right circle clockwise, then the left circle clockwise, then the left circle counterclockwise, the right circle counterclockwise, and ending at the point again. It's easier to see why the pushout is a torus if we imagine the figure 8 with this quotient diagram:

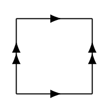

If you identify the opposite edges of a square, you get a figure 8. And gluing a circle along the edges of the figure 8 exactly consists of filling in this square, which results in a torus. ◊

This leads us to a definition and a new subsection.

5.2. CW-COMPLEXES

It's like that old joke: "Oh, you know those algebraic topologists; every other one has a massive CW-complex."[10]

DEFINITION 5.8. A *CW-complex* is a topological space X together with a so-called filtration by subspaces:

$$\emptyset = \mathrm{Sk}_{-1}(X) \subseteq \mathrm{Sk}_0(X) \subseteq \mathrm{Sk}_1(X) \subseteq \cdots \subseteq X$$

such that $X = \bigcup_{i=0}^{\infty} \mathrm{Sk}_i(X)$ and each $\mathrm{Sk}_k(X)$ is obtained from $\mathrm{Sk}_{k-1}(X)$ by attaching k-cells. Moreover, X must have the weak topology obtained from this filtration: $A \subseteq X$ is open if and only if $A\big|_{\mathrm{Sk}_k(X)}$ is open in $\mathrm{Sk}_k(X)$ for every k.

The C in CW-complex stands for cell, and the W stands for weak topology. The set $\mathrm{Sk}_k(X)$ is called the *k-skeleton* of X.

[10] ... as in the style of Freud. Get it?

Example 5.9. The torus is a CW-complex with the filtration

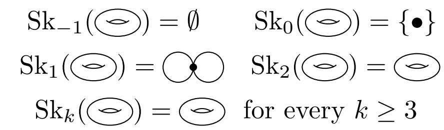

◊

Example 5.10. Every sphere is a CW-complex formed by attaching one n-cell to a point. Its skeleta are

$$\mathrm{Sk}_k(S^n) = \begin{cases} \emptyset & \text{if } k = -1 \\ \bullet & \text{if } 0 \le k \le n-1 \\ S^n & \text{if } k \ge n. \end{cases}$$

◊

Example 5.11. This CW-structure is not unique. The n-sphere also has a CW-complex structure made by attaching two k-spheres for every $0 \le k \le n$. For example, we can construct S^2 like this:

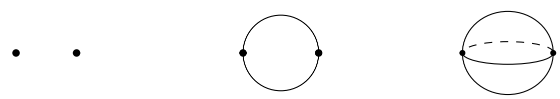

◊

In what follows, we say "a CW-complex X." when X is a topological space, which is strictly speaking bad form; but the skeletal CW-structure will be clear from context.

Definition 5.12. A CW-complex is called *CW-complex* if its skeleta eventually stabilize, if there is an $n \in \mathbb{N}$ such that $\mathrm{Sk}_n(X) = X$. The minimum such n is called the *dimension* of X. A CW-complex is *finite* if it is finite dimensional and each skeleton is obtained by adding only finitely many cells.

If we continue the second CW-construction of S^n indefinitely, we get S^∞, the CW-complex obtained by adding two n-cells at every step (and never stopping). It can be represented in $\mathbb{R}^{\mathbb{N}}$ by

$$\{(x_1, x_2, \dots) \in \mathbb{R}^{\mathbb{N}} : \sum_{i=1}^{\infty} x_i^2 = 1 \text{ and all but finitely many coordinates are } 0\}$$

It turns out this space, even though it's infinite-dimensional, isn't actually very interesting.

Proposition 5.13. *S^∞ is homotopy equivalent to a point.*

Proof. We define the shift operator $T \colon S^\infty \to S^\infty$ by $(x_1, x_2, \dots) \mapsto (0, x_1, x_2, \dots)$. First, T is homotopic to the identity via the homotopy

$$h_1(\mathbf{x}, t) = \frac{t\mathbf{x} + (1-t)T(\mathbf{x})}{\|t\mathbf{x} + (1-t)T(\mathbf{x})\|}.$$

(To verify that the denominator is never zero, consider the first nonzero coordinate of x if $t \neq 0$; if $t = 0$, then the denominator is also nonzero.) Also, T is homotopic to the map sending every point to $\mathbf{z} = (1, 0, 0, \dots)$ via

$$h_2(\mathbf{x}, t) = \frac{t\mathbf{z} + (1-t)T(\mathbf{x})}{\|t\mathbf{z} + (1-t)T(\mathbf{x})\|}.$$

So the identity is homotopic to the point morphism: S^∞ is homotopy equivalent to a point. ☐

We needed the middle step of T because the "homotopy" with the same style that you might define from the identity to the point map,

$$h_2(\mathbf{x}, t) = \frac{t\mathbf{z} + (1-t)\mathbf{x}}{\|t\mathbf{z} + (1-t)\mathbf{x}\|},$$

has a denominator that is sometimes 0 (in particular, when $\mathbf{x} = -\mathbf{z}$ and $t = 1/2$.

Anyway, that's mostly a curiosity; here's an infinite-dimensional space that's not trivial.

DEFINITION 5.14. The *real projective space* $\mathbb{RP}^n$ is the quotient of S^n by the equivalence relation $\{x \sim -x : x \in S^n\}$.

It turns out that $\mathbb{RP}^0$ is a point and $\mathbb{RP}^1$ is homeomorphic to S^1, but $\mathbb{RP}^2$ is genuinely different from S^2. One way to think of it is as the set of lines through the origin in $\mathbb{R}^3$.

The projective space $\mathbb{RP}^n$ can be given a CW-structure with one 0-cell, one 1-cell, ..., and one n-cell. Each cell is attached via a "double cover" of S^{n-1} by itself. As we extended S^n to S^∞, we can also extend $\mathbb{RP}^n$ to $\mathbb{RP}^\infty$. This, it turns out, is not homeomorphic to a point, though it will take a while to prove this.

POP QUIZ. What's the next thing we're going to do?
(a) Abuse notation
(b) Draw a commutative diagram
(c) Form a category
(d) Form a chain complex

They're all likely options in this course, but this time, it's (c).

DEFINITION 5.15. If X and Y are both CW-complexes, a continuous map $f\colon X \to Y$ is called *cellular* if $f\big(\mathrm{Sk}_k(X)\big) \subseteq \mathrm{Sk}_k(Y)$ for every $k \in \mathbb{N}_0$.

The category CWComp has all CW-complexes as objects, and the morphisms from X to Y are the set of cellular maps. For example, the inclusion map $S^{n-1} \hookrightarrow S^n$ is a cellular map if we equip the n-spheres with the CW-structure in Example 5.11, but not if the spheres have the CW-structure in Example 5.10. This gives an indication as to why a slightly less efficient CW-structure is sometimes beneficial.

5.3. HOMOLOGY GROUPS OF CW-COMPLEXES

Fix a CW-complex X, and suppose the pushout that attches n-cells indexes the $(n-1)$-spheres by I_n. We'll use X_n as shorthand for $\mathrm{Sk}_n(X)$.

Since $X_n \subseteq X_{n+1}$, we get (when $q \geq 1$) a long exact sequence

$$\cdots \to H_{q+1}(X_{n+1}/X_n) \to H_q\big(X_{n-1}\big) \to H_q\big(X_n\big) \to H_q\big(X_n/X_{n-1}\big) \to \cdots \tag{1}$$

So we need to understand X_n/X_{n-1}. To do this, think of the pushout

$$\begin{array}{ccc} \bigsqcup_{i\in I_n} S^{n-1} & \longrightarrow & X_{n-1} \\ \downarrow & & \downarrow \\ \bigsqcup_{i\in I_n} D^n & \longrightarrow & X_n. \end{array}$$

If X_{n-1} is collapsed to a point, then so are the boundaries of each of the n-spheres that are attached. So what we end up with is a *bouquet of spheres*, a collection of spheres joined to each other at a single point:

$$\bigsqcup_{i\in I_n} S^n \Big/ \bigsqcup_{i\in I_n} \bullet.$$

This is usually denoted by $\bigvee_{i\in I_n} S^n$; the symbol is the wedge sum. To determine the homology of this space, we'll look at the homology of $H_q(\bigvee_{i\in I_n} S^n, \bullet) \cong H_q(\bigsqcup_{i\in I_n} S^n, \bigsqcup_{i\in I_n} \bullet)$. For this latter homology group, we can analyze the exact sequence

$$H_q\Big(\bigsqcup_{i\in I_n} \bullet\Big) \to H_q\Big(\bigsqcup_{i\in I_n} S^n\Big) \to H_q\Big(\bigsqcup_{i\in I_n} S^n, \bigsqcup_{i\in I_n} \bullet\Big) \to H_{q-1}\Big(\bigsqcup_{i\in I_n} \bullet\Big) \to H_{q-1}\Big(\bigsqcup_{i\in I_n} S^n\Big).$$

If $q \neq 0, n$, then the middle homology group is flanked by two zero groups, so it is itself 0. If $q = n$, then the homology groups of $\bigsqcup_{i \in I_n} \bullet$ are 0, so the middle homology group is isomorphic to $H_q(\bigsqcup_{i \in I_n} S^n) \cong \oplus_{i \in I_n} \mathbb{Z}$. Therefore, we've determined that

$$H_q\Big(\bigvee_{i \in I_n} S^n, \bullet\Big) \cong \begin{cases} \bigoplus_{i \in I_n} \mathbb{Z} & \text{if } q = n \\ 0 & \text{otherwise.} \end{cases} \tag{2}$$

Plugging in this fact into (1), we get the following result.

THEOREM 5.16. *If X is a CW-complex and $q \neq n, n-1$, then $H_q\big(\mathrm{Sk}_{n-1}(X)\big) \cong H_q\big(\mathrm{Sk}_n(X)\big)$.*

In particular, the homology groups of the skeleta eventually stabilize: If $n \geq q+1$, then $H_q(X_n) = H_q(X_{q+1})$.

EXERCISE 5.17. Prove this.

EXERCISE 5.18. Show that $H_q(X_n) = 0$ if $n \leq q-1$. [HINT: Use the fact that X_0 is a collection of discrete points, so $H_q(X_0) = 0$.]

Okay, so the homologies of the skeleta stabilize when $n > q$. But what about the homologies of X? This is what we really care about.

PROPOSITION 5.19. *$H_q(X_k) \cong H_q(X)$ whenever $k > q$.*

Proof sketch. Something something compactness. □

So these relative homology groups $H_n(X_n, X_{n-1})$ hold a lot of information about X. As algebraic topologists are wont to do, we now form a chain complex with them.

DEFINITION 5.20. The *cellular n-chains* of a CW-complex X are the groups $H_n(X_n, X_{n-1})$.

How do we define a map $C_n(X) \to C_{n-1}(X)$? As usual, we rely on the long exact sequence. We have a sequence of maps

$$C_n(X) = H_n(X_n, X_{n-1}) \xrightarrow{\partial_n} H_{n-1}(X_{n-1}) \to H_{n-1}(X_{n-1}, X_{n-2}) = C_{n-1}(X),$$

where the second map is induced by the inclusion $(X_{n-1}, \emptyset) \hookrightarrow (X_{n-1}, X_{n-2})$. We denote the composition of these maps by d_n.

THEOREM 5.21. *$C_\bullet(X)$ and the maps d_n form a chain complex, called the* cellular chain complex *of X. Moreover, $H_n\big(C_\bullet(X)\big) \cong H_n(X)$ for every $n \in \mathbb{Z}$.*

The proof consists of constructing a large diagram and inspecting it. We'll do that here, but feel free to skip to the next section to see how this theorem is applied—so it's clear why the statement is even useful—before reading through the proof.

Proof of Theorem 5.21. As promised, first we get a great big diagram:

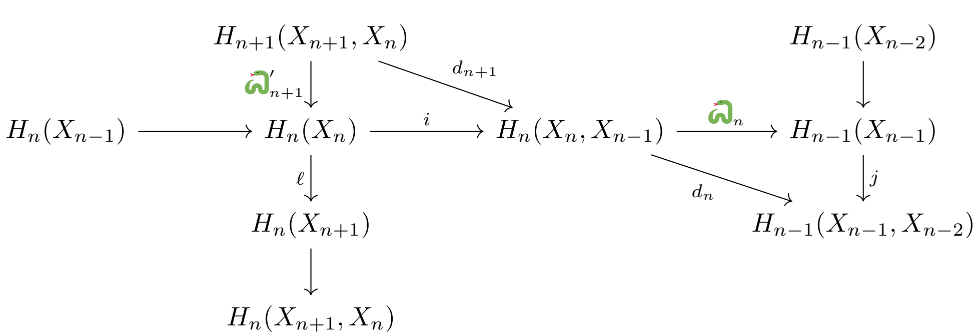

The columns and the rows are each from a homology long exact sequence. First up, the diagram is commutative, $d_n \circ d_{n+1} = j \circ \partial_n \circ i \circ \partial'_{n+1}$; since the row is exact, $\partial_n \circ i = 0$, so $d_n \circ d_{n+1} = 0$. This means that $\big(C_\bullet(X), d\big)$ is a chain complex.

Next, we show that $\ker d_n / \operatorname{im} d_{n+1} \cong H_n(X)$. First, $\ker d_n = \ker(j \circ \partial_n)$. But $H_{n-1}(X_{n-2}) = 0$, so j is injective; that means that $\ker d_n = \ker \partial_n$. Exactness of the row shows that $\ker \partial_n = \operatorname{im} i$. But i is injective because $H_n(X_{n-1}) = 0$, so $\operatorname{im} i \cong H_n(X)$. Now for $\operatorname{im} d_{n+1}$. Since i is injective, $\operatorname{im} d_{n+1} = \operatorname{im}(i \circ \partial'_{n+1}) = i(\operatorname{im} \partial'_{n+1})$.

Therefore, (keeping in mind that i is injective), we have that

$$\ker d_n / \operatorname{im} d_{n+1} \cong H_n(X) / \operatorname{im} \partial'_{n+1}.$$

Since $H_n(X_{n+1}, X_n) = 0$ by (2), this quotient is isomorphic to $H_n(X_{n+1})$. (See the following exercise.) But from Proposition 5.19, $H_n(X_{n+1}) \cong H_n(X)$, which finishes the proof. $\square$

EXERCISE 5.22. Suppose that $A \xrightarrow{f} B \to C \to 0$ is an exact sequence of abelian groups. Show that $B/\operatorname{im} f \cong C$.

5.4. USING THE CELLULAR CHAIN COMPLEX

Let's start with our standard first step: the sphere.[11]

EXAMPLE 5.23. Equip S^2 with the CW-structure from Example 5.10. Then

$$C_n(S^2) \cong \begin{cases} \mathbb{Z} & \text{if } n = 0, 2 \\ 0 & \text{otherwise.} \end{cases}$$

So the long exact sequence

$$\cdots \to C_4(X) \to C_3(X) \to C_2(X) \to C_1(X) \to C_0(X) \to 0$$

becomes

$$\cdots \to 0 \to 0 \to \mathbb{Z} \to 0 \to \mathbb{Z} \to 0.$$

The image of every map must be 0, and the kernel of every map is the entire group, so it's easy to read off that

$$H_n(X) \cong H_n\big(C_\bullet(X)\big) \cong \begin{cases} \mathbb{Z} & \text{if } n = 0, 2 \\ 0 & \text{otherwise.} \end{cases} \qquad \lozenge$$

In this example, we didn't even need to know anything about the maps $d_n \colon C_n(X) \to C_{n-1}(X)$, but in general we will. Since these maps are defined in terms of the mysterious snake map, this seems rather difficult. To work with them, we'll need one more tool.

DEFINITION 5.24. Any continuous map $f \colon S^n \to S^n$ induces a map $f^* \colon H_n(S^n) \to H_n(S^n)$ of an infinite cyclic group, and this map must be of the form $f^*(x) = dx$ for some $d \in \mathbb{Z}$. This integer d is called the *degree* of f.

For example, identity map $\iota \colon S^n \to S^n$ induces the identity map $H_n(S^n) \to H_n(S^n)$, so $\deg(\iota) = 1$. The map sending S^n to a point on S^n has degree 0.

Degree is useful because it allows us to actually calculate the chain maps d_n. Each of the groups $C_n(X)$ is free abelian, and we can choose a generating set where each element corresponds to a sphere in the bouquet X_n/X_{n-1}. For each sphere α, let e^n_α denote this element. This is the basis we'd like to calculate in, but it's unclear how the snake map will rearrange this basis. The degree actually gives us the answer.

[11] This is sort of a cheat, since we needed to know the homology of the sphere to prove things about cellular chain complexes, but that's no reason to quibble.

THEOREM 5.25. *If $n > 1$, the map $d_n : C_n(X) \to C_{n-1}(X)$ is given by*

$$d_n(e^n_\alpha) = \sum_\beta d_{\alpha\beta} e^{n-1}_\beta,$$

where the sum ranges over all spheres β in the bouquet X_{n-1}/X_{n-2} and $d_{\alpha\beta}$ is the degree of the composite map $S^{n-1} \to X^{n-1} \to S^{n-1}_\beta$: The first arrow is the attaching map of the boundary of α (from the pushout), and the second arrow is the quotient map that collapses $X_{n-1} \setminus \beta$ to a point.

I won't prove this. It consists of drawing a big diagram and chasing through it, and I wouldn't write the proof any more clearly than is written elsewhere. See, for example, these notes on cellular homology; or, if for some reason you'd like a terser version of the same proof, see Hatcher's textbook.

Instead, it's important to see how to use it. What Theorem 5.25 says is that to determine $d_n(e^n_\alpha)$, we just need to look at how the attaching map of α interacts with each $(n-1)$-cell.

Let's see how to use it. One important consequence is that if two n-cells α_1 and α_2 are attached to X using the same attaching map, then $d_n(\alpha_1) = d_n(\alpha_2)$. We can use that, for example, to calculate the degree of slightly more complex maps.

EXAMPLE 5.26. What is the degree of the reflection map over some hyperplane through the origin? Let's start with S^1 and a reflection across the x-axis. If we equip S^1 with the CW-structure from Example 5.11, the reflection map is cellular, so it induces a map $C_\bullet(S^1) \to C_\bullet(S^1)$ of chain complexes. Let's say that the 0-cells of S^1 are x and y and the 1-cells are u and v. Then $C_0(S^1) \cong \mathbb{Z}\{x, y\}$ and $C_1(S^1) \cong \mathbb{Z}\{u, v\}$, so we get a diagram like this:

$$\begin{array}{ccccc} \mathbb{Z}\{u,v\} & \xrightarrow{d_1} & \mathbb{Z}\{x,y\} & \longrightarrow & 0 \\ \downarrow r & & \downarrow \iota & & \\ \mathbb{Z}\{u,v\} & \xrightarrow{d_1} & \mathbb{Z}\{x,y\} & \longrightarrow & 0 \end{array}$$

The reflection is an identity on the 0-skeleton of S^1 and interchanges u and v. So the map ι is the identity, and $r(u) = v$, $r(v) = u$.

Since u and v have the same attaching map, $d_1(u) = d_1(v)$, so $u - v \in \ker(d_1)$. This means that $u - v$ represents an equivalence class in $H_n(C_n(X))$. Setting $x = u - v$, we have $r(x) = -x$, so the degree of r is -1. (We only need to check the multiplying factor for one element in $H_n(C_n(X))$, since it is the same for all elements.) ◊

EXERCISE 5.27. Show that the degree of reflection across a hyperplane through the origin is -1 for every dimension. (You should just follow the argument in Example 5.26.)

To get even more degree calculations for free, we can see that it plays nicely with function composition.

LEMMA 5.28. $\deg(f \circ g) = \deg(f)\deg(g)$.

Proof. If f and g induce maps of a cyclic group that consist of multiplying by $\deg(f)$ and $\deg(g)$, respectively, then $f \circ g$ consists of multiplying by $\deg(f)\deg(g)$. □

COROLLARY 5.29. *If f is a homeomorphism, then* $\deg(f) = \pm 1$.

Proof. The degree of the the identity function is 1, so $1 = \deg(f \circ f^{-1}) = \deg(f)\deg(f^{-1})$. □

One more example of degree:

EXAMPLE 5.30. The *antipodal* map on S^n is the one that sends $x \to -x$. Each coordinate of S^n has the form $(x_1, \ldots, x_{n+1})$, and the antipodal map is the composition of a reflection for each coordinate, so its degree is $(-1)^n$. ◊

Now let's use CW homology to compute the homologies of the torus.

EXAMPLE 5.31. From Example 5.7, the torus can be realized as a CW-complex with one 0-cell x, two 1-cells u and v, and one 2-cell A. Its cellular long exact sequence is therefore

$$\cdots \to 0 \to \mathbb{Z}\{A\} \to \mathbb{Z}\{u, v\} \to \mathbb{Z}\{x\} \to 0.$$

Since $\circledcirc$ is connected, we know that $H_0\big(C_0(\circledcirc)\big) \cong \mathbb{Z}$, so $\operatorname{im}(d_1) = 0$. Therefore $\ker(d_1) = \mathbb{Z}\{u, v\}$. Recall that A is attached by following u, then v, then the reverse of u, and then the reverse of v. Symbolically, we might write the attaching map of A as $uv\overleftarrow{u}\,\overleftarrow{v}$.

Now let's look at $d_2(A)$ using Theorem 5.25. If we compose the attaching map with the map that quotients out everything but u, then we get the map sending S^1 to $u\overleftarrow{u}$. This map is homotopic to the map sending S^1 to a single point, so $u\overleftarrow{u} = 0$ in $H_1\big(C_\bullet(\circledcirc)\big)$. The degree of this map is 0; similarly, the degree $d_{A,v} = 0$. So $d_2(A) = 0u + 0v = 0$, which means that $H_1\big(C_\bullet(\circledcirc)\big) = \mathbb{Z}\{u, v\}$. Finally, the kernel of d_2 is $\mathbb{Z}\{A\}$, and the image of d_3 is $\{0\}$, so $H_2\big(C_\bullet(\circledcirc)\big) = \mathbb{Z}\{A\}$. In short, then,

$$H_n(\circledcirc) \cong H_n\big(C_\bullet(\circledcirc)\big) \cong \begin{cases} \mathbb{Z} & \text{if } n = 0, 2 \\ \mathbb{Z}^2 & \text{if } n = 1 \\ 0 & \text{otherwise,} \end{cases}$$

just as we computed before. ◊

Now let's do a new homology calculation.

5.5. HOMOLOGY OF REAL PROJECTIVE SPACE

Remember that $\mathbb{RP}^n$ is defined to be $S^n/\{x \sim -x\}$: the sphere quotiented by the antipodal map. We can also think of this space as $D^n/\{x \sim -x : x \in S^{n-1}\}$. In this way, $\mathbb{RP}^n$ is obtained from $\mathbb{RP}^{n-1}$ by attaching a single n-cell via the map $S^{n-1} \mapsto S^{n-1}/\{x \sim -x\}$.

If we give S^n the CW-structure of Example 5.11, then the quotient map $S^n \to S^n/\{x \sim -x\}$ is a cellular map from S^n to $\mathbb{RP}^n$. For $n = 2$, we get a diagram like this:

$$\begin{array}{ccccccccc} \cdots & \longrightarrow & \mathbb{Z}\{a_2, b_2\} & \xrightarrow{d_2^S} & \mathbb{Z}\{a_1, b_1\} & \xrightarrow{d_1^S} & \mathbb{Z}\{a_0, b_0\} & \longrightarrow & 0 \\ & & \Big\downarrow{\scriptstyle f_2} & & \Big\downarrow{\scriptstyle f_1} & & \Big\downarrow{\scriptstyle f_0} & & \\ \cdots & \longrightarrow & \mathbb{Z}\{c_2\} & \xrightarrow{d_2^{\mathbb{RP}}} & \mathbb{Z}\{c_1\} & \xrightarrow{d_1^{\mathbb{RP}}} & \mathbb{Z}\{c_0\} & \longrightarrow & 0 \end{array}$$

where a_i, b_i are the i-cells of S^2 and c_i are the i-cells of $\mathbb{RP}^2$. We know that $d_i^S(a_i) = d_i^S(b_i)$ for every $i \in \{0, 1, 2\}$, since a_i and b_i have the same attaching map. So the kernel of d_1^S is generated by $a_1 - b_1$. We know that $H_1(S^2) = \{0\}$, so $d_2^S(a_2) = d_2^S(b_2)$ must generate the whole group, so $d_2^S(a_2) = \pm(a_1 - b_1)$.

Now, what about these downward maps? For each pair of n-cells, the quotient map is the identity on one of them; arbitrarily, let's say it's the identity on a_i; then $f(a_i) = c_i$. What is the action on b_i? The antipodal map α sends b_i to $(-1)^i a_i$—the multiplicative factor is the degree of the antipodal map on the boundary of b_i. So $f_i(\alpha(b_i))) = (-1)^i c_i$. But the quotient map $S^n \to S^n/\{x \sim -x\}$ is invariant under antipodal interchange, so $f(\alpha(b_i)) = f(b_i)$.

Now we're ready to calculate $d_i^{\mathbb{RP}}$:

$$\begin{aligned} d_2^{\mathbb{RP}}(c_2) &= d_2^{\mathbb{RP}}\big(f_2(a_2)\big) \\ &= f_1\big(d_2^S(a_2)\big) \\ &= f_1\big(\pm(a_1 - b_1)\big) \\ &= \pm\big(c_1 - (-c_1)\big) \\ &= \pm 2c_1. \end{aligned}$$

A similar calculation shows that $d_1^{\mathbb{RP}}(c_1) = 0$. So we can write the cellular chain sequence of $\mathbb{RP}^2$ as

$$\cdots \to 0 \to \mathbb{Z} \xrightarrow{\pm 2} \mathbb{Z} \xrightarrow{0} \mathbb{Z} \to 0.$$

This makes it easy to see the homology:

$$H_n(\mathbb{RP}^2) \cong \begin{cases} \mathbb{Z} & \text{if } n = 0 \\ \mathbb{Z}/2\mathbb{Z} & \text{if } n = 1 \\ 0 & \text{otherwise.} \end{cases}$$

This is *super* weird. No. It's borderline nonsensical. We introduced homology as measuring the number of holes: $\mathbb{Z}^k$ means k holes in a certain dimension: easy-peasy. But what in the world does it mean to say you have one hole modulo even multiples? Hatcher has little to say about it when it comes up; neither does Miller in his notes. Fortunately, there's StackExchange, which gives a little bit of intuition. But perhaps the best thing is to consider homology as simply an invariant. In the wise words of StackExchange user D Wiggles,

> Here's another perspective. On some level, it doesn't matter what homology measures. The point is that it takes something very hard (topology) and turns it into something easy (abelian groups). If someone "hands you" two topological spaces, you basically have a useless pile of garbage. Great: Now you know what all the open sets are . . . How would you ever tell two such things apart? The answer is homology. It's easy to compute and often lets you answer a hard question with relative ease.

5.6. GEOMETRIC REALIZATION

Any semisimplicial set $\mathcal{X}$ can be realized as a CW-complex; X_0 is the set of points, X_1 is the set of edges, and so on, and the maps d_i provide the attaching maps. This is called the *geometric realization* of $\mathcal{X}$. Here's why it's important: Every triangulation of a topological space X is a semisimplicial set $\mathcal{X}$, and it would be nice if we could calculate $H_n(X)$ using the (probably much nicer) chain complex $\mathcal{X}$. This is indeed the case.

THEOREM 5.32. *If $\mathcal{X}$ is a semisimplicial set whose geometric realization is homeomorphic to $X \in$ Top, then $H_n(X) = H_n(\mathcal{X})$ for every $n \in \mathbb{Z}$.*

The reason for this is simple: The cellular chain complex $C_n(X)$ is the exact same as the chain complex $S_n(\mathcal{X})$. We know that the homology computed from the cellular chain complex is the same as the homology computed from the singular chain complex, so the homology groups must be the same.

If you can find a finite triangulation of a space, this theorem makes calculating homology groups quite simple. We won't use it for a while—not until Example 7.18, where we calculate the cohomology ring of the Klein bottle—but it's good to know, regardless.

5.7. EULER CHARACTERISTIC

Here's an invariant that's simpler to compute than homology but still can distinguish some spaces.

DEFINITION 5.33. If X is a finite CW-complex, the *Euler characteristic* of X is $\chi(X) = \sum_k (-1)^k n_k$, where n_k is the number of k-cells in X.

The good thing about the Euler characteristic is that it's very easy to compute. But it seems to depend on the particular CW-structure we pick for X, which would make it rubbish as an invariant. Luckily, that's not the case. Here's why.

THEOREM 5.34. $\chi(X) = \sum_k (-1)^k \operatorname{rank}\big(H_k(X)\big)$, *where the rank of a group G is the size of the largest free abelian group it contains.*[12]

This means that the Euler characteristic depends only on the homotopy type of the space. In particular, it doesn't depend on the choice of CW-structure.

The use of the Euler characteristic is in how extraordinarily easy it is to compute. So before proving Theorem 5.34, let's see the Euler characteristic in action.

EXAMPLE 5.35. The sphere has a CW-structure with one n-cell and one 0-cell, so $\chi(S^n) = (-1)^n + 1$. ♢

EXAMPLE 5.36. The torus has a CW-structure with one 2-cell, two 1-cells, and one 0-cell, so $\chi($$) = 1 - 2 + 1 = 0$. ♢

So with this very short calculation, we can see that S^2 and  are not homeomorphic or even homotopy equivalent. On the other hand, the Euler characteristic can't distinguish S^1 from S^3 from 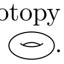. Nevertheless, it's a good first invariant to check because it's so quick.

Sketch of Theorem 5.34. First, do the exercise that follows this proof. We write Z_k for $\ker(d_k)$ and B_k for $\operatorname{im}(d_k)$ in the cellular chain sequence. Then

$$0 \to Z_k \to C_k \to B_k \to 0$$

is exact (since $C_k/\ker d_k \cong \operatorname{im} d_k$), as is

$$0 \to B_{k+1} \to Z_k \to H_k(X) \to 0$$

(since $H_k(X) \cong Z_k/B_{k+1}$). Combining, we get that

$$n_k = \operatorname{rank}\big(C_k(X)\big) = \operatorname{rank}(Z_k) + \operatorname{rank}(B_k) = \operatorname{rank}(B_k) + \operatorname{rank}(B_{k+1}) + \operatorname{rank}\big(H_k(X)\big).$$

So in the formula for the Euler characteristic, $\operatorname{rank}(B_k)$ and $\operatorname{rank}(B_{k+1})$ both appear in two terms, but with opposite signs. So all that remains after cancelling is $\sum_k (-1)^k \operatorname{rank}\big(H_k(X)\big)$. □

EXERCISE 5.37. If $0 \to A \to B \to C \to 0$ is an exact sequence of finitely generated abelian groups, then $\operatorname{rank}(A) + \operatorname{rank}(C) = \operatorname{rank}(B)$. Prove this.

EXERCISE 5.38. Fill in the gaps in the proof of Theorem 5.34.

Our next goal is to develop some invariants that are more powerful than Euler characteristic but easier to compute than homology.

6. HOMOLOGY WITH COEFFICIENTS

6.1. WHAT ARE COEFFICIENTS?

For now and evermore, R is a commutative ring with unit. For our purposes, standard examples are $\mathbb{Z}$, $\mathbb{Q}$, $\mathbb{R}$, and $\mathbb{Z}_p = \mathbb{F}_p = \mathbb{Z}/p\mathbb{Z}$ (depending on your preferred notation). But you really can use anything here: $SL_n(\mathbb{Z})$, $\prod_{n=2}^{\infty} \mathbb{Z}_n$, and so on. But we won't really deal with these more exotic rings. Anyway.

[12] The Fundamental Theorem of Finitely Generated Abelian Groups states that every finitely generated abelian group G is isomorphic to $\mathbb{Z}^r \oplus H$ for some finite abelian group H. The number r is the *rank* of G.

DEFINITION 6.1. If $\mathcal{X}$ is a semisimplicial set, then $S_k(\mathcal{X}; R)$ denotes the free R-module generated by X_k.[13] The chain complex $S_\bullet(\mathcal{X}; R)$ is obtained from these modules by defining the boundary operators exactly as in Definition 1.9. If X is a topological space, we abbreviate $S_\bullet\big(\operatorname{Sing}_n(X); R\big)$ by $S_\bullet(X; R)$. Given a chain complex $S_\bullet(\mathcal{X}; R)$, the *nth homology of X with coefficients in R* is the R-module $H_n(\mathcal{X}; R) := \ker \partial_n / \operatorname{im} \partial_{n+1}$.

So, for example, $S_*(X; \mathbb{Z})$ is just the regular chain complex of X that we've been studying all along. In fact, everything that we've done so far could be repeated, essentially word-for-word, with *any* commutative ring in place of the integers. In particular:

- $H_k(X; R)$ is a homotopy invariant for every $k \in \mathbb{N}_0$ and commutative ring R.
- All of the Eilenberg–Steenrod axioms hold except the dimension axiom, which is replaced by

$$H_n(\bullet; R) \cong \begin{cases} R & \text{if } 0 \\ 0 & \text{otherwise.} \end{cases}$$

- All of our tools—like the long homology sequence, Mayer-Vietoris, and cellular chain complexes—still work.

The second bullet point indicates that what we have now are *extraordinary homology theories*. Even though they're extraordinary, if you know the chain complex, they're really not much different to compute.

EXAMPLE 6.2. Consider the topological space $\mathbb{RP}^2$. We know that its cellular chain complex $C_\bullet(\mathbb{RP}^2)$ looks like this:

$$\cdots \longrightarrow 0 \longrightarrow \mathbb{Z} \xrightarrow[\pm 2]{\partial_2} \mathbb{Z} \xrightarrow[0]{\partial_1} \mathbb{Z} \xrightarrow{\partial_0} 0.$$

So the cellular chain complex $C_\bullet(\mathbb{RP}^2; \mathbb{Q})$ looks like this:

$$\cdots \longrightarrow 0 \longrightarrow \mathbb{Q} \xrightarrow[\pm 2]{\partial_2} \mathbb{Q} \xrightarrow[0]{\partial_1} \mathbb{Q} \xrightarrow{\partial_0} 0.$$

As before, it's easy to read off the homology from this: the kernel of ∂_0 is all of $\mathbb{Q}$, and the image of ∂_1 is 0, so $H_0(\mathbb{RP}^2; \mathbb{Q}) \cong \mathbb{Q}$. And $\ker(\partial_1) = \mathbb{Q}$ with $\operatorname{im}(\partial_2) = \mathbb{Q}$, so $H_1(\mathbb{RP}^2; \mathbb{Q}) \cong 0$. Similarly, $H_2(\mathbb{RP}^2; \mathbb{Q}) \cong 0$.

We can compare the homologies with coefficients in $\mathbb{Z}$ or $\mathbb{Q}$:

$$H_n(\mathbb{RP}^2; \mathbb{Z}) \cong \left\{\begin{matrix} \mathbb{Z} & \text{if } n = 0 \\ \mathbb{Z}_2 & \text{if } n = 1 \\ 0 & \text{otherwise} \end{matrix}\right\} \qquad H_n(\mathbb{RP}^2; \mathbb{Q}) \cong \left\{\begin{matrix} \mathbb{Q} & \text{if } n = 0 \\ 0 & \text{otherwise.} \end{matrix}\right\}$$

So homology with $\mathbb{Q}$ somehow picks out less than homology with $\mathbb{Z}$ does. On the other hand, if you calculate out the sequence $C_\bullet(\mathbb{RP}^2; \mathbb{Z}_2)$, you'll find that

$$H_n(\mathbb{RP}^2; \mathbb{Z}_2) \cong \left\{\begin{matrix} \mathbb{Z}_2 & \text{if } n = 0, 1, 2 \\ 0 & \text{otherwise.} \end{matrix}\right\}$$

Somehow homology with coefficients in $\mathbb{Z}_2$ can see that $\mathbb{RP}^2$ is 2-dimensional, while homology with coefficients in $\mathbb{Z}$ cannot. ◊

It seems, from this example, that $H_n(X; R)$ gives us genuinely new information: $\mathbb{RP}^2$ has a nonzero homology group when we consider coefficients in $\mathbb{Z}_2$. It turns out, though, that $H_n(X; R)$ can always be determined by the homology groups $H_n(X; \mathbb{Z})$.

[13] What's an R-module? It's just a vector space, but over a ring instead of a field. In fact, if R is a field, then R-modules just vector spaces over R. You're familiar with another set of modules: $\mathbb{Z}$-module are nothing more and nothing less than abelian groups. See more here. A free R-module is isomorphic to $\oplus_{i \in I} R$ for some index set I.

However, the relationship between these homology groups is complex. For some topological spaces, we know the homology groups with coefficients in $\mathbb{Q}$ or $\mathbb{Z}_2$, but not in $\mathbb{Z}$. One such topological space is the space of k unordered points in the torus

$\mathrm{config}_k($  $) := \{(x_1, \ldots, x_k) \in$  $^k : x_i \neq x_j \text{ if } i \neq j\}/S_k$

(The quotient by S_k is a quotient by the action of the symmetric group.) For some reason, some physicists really care about this space, and the homology groups will tell them something about it. The homology groups are completely known—but only with coefficients in $\mathbb{Q}$ or $\mathbb{Z}_2$, not even in $\mathbb{Z}_3$, let alone $\mathbb{Z}$.

Also, if you have a big (but finite) simplicial set, then a computer can calculate the homology of it for you. It's been shown that these calculations are faster over $\mathbb{Q}$ than over $\mathbb{Z}$. Practically speaking, then, homology with coefficients is sometimes necessary to actually compute the groups, either by hand or by computer. Moreover, it's usually just easier to compute homology groups over $\mathbb{Z}_2$ than over $\mathbb{Z}$. The little niggle we had over the sign when we calculated $H_n(\mathbb{RP}^2)$ in Section 5.5—something that can become quite a headache in more complicated calculations—completely disappears if we're working in $\mathbb{Z}_2$, since $+1 = -1$.

So that's why we want to study homology with coefficients. And anyway, we still have to prove exactly *how* homology with $\mathbb{Z}$-coefficients determines all the rest.

We'll start with an algebraic prelude.

6.2. THE CATEGORICAL ALGEBRA OF R-MODULES

⚠ I'm going to be a terse with the commutative algebra background here. If you're not familiar with modules over a ring, then consider reading Section 3 in these notes; that should get you up to speed.

6.2.1. The category of R-modules

DEFINITION 6.3. The category Mod_R has as objects the collection of R-modules, and $\mathrm{Hom}_{\mathrm{Mod}_R}(M, N)$ is the collection of R-linear maps from M to N (that is, the maps $f\colon M \to N$ such that $f(x+y) = f(x) + f(y)$ and $f(rx) = rf(x)$ for every $x, y \in M$ and $r \in R$.)

So we have another category. This one's a little different, because its Hom-sets aren't just sets. Given two module homomorphisms $f, g \in \mathrm{Hom}_{\mathrm{Mod}_R}(M, N)$, there is another map $h \in \mathrm{Hom}_{\mathrm{Mod}_R}(M, N)$ such that

$$h(x) = f(x) + g(x)$$

for every $x \in M$. Moreover, if $r \in R$, then there is another map $\tilde{h} \in \mathrm{Hom}_{\mathrm{Mod}_R}(M, N)$ such that

$$\tilde{h}(x) = r \cdot f(x)$$

for every $x \in M$. These operations actually make $\mathrm{Hom}_{\mathrm{Mod}_R}(M, N)$ into an R-module. When we consider this set as an R-module, we'll write $\underline{\mathrm{Hom}}_{\mathrm{Mod}_R}(M, N)$; when we want to emphasize only the set structure, we'll omit the underline.

Let's go deeper. For a fixed R-module M, we can define a map $N \mapsto \underline{\mathrm{Hom}}_{\mathrm{Mod}_R}(M, N)$ on the category of R-modules. Any map $f\colon N_1 \to N_2$ induces a map $\underline{\mathrm{Hom}}_{\mathrm{Mod}_R}(M, N_1) \to \underline{\mathrm{Hom}}_{\mathrm{Mod}_R}(M, N_2)$ given by $g \mapsto f \circ g$. The map we just defined is a functor $\mathrm{Mod}_R \to \mathrm{Mod}_R$, denoted $\underline{\mathrm{Hom}}_{\mathrm{Mod}_R}(M, -)$.

If we switch the positions and fix an R-module N, then $\underline{\mathrm{Hom}}_{\mathrm{Mod}_R}(-, N)$ is almost a functor; it switches the arrows. The map $f\colon M_1 \to M_2$ induces the map $\underline{\mathrm{Hom}}_{\mathrm{Mod}_R}(M_2, N) \to \underline{\mathrm{Hom}}_{\mathrm{Mod}_R}(M_1, N)$ given by $g\colon g \circ f$. This means that $\underline{\mathrm{Hom}}_{\mathrm{Mod}_R}(-, N)$ is a contravariant functor, or in other words, a functor $\mathrm{Mod}_R^{\mathrm{op}} \to \mathrm{Mod}_R$.

Putting the two together, we get a functor

$$\underline{\mathrm{Hom}}_{\mathrm{Mod}_R}(-,-)\colon\ \mathrm{Mod}_R^{\mathrm{op}} \times \mathrm{Mod}_R \to \mathrm{Mod}_R\,.$$

In every category, $\mathrm{Hom}_{\mathcal{C}}(-,-)\colon \mathcal{C}^{\mathrm{op}} \times \mathcal{C} \to \mathsf{Set}$ is a functor. An *internal Hom* is a functor $\underline{\mathrm{Hom}}_{\mathcal{C}}(-,-)\colon \mathcal{C}^{\mathrm{op}} \times \mathcal{C} \to \mathcal{C}$. In short, then, the category Mod_R has an internal Hom.

6.2.2. Tensor products

Every category with an internal Hom has a type of product structure on its objects. Let's begin with the easiest case.

Example 6.4. In the category Set, the regular Hom functor is also an internal Hom, simply because the image Hom is already a set. There is what's called a *currying isomorphism*

$$\underline{\mathrm{Hom}}_{\mathsf{Set}}(A \times B, C) \cong \underline{\mathrm{Hom}}_{\mathsf{Set}}\big(A, \underline{\mathrm{Hom}}_{\mathsf{Set}}(B,C)\big),$$

where $A \times B$ is the usual product of sets. This isomorphism maps the element $f\colon A \times B \in C$ to the map $g\colon A \to \underline{\mathrm{Hom}}_{\mathsf{Set}}(B,C)$ given by $g(a) = \big(b \mapsto f(a,b)\big)$. Furthermore, this isomorphism is natural in each of the objects A, B, and C. To unpack that, if we denote the currying isomorphism by 🌶, then whenever $f\colon A \to D$, the following square commutes.

$$\begin{array}{ccc} \underline{\mathrm{Hom}}_{\mathsf{Set}}(A \times B, C) & \xrightarrow{\text{🌶}} & \underline{\mathrm{Hom}}_{\mathsf{Set}}\big(A, \underline{\mathrm{Hom}}_{\mathsf{Set}}(B,C)\big) \\ {\scriptstyle \underline{\mathrm{Hom}}(-,B)(f\times 1_B)}\big\downarrow & & \big\downarrow{\scriptstyle \underline{\mathrm{Hom}}(-,\underline{\mathrm{Hom}}(B,C))(f)} \\ \underline{\mathrm{Hom}}_{\mathsf{Set}}(D \times B, C) & \xrightarrow{\text{🌶}} & \underline{\mathrm{Hom}}_{\mathsf{Set}}\big(D, \underline{\mathrm{Hom}}_{\mathsf{Set}}(B,C)\big) \end{array}$$

Similar conditions hold for B and C. ◊

The *tensor product* is a construct that generalizes this currying isomorphism. Given two R-modules M and N, their tensor product is a new R-module $M \otimes_R N$ that satisfies the following property.

Proposition 6.5. *There is a currying isomorphism*

$$\underline{\mathrm{Hom}}_{\mathrm{Mod}_R}(A \otimes_R B, C) \cong \underline{\mathrm{Hom}}_{\mathrm{Mod}_R}\big(A, \underline{\mathrm{Hom}}_{\mathrm{Mod}_R}(B,C)\big)$$

which is natural in A, B, and C.

A map $f\colon A \times B \to C$ of R-modules is called *bilinear* if
- $f(x+z,y) = f(x,y) + f(z,y)$,
- $f(x,y+w) = f(x,y) + f(x,w)$, and
- $f(r \cdot x, y) = r \cdot f(x,y) = f(x, r \cdot y)$

for every $x, z \in A$ and $y, w \in B$. You can check that the set $\mathrm{Hom}_{\mathrm{Mod}_R}\big(A, \mathrm{Hom}_{\mathrm{Mod}_R}(B,C)\big)$ is "the same" as the collection of bilinear maps $A \oplus B \to C$, via the bijection that sends $f\colon A \times B \to C$ to the map $g\colon A \to \mathrm{Hom}_{\mathrm{Mod}_R}(B,C)$ defined by $g(a) = \big(b \mapsto f(a,b)\big)$. Among these maps, the tensor product enjoys a certain universal property.

Proposition 6.6. *There is a map $\varphi\colon A \times B \to A \otimes_R B$ such that: If f is any bilinear map $A \times B \to D$ of R-modules, then there is a unique R-module homomorphism $A \otimes_R B \to D$ where the following diagram commutes:*

$$\begin{array}{ccc} A \times B & \xrightarrow{\varphi} & A \otimes_R B \\ & {\scriptstyle f}\searrow & \downarrow \\ & & D \end{array}$$

This is the way to think of tensor products abstractly. Sometimes it's useful to think of them concretely, which is when the construction of the tensor product comes in useful. You can read more here, but the upshot is that the tensor product $A \otimes_R B$ is generated by elements denoted $a \otimes b$ for each $a \in A$ and $b \in B$. It is the "largest" R-module generated by these elements such that these elements are bilinear. That means that:

$$\begin{aligned} (a_1 + a_2) \otimes b &= a_1 \otimes b + a_2 \otimes b \\ a \otimes (b_1 + b_2) &= a \otimes b_1 + a \otimes b_2 \\ (ra) \otimes b &= r(a \otimes b) = a \otimes (rb) \end{aligned} \tag{3}$$

⚠ It's important to note that while the elements $a \otimes b$ generate $A \otimes_R B$, not every element is necessarily of this form—the tensor product contains finite R-linear combinations of these elements, as well.

EXAMPLE 6.7. The element $0 \otimes 0$ is the identity element of $A \otimes_R B$. This is because

$$0 \otimes 0 = (0 \cdot 0) \otimes 0 = 0 \cdot (0 \otimes 0).$$

Moreover, $a \otimes 0 = 0 \otimes b = 0 \otimes 0$ for every $a \in A$ and $b \in B$ by doing the same trick of pulling out a 0. ♢

EXAMPLE 6.8. Suppose that p and q are distinct primes. What is $\mathbb{Z}_p \otimes_{\mathbb{Z}} \mathbb{Z}_q$? Well, we can choose an element $n \in \mathbb{Z}$ such that $np \cong 1 (\text{mod } q)$. So by the bilinearity of the tensor product, we have

$$a \otimes b = a \otimes (npb) = (pa) \otimes (nb) = 0 \otimes (nb) = 0.$$

So every generating element of $\mathbb{Z}_p \otimes_{\mathbb{Z}} \mathbb{Z}_q$ is 0, which means that the entire $\mathbb{Z}$-module is zero. ♢

EXAMPLE 6.9. What is $\mathbb{Z}_4 \otimes_{\mathbb{Z}} \mathbb{Z}_2$? There's only one element, $1 \otimes 1$, that doesn't have 0 in one coordinate, so this is the only possible nonzero generator. We know that $2 \cdot (1 \otimes 1) = 1 \otimes 2 = 1 \otimes 0 = 0$, so $1 \otimes 1$ is either 0 or has order 2. This means that $\mathbb{Z}_4 \otimes_{\mathbb{Z}} \mathbb{Z}_2$ is either isomorphic to the 0 group or isomorphic to $\mathbb{Z}_2$. If it were isomorphic to the zero group, then we could write, using Proposition 6.5 with $C = \mathbb{Z}_2$, that

$$\underline{\text{Hom}}_{\text{Mod}_R}(0, \mathbb{Z}_2) \cong \underline{\text{Hom}}_{\text{Mod}_R}\big(\mathbb{Z}_4, \underline{\text{Hom}}_{\text{Mod}_R}(\mathbb{Z}_2, \mathbb{Z}_2)\big).$$

But the left side has only one element and the right side has at least 2 (since there are two homomorphisms $\mathbb{Z}_2 \to \mathbb{Z}_2$). So $\mathbb{Z}_4 \otimes_{\mathbb{Z}} \mathbb{Z}_2$ is not the zero group, meaning $\mathbb{Z}_4 \otimes_{\mathbb{Z}} \mathbb{Z}_2 \cong \mathbb{Z}_2$. ♢

This can be extended to a general calculation:

PROPOSITION 6.10. *$\mathbb{Z}/m\mathbb{Z} \otimes_{\mathbb{Z}} \mathbb{Z}/n\mathbb{Z} \cong \mathbb{Z}/\gcd(m,n)\mathbb{Z}$.*

Here are some more properties of the tensor product

PROPOSITION 6.11. *$R \otimes_R B \cong B$.*

Sketch. The map $f(r \otimes b) = rb$ on the generators of $R \otimes_R B$ extends to an isomorphism. □

This proof sketch brushes one subtlety under the rug: When a map is defined on the generators of $A \otimes_R B$, you need to check that it's well-defined, because each generator has multiple names which follow from the generating relations in (3). In this case, it's all good, and you can in general prove that a map on the generators is well-defined by using Proposition 6.6.

PROPOSITION 6.12. *If M, N, and P are A-modules, then*

1. *$M \otimes N \cong N \otimes M$ (commutativity)*
2. *$(M \otimes N) \otimes P \cong M \otimes (N \otimes P)$ (associativity)*

3. $(M \oplus N) \otimes P \cong (M \otimes P) \oplus (N \otimes P)$ *(distributivity)*

Now we have everything we need to calculate the direct sum of any two finitely generated $\mathbb{Z}$-modules. Every finitely generated abelian group has the form $\mathbb{Z}^r \oplus A$, where A is the direct sum of finite cyclic groups. So to calculate the tensor product, you can just distribute the tensor of the direct sums and then apply Propositions 6.10 and 6.11.

6.2.3. Exactness properties of the Hom functor and tensor product

There are some results worth stating:

PROPOSITION 6.13 (Hom$(M, -)$ is left exact). *The sequence*

$$0 \to N_1 \xrightarrow{f} N_2 \xrightarrow{g} N_3$$

of R-modules is exact if and only if the sequence

$$0 \to \mathrm{Hom}(M, N_1) \xrightarrow{\bar{f}} \mathrm{Hom}(M, N_2) \xrightarrow{\bar{g}} \mathrm{Hom}(M, N_3)$$

is exact for every R-module M.

PROPOSITION 6.14 (Hom$(-, N)$ is left exact). *The sequence*

$$M_1 \xrightarrow{f} M_2 \xrightarrow{g} M_3 \to 0$$

of R-modules is exact if and only if the sequence

$$0 \to \mathrm{Hom}(M_3, N) \xrightarrow{\bar{g}} \mathrm{Hom}(M_2, N) \xrightarrow{\bar{f}} \mathrm{Hom}(M_1, N)$$

is exact for every R-module N.

PROPOSITION 6.15 (Tensor product is right exact). *If the sequence*

$$M_1 \xrightarrow{f} M_2 \xrightarrow{g} M_3 \to 0$$

of R-modules is exact, then the sequence

$$M_1 \otimes_R N \xrightarrow{f \otimes 1_N} M_2 \otimes_R N \xrightarrow{g \otimes 1_N} M_3 \otimes_R N \to 0$$

is exact for any R-module N. (1_N is the identity map on N.)

The proofs of the first two consist of unravelling the definitions; nothing interesting occurs. The third statement can be proved that way, but it can also be proved in a more clever way by utilizing the correspondence between tensor product and bilinear maps; see this proof.

It is *not* true that the tensor product is left exact. For example, $0 \to \mathbb{Z} \xrightarrow{\times p} \mathbb{Z}$ is exact, but it is not exact after tensoring (over $\mathbb{Z}$) with $\mathbb{Z}_p$. But if M is a *free* R-module, then it is left exact.

EXERCISE 6.16. Suppose that M is a free R-module. Show that if $f\colon N_1 \to N_2$ is an injective R-module map, then $f \otimes 1_M\colon N_1 \otimes_R M \to N_2 \otimes_R M$ is also injective.

That's enough straight-up algebra for now. Don't worry; we'll be back for more soon.

6.3. QUASI-ISOMORPHISMS AND FREE RESOLUTIONS

When we're doing computations, in the end we don't really care too much about the specific chain complex that we've constructed; our focus is on its homology. This next definition is one way to loosen the definition of equivalence for chain complexes with this in mind.

DEFINITION 6.17. A map $f\colon C_\bullet \to D_\bullet$ of R-module chain complexes is called a *quasi-isomorphism* if $H_n(f)\colon H_n(C_\bullet) \to H_n(D_\bullet)$ is an isomorphism for every $n \in \mathbb{Z}$.

In practice, algebraic topologists (and homological algebraists more generally) only care about chain complexes up to quasi-isomorphism.

EXAMPLE 6.18. In the diagram

$$\begin{array}{ccccccccccc} \cdots & \longrightarrow & 0 & \longrightarrow & \mathbb{Z} & \xrightarrow{\times n} & \mathbb{Z} & \longrightarrow & 0 & \longrightarrow & \cdots \\ & & \downarrow{\scriptstyle f_2} & & \downarrow{\scriptstyle f_1} & & \downarrow{\scriptstyle f_0} & & \downarrow{\scriptstyle f_{-1}} & & \\ \cdots & \longrightarrow & 0 & \longrightarrow & 0 & \longrightarrow & \mathbb{Z}_n & \longrightarrow & 0 & \longrightarrow & \cdots \end{array}$$

the map f defined by $f_i = 0$ if $i \neq 0$ and $f_0(x) = x \bmod p$ is a quasi-isomorphism. ◊

EXERCISE 6.19. Check this. First, does the diagram commute? (This ensure that f is actually a chain map.) Second, what is $H_n(f)$?

We'd like to codify this equivalence of quasi-isomorphism into a category in which two chain complexes are actually isomorphic if they are quasi-isomorphic in the category of chain complexes of R-modules. Such a category exists; it is called the *derived category* (of chain complexes of R-modules) and is denoted $D(R)$. You can think of it as being formed from the category of R-module chain complexes by adding formal inverse morphisms for every quasi-isomorphism. This idea runs into set-theoretic issues: Since you also have to include the composition of any two morphisms, it's not clear that what you get at the end of this process is a category; some collections of morphisms between two objects might be "too big" to be a set.

For this reason, none of the proofs here will rely on the existence of the derived category, but its existence provides a better understanding and motivation of some of the definitions to come.

DEFINITION 6.20. A *free resolution* of an R-module M is a chain complex $C_\bullet$ of free R-modules together with a quasi-isomorphism to the chain complex that has 0 everywhere except at degree 0, where it has M.

In the derived category, free resolutions are isomorphic to the module M, so they're a bit nicer to work with. Example 6.18 is, for example, a free resolution of $\mathbb{Z}_n$.

EXERCISE 6.21. Show that any finitely generated abelian group (considered as a $\mathbb{Z}$-module) has a free resolution with only two nonzero terms. (You'll need to use the fundamental theorem of finitely generated abelian groups.)

In R is a field, then every R-module is already free, so M itself (or, more properly, the chain complex $M_\bullet := \cdots \to 0 \to M \to 0 \to \cdots$) is already a free resolution of itself.

EXAMPLE 6.22. Chain resolutions are not unique. For example, here is a different resolution of $\mathbb{Z}_n$:

$$\begin{array}{ccccccccccc} \cdots & \longrightarrow & 0 & \longrightarrow & \mathbb{Z}\oplus\mathbb{Z} & \xrightarrow{(a,b)\mapsto(a,nb)} & \mathbb{Z}\oplus\mathbb{Z} & \longrightarrow & 0 & \longrightarrow & \cdots \\ & & \downarrow{\scriptstyle \epsilon_2} & & \downarrow{\scriptstyle \epsilon_1} & & \downarrow{\scriptstyle \epsilon_0} & & \downarrow{\scriptstyle \epsilon_{-1}} & & \\ \cdots & \longrightarrow & 0 & \longrightarrow & 0 & \longrightarrow & \mathbb{Z}_n & \longrightarrow & 0 & \longrightarrow & \cdots \end{array}$$

Or one with three terms:

$$\begin{array}{ccccccccccccc} \cdots & \longrightarrow & 0 & \longrightarrow & \mathbb{Z} & \xrightarrow{1\mapsto(1,0)} & \mathbb{Z}\oplus\mathbb{Z} & \xrightarrow{(a,b)\mapsto(0,nb)} & \mathbb{Z}\oplus\mathbb{Z} & \longrightarrow & 0 & \longrightarrow & \cdots \\ & & \downarrow{\scriptstyle \epsilon_3} & & \downarrow{\scriptstyle \epsilon_2} & & \downarrow{\scriptstyle \epsilon_1} & & \downarrow{\scriptstyle \epsilon_0} & & \downarrow{\scriptstyle \epsilon_{-1}} & & \\ \cdots & \longrightarrow & 0 & \longrightarrow & 0 & \longrightarrow & 0 & \longrightarrow & \mathbb{Z}_n & \longrightarrow & 0 & \longrightarrow & \cdots \end{array}$$

◊

Now we get to something Fundamental.

THEOREM 6.23 (Fundamental lemma of homological algebra). *Suppose M and N are R-modules and $F_\bullet$ and $G_\bullet$ are free resolutions of M and N, respectively. For any R-module homomorphism $f\colon M \to N$, there is a chain map $f^*\colon F_\bullet \to G_\bullet$ such that $f = H_0(f^*)\colon M \to N$. Moreover, this map is unique up to chain homotopy.*

The chain map f^* is called a *lift* of f.

The existence part of this proof is a classic "Who's going to stop me?" argument: You pick the maps $f_1^*, f_2^*, \dots$ inductively, ensuring that the diagram commutes at each step. For example: start with the diagram

$$\begin{array}{ccc} F_0 & \xrightarrow{\epsilon_M} & M \\ {\scriptstyle f_0}\downarrow & & \downarrow{\scriptstyle f} \\ G_0 & \xrightarrow{\epsilon_N} & N \end{array}$$

where $H_0(\epsilon_M)$ is an isomorphism $H_0(F_\bullet) \to M$ and $H_0(\epsilon_N)$ is an isomorphism $H_0(G_\bullet) \to N$. (These are the maps from the free resolutions.) First, note that this implies that ϵ_N is surjective; if it weren't, then $H_0(E_\bullet) \subsetneq N$. Therefore, we define f_0 like this: F_0 is a free R-module on some set S. For each s, choose an element $x_s \in G_0$ such that $f\epsilon_M(s) = \epsilon_N(s_x)$; then define $f_0(s) = t$. Booyah, a well-defined function that makes the square commute.

Then you extend the diagram like this:

$$\begin{array}{ccccccc} F_1 & \longrightarrow & \ker(\epsilon_M) & \longrightarrow & F_0 & \xrightarrow{\epsilon_M} & M \\ {\scriptstyle f_1}\downarrow & & \downarrow{\scriptstyle g_0} & & \downarrow{\scriptstyle f_0} & & \downarrow{\scriptstyle f} \\ G_1 & \longrightarrow & \ker(\epsilon_N) & \longrightarrow & G_0 & \xrightarrow{\epsilon_N} & N \end{array}$$

The leftmost horizontal maps are surjective; this comes from the commutativity of the free resolution diagram. The map g_0 comes from the commutativity of the rightmost square in this diagram. Now we pull the same trick: F_1 is free on some set, yada yada yada. And so on and so forth. (These notes provide a more detailed proof.)

That any two maps are chain homotopic is a bit harder to prove, but it should be believable, because any two of these chain maps are the same under the homology functor.

6.4. TENSOR PRODUCT OF CHAIN COMPLEXES

Here is a funky definition:

DEFINITION 6.24. If $C_\bullet$ and $D_\bullet$ are two chain complexes of R-modules, then their tensor product, denoted $C_\bullet \otimes D_\bullet$, has the module $(C_\bullet \otimes D_\bullet)_n = \bigoplus_{p+q=n} C_p \otimes_R D_q$ in degree n and the boundary maps $\partial_n(c_p \otimes d_q) = (\partial c_p) \otimes d_q + (-1)^p c_p \otimes (\partial d_q)$, where $c_p \in C_p$ and $d_q \in D_q$.

Eventually, we'll prove that the tensor product of two chain complexes associated to topological spaces is the chain complex of the product space. If X and Y are CW-complexes, each n-cell in $X \times Y$ is the products of a p-cell in X and an $(n-p)$-cell in Y; this is a partial explanation for the definition of the degree-n term in $C_\bullet \otimes D_\bullet$. The definition of the boundary map on $C_\bullet \otimes D_\bullet$ is just weird; we'll discuss that later, after we've built up some machinery.

EXAMPLE 6.25. There's one very important case of a tensor product: Let $R_\bullet$ denote the chain complex $\cdot \to 0 \to M \to 0 \to \cdots$ that has the ring R in degree 0 and the zero module everywhere else. For any chain complex $C_\bullet$ of $\mathbb{Z}$-modules, we have $(C_\bullet \otimes_{\mathbb{Z}} R_\bullet)_n = C_n \otimes R$. This tensor product is actually an R-module, and with a little investigation you'll find that $H_n(C_\bullet \otimes_{\mathbb{Z}} R_\bullet) = H_n(C_\bullet; R)$. $\lozenge$

This tensor product doesn't come from an internal Hom, which makes it seem a little ad hoc. It turns out that the derived category $D(R)$ also has a tensor product, and this one does come from an internal Hom. This tensor product is denoted $\otimes^{\mathbb{L}}_R$, where for historical reasons, the $\mathbb{L}$ stands for "left" (this is simply because when you write out the statement of the currying isomorphism, the tensor product appears on the left). It turns out that if $C_\bullet$ or $D_\bullet$ is a chain complex of free R-modules, then

$$C_\bullet \otimes^{\mathbb{L}}_R D_\bullet \cong C_\bullet \otimes_R D_\bullet.$$

Now, we don't really care so much about the actual chain complex $C_\bullet \otimes^{\mathbb{L}}_R D_\bullet$; what we care about is its homology.

Definition 6.26. If M and N are R-modules, then $\mathrm{Tor}_i^R(M,N) := H_i(M_\bullet \otimes^{\mathbb{L}}_R N_\bullet)$.

Now, I said we wouldn't do anything that assumed the existence of the derived category, and it seems I just violated that rule. But fear not! We can show that this definition is well-defined without using the derived category.

Proposition 6.27. *If M and N are R-modules and $F_\bullet$ and $F'_\bullet$ are two free resolutions of M, then $H_n(F_\bullet \otimes_R N_\bullet) \cong H_n(F'_\bullet \otimes_R N_\bullet)$.*

With this in hand, we can instead *define* $\mathrm{Tor}_i^R(M,N) = H_i(F_\bullet \otimes_R N_\bullet)$ for any free resolution $F_\bullet$ of M; this is well-defined up to isomorphism and doesn't require any mention of the derived category.

Warning: This proof isn't very enlightening.

Proof of Proposition 6.27. We may use the fundamental lemma to obtain two lifts of the map $1_M\colon M \to M$; call them $f\colon F_\bullet \to F'_\bullet$ and $g\colon F'_\bullet \to F_\bullet$. Now, $g \circ f\colon F_\bullet \to F_\bullet$ is a lift of 1_M; but so is $1_{F_\bullet}$, so these two maps are chain homotopic. Similarly, $f \circ g$ is chain homotopic to $1_{F'_\bullet}$.

Now we have maps $f \otimes 1_{N_\bullet}\colon F_\bullet \otimes N_\bullet \to F'_\bullet \otimes N_\bullet$ and $g \otimes 1_{N_\bullet} F'_\bullet \otimes N_\bullet \to F_\bullet \otimes N_\bullet$. Their compositions are $(g \otimes 1_{N_\bullet}) \circ (f \otimes 1_{N_\bullet}) = (g \circ f) \otimes 1_{N_\bullet}$, which is chain homotopic to $1_{F_\bullet} \otimes 1_{N_\bullet}$. Similarly, $(f \otimes 1_{N_\bullet}) \circ (g \otimes 1_{N_\bullet})$ is chain homotopic to $1_{F'_\bullet} \otimes 1_{N_\bullet}$. This means that, after applying the homology functors, $f \otimes 1_{N_\bullet}$ and $g \otimes 1_{N_\bullet}$ are inverse functions; so $F_\bullet \otimes N$ and $F'_\bullet \otimes N$ have the same homology groups. □

This proof isn't so important. What is important is that you can calculate Tor modules. So let's calculate.

Example 6.28. What is $\mathrm{Tor}_i^{\mathbb{Z}}(\mathbb{Z}_2, \mathbb{Z}_4)$? First, choose a free resolution of $\mathbb{Z}_2$, for example $C_\bullet = \cdot \to 0 \to \mathbb{Z} \xrightarrow{\times 2} \mathbb{Z} \to 0 \to \cdots$. Then we can form the chain complex $C_\bullet \otimes_{\mathbb{Z}} (\mathbb{Z}_4)_\bullet$:

$$\cdots \to 0 \to \mathbb{Z} \otimes_{\mathbb{Z}} \mathbb{Z}_4 \xrightarrow{a \otimes b \mapsto 2a \otimes b} \mathbb{Z} \otimes \mathbb{Z}_4 \to 0 \to \cdots.$$

But we know that $\mathbb{Z} \otimes_{\mathbb{Z}} \mathbb{Z}_4 \cong \mathbb{Z}_4$ via the isomorphism $a \otimes b \mapsto ab$ (see Proposition 6.11). So this chain complex is isomorphic to

$$\cdots \to 0 \to \mathbb{Z}_4 \xrightarrow{\times 2} \mathbb{Z}_4 \to 0 \to \cdots.$$

Now we can just read off the homology of this chain complex. The 0th homology is $\mathbb{Z}_4/2\mathbb{Z}_4 \cong \mathbb{Z}_2$, and the 1st homology is $2\mathbb{Z}_4/\{0\}$. So $\mathrm{Tor}_0^{\mathbb{Z}}(\mathbb{Z}_2, \mathbb{Z}_4) \cong \mathrm{Tor}_1^{\mathbb{Z}}(\mathbb{Z}_2, \mathbb{Z}_4) \cong \mathbb{Z}_2$. The higher Tor groups are 0. ♢

Exercise 6.29. Show that $\mathrm{Tor}_i^{\mathbb{Z}}(\mathbb{Z}_2, \mathbb{Z}_3) = 0$ for every i.

Example 6.30. If R is a field, then $\mathrm{Tor}_i^R(M,N) = 0$ for every $i > 0$. This is because $M_\bullet$ is already a free resolution of M, and $M_\bullet \otimes_R N_\bullet \cong (M \otimes_R N)_\bullet$. So $\mathrm{Tor}_i^R(M,N)$ is the ith homology group of the chain complex

$$\cdots \to 0 \to M \otimes_R N \to 0 \to \cdots$$

Similarly, if M is a finitely generated $\mathbb{Z}$-module (i.e., a finitely generated abelian group), then $\operatorname{Tor}_i^{\mathbb{Z}}(M, N) = 0$ if $i > 1$. This is because M always has a free resolution with at most two nonzero terms. (See Exercise 6.21.) ◊

A final note: You can modify the proof of Proposition 6.27 to show that if $G_\bullet$ and $G'_\bullet$ are two free resolutions of N, then

$$H_n(M_\bullet \otimes_R G_\bullet) \cong H_n(M_\bullet \otimes_R G'_\bullet).$$

In fact, you can extend the proof to show that

$$\operatorname{Tor}_i^R(M, N) \cong H_i(M \otimes_R G_\bullet) \cong H_i(F_\bullet \otimes_R G_\bullet)$$

for any two free resolutions $F_\bullet$ and $G_\bullet$ of M and N, respectively.

6.5. UNIVERSAL COEFFICIENT THEOREM

The homology coefficients can actually be even more general than commutative rings.

DEFINITION 6.31. Let M be an abelian group and $\mathcal{X}$ a semisimplicial set. The chain complex of $\mathcal{X}$ with coefficients in M is defined as $S_\bullet(\mathcal{X}; M) = S_\bullet(\mathcal{X}; \mathbb{Z}) \otimes_{\mathbb{Z}} M_\bullet$. If $\mathcal{X} = \operatorname{Sing}(X)$, we write simply $S_\bullet(X; M)$. Similarly, the homology of $\mathcal{X}$ with coefficients in M is $H_n(\mathcal{X}; M) = H_n\big(S_\bullet(\mathcal{X}; M)\big)$.

A few comments are in order. First, if R is a commutative ring, then $S_\bullet(\mathcal{X}; \mathbb{Z}) \otimes_{\mathbb{Z}} R_\bullet = S_\bullet(\mathcal{X}; R)$ as it was defined before, except here we're only considering it as an abelian group. Also, since $S_\bullet(\mathcal{X}; \mathbb{Z})$ is by definition a free group, the tensor product $\otimes_{\mathbb{Z}}$ is the same as the tensor product $\otimes_{\mathbb{Z}}^{\mathbb{L}}$ in the derived category.

Now, the punchline: The homology with $\mathbb{Z}$-coefficients completely determines the homology groups with any other coefficients.

THEOREM 6.32 (Universal coefficient theorem). *If $C_\bullet$ is a chain complex of free $\mathbb{Z}$-modules and M is an abelian group, then*

$$H_q(C_\bullet \otimes_{\mathbb{Z}} M) \cong \big(H_q(C_\bullet; \mathbb{Z}) \otimes_{\mathbb{Z}} M\big) \oplus \operatorname{Tor}_1^{\mathbb{Z}}\big(H_{q-1}(C_\bullet), M\big).$$

Specializing to semisimplicial sets, we get

$$H_q(\mathcal{X}; M) \cong \big(H_q(\mathcal{X}) \otimes_{\mathbb{Z}} M\big) \oplus \operatorname{Tor}_1^{\mathbb{Z}}\big(H_{q-1}(C_\bullet), M\big).$$

If M is a free abelian group, then $\operatorname{Tor}_1^{\mathbb{Z}}(H_{q-1}(C_\bullet), M) = 0$, because $M_\bullet$ is already a free resolution of M. So in this case, the formula simplifies even more to $H_q(\mathcal{X}; M) \cong H_q(\mathcal{X}) \otimes_{\mathbb{Z}} M$.

This formula is a little odd, so before we prove it, let's see how to use it.

EXAMPLE 6.33. We can go back to our old friend S^2 first. Theorem 6.32 tells us that

$$H_2(S^2; \mathbb{Z}_2) \cong \big(H_2(S^2) \otimes_{\mathbb{Z}} \mathbb{Z}_2\big) \oplus \operatorname{Tor}_1^{\mathbb{Z}}\big(H_1(S^2) \otimes_{\mathbb{Z}} \mathbb{Z}_2\big).$$

Now, $H_1(S^2) = 0$ and $H_2(S^2) \cong \mathbb{Z}$. So the right term of the direct product is 0 and the left term is $\mathbb{Z} \otimes_{\mathbb{Z}} \mathbb{Z}_2 \cong \mathbb{Z}_2$. Therefore $H_2(S^2; \mathbb{Z}^2) \cong \mathbb{Z}_2$, just as we expected. ◊

EXAMPLE 6.34. Next, let's calculate $H_2(\mathbb{RP}^2; \mathbb{Z}_2)$. This time, $H_2(\mathbb{RP}^2) \cong 0$ and $H_1(\mathbb{RP}^2) \cong \mathbb{Z}_2$. So

$$\begin{aligned} H_2(\mathbb{RP}^2; \mathbb{Z}_2) &\cong \big(H_2(\mathbb{RP}^2) \otimes_{\mathbb{Z}} \mathbb{Z}_2\big) \oplus \operatorname{Tor}_1^{\mathbb{Z}}\big(H_1(\mathbb{RP}^2) \otimes_{\mathbb{Z}} \mathbb{Z}_2\big) \\ &\cong (0 \otimes_{\mathbb{Z}} \mathbb{Z}_2) \oplus \operatorname{Tor}_1^{\mathbb{Z}}(\mathbb{Z}_2, \mathbb{Z}_2). \end{aligned}$$

You can calculate that $\operatorname{Tor}_1^{\mathbb{Z}}(\mathbb{Z}_2, \mathbb{Z}_2) \cong \mathbb{Z}_2$, so $H_2(\mathbb{RP}^2; \mathbb{Z}_2) \cong \mathbb{Z}_2$. ◊

⚠ The isomorphism in Theorem 6.32 is not natural in the terms $C_\bullet$ and M. For example, the antipodal map $S^2 \to \mathbb{RP}^2$ induces an isomorphism $H_2(S^2;\mathbb{Z}_2) \to H_2(\mathbb{RP}^2;\mathbb{Z}_2)$; but it does *not* induce an isomorphism $H_2(S^2)\otimes \mathbb{Z}_2 \to H_2(\mathbb{RP}^2)\otimes\mathbb{Z}_2$, simply because no map can—these are different groups. (The map $\operatorname{Tor}_1^{\mathbb{Z}}\big(H_1(S^2),\mathbb{Z}_2\big) \to \operatorname{Tor}_1^{\mathbb{Z}}\big(H_1(\mathbb{RP}^2),\mathbb{Z}_2\big)$ is also not an isomorphism.)

Proof of Theorem 6.32. First, if M is a free module, then

$$H_q(C_\bullet)\otimes_{\mathbb{Z}} M \cong \bigoplus_{1\le i\le \operatorname{rank}(M)} H_q(C_\bullet) \cong H_q(C_\bullet\otimes M).$$

We noted above that $\operatorname{Tor}_i^{\mathbb{Z}}(C_\bullet, M) = 0$ if M is free, so this proves the theorem in this case.

Otherwise, let $F_\bullet = \cdots \to 0 \to F_1 \to F_0 \to 0$ be a free resolution of M. The sequence

$$0 \to F_1 \xrightarrow{\alpha} F_0 \to M \to 0$$

is exact; since each of the groups in $C_\bullet$ is free, Proposition 6.15 and Exercise 6.16 tell us that the sequence

$$0 \to C_\bullet\otimes F_1 \to C_\bullet\otimes F_0 \to C_\bullet\otimes M \to 0$$

is exact. This is a short exact sequence of chain complexes, so it gives us a long exact sequence in homology:

$$\cdots \to H_q(C_\bullet\otimes F_1) \xrightarrow{f} H_q(C_\bullet\otimes F_0) \to H_q(C_\bullet\otimes M) \longrightarrow H_{q-1}(C_\bullet\otimes F_1) \xrightarrow{g} H_{q-1}(C_\bullet\otimes F_0) \to \cdots$$

Each long exact sequence can be cut into short exact sequences; so we get the short exact sequence

$$0 \to \operatorname{coker}(f) \to H_q(C_\bullet\otimes M) \to \ker(g) \to 0.$$

The cokernel of f is

$$H_q(C_\bullet\otimes F_1)/f\big(H_q(C_\bullet\otimes F_0)\big) \cong \big(H_q(C_\bullet)\otimes F_1\big)/\big(H_q(C_\bullet)\otimes\alpha(F_0)\big)$$

since both F_0 and F_1 are free. It's a general fact about tensor products that if $L_1\subseteq L_2$ and N are all R-modules, then $(N\otimes L_2)/(N\otimes L_1)\cong N\otimes(L_2/L_1)$. In this case, we get

$$\operatorname{coker}(f)\cong H_q(C_\bullet)\otimes F_0/f(F_1)\cong H_q(C_\bullet)\otimes M.$$

The kernel of g is isomorphic to the kernel of the map $1\otimes\alpha\colon H_{q-1}(C_\bullet)\otimes F_1 \to H_{q-1}(C_\bullet)\otimes F_0$. But this is by definition $\operatorname{Tor}_1^{\mathbb{Z}}\big(H_{q-1}(C_\bullet), M\big)$.

Putting this together, we have a short exact sequence

$$0 \to H_q(C_\bullet)\otimes M \to H_q(C_\bullet\otimes M) \to \operatorname{Tor}_1^{\mathbb{Z}}\big(H_{q-1}(C_\bullet), M\big) \to 0.$$

One can show that this exact sequence splits, so that the middle term is the direct sum of the outer two. □

The universal coefficients theorem can be strengthened:

THEOREM 6.35 (Universal Coefficients Theorem 2.0). *If R is a principal ideal domain and $C_\bullet$ and $D_\bullet$ are chain complexes of R-modules where every R-module in $C_\bullet$ is free, then there is a short exact sequence*

$$0 \to \bigoplus_{p+q=n}\big(H_p(C_\bullet)\otimes_R H_q(D_\bullet)\big) \to H_n(C_\bullet\otimes_R D_\bullet) \to \bigoplus_{p+q=n-1}\operatorname{Tor}_1^R\big(H_p(C_\bullet), H_q(D_\bullet)\big) \to 0$$

that splits, but not naturally. In particular, if every homology group is free, then the Tor term vanishes, and

$$H_n(C_\bullet\otimes_F D_\bullet)\cong\bigoplus_{p+q=n}\big(H_p(C_\bullet)\otimes_F H_q(D_\bullet)\big).$$

This theorem takes some more work. One nice way to do this is to use the derived category. You first show that the sequence

$$\cdots \xrightarrow{0} H_2(D_\bullet) \xrightarrow{0} H_1(D_\bullet) \xrightarrow{0} H_0(D_\bullet) \to 0$$

is quasi-isomorphic to $D_\bullet$. This is plausible, since all the homology groups are the same, but that's not enough: You need to find an actual sequence of quasi-isomorphisms that sends $D_\bullet$ to this new sequence. But once you do this, UCT2.0 just becomes a direct sum of applications of UTC1.0. Pretty nice.

Now we get a second payoff—the homology of the product of two spaces.

THEOREM 6.36 (Eilenberg–Zilber). *For any two topological spaces X and Y and any commutative ring R, the chain complex $S_\bullet(X \times Y; R)$ is quasi-isomorphic to $S_\bullet(X; R) \otimes_R S_\bullet(Y; R)$.*

First, an example.

EXAMPLE 6.37. What is the homology of $\mathbb{RP}^2 \times \mathbb{RP}^2$ over $\mathbb{F}_2$? It's EZ with Eileberg-Zilber.[14] We know the homology groups of $\mathbb{RP}^2$ cold, and EZ says

$$H_n(\mathbb{RP}^2 \times \mathbb{RP}^2; \mathbb{F}_2) \cong \bigoplus_{p+q=n} H_p(\mathbb{RP}^2; \mathbb{F}_2) \otimes_{\mathbb{Z}} H_q(\mathbb{RP}^2; \mathbb{F}_2).$$

Plugging in the homology groups and recalling that $\mathbb{F}_2 \otimes \mathbb{F}_2 \cong \mathbb{F}_2$, we get the table

n	$H_n(\mathbb{RP}^2 \times \mathbb{RP}^2)$
0	$\mathbb{F}_2$
1	$\mathbb{F}_2 \oplus \mathbb{F}_2$
2	$\mathbb{F}_2 \oplus \mathbb{F}_2 \oplus \mathbb{F}_2$
3	$\mathbb{F}_2 \oplus \mathbb{F}_2$
4	$\mathbb{F}_2$
5	0

◊

EXERCISE 6.38. Do the computations to verify that the table in Example 6.37 is correct.

We won't prove the Eilenberg–Zilber theorem; you can see these notes, or surely many other places, for the full proof. Instead, we'll focus on the canonical map that provides an isomorphism.

The *Alexander–Whitney map* $\alpha \colon S_\bullet(X \times Y; R) \to S_\bullet(X; R) \otimes_R S_\bullet(Y; R)$ is defined as follows. Since α is a chain map, there is a homomorphism of R-modules

$$\alpha_n \colon S_n(X \times Y; R) \to \bigoplus_{p+q=n} S_p(X; R) \otimes_R S_q(Y; R).$$

To define these maps, we will define auxiliary maps $\alpha_{p,q} \colon S_{p+q}(X \times Y; R) \to S_p(X; R) \otimes_R S_q(Y; R)$ for each $p, q \geq 0$; then we can set $\alpha(c) = \bigoplus_{p+q=n} \alpha_{p,q}(c)$.

Remember that $\alpha_{p,q}$ will be defined on the free abelian group generated by the $(p+q)$-simplices. To define $\alpha_{p,q}$, we only need to specify $\alpha_{p,q}(\sigma)$ for every map $\sigma \colon \Delta^{p+q} \to X \times Y$. We do this simply by restriction. Let $\alpha|_k$ denote the restriction of α to the first k coordinates of Δ^{p+q} and $\alpha|^k$ denote the restriction to the last k coordinates of Δ^{p+q}. (These are called the *front* and *back* k-simplices of Δ^{p+q}, respectively.) Also, let π_1 denote the projection onto the first coordinate of σ and π_2 denote the projection onto the second coordinate of σ. We can then define $\alpha_{p,q} = (\pi_1 \circ \alpha)|_p \otimes (\pi_2 \circ \alpha)|^q$. This completely determines the Alexander-Whitney map.

[14] get it?

7. COHOMOLOGY

7.1. HOMOLOGY & CO.

Remember that in the currying isomorphism, there is a sort of duality between the tensor product and the Hom functor. Cohomology exploits this duality to create a parallel theory that marches in the opposite direction, reversing each arrow in homology. It starts with a definition.

DEFINITION 7.1. For any semisimplicial set $\mathcal{X}$ and abelian group M, we define $S^n(\mathcal{X}; M) = \underline{\mathrm{Hom}}_{\mathsf{Ab}}\big(S_n(\mathcal{X}), M\big)$. An element of $S^n(\mathcal{X}; M)$ is called a *singular n-cochain.*

As usual, if X is a topological space, we write $S^n(X; M)$ for $S^n\big(\mathrm{Sing}_n(X); M\big)$. A homomorphism $S_n(\mathcal{X}) \to M$ is defined by a function $\mathcal{X} \to M$ and vice versa. We may also define, for any commutative ring R and R-module M, the set $S^n(\mathcal{X}; M) = \underline{\mathrm{Hom}}_{\mathrm{Mod}_R}\big(S_n(\mathcal{X}), M\big)$. We'll switch back an forth between these: Abelian groups are easier to think about, but everything we say about them works in the more general case of modules over a commutative ring.

As you might expect, these groups assemble themselves into a chain complex.

DEFINITION 7.2. The cohomology boundary map $\partial^n \colon S^n(\mathcal{X}; M) \to S^{n+1}(\mathcal{X}; M)$ is defined by $\partial^n f(\sigma) = (-1)^{n+1} f(\partial_{\mathcal{X}}\sigma)$, where $\partial_{\mathcal{X}}$ is the boundary map in $\mathcal{X}$.

Wait a minute! That map goes in the wrong direction! I told you the arrows would be reversed, didn't I? To make things fit into the category theory and algebra of chain complexes, (co)homological algebraists will sometimes define an actual chain complex $S^*(\mathcal{X}; M)$, where the degree $-k$ part is $S^k(\mathcal{X}; M)$. Then ∂^n is a degree-lowering map from degree $-n$ to degree $-n-1$, as it "should be."

But there's some other monkey business going on: What's up with the sign $(-1)^{n+1}$? If you're only going to do basic cohomology, the sign is irrelevant: $\ker(\partial^n)$ and $\mathrm{im}(\partial^n)$ are the same whether we use the sign or not. But if you want cohomology to interact nicely with homology, then the sign is necessary. Pages 81–82 of Miller's notes provide a good justification for why you might want to include the sign.

⚠ Be wary! Hatcher omits the sign. If you're comparing with his textbook, some things might correspondingly differ.

These issues aside, we do in fact get what we want:

$$\partial^{n+1} \circ \partial^n f(\sigma) = (-1)^{n+1}(-1)^{n+2} f(\partial^2\sigma) = 0.$$

So $S^*(\mathcal{X}; M)$ is a chain complex. We write $H^q(\mathcal{X}; M)$ for the qth *cohomology group* of this chain complex:

$$H^q(\mathcal{X}; M) = H_{-q}\big(S^*(\mathcal{X}; M)\big).$$

Explicitly, $H^q(\mathcal{X}; M) = \ker(\partial^q)/\,\mathrm{im}(\partial^{q-1})$.

Let's do a first example that parallels Exercise 3.1. What is the 0th cohomology group of a topological space X? Well, ∂^{-1} is the zero map, so its image is 0 and $H^0(X; M) \cong \ker(\partial^0)$. Every element in $S^0(X; M)$ is just a function $\mathrm{Sing}_0(X) \to M$, or, in other words, a (not necessarily continuous) function $X \to M$. Say $f \in S^0(X; M)$. To evaluate $\partial^0(f)$, which lies in $S^1(X; M)$, pick a 1-simplex $\sigma \colon \Delta^1 \to X$ and calculate

$$(\partial^0 f)(\sigma) = -f(\partial\sigma) = -f\big(\sigma(e_1) - \sigma(e_0)\big) = f(e_0) - f(e_1).$$

If $f \in \ker \partial^0$, then f is constant on every path component of X and vice versa. So:

PROPOSITION 7.3. *If X is a topological space, then $H^0(X;M)$ is isomorphic to the set of functions $\pi_0(X) \to M$, where $\pi_0(X)$ is the collection of path-connected components of X. In other words, $H^0(X;M) \cong \prod_{i\in\pi_0(X)} M$.*

This means that, if X has only finitely many path components, then $H^0(X) \cong H_0(X)$. This is not true if $\pi_0(X)$ is infinite, since $H^0(X)$ is then a direct product and has no countable basis, while $H_0(X)$ is a direct sum with a countable basis.

7.2. THE (N)EXT FUNCTOR

The dual functor to Hom is the Ext functor.

DEFINITION 7.4. Suppose that M and N are R-modules and $F_\bullet$ is a free resolution of M. Applying the functor $\mathrm{Hom}(-,N)$ to this sequence, we get a cochain complex

$$0 \to \mathrm{Hom}(F_0,N) \to \mathrm{Hom}(F_1,N) \to \mathrm{Hom}(F_2,N) \to \cdots$$

The *Ext functor* $\mathrm{Ext}^n_R(M,N)$ is the nth cohomology group of this sequence: $\mathrm{Ext}^n_R(M,N) = H - -n\big(\mathrm{Hom}(F_\bullet,N)\big)$.

Importantly, the definition of Ext is independent of the choice of free resolution of M. This is proven in pretty much the same way as it was for the Tor functor, using the fundamental theorem of homological algebra. If M is a free R-module, then $\mathrm{Ext}^n_R(M,N) = 0$ for every $n > 0$, since $M_\bullet$ is already a free resolution of M. In particular, if R is a field, then M is necessarily free.

EXAMPLE 7.5. To calculate $\mathrm{Ext}^1_{\mathbb{Z}}(\mathbb{Z}_2,\mathbb{Z}_2)$, first take a free resolution $0 \to \mathbb{Z} \xrightarrow{\times 2} \mathbb{Z} \to 0$ of $\mathbb{Z}_2$. The only maps in $\mathrm{Hom}(\mathbb{Z},\mathbb{Z}_2)$ is $1 \mapsto 1$ and $1 \mapsto 0$, so the cochain complex looks like $0 \to \mathbb{Z}_2 \xrightarrow{f} \mathbb{Z}_2 \to 0$. But what is this map f^*? It's induced by the map $\times 2$ in the free resolution. If $g \in \mathrm{Hom}(\mathbb{Z},\mathbb{Z}_2)$ with $g(1) = x$, then $f^* \circ g$ is a map $\mathbb{Z} \to \mathbb{Z}_2$ given by $1 \mapsto x \mapsto 2x = 0$. So f^* is the zero map, and the cochain complex looks like

$$0 \to \mathbb{F}_2 \xrightarrow{\times 0} \mathbb{F}_2 \to 0 \to \cdots$$

Thus $\mathrm{Ext}^n_{\mathbb{Z}}(\mathbb{Z}_2,\mathbb{Z}_2) = \mathbb{F}_2$ if $n = 0, 1$ and $\mathrm{Ext}^n_{\mathbb{Z}}(\mathbb{Z}_2,\mathbb{Z}_2) = 0$ otherwise. ◊

This funky functor plays the role of Tor in the dual universal coefficients theorem.

THEOREM 7.6 (Cohomology UCT). *If $C_\bullet$ is a chain complex of free $\mathbb{Z}$-modules and M is an abelian group, then*

$$H^q(C_\bullet;M) \cong \underline{\mathrm{Hom}}_{\mathbb{Z}}\big(H_q(C_\bullet),M\big) \oplus \mathrm{Ext}^1_{\mathbb{Z}}\big(H_{q-1}(C_\bullet),M\big).$$

You can see a proof in Miller's notes. Here is one important take-away. In the UCT, only the homology groups of the chain complexes are used, so: If $C_\bullet$ and $D_\bullet$ are quasi-isomorphic, then their cohomology groups are the same! So instead of using the cumbersome singular chain complex of a topological space, we can use, for example, the cellular chain complex or a semisimplicial set derived from a triangulation. Let's see a few examples.

EXAMPLE 7.7. First up, our old pal $\mathbb{RP}^2$: Let's calculate $H^*(\mathbb{RP}^2;\mathbb{Z})$. We'll use the cellular chain complex $C_\bullet = 0 \to \mathbb{Z} \xrightarrow{\times 2} \mathbb{Z} \xrightarrow{\times 0} \mathbb{Z} \to 0$. Each element in $\underline{\mathrm{Hom}}(\mathbb{Z},\mathbb{Z})$ is completely determined by the image of 1, so $\underline{\mathrm{Hom}}(\mathbb{Z},\mathbb{Z}) \cong \mathbb{Z}$. So the cochain complex looks like $0 \to \mathbb{Z}\{f_0\} \xrightarrow{\partial^0} \mathbb{Z}\{f_1\} \xrightarrow{\partial^1} \mathbb{Z}\{f_2\} \to 0$, where $f_i\colon 1 \mapsto 1$ and $S^n(\mathbb{RP}^2;\mathbb{Z}) = \mathbb{Z}\{f_n\}$ for $n = 0,1,2$; and $S^n(\mathbb{RP}^2;\mathbb{Z}) = 0$ otherwise. We can calculate the differential maps like this: $\partial^0 f_0$ is a map from the degree-1 term of $C_\bullet$ to $\mathbb{Z}$ defined by

$$(\partial^0 f_0)(x) = -f_1(\partial_1 x) = -f_1(0) = 0,$$

so ∂_0 is the zero map. Similarly,

$$(\partial^1 f_1)(x) = f_2(\partial_2 x) = 2f_2(x),$$

so ∂^1 is the doubling map. The cohomology chain complex therefore looks like $0 \to \mathbb{Z} \xrightarrow{\times 0} \mathbb{Z} \xrightarrow{\times 2} \mathbb{Z} \to 0$, so we have

$$H^q(\mathbb{RP}^2;\mathbb{Z}) \cong \begin{cases} \mathbb{Z} & \text{if } q = 0 \\ \mathbb{Z}_2 & \text{if } q = 2 \\ 0 & \text{otherwise.} \end{cases} \qquad \Diamond$$

Exercise 7.8. Calculate $H^q(\mathbb{RP}^2;\mathbb{F}_2)$ using the same method as above.

Now an example using the cohomology UCT.

Example 7.9. The universal coefficient theorem tells us that

$$H^1(\mathbb{RP}^2;\mathbb{F}_2) = \underline{\mathrm{Hom}}\big(H_1(\mathbb{RP}^2;\mathbb{Z}),\mathbb{F}_2\big) \oplus \mathrm{Ext}^1_{\mathbb{Z}}(\mathbb{Z},\mathbb{F}_2).$$

Since $\mathbb{Z}$ is free, the Ext term vanishes. So $H^1(\mathbb{RP}^2;\mathbb{F}_2)$ is isomorphic to $\underline{\mathrm{Hom}}(\mathbb{F}_2,\mathbb{F}_2) \cong \mathbb{F}_2$. $\Diamond$

Exercise 7.10. Calculate $H^2(\mathbb{RP}^2;\mathbb{F}_2)$ using the UCT and Example 7.5.

7.3. The cohomology ring

We can assemble all the cohomology groups into one big monster:

$$H^*(X;R) := \bigoplus_{n\geq 0} H^n(X;R).$$

Each term $H^n(X;R)$ is an R-module, so the whole direct sum is an R-module, as well. This kind of direct sum is called a *graded R-module*. The goal of this section is to define a product that turns this structure into a ring. Up to this point, we've basically just rephrased homology with coefficients as cohomology. Introducing this product will be something new, and it will allow us to distinguish even more spaces. Let's get started.

Lemma 7.11. *If $C_\bullet$ and $D_\bullet$ are chain complexes of R-modules, then there is a map*

$$S^\bullet(C_\bullet;R) \otimes S^\bullet(D_\bullet;R) \longrightarrow S^\bullet(C_\bullet \otimes D_\bullet;R)$$

defined by sending the tensor product of $f\colon C_p \to R$ and $g\colon C_q \to R$ to the map $h\colon C_\bullet \otimes D_\bullet \to R$ defined by

$$h(c_p \otimes d_q) = \begin{cases} (-1)^{pq} f(c_p)g(d_q) & \text{if } c_p \in C_p \text{ and } d_q \in D_q \\ 0 & \text{otherwise.} \end{cases}$$

If $H_n(C_\bullet)$ is finitely generated for every n, then this map is an isomorphism.

Theorem 7.12 (Cohomology cross product). *If X and Y are topological spaces, there is a map*

$$H^*(X;R) \otimes H^*(Y;R) \longrightarrow H^*(X \times Y;R)$$

of graded R-modules that is an isomorphism if $H_n(X;R)$ is free and finitely generated for each $n \in \mathbb{Z}$.

Since $H^*(X;R)$ and $H^*(Y;R)$ are both graded rings (in particular, direct products), their tensor product is graded, with the part in degree n being

$$\bigoplus_{p+q=n} H^q(X;R) \otimes_R H^p(X;R).$$

So to prove the statement, we need to prove it for each degree: When $H^n(X;R)$ is free and finitely generated, this says

$$H^n(X \times Y; R) \cong \bigoplus_{p+q=n} H^p(X;R) \otimes_R H^q(Y;R),$$

which is a kind of coEilenberg–Zilber theorem.

Proof of Theorem 7.12. First, apply the UCT to get a map

$$\big(H^*(X;R) \otimes_R H^*(Y;R)\big)_n = \bigoplus_{p+q=n} H_{-p}\big(S^\bullet(X);R\big) \otimes H_{-q}\big(S^\bullet(Y);R\big) \longrightarrow H_{-n}\big(S^\bullet(X) \otimes S^\bullet(Y);R\big).$$

If $H_n(X)$ is free for every $n \in \mathbb{Z}$, then the Tor terms vanish in the UCT, so this map is an isomorphism. Lemma 7.11 provides a map

$$H_{-n}\big(S^\bullet(X) \otimes S^\bullet(Y);R\big) \longrightarrow H_{-n}\big(\underline{\mathrm{Hom}}(S_\bullet(X) \otimes S_\bullet(Y), R)\big)$$

which is an isomorphism if the $H_n(X)$ are finitely generated. The Alexander–Whitney map induces a map

$$H_{-n}\big(\underline{\mathrm{Hom}}(S_\bullet(X) \otimes S_\bullet(Y), R)\big) \longrightarrow H_{-n}\big(\underline{\mathrm{Hom}}(S_\bullet(X \times Y), R\big);$$

this is always an isomorphism. (The arrow is in the opposite direction of the Alexander–Whitney map because cohomology is contravariant.) Finally, the definition of cohomology gives us

$$H_{-n}\big(\underline{\mathrm{Hom}}(S_\bullet(X \times Y), R\big) \cong H_{-n}\big(S^\bullet(X \times Y);R\big) \cong H^n(X \times Y;R).$$

Composing all these maps proves the theorem. ☐

Given a topological space X, we'll let Δ denote the diagonal map $X \to X \times X$ that sends $x \mapsto (x,x)$. Since cohomology is contravariant, the diagonal map induces a map $H^*(X \times X;R) \to H^*(X;R)$. With this, we can define a product structure on $H^*(X;R)$.

DEFINITION 7.13. The *cup product multiplication* on $H^*(X;R)$ is the composite map

$$H^*(X;R) \otimes_R H^*(X;R) \longrightarrow H^*(X \times X;R) \xrightarrow{\Delta^*} H^*(X;R).$$

The image of $x \otimes y$ under this map is denoted $x \smile y$.

An element $x \in H^p(X;R)$ is called *homogeneous of degree* p, and we write $|x| = p$. Every element of $H^*(X;R)$ is a finite sum of homogeneous elements. If x and y are homogeneous elements of degree p and q, respectively, then their cup product $x \smile y$ is a homogeneous element of degree $p+q$. (You can verify this by tracing through the maps in the proof of Theorem 7.12.)

So how do we actually compute the cup product of two elements? Each element $\bar{f} \in H^p(X;R)$ is actually an equivalence class of cocycles: functions $f\colon \mathrm{Sing}_p(X) \to R$ that satisfy $\partial^p f = 0$, that is, $f(\partial_{p+1}\sigma) = 0$ for every $\sigma\colon \Delta^{p+1} \to X$. (In short: equivalence classes of elements in $\ker \partial^p$.) If $\bar{g} \in H^q(X;R)$, then we can compute $\bar{f} \smile \bar{g}$ as follows: First fix any functions $f\colon \mathrm{Sing}_p(X) \to R$ and $g\colon \mathrm{Sing}_q(X) \to R$ in the equivalence classes $\bar{f}$ and $\bar{g}$, respectively. Then $\bar{f} \smile \bar{g}$ is the equivalence class of the function $f \smile g\colon \mathrm{Sing}_{p+q}(X) \to R$ defined by

$$(f \smile g)(\sigma) = (-1)^{pq} f(\sigma|_p) g(\sigma|^q)$$

for every $\sigma\colon \Delta^{p+q} \to X$. (Recall that $\sigma|_p$ is the restriction of σ to the front p face and $\sigma|^q$ is the restriction to the back q face.)

Where does this formula come from? The Alexander–Whitney map along with Lemma 7.11. You can trace through the maps, if you want, to see how this comes out. The point is, this gives a method to explicitly calculate the cup product of two elements. It's a bit arduous, and we'll develop tools to avoid doing it, but it is possible.

At this point, I will state some facts about the cup product. Some of them are easy to prove; some of them are harder.

1. There is an identity element $1 \in H^0(X;R)$ such that $1 \smile x = x \smile 1 = x$ for every $f \in H^*(X;R)$.
2. The cup product is associative.
3. If x and y are homogeneous elements of $H^*(X;R)$ with degrees p and q, respectively, then $x \smile y = (-1)^{pq} y \smile x$.

Of these, the last is definitely the most surprising. The cup product is not commutative; rather, it's what's called *graded-commutative*.

Let's start with the easiest one: identity.

PROPOSITION 7.14. *Let* $\mathbf{1}\colon \mathrm{Sing}_0(X) \to R$ *be the function that sends each simplex* $\sigma^0 \to X$ *to* $1 \in R$. *The image* $\bar{\mathbf{1}}$ *of* $\mathbf{1}$ *in cohomology is the identity element for the cup product on* $H^*(X;R)$.

Proof. This follows directly from our formula for the cup product: If $\bar{f} \in H^p(X;R)$, then

$$(f \smile \mathbf{1})(\sigma) = (-1)^{p\cdot 0} f(\sigma|_p) 1(\sigma|^0) = f(\sigma),$$

since the front p-face of $\sigma\colon \Delta^p \to X$ is σ itself. The calculation for $\mathbf{1} \smile f$ is similar. □

Next: associativity.

PROPOSITION 7.15. *The cup product on* $H^*(X;R)$ *is associative.*

Proof. Suppose that $\bar{f}, \bar{g}, \bar{h} \in H^*(X;R)$ are homogeneous of degree p, q, and r, respectively. Then $(\bar{f} \smile \bar{g}) \smile \bar{h}$ is the equivalence class of $(f \smile g) \smile h$, which we can calculate. Let $\sigma\colon \Delta^{p+q+r} \to X$ and $\sigma|_{{}_q}$ denote restriction to the "middle" q-face of σ (that is, to the coordinates $p, p+1, \ldots, p+q$). Then

$$\begin{aligned}\big((f \smile g) \smile h\big)(\sigma) &= (-1)^{(p+q)r} (f \smile g)(\sigma|_{p+q})\, h(\sigma|^r) \\ &= (-1)^{pr+qr} (-1)^{pq} f(\sigma|_p)\, g(\sigma|_{{}_q})\, h(\sigma|^r).\end{aligned}$$

Expanding out the product $f \smile (g \smile h)$ gives the same result, so the cup product is associative. □

The fact that $x \smile y = (-1)^{pq} y \smile x$ whenever $|x| = p$ and $|y| = q$ is harder to prove. So I'll just authoritatively state it as a proposition.

PROPOSITION 7.16. *If* $x, y \in H^*(X;R)$ *are homogeneous elements of degree* p *and* q, *respectively, then* $x \smile y = (-1)^{pq} y \smile x$.

Graded commutativity has the odd consequence that if x is a homogeneous element with odd degree, then $x \smile x = -(x \smile x)$; that is, $2(x \smile x) = 0$. So whenever 2 is not a zero-divisor in R and x is a homogeneous element with odd degree, then $x \smile x = 0$.

Putting everything we've done together, we find that $H^*(X;R)$, with the cup product, is what's called a graded commutative ring.

Let's close with a few example computations.

EXAMPLE 7.17. We will compute the cohomology ring of S^n (with $n \geq 1$). First, what are the cohomology groups? The cellular chain complex for S^n has $\mathbb{Z}$ in degrees n and 0, so $S^\bullet(S^n)$ (what a notation clash!) has $\mathbb{Z}$ in degrees 0 and $-n$. (Check this!) So $H^q(S^n) = \mathbb{Z}$ if $q = 0, n$ and $H^q(S^n) = 0$ for $q \neq 0, n$. There's actually not much to compute in this example! We can write

$$H^*(S^n;\mathbb{Z}) = \begin{cases} \mathbb{Z}\{\mathbf{1}\} & \text{in degree } 0 \\ \mathbb{Z}\{x\} & \text{in degree } n \\ 0 & \text{in all other degrees.} \end{cases}$$

Here, x is just some generator for $H^n(S^n)$. Since $\mathbf{1}$ is the identity, we just need to compute $x \smile x$ to know the full product structure of $H^*(S^n)$. But $x \smile x$ is a homogeneous element of degree $2n$, and there's only one of those: $x \smile x = 0$. As a ring, then, $H^*(S^n) \cong \mathbb{Z}[x]/x^2$, where $|x| = n$. ◊

EXAMPLE 7.18. Now let's do something a little more involved: the cohomology ring of the Klein bottle K with $\mathbb{F}_2$ coefficients. We can first take a simplicial realization of the bottle like this:

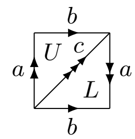

The chain complex for this geometric realization looks like

$$\cdots \to 0 \to \mathbb{F}_2\{U, L\} \to \mathbb{F}_2\{a, b, c\} \to \mathbb{F}_2\{v\} \to 0$$

with

$$\begin{aligned} \partial U &= b - c + a = a + b + c \\ \partial L &= a - b + c = a + b + c \\ \partial a = \partial b = \partial c &= 0. \\ \partial v &= 0. \end{aligned}$$

(Remember that the ring of coefficients is $\mathbb{F}_2$, where $x = -x$.) Dualizing, we get the cohomological chain complex

$$0 \to \mathbb{F}_2\{\delta_v\} \to \mathbb{F}_2\{\delta_a, \delta_b, \delta_c\} \to \mathbb{F}_2\{\delta_U, \delta_L\} \to 0 \to \cdots$$

where $\delta_a\colon \mathbb{F}_2\{a, b, c\} \to \mathbb{F}_2$ is the group homomorphism that satisfies $\delta_a(n_1a + n_2b + n_3c) = n_1$; and similar definitions for the others. (These are the group homomorphisms induced by the characteristic functions.) Their boundaries are defined as follows. First, $\partial^2(U) = \partial^2(L) = 0$, since ∂^2 is necessarily the zero map. Next simplest, $\partial^0\delta_v$ is an element in $\mathbb{F}_2\{\delta_a, \delta_b, \delta_c\}$; more to the point, it's a function $\mathbb{F}_2\{a, b, c\} \to \mathbb{F}_2$. To figure it out, we simply evaluate it at a, b, and c:

$$\partial^0\delta_v(a) = \delta_v(\partial_1 a) = \delta_v(0) = 0.$$

Similar calculations will tell you that $\partial^0\delta_v(b) = \partial^0\delta_v(c) = 0$, so $\partial^0\delta_v = 0$, the zero map. Finally, the middle maps. We apply the same procedure:

$$\partial^1\delta_a(U) = \delta_a(\partial_2 U) = \delta_a(b - c + a) = 1$$

and

$$\partial^1\delta_a(L) = \delta_a(\partial_2 L) = \delta_a(a - b + c) = 1.$$

So $\partial^1\delta_a = \delta_U + \delta_L$. If you do similar calculations, you'll find that $\partial^1\delta_a = \partial^1\delta_b = \partial^1\delta_c$.

Now we can calculate the cohomology groups of the Klein bottle. The kernel of ∂^2 is $\mathbb{F}_2\{U, L\}$, while the image of ∂^1 is $\delta_U + \delta_L$. Denoting by $\overline{\delta_U}$ the image of δ_U in the cohomology group, we have $H^2(K; \mathbb{F}_2) = \mathbb{F}_2\{\overline{\delta_U}\}$.

Moving on, the kernel of ∂^1 is generated by $\delta_a + \delta_b$ and $\delta_a + \delta_c$ (remember that we have $\mathbb{F}_2$ coefficients), while the image of ∂^0 is $\{0\}$. Letting $\overline{\delta_{a,b}}$ and $\overline{\delta_{b,c}}$ denote the images of $\delta_a + \delta_b$ and $\delta_b + \delta_c$ in the cohomology group, we have $H^1(K; \mathbb{F}_2) = \mathbb{F}_2\{\overline{\delta_{a,b}}, \overline{\delta_{b,c}}\}$.

Finally, the kernel of ∂^0 is generated by δ_v, so $H^0(K; \mathbb{F}_2) = \mathbb{F}_2\{\overline{\delta_v}\}$. Putting this all together, we get

$$H^*(K; \mathbb{F}_2) \cong \begin{cases} \mathbb{F}_2\{\overline{\delta_v}\} & \text{in degree } 0 \\ \mathbb{F}_2\{\overline{\delta_{a,b}}, \overline{\delta_{b,c}}\} & \text{in degree } 1 \\ \mathbb{F}_2\{\overline{\delta_U}\} & \text{in degree } 2 \\ 0 & \text{otherwise.} \end{cases}$$

Now we need to calculate the cup products. Some are easy: There are only two elements in $\mathbb{F}_2\{\overline{\delta_v}\}$: the zero element and $\overline{\delta_v}$. Since the identity element of $H^*(K; \mathbb{F}_2)$ must be in here and it can't be 0, it must be $\overline{\delta_v}$. Also, $\overline{\delta_U} \smile \overline{\delta_U} = 0$ for degree reasons; as does $\overline{\delta_{a,b}} \smile \overline{\delta_U}$ and $\overline{\delta_{b,c}} \smile \overline{\delta_U}$.

The remaining interesting products are those between homogeneous degree-1 elements. For these, we can just work through the calculation, which are somewhat easier, since in $\mathbb{F}_2$, we don't have to worry about signs. To determine $\overline{\delta_{a,b}} \smile \overline{\delta_{a,b}}$, we calculate

$$\begin{aligned}\big((\delta_a + \delta_b) \smile (\delta_a + \delta_b)\big)(U) &= (\delta_a + \delta_b)(U|_1)\,(\delta_a + \delta_b)(U|^1) \\ &= (\delta_a + \delta_b)(a)\,(\delta_a + \delta_b)(b) \\ &= 0\end{aligned}$$

and

$$\begin{aligned}\big((\delta_a + \delta_b) \smile (\delta_a + \delta_b)\big)(L) &= (\delta_a + \delta_b)(L|_1)\,(\delta_a + \delta_b)(L|^1) \\ &= (\delta_a + \delta_b)(c)\,(\delta_a + \delta_b)(a) \\ &= 1.\end{aligned}$$

So $\overline{\delta_{a,b}} \smile \overline{\delta_{a,b}} = \overline{\delta_U}$. You can similarly calculate that

$$\begin{aligned}\overline{\delta_{a,b}} \smile \overline{\delta_{b,c}} &= \overline{\delta_U} \\ \overline{\delta_{b,c}} \smile \overline{\delta_{b,c}} &= 0.\end{aligned}$$

Together with graded commutativity (which implies that $\overline{\delta_{b,c}} \smile \overline{\delta_{a,b}} = \overline{\delta_U}$, as well), this completely determines the structure of the ring $H^*(K;\mathbb{F}_2)$. ◊

EXERCISE 7.19. Determine the ring structure of $H^*($$;\mathbb{F}_2)$.

7.4. COHOMOLOGY OF COMPOSITE SPACES

The motto of this section is: Basically everything that we've proved for homology is true for cohomology with the same proofs.

So, taking that as an axiom, let's state some results.

THEOREM 7.20. *$H^*(X \sqcup Y;R) \cong H^*(X;R) \oplus H^*(Y;R)$ as graded rings.*

If X and Y are *pointed spaces*, meaning that they have distinguished points $x \in X$ and $y \in Y$, then their *wedge product* $X \vee Y$ is defined by gluing X and Y together at a point: $(X \sqcup Y)/(x \sim y)$. We determined the cohomology of the wedge product of spheres; in general,

$$H_n(X \sqcup Y;R) \cong \begin{cases} H_n(X \vee Y;R) & \text{if } n > 0 \\ H_0(X \vee Y;R) \oplus \mathbb{Z} & \text{if } n = 0.\end{cases}$$

The case for cohomology is similar:

THEOREM 7.21. *If (X,x) and (Y,y) are pointed spaces, then the surjective map $X \sqcup Y \to X \vee Y$ induces a ring homomorphism $H^*(X \vee Y) \to H^*(X \sqcup Y)$ that is an isomorphism in positive degrees and injective in degree 0.*

Of course, we can always calculate $H^0(X \sqcup Y;R)$ explicitly using Proposition 7.3.

EXAMPLE 7.22. The cohomology ring of $S^1 \vee S^2$ (which is independent of the chosen base point) is

$$H^*(S^1 \vee X^2) \cong \begin{cases} \mathbb{Z}\{\mathbf{1}\} & \text{in degree } 0 \\ \mathbb{Z}\{x\} & \text{in degree } 1 \\ \mathbb{Z}\{y\} & \text{in degree } 2 \\ 0 & \text{otherwise.}\end{cases}$$

Since $H^*(S^1 \vee S^2)$ is a subring of $H^*(S^1 \sqcup S^2) \cong H^*(S^1) \oplus H^*(S^2)$ (this is Theorem 7.21), we can determine the cup products in $H^*(S^1 \vee S^2)$ by the cup products in the product ring. So, for example, $x \smile x = 0$ in the ring $H^*(S^1)$, so $(x, 0) \smile (x, 0) = (0, 0)$ in the ring $H^*(S^1) \oplus H^*(S_2)$. Similarly, $y \smile y = 0$. Because $x \smile y$ is a homogeneous element of degree 3, we know $x \smile y = 0$. This determines all products. So, as a graded ring, $H^*(S^1 \vee S^2) \cong \mathbb{Z}[x]/(x^2) \oplus \mathbb{Z}[y]/(y^2)$ with $|x| = 1$ and $|y| = 2$. ◊

And now: How might we calculate the cohomology ring of a cross product of two space?

THEOREM 7.23. *The cohomology cross product is a homomorphism of graded rings with the multiplication*

$$(a_1 \otimes b_1) \smile (a_2 \otimes b_2) = (-1)^{|b_1||a_2|}\big((a_1 \smile a_2) \otimes (b_1 \smile b_2)\big)$$

in the ring $H^*(X; R) \otimes_R H^*(Y; R)$. This map is an isomorphism if $H_n(X; R)$ is free and finitely generated for every $n \in \mathbb{Z}$.

Let's see this in action.

EXAMPLE 7.24. I hope you've done Exercise 7.19; now we'll compute it another way, using the equation  $= S^1 \times S^1$. Since $H_n(S^1)$ is free and finitely generated,

$H^*($$)$ $$\cong H^*(S^1) \otimes H^*(S^1) \cong \begin{cases} \mathbb{Z}\{\mathbf{1} \otimes \mathbf{1}\} & \text{in degree } 0 \\ \mathbb{Z}\{\mathbf{1} \otimes x, x \otimes \mathbf{1}\} & \text{in degree } 1 \\ \mathbb{Z}\{x \otimes x\} & \text{in degree } 2 \\ 0 & \text{otherwise.} \end{cases}$$

The products are easy to calculate with Theorem 7.23. For example:

$$(\mathbf{1} \otimes x) \smile (x \otimes \mathbf{1}) = (-1)^{|x| \cdot |x|}\big((\mathbf{1} \smile x) \otimes (x \smile \mathbf{1})\big) = -(x \otimes x).$$

You can calculate the rest of the products, should you wish. ◊

7.5. COHOMOLOGY RING AS AN INVARIANT

Why cohomology?

This is an excellent question—wasn't homology, plain and simple, complicated enough? In some sense, yes: Everything in cohomology is determined by homology. Strictly speaking, even the product structure of the cohomology ring comes from the homology structure; it's just that in homology, what you have is a "coproduct." and it's pschologically much easier to deal with rings and algebras than co-rings and coalgebras. Plus, we get to immediately import everything from algebra when we work with cohomology rings; we don't have to build up an entire dual theory.

Either way you spin it—cohomological algebras or homologial coalgebras—the product structure that you get is genuinely useful: It can distinguish more spaces than simple homology groups can. Let's see an example.

Attach two copies of S^1 at two different points of the sphere S^2 to obtain the space 🐭 $= (S^2 \vee S^1) \vee S^1$. You can use Theorem 7.21 to calculate that

$$H^*(🐭) \cong \begin{cases} \mathbb{Z}\{\mathbf{1}\} & \text{in degree } 0 \\ \mathbb{Z}\{y_1, y_2\} & \text{in degree } 1 \\ \mathbb{Z}\{x\} & \text{in degree } 2 \\ 0 & \text{otherwise.} \end{cases}$$

But we've also calculated the cohomology of the torus, and $H^n($🐭$) = H^n($$)$ for every $n \in \mathbb{Z}$. The cohomology groups of these spaces, then, cannot distinguish them. (Neither can the homology groups.) Are they actually different spaces, up to homotopy? It seems like not—there are two points in 🐭 that, if removed, disconnect the space, whereas 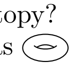 has no such point. But that only shows that they're not homeomorphic.

In fact, the ring structure of H^*(🐭) differs from H^*(), and this shows that 🐭 and  are *not* homotopy equivalent. In H^*(🐭), the square of every homogeneous element of degree 1 is 0. This is because H^*(🐭) is a subring of $H^*(S^2) \oplus H^*(S^1) \oplus H^*(S^1)$, and this inclusion is an isomorphism in positive degrees. On the other hand, the squares of homogeneous elements in degree 1 of H^*() are nonzero. So these rings are not isomorphic, and the spaces are not homotopy equivalent.

7.6. POINCARÉ DUALITY

Elements of $H^q(X;R)$ are (equivalence classes of) functions that map q-simplices to R, while elements of $H_q(X;R)$ are (equivalence classes of) sums of q-simplices. This suggests the possibility of an action of the cohomology groups on the homology groups: We can map the pair $(f,c) \in H^q(X;R)\times H_q(X;R)$ to $f(c)$. This is a bilinear map, so it induces a map $H^q(X;R)\otimes H_q(X;R) \to R$ given by $f \otimes c \mapsto f(c)$. Well, as long as it's well-defined. We need to make sure that the output of this function doesn't depend on the choice of representative from each equivalence class.

LEMMA 7.25. *Suppose that* $A \xrightarrow{\partial_A} B \xrightarrow{\partial_B} C$ *is a chain complex and* $\mathrm{Hom}(C;R) \xrightarrow{\partial_B^\vee} \mathrm{Hom}(B;R) \xrightarrow{\partial_A^\vee} \mathrm{Hom}(A;R)$ *is its dual. The map*

$$\varphi\colon \frac{\ker \partial_A^\vee}{\operatorname{im} \partial_B^\vee} \otimes \frac{\ker \partial B}{\operatorname{im} \partial A} \to R$$

given by $\varphi([f]\otimes[c]) = f(c)$ *is well-defined.*

Proof. We define $\varphi'\colon \ker \partial_A^\vee \otimes \ker \partial_B \to R$ by $\varphi'(f\otimes c) = f(c)$. Since φ' is bilinear, this is well-defined. We now need to prove that φ' is constant on equivalence classes $[f]$ and $[c]$.

First, for any $f \in \ker \partial_A^\vee$ and $b \in \operatorname{im} \partial_A$, we have $f(b) = 0$; so $\varphi'(f\otimes c) = \varphi'(f \otimes d)$ when $[c] = [d]$. To show that φ' is constant on equivalence classes of $\operatorname{im} \partial_B^\vee$, pick a function $g \in \operatorname{im} \partial_B^\vee$. This means that $g = h\partial_B$ for some $h\colon C \to \mathbb{F}_2$. In particular, for each $c \in \ker \partial_B$, we have $g(c) = h\partial_B(c) = 0$. So $\varphi'(g \otimes c) = 0$; therefore φ' is constant on cosets of $\operatorname{im} \partial_B^\vee$. ☐

This map is called the *Kronecker pairing*. Remember that the currying isomorphism associates every map $A \otimes B \to C$ with a map $A \to \mathrm{Hom}(B, C)$. (The latter is called the *adjoint map*.) The corresponding map for the Kronecker pairing is $\psi\colon H^q(X;R) \to \mathrm{Hom}_{\mathrm{Mod}_R}\big(H_q(X;R), R\big)$ where $\varphi(f)$ is the map $c \mapsto f(c)$.

DEFINITION 7.26. A *pairing* of two finitely generated R-modules M and N is a map $M \otimes_R N \to R$. The pairing is called *perfect* if the adjoint map is an isomorphism.

Given a pairing $\varphi\colon M \otimes N \to R$, algebraic topologists often write $\langle x, y\rangle$ for $\varphi(x \otimes y)$. If $M \otimes_R N \to R$ is a perfect pairing, then M is isomorphic to $\mathrm{Hom}_{\mathrm{Mod}_R}(N, R)$, the dual of N. Poinaré duality is a statement about pairings of the cohomology and homology of manifolds. So we really should define a manifold first.

DEFINITION 7.27. A *n-dimensional manifold*, or n-manifold for short, is a Hausdorff topological space[15] in which every point has a neighborhood that is homeomorphic to $\mathbb{R}^n$.

So $\mathbb{R}^n$ is definitely an n-manifold. But so are S^{n-1} and $D^n \setminus S^n$. (The closed unit disc is *not* a manifold.) The torus is a 2-manifold, as is $\mathbb{RP}^2$. The wedge product $S^2 \vee S^2$ is not a manifold, since the point where the spheres are glued together has no neighborhood homeomorphic to $\mathbb{R}^2$, even though every other point does.

Now we sneak into the halls of topology and steal the fact that

[15] A topological space X is *Hausdorff* if every pair of points can be separated by open sets. Specifically: For every pair of distinct points $x, y \in X$ there are open sets U and V with $x \in U$ and $y \in V$ such that $U \cap V = \emptyset$.

THEOREM 7.28. *Every compact manifold is homotopy equivalent to a finite CW-complex.*

Since it's stolen, we don't have to prove it.

This means that the (co)homology groups of a compact manifold are always finitely generated. Now we can state Poincaré duality.

THEOREM 7.29 (Poincaré duality). *If M is a compact n-manifold, then there is a unique element $[m] \in H_n(M; \mathbb{F}_2)$ such that, for all integers p and q with $p + q = n$, the composition*

$$H^p(M; \mathbb{F}_2) \otimes_{\mathbb{F}_2} H^q(M; \mathbb{F}_2) \xrightarrow{\smile} H^n(M; \mathbb{F}_2) \xrightarrow{\langle \cdot, [m] \rangle} \mathbb{F}_2$$

is a perfect pairing.

You might wonder why I've abruptly switched from an arbitrary ring to the very specific ring $\mathbb{F}_2$. A version of Theorem 7.29 is true for arbitrary rings, but it takes much more machinery to set up and state. (Though it might be considered a more "true" expression of the duality.) If all you're interested in is the $\mathbb{F}_2$-cohomology groups of a compact manifold, then Poincaré duality can make this quite easy.

EXAMPLE 7.30. Suppose that M is a 3-manifold with $H^0(M; \mathbb{F}_2) \cong \mathbb{F}_2$ and $H^1(M; \mathbb{F}_2) \cong \mathbb{F}_2 \oplus \mathbb{F}_2$. Then using the fact that $2 + 1 = 3$ and Theorem 7.29, we have

$$H^2(M; \mathbb{F}_2) \cong \underline{\mathrm{Hom}}\big(H^1(M; \mathbb{F}_2), \mathbb{F}_2\big) \cong \mathbb{F}_2 \oplus \mathbb{F}_2.$$

And $0 + 3 = 3$, so

$$H^3(M; \mathbb{F}_2) \cong \underline{\mathrm{Hom}}\big(H^0(M; \mathbb{F}_2), \mathbb{F}_2\big) \cong \mathbb{F}_2.$$

Since $H^k(M; \mathbb{F}_2) \cong 0$ for every $k < 0$, we can use the same trick to conclude that $H^q(M; \mathbb{F}_2) = \{0\}$ for every $q > 3$. So, using only the first two cohomology groups, we can to compute

$$H^q(M; \mathbb{F}_2) \cong \begin{cases} \mathbb{F}_2 & \text{if } q = 0 \\ \mathbb{F}_2 \oplus \mathbb{F}_2 & \text{if } q = 1 \\ \mathbb{F}_2 \oplus \mathbb{F}_2 & \text{if } q = 2 \\ \mathbb{F}_2 & \text{if } q = 3 \\ 0 & \text{otherwise.} \end{cases} \qquad \Diamond$$

Poincaré duality can be used to extract information about the ring structure (not just the groups structure), as well.

EXAMPLE 7.31. Let's calculate $H^*(\mathbb{RP}^2; \mathbb{F}_2)$. We know from earlier calculations (or you could calculate right now, using the cellular chain complex) that $H^q(\mathbb{RP}^2; \mathbb{F}_2) \cong \mathbb{F}_2$ for $q = 0, 1, 2$ and $H^q(\mathbb{RP}^2; \mathbb{F}_2) \cong \{0\}$ otherwise. So

$$H^*(\mathbb{RP}^2; \mathbb{F}_2) \cong \begin{cases} \mathbb{F}_2\{\mathbf{1}\} & \text{in degree } 0 \\ \mathbb{F}_2\{x\} & \text{in degree } 1 \\ \mathbb{F}_2\{y\} & \text{in degree } 2 \\ 0 & \text{all other degrees} \end{cases}$$

since $x \smile y = y \smile y = 0$ (by considering degree), we only need to calculate $x \smile x$.

Conveniently, $\mathbb{RP}^2$ is a 2-manifold and $1 + 1 = 2$. (It's almost like this was planned.) Poincaré duality says that the map

$$H^1(\mathbb{RP}^2; \mathbb{F}_2) \otimes_{\mathbb{F}_2} H^1(\mathbb{RP}^2; \mathbb{F}_2) \xrightarrow{\smile} H^2(\mathbb{RP}^2; \mathbb{F}_2) \xrightarrow{\langle \cdot, [m] \rangle} \mathbb{F}_2$$

is a perfect pairing for some $[m] \in H_n(\mathbb{RP}^2; \mathbb{F}_2)$. In other words, this composition induces an isomorphism $H^1(\mathbb{RP}^2; \mathbb{F}_2) \cong \underline{\mathrm{Hom}}\big(H^1(\mathbb{RP}^2; \mathbb{F}_2), \mathbb{F}_2\big)$ via the currying isomorphism. If the cup product in this composition is the zero map, then the composition is the zero map, which the currying isomorphism takes to, you guessed it, the zero map—which is definitely not an isomorphism. So $\smile : H^1(\mathbb{RP}^2;\mathbb{F}_2) \otimes_{\mathbb{F}_2} H^1(\mathbb{RP}^2;\mathbb{F}_2) \to H^2(\mathbb{RP}^2;\mathbb{F})$ cannot be the zero map. Since $H^1(\mathbb{RP}^2;\mathbb{F}_2) \otimes_{\mathbb{F}_2} H^1(\mathbb{RP}^2;\mathbb{F}_2)$ has only two elements, 0 and $x \otimes x$, this means that $x \smile x \neq 0$. The only other option is that $x \smile x = y$.

And there we have it: The multiplicative structure of $H^*(\mathbb{RP}^2;\mathbb{F}_2)$ is given by $x \smile x = y$ and $x \cup y = y \cup y = 0$. In other words, $H^*(\mathbb{RP}^2;\mathbb{F}_2) \cong \mathbb{F}_2[x]/(x^3)$, where x has degree 1. ♢

In Example 7.22, we computed the cohomology ring of $S^1 \vee S^2$. Looking back, we can see that $H^q(\mathbb{RP}^2;\mathbb{F}_2) \cong H^q(S^1 \vee S^2;\mathbb{F}_2)$ for every $q \in \mathbb{Z}$, so cohomology groups with $\mathbb{F}_2$ coefficients cannot distinguish these two spaces. The ring structure, on the other hand, can. The nonzero degree-1 element in $H^*(S^1 \vee S^2;\mathbb{F}_2)$ squares to 0, while the nonzero degree-1 element in $H^*(\mathbb{RP}^2;\mathbb{F}_2)$ does not. (Of course $S^1 \vee S^2$ isn't even a manifold because of the wedge point.)

Example 7.32. You can calculate that

$$H^*(\mathbb{RP}^3;\mathbb{F}_2) \cong \begin{cases} \mathbb{F}_2\{\mathbf{1}\} & \text{in degree 0} \\ \mathbb{F}_2\{x\} & \text{in degree 1} \\ \mathbb{F}_2\{y\} & \text{in degree 2} \\ \mathbb{F}_2\{z\} & \text{in degree 3} \\ 0 & \text{all other degrees.} \end{cases}$$

The embedding $\mathbb{RP}^2 \hookrightarrow \mathbb{RP}^3$ can be chosen to be a cellular map of CW-complexes, so it induces a surjective homomorphism of graded rings $H^*(\mathbb{RP}^3;\mathbb{F}^2) \to H^*(\mathbb{RP}^2;\mathbb{F}_2)$ which sends $\mathbf{1}$, x, and y to $\mathbf{1}$, x, and y, respectively, and z to 0 (in degree 3). A map of this form is an isomorphism in degrees 0, 1, and 2, so this tells us that $x \smile x = y$ in $H^*(\mathbb{RP}^3;\mathbb{F}_2)$. (Alternatively, you can go back to the original way we calculated the cup product as

$$(f \smile g)(\sigma) = (-1)^{pq} f(\sigma|_p) g(\sigma|^q)$$

and note that the equivalence classes in $H^q(\mathbb{RP}^2;\mathbb{F}_2)$ and $H^1(\mathbb{RP}^3;\mathbb{F}_2)$ are the same for $q = 0, 1, 2$ for the same reason.) We also know that $y \smile y = x \smile z = 0$ (because of degree), so we need to calculate $x \smile y$. Just apply the same reasoning as in the last example: Since $1 + 2 = 3$ and $\mathbb{RP}^3$ is a 3-mainfold, Poincaré duality tells us that $x \smile y \neq 0$; therefore $x \smile y = z$.

So $x^2 = y$ and $x^3 = z$, which means that $H^*(\mathbb{RP}^3;\mathbb{F}_2) \cong \mathbb{F}_2[x]/(x^4)$. You can iterate this reasoning to prove that $H^*(\mathbb{RP}^n;\mathbb{F}_2) \cong \mathbb{F}_2[x]/(x^{n+1})$ for every $n \in \mathbb{N}$. ♢

This calculation of the cohomology rings of $\mathbb{RP}^n$ can be used to prove the Borsuk-Ulam theorem, which has applications across topology and combinatorics.

Theorem 7.33 (Borsuk-Ulam). *For every continuous function $f\colon S^n \to \mathbb{R}^n$, there is a point $x \in S^n$ for which $f(x) = f(-x)$.*

An alternative (but equivalent) statement is that for any *odd* continuous function $f\colon S^n \to \mathbb{R}^n$ (meaning that $f(-x) = -f(x)$), there is a point $x \in S^n$ for which $f(x) = 0$. For a proof of the Borsuk-Ulam theorem, see Miller's notes.