MA3H6: Algebraic Topology

Martin Gallauer

January 12, 2024

Abstract

These are lecture notes for the course MA3H6 (Algebraic Topology) taught at the University of Warwick. The topics for this course are selected from Chapter 2 of Hatcher's "Algebraic Topology". Much of this selection and indeed the content of this course are due to the lecturers in previous years, including most recently Chris Lazda and John Greenlees. Many pictures used are also due to Chris Lazda.

Please send any corrections to martin.gallauer@warwick.ac.uk or post them on the forum for this course.

Contents

1 Introduction

1.1 Goals

Commentary 1.1. In very rough terms, Algebraic Topology is a set of tools for *translating questions in Topology to questions in Algebra*:

$$\text{Topology} \rightsquigarrow \text{Algebra}$$

In slightly more precise terms, it constructs algebraic invariants of topological spaces,

$$X \longmapsto A(X).$$

Here, "algebraic" refers to the fact that $A(X)$ is a group, a vector space, etc. And "invariant" indicates that if two spaces are homeomorphic (or just homotopy equivalent), then the associated algebraic objects are isomorphic:

$$X \simeq Y \qquad \Rightarrow \qquad A(X) \cong A(Y) \tag{1.2}$$

Example 1.3. Given a topological space X, let $\mathbb{Z} \cdot \pi_0(X)$ denote the free abelian group on the set of path-connected components of X. Of course, if $f : X \xrightarrow{\approx} Y$ is a homeomorphism then it induces a bijection $\pi_0(X) \xrightarrow{\sim} \pi_0(Y)$ and hence an isomorphism $\mathbb{Z} \cdot \pi_0(X) \xrightarrow{\cong} \mathbb{Z} \cdot \pi_0(Y)$. The same is true if f is assumed to be a homotopy equivalence only.

We will encounter this algebraic invariant later on (Corollary 3.15) and identify it with the zeroth homology group $\mathrm{H}_0(-)$.

Example 1.4. You will have encountered the fundamental group $\pi_1(-)$ of a topological space before. This is an important algebraic invariant although you will remember that it takes as input not a topological space but rather a *pointed* topological space.

Recall that $\pi_1(X, x)$ is the set of homotopy classes of pointed maps $(S^1, *) \to (X, x)$, that is, of loops in X based at x. The group structure arises from concatenation of loops.

Example 1.5. The previous example admits a straightforward generalization. Given a pointed space (X, x) and $n \geq 1$, we denote by $\pi_n(X, x)$ the set of homotopy classes of pointed maps $(S^n, *) \to (X, x)$. There is a group structure on $\pi_n(X, x)$ generalizing the one of the fundamental group. The resulting groups are called the *homotopy groups of* (X, x).

For example,

$$\pi_n(S^l) \cong \begin{cases} \mathbb{Z} & l = n \\ 0 & l > n \end{cases} \tag{1.6}$$

Commentary 1.7. Algebraic invariants may be used to tell topological spaces apart. By simple contraposition of (1.2), if $A(X) \not\cong A(Y)$ then $X \not\simeq Y$. For this to be a useful strategy, we would like that:

1. it is 'easy' to compute $A(-)$ and to tell such algebraic objects apart;
2. the algebraic invariant is 'fine' enough in that $A(X) \not\cong A(Y)$ actually happens 'often'.

Let us see how the algebraic invariants considered in Examples 1.3 to 1.5 stack up against these desiderata.

Example 1.8. The invariant $\mathbb{Z} \cdot \pi_0(-)$ performs well on the first but not so much on the second. For example, it does not distinguish between any connected manifolds.

Example 1.9. We already mentioned ("pointed out") the issue that the fundamental group—and in fact, all homotopy groups—depends on a base point. Moreover, it is in general a non-abelian group, which can make it quite tricky to tell whether or not two fundamental groups are isomorphic. Similarly, the van Kampen theorem, while it promises a powerful tool to compute $\pi_1(-)$, involves amalgated sums and is algebraically a bit awkward to put to use.

On the positive side, the fundamental group is a complete invariant for compact surfaces. On the negative side, it is not a fine invariant for higher dimensional spaces. In fact, for a CW-complex X, $\pi_1(X, x) \cong \pi_1(X^2, x)$ depends on the 2-skeleton $X^2 \subseteq X$ only. For illustration:

- $\pi_1([0,1]) = *, \qquad \pi_1(S^1) \cong \mathbb{Z}, \qquad \pi_1(S^2) = 0;$
- π_1 can show that $\mathbb{R}^2$ and $\mathbb{R}^3$ are not homeomorphic: indeed, $\pi_1(\mathbb{R}^2 \backslash \{0\}) \cong \pi_1(S^1) \cong \mathbb{Z}$ while $\pi_1(\mathbb{R}^3 \backslash \{0\}) \cong \pi_1(S^2) = 0$;

- however, π_1 cannot tell $\mathbb{R}^3$ and $\mathbb{R}^4$ apart in this manner.

Example 1.10. As seen in Example 1.5, the higher homotopy groups combined remove this shortcoming.[1] They are powerful invariants, able to tell apart many topological spaces of interest. Unfortunately, however, they are very difficult to compute. For example, it is not an accident that (1.6) misses many cases (namely, $l < n$): Most of these are unknown!

Remark 1.11. It seems from these examples that we cannot hope to simultaneously satisfy both desiderata of Commentary 1.7. Either an invariant is easy to employ but not very powerful, or it is able to tell many spaces apart but difficult to actually apply. We therefore see something at play that is sometimes called the law of conservation of difficulty. This isn't a law of nature but rather a rule of thumb: in order to establish something difficult, some difficult mathematics has to be done somewhere in the course of the proof.

Commentary 1.12. In this course, we will introduce and study another infinite sequence of algebraic invariants. These are called *homology groups* and are denoted by $\mathrm{H}_n(-)$, for $n \geq 0$. Along the spectrum 'easy to employ'—'powerful' they lie somewhere in the middle.[2] For example, as we will see, they are abelian groups and easier to compute than the homotopy groups. (The homology groups of spheres are all known, and in fact we will compute them, see Corollary 4.13.) On the other hand, they contain typically less information than those.

Commentary 1.13. The computability of homology groups comes at a price, however. Even defining them is involved and will take us a week or so. And establishing some of their fundamental properties will require substantial effort.

1.2 The idea of homology

Example 1.14. Consider the following CW complex X:

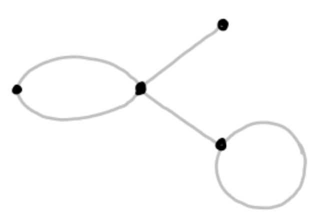

We may view the characteristic maps $\Phi_\alpha : [-1, 1] = D^1 \to X$ as defining paths ℓ_α in X from $\Phi_\alpha(-1)$ to $\Phi_\alpha(1)$.

[1]The homotopy groups π_n for $n \geq 2$ are also abelian.

[2]Arguably, in the grand scheme of things, they are situated more towards the left end of the spectrum.

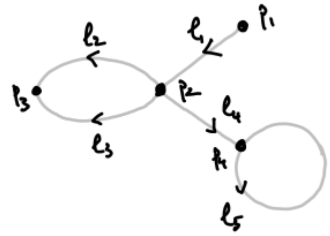

Then every path in X, say starting and ending at p_4, is homotopic to a concatenation of these basic paths ℓ_α or their inverses. Examples:

- $(\ell_5)^{-2}$, the path starting at p_4, running twice around the circle in clockwise direction, and ending at p_4
- $(\ell_4)^{-1}\ell_2(\ell_3)^{-1}\ell_4$, the path running from p_4 to p_2, around the ellipse counterclockwise and returning to p_4

Commentary 1.15. In defining the fundamental group $\pi_1(X, p_4)$ we take into account the order in which these basic paths are concatenated. As mentioned above (Commentary 1.12), homology groups will be abelian so the following definition is a first step in the direction of homology.

Definition 1.16. A 1-*chain* in X is a formal linear combination of 1-cells

$$n_1\ell_1 + n_2\ell_2 + n_3\ell_3 + n_4\ell_4 + n_5\ell_5$$

with $n_\alpha \in \mathbb{Z}$. In other words, the group of 1-chains in X is the free abelian group $\oplus_{\alpha=1}^{5}\mathbb{Z}\cdot\ell_\alpha \cong \mathbb{Z}^5$ on the basis $\{\ell_\alpha\}$.

Commentary 1.17. Some 1-chains in X may be thought of as paths. For example, $-2\cdot\ell_5$ may be thought of as the first path considered in Example 1.14. Others, like $\ell_2+\ell_5$, don't admit such an interpretation. In any case we should remember that the order of 'concatenation' is immaterial as far as 1-chains are concerned.

For a more accurate description of the relation between paths and 1-chains in X see Commentary 1.25.

Example 1.18. The two loops $\ell_2\ell_3^{-1}$ (based at p_2) and $\ell_3^{-1}\ell_2$ (based at p_3) are not distinguished anymore when viewed as 1-chains. Indeed, the two linear combinations

$$\ell_2 - \ell_3 \text{ and } -\ell_3 + \ell_2 \tag{1.19}$$

are *equal*. So, abelianization has the consequence of 'forgetting about base-points'.

Commentary 1.20. We may now ask when a 1-chain represents a loop in X. (Such 1-chains will be called *1-cycles* below.) In analogy with paths, the answer should be: if it enters and exits every 0-cell the same number of times. We can formalize this as follows.

Definition 1.21. A 0-*chain* in X is a formal linear combination of 0-cells:

$$m_1p_1 + m_2p_2 + m_3p_3 + m_4p_4$$

with $m_\beta \in \mathbb{Z}$. In other words, the group of 0-chains is the free abelian group $\oplus_{\beta=1}^4 \mathbb{Z} \cdot p_\beta \cong \mathbb{Z}^4$ on the basis $\{p_\beta\}$.

Definition 1.22. We define the *boundary operator* ∂_1 that takes a basic 1-chain ℓ_α to the formal difference between start and end point of ℓ_α, viewed as a 0-chain in X. In other words, it is the group homomorphism given on generators by:

$$\begin{aligned} \partial_1 : \ \oplus_\alpha \mathbb{Z} \cdot \ell_\alpha &\longrightarrow \oplus_\beta \mathbb{Z} \cdot p_\beta \\ \ell_1 &\longmapsto p_2 - p_1 \\ \ell_2 &\longmapsto p_3 - p_2 \\ \ell_3 &\longmapsto p_3 - p_2 \\ \ell_4 &\longmapsto p_4 - p_2 \\ \ell_5 &\longmapsto p_4 - p_4 = 0 \end{aligned}$$

A 1-*cycle* in X is a 1-chain in the kernel of ∂_1.

Example 1.23.

- For any $n \in \mathbb{Z}$, the 1-chain $n \cdot \ell_5$ is a 1-cycle since $\partial_1(n \cdot \ell_5) = n \cdot \partial_1(\ell_5) = n \cdot 0 = 0$.
- The 1-chain $\ell_2 - \ell_3 = -\ell_3 + \ell_2$ of (1.19) is a 1-cycle.

Remark 1.24. With the given basis elements for 1- and 0-chains, the boundary operator $\partial_1 : \mathbb{Z}^5 \to \mathbb{Z}^4$ is represented by the matrix

$$\begin{pmatrix} -1 & 0 & 0 & 0 & 0 \\ 1 & -1 & -1 & -1 & 0 \\ 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \end{pmatrix}$$

hence $\ker(\partial_1) = \mathbb{Z} \cdot (\ell_2 - \ell_3) \oplus \mathbb{Z} \cdot \ell_5 \cong \mathbb{Z}^2$. This is what we will later call the first homology group of X and denote by $\mathrm{H}_1(X)$.

Commentary 1.25. The fundamental group of X is (for example, by the van Kampen Theorem) the free group on the two generators $\ell_2(\ell_3)^{-1}$ and ℓ_5. We therefore see in this example that the morphism

$$\begin{aligned} \mathbb{Z} * \mathbb{Z} \cong \pi_1(X) &\longrightarrow \mathrm{H}_1(X) \cong \mathbb{Z}^2 \\ \ell_2(\ell_3)^{-1} &\longmapsto \ell_2 - \ell_3 \\ \ell_5 &\longmapsto \ell_5 \end{aligned}$$

identifies the first homology group with the *abelianization* of the fundamental group. We will see later (Theorem 3.28) that this is a general phenomenon.

Example 1.26. Let Y be the CW complex obtained from X by filling in the ellipse with a new 2-cell:

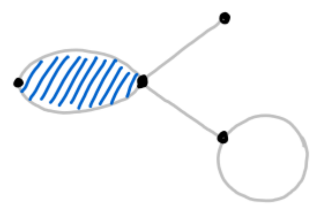

More precisely, we choose the attaching map

$$\phi_1 : S^1 \to X \tag{1.27}$$

described by the loop $\ell_2(\ell_3)^{-1}$ based at p_2. At the level of π_1, attaching this new 2-cell d_1 has the effect of killing the loop $\ell_2(\ell_3)^{-1}$.

The following definition will not come as a surprise.

Definition 1.28. The group of 2-*chains* in Y is the free abelian group on d_1. The boundary operator ∂_2 on 2-chains is given by:

$$\begin{aligned} \partial_2 : \ \mathbb{Z} \cdot d_1 &\longrightarrow \oplus_\alpha \mathbb{Z} \cdot \ell_\alpha \\ d_1 &\longmapsto \ell_2 - \ell_3 \end{aligned}$$

Commentary 1.29. Here the image of d_1 should be thought of as the 1-cycle described by the loop used to define the attaching map ϕ.

Remark 1.30. If we define the first homology group of Y as the quotient

$$\mathrm{H}_1(Y) = \frac{\ker(\partial_1 : \{1\text{-chains}\} \to \{0\text{-chains}\})}{\mathrm{img}(\partial_2 : \{2\text{-chains}\} \to \{1\text{-chains}\})}$$

then we find that $\mathrm{H}_1(Y) \cong \mathbb{Z}$ is generated by the 1-cycle ℓ_5.

This is another instance of the phenomenon mentioned in Commentary 1.25. Indeed, $\pi_1(Y) \cong \pi_1(S^1) \cong \mathbb{Z}$ is already abelian hence equal to its abelianization.

Example 1.31. Suppose we change the attaching map to $\phi' : S^1 \to X$ described by the loop $\ell_2(\ell_3)^{-1}\ell_2(\ell_3)^{-1}$ to get another CW complex Y'. Then the boundary map sends the unique 2-cell d_1' to $2(\ell_2 - \ell_3)$ so that

$$\mathrm{H}_1(Y') \cong \mathbb{Z} \oplus \mathbb{Z}/2\mathbb{Z}.$$

Example 1.32. We consider one last modification. Let Z be the CW complex obtained from Y by adding a second 2-cell d_2 via the same attaching map (1.27).

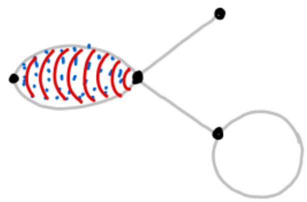

Note that attaching this new 2-cell does not have any effect on the fundamental group so we expect that $\mathrm{H}_1(Z) \cong \mathrm{H}_1(Y) \cong \mathbb{Z}$.

The group of 2-chains in Z is the free abelian group $\mathbb{Z} \cdot d_1 \oplus \mathbb{Z} \cdot d_2$ on the basis $\{d_1, d_2\}$ and the boundary operator takes both of these generators to the 1-cycle $\ell_2 - \ell_3$. This means that $\mathrm{H}_1(Z) = \frac{\ker(\partial_1)}{\mathrm{img}(\partial_2)}$ is indeed the same as $\mathrm{H}_1(Y) \cong \mathbb{Z}$. On the other hand, the kernel of the boundary operator $\partial_2 \colon \{2\text{-chains}\} \to \{1\text{-chains}\}$ is non-trivial: it is infinite cyclic generated by $d_1 - d_2$. This is what we will eventually call the second homology group $\mathrm{H}_2(Z) = \ker(\partial_2) \cong \mathbb{Z}$.

Commentary 1.33. What is the upshot of these examples?

- You will now probably be in a position to guess how the definition of homology groups of any CW complex X will look like. Namely, one starts by defining the group of n-chains in X as the free abelian group on the n-cells e^n_α in X. The boundary operator $\partial_n \colon \{n\text{-chains}\} \to \{(n-1)\text{-chains}\}$ takes a generator e^n_α to an $(n-1)$-cycle described in terms of the attaching map for e^n_α. (At this point of the discussion it is admittedly not clear at all yet how this description looks like in general.) The nth homology group $\mathrm{H}_n(X)$ is then the quotient $\ker(\partial_n)/\mathrm{img}(\partial_{n+1})$.

  More precisely, this is what we will call the *cellular homology* of X in Section 8.3.
- The examples also suggest one way to think about the homology of a space X. Namely, roughly, the nth homology group $\mathrm{H}_n(X)$ measures 'how many' n-dimensional holes there are in X.

Commentary 1.34. The idea of using algebraic structures to distinguish non-homeomorphic topological spaces was first systematically developed by Henri Poincaré in his series of papers on "Analysis Situs" (1899–1904). For this reason these papers are commonly taken as the birth of Algebraic Topology. And quite an impressive birth they were! Poincaré introduced not only the fundamental group and a form of simplicial homology (see section 2 below), but also formulated what is now known as *Poincaré duality* (see the fourth-year module MA4J7) and the *Poincaré Conjecture* (now a celebrated theorem), among many other things.

2 Simplicial homology

As mentioned in the introduction, the definition of the homology groups of a topological space is rather involved. For this reason we start in this section with a simplified version, called *simplicial homology*, that takes as input not a topological space but rather a Δ-*complex*. In the next section 3 we go on to study the generalized version applicable to all topological spaces.

2.1 Δ-complexes

Δ-complexes are topological spaces built out of simplices: points, lines, triangles, tetrahedrons, etc. To make this precise, let us introduce:

Definition 2.1. The *(standard) n-simplex* $\Delta^n \subseteq \mathbb{R}^{n+1}$ is the subspace

$$\Delta^n := \left\{ (x_0, \dots, x_n) \mid \sum_{i=0}^{n} x_i = 1,\ x_i \geq 0 \right\}.$$

More generally, any homeomorphic space will also be called an n-simplex. Its vertices will often be denoted $v_0, v_1, \dots, v_n$. These correspond to the points $(1, 0, \dots, 0)$, $(0, 1, 0, \dots, 0)$, …, $(0, \dots, 0, 1)$ in the standard n-simplex.

Example 2.2.

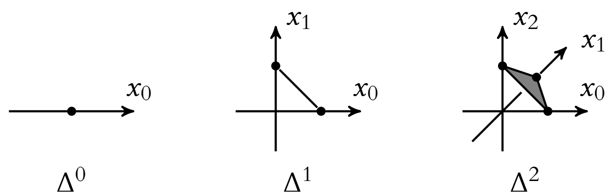

Definition 2.3. The jth *face* of the n-simplex is the subspace $\partial_j \Delta^n \subseteq \Delta^n$ of points whose jth coordinate vanishes. That is,

$$\partial_j \Delta^n = \left\{ (x_0, \dots, x_n) \mid \sum_{i=1}^{n} x_i = 1, x_i \geq 0\, \forall i, x_j = 0 \right\}.$$

Note that $\partial_j \Delta^n$ is canonically an $(n-1)$-simplex with vertices $v_0, \dots, \widehat{v_j}, \dots, v_n$.

Example 2.4. Here, the three faces of the standard 2-simplex are shown in color:

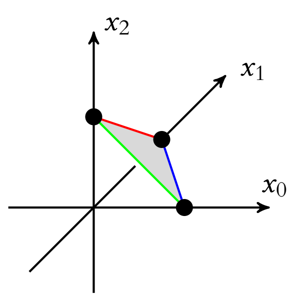

Note that the jth face is the one *opposite* the jth vertex.

Definition 2.5. The *boundary* of Δ^n is the union of its $(n+1)$ faces, and will be denoted by $\partial \Delta^n$.

Example 2.6. Since the unique face $\partial_0 \Delta^0$ is the empty set we have $\partial \Delta^0 = \emptyset$.

Exercise 2.7. Show that the n-simplex Δ^n is homeomorphic to the n-ball

$$D^n = \left\{ (y_1, \dots, y_n) \mid \sum_{i=1}^{n} y_i^2 \leq 1 \right\} \subset \mathbb{R}^n$$

and its boundary $\partial\Delta^n$ homeomorphic to the sphere S^{n-1}.

Definition 2.8. A Δ-*complex* is a topological space obtained inductively as follows:

1. Start with a (discrete) collection of 0-simplices, that is, points. This is the 0-skeleton X^0.
2. Inductively, the n-skeleton X^n is obtained from X^{n-1} by attaching n-simplices Δ^n_α whereby each face $\partial_i\Delta^n_\alpha$ gets identified with an $(n-1)$-simplex Δ^{n-1}_β in X^{n-1}.
3. If $X = X^k$ and k is minimal with this property then X is of dimension k. More generally we can have $X = \cup_n X^n$ in which case a subspace $U \subseteq X$ is open iff $U \cap X^n \subseteq X^n$ is open for all n.

Example 2.9. Start with a single point $X^0 = \Delta^0$. Attach now a single 1-simplex Δ^1, in the only possible way. Namely, both boundary points are identified with the 0-skeleton:

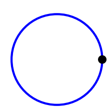

In other words, we get a circle.

Example 2.10. Now we start with two points $X^0 = \{p, q\}$. We then attach two 1-simplices a, b as follows:

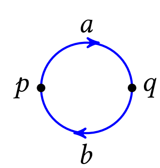

The arrows indicate that the 'line a runs from p to q' (and similarly for b). More precisely, this means that the 0th face of a (the end point) gets identified with q, and the 1st face (the start point) with p.

Note that we get again the circle. This is something to keep in mind: Δ-complexes can typically be built out of simplices in different ways.

Example 2.11. Moving to 2-dimensional spaces, the filled-in square is a Δ-complex, as depicted:

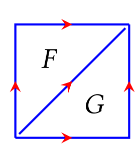

It has four 0–simplices, five 1–simplices and two 2–simplices F and G. Note how the arrows specify an ordering of the vertices of each 2–simplex, and therefore a homeomorphism with Δ^2. In particular, there is no ambiguity about which face is which. (For example, check that the first face of F is the diagonal. Or that the 0th face of G is the right vertical line.)

Example 2.12. The torus $\mathbb{T} = S^1 \times S^1$ is obtained by identifying opposite edges of a square:

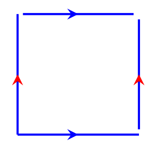

It is a Δ–complex, as follows:

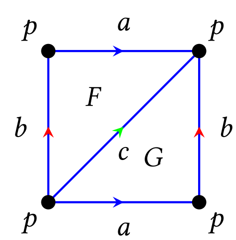

Here there is a single 0–simplex p, three 1–simplices a, b and c, and two 2–simplices F and G.

Exercise 2.13. See the first exercise sheet for further examples.

Remark 2.14. A Δ–complex is essentially just a combinatorial datum. Namely, it is determined (up to homeomorphism) by the sets of n–simplices S_n, $n \geq 0$, together with the attaching rules. These are 'face' maps $d_i^n \colon S_n \to S_{n-1}$ (for $0 \leq i \leq n$) specifying that $\partial_i \Delta_\alpha^n$ gets identified with $\Delta^{n-1}_{d_i^n(\alpha)}$.

These maps are not arbitrary but satisfy the relation (whenever $i < j$)

$$d_i^{n-1} \circ d_j^n = d_{j-1}^{n-1} \circ d_i^n, \tag{2.15}$$

as you can easily check. Such a combinatorial datum $S = (S_\bullet, d_\bullet^\bullet)$ is called a Δ*-set* or *semi-simplicial set*.

Given a Δ–set S we denote the associated Δ–complex by $|S|$. (It is sometimes called its geometric realization.) To be pedantic (but one rarely is), a Δ–complex is really a topological space X together with a homeomorphism $X \approx |S|$ for some Δ–set S. The latter could then be called a Δ*-complex structure* on X. As we saw in Examples 2.9 and 2.10, a topological space can admit distinct Δ–complex structures.

Example 2.16. In the case of the torus (Example 2.12) we found

$$S_0 = \{p\}, \qquad S_1 = \{a, b, c\}, \qquad S_2 = \{F, G\}$$

with face maps that can be read off the picture. For example,

$$d_0^2(F) = a, \quad d_0^2(G) = b; \qquad d_i^1(a) = d_i^1(b) = d_i^1(c) = p \text{ for } i = 0, 1.$$

Commentary 2.17 (unimportant)**.** You might have encountered simplicial complexes or CW-complexes before. Every simplicial complex is automatically a Δ-complex, and every Δ-complex is automatically a CW-complex. So, the relation between these is as follows:

$$\{\text{simplicial complexes}\} \subseteq \{\Delta\text{-complexes}\} \subseteq \{\text{CW-complexes}\}$$

It turns out (we will not prove nor use that) that a topological space admits the structure of a Δ-complex iff it admits the structure of a simplicial complex (albeit with different simplices in general). So, in some sense, the first inclusion above is an equality.[3] However, this is not true for the second inclusion.

2.2 Homology

Commentary 2.18. In Section 1.2 we defined homology in the context of some simple CW complexes. A key ingredient was the definition of certain boundary operators. It turns out that moving from cells to simplices and from CW complexes to Δ-complexes makes things even more transparent. Let us now see how.

- Start with the standard 1-simplex:

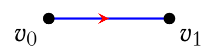

  Taking a cue from Definition 1.22 we view its 'oriented' boundary as the formal difference between head and tail:
$$v_1 - v_0.\,^{4}$$

- Similarly when passing to 2-simplices,

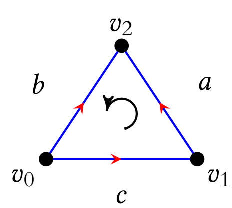

  we define its 'oriented' boundary as
$$a - b + c.\,^{5}$$

[3]You might now wonder why we don't use simplicial complexes instead. The reason is that typically many more simplices are required than for Δ-complexes, making the homology computations below more involved.

[4]The difference $v_0 - v_1$ would work equally well.

[5]Again, we have made a choice here. Orienting the triangle in a clockwise fashion works equally and would have given the negative of this expression: $-a + b - c$

- The general formula for an n-simplex s that emerges is then

$$\partial_n(s) = \sum_{i=0}^{n}(-1)^i d_i^n(s), \tag{2.19}$$

  the alternating sum of its faces.

In fact, this whole section is nothing but an elaboration of section 1.2 in the context of Δ-complexes. With the little input given now, you should be able to define the (simplicial) homology groups of a Δ-complex. Try it before reading on!

Definition 2.20. Let S be a Δ-set.

1. The group of *n-chains* in S is the free abelian group on S_n. It is denoted by $\Delta_n(S)$.
2. The boundary operator $\partial_n \colon \Delta_n(S) \to \Delta_{n-1}(S)$ is the homomorphism given on the generators $s \in S_n$ by the formula in (2.19). (By convention, $\Delta_{-1}(S) = 0$ and ∂_0 is the zero map.)

Example 2.21. Recall from Example 2.9 the 'minimal' Δ-complex structure on S^1. The only possibly interesting boundary operator is $\partial_1 : \mathbb{Z}\ell = \Delta_1(S) \to \Delta_0(S) = \mathbb{Z}p$. It takes the generator ℓ to $\partial_1(\ell) = d_0(\ell) - d_1(\ell) = p - p = 0$ and hence is the zero map.

Example 2.22. Consider again the Δ-complex structure on the torus from Example 2.12. For instance, we find that

$$\partial_2(F - G) = \partial_2(F) - \partial_2(G) = (b - c + a) - (a - c + b) = 0$$

so we have found a non-trivial 2-chain in the kernel of ∂_2.

Definition 2.23. Let S be a Δ-set. We define the following subgroups of $\Delta_n(S)$, for $n \geq 0$:

- the group of *n-cycles* $Z_n(S) := \ker(\partial_n)$
- the group of *n-boundaries* $B_n(S) := \operatorname{img}(\partial_{n+1})$
- the *nth simplicial homology group* is the quotient:

$$\mathrm{H}_n(S) := \frac{Z_n(S)}{B_n(S)}$$

  Elements of this group are called *homology classes*.

Commentary 2.24. For the last expression to make sense we need to know that $B_n(S) \subseteq Z_n(S)$. We will indeed prove this below in Lemma 2.34, but we will also see it the next couple of examples directly. (In section 1.2 we also defined the homology as the quotient of cycles by boundaries. In each case, it could be easily verified that this made sense, to wit, that the n-boundaries are all n-cycles.)

Example 2.25. Let S be the Δ-set from Example 2.9. We saw in Example 2.21 that the boundary operator $\partial_1 = 0$ vanishes. It follows that:

$$Z_1(S) = Z_0(S) = \mathbb{Z}, \qquad B_n(S) = 0 \quad \forall n,$$

and therefore $\mathrm{H}_0(S) = \mathrm{H}_1(S) = \mathbb{Z}$ and all other homology groups 0.

Example 2.26. Let T be the Δ-set from Example 2.10. We may represent the situation as follows:

$$\cdots \to 0 \to \mathbb{Z}a \oplus \mathbb{Z}b \xrightarrow[\partial_1]{\begin{pmatrix} -1 & 1 \\ 1 & -1 \end{pmatrix}} \mathbb{Z}p \oplus \mathbb{Z}q \to 0$$

with all other n-chains trivial. We deduce that

$$\begin{aligned} B_0(T) = \mathbb{Z}(p-q) \quad &\subset \quad \mathbb{Z}p \oplus \mathbb{Z}q = Z_0(T) \\ B_1(T) = 0 \quad &\subset \quad \mathbb{Z}(a+b) = Z_1(T) \end{aligned}$$

and therefore $H_0(T) = H_1(T) = \mathbb{Z}$ and all other homology groups 0.

Remark 2.27. We just saw that the two Δ-sets S and T of Examples 2.25 and 2.26 have the same homology groups. This is not a coincidence (cf. also exercise 1.5 on sheet 1). We showed in Examples 2.9 and 2.10 that their geometric realization is the same, namely the 1-sphere. And it turns out, although it isn't clear at all at this point, that the homology is an invariant of the geometric realization. We will establish this in Corollary 8.5 and thereby justify the following definition.

Definition 2.28. If X is a topological space with a Δ-complex structure $|S| \approx X$ we define its *nth simplicial homology group* to be

$$\mathrm{H}_n^{\Delta}(X) := \mathrm{H}_n(S).$$

Commentary 2.29. Our informal discussion of homology, Commentary 1.33, lets us think about these groups as follows. An n-cycle is an n-chain that could be the boundary of an $(n+1)$-chain. Therefore, a nonzero homology class detects an n-dimensional hole: something that could be but isn't a boundary.

Example 2.30. So, we have seen that $\mathrm{H}_n^{\Delta}(S^1) = \begin{cases} \mathbb{Z} & : n = 0, 1, \\ 0 & : \text{else.} \end{cases}$

Example 2.31. Let us compute the simplicial homology of the torus, see Example 2.12. We may represent the situation as follows:

$$\cdots \to 0 \to \mathbb{Z}F \oplus \mathbb{Z}G \xrightarrow[\partial_2]{\begin{pmatrix} 1 & 1 \\ 1 & 1 \\ -1 & -1 \end{pmatrix}} \mathbb{Z}a \oplus \mathbb{Z}b \oplus \mathbb{Z}c \xrightarrow[\partial_1]{(0\ 0\ 0)} \mathbb{Z}p \to 0$$

We see that the 2-cycle from Example 2.22 generates the 2-cycles and $\mathrm{H}_2^{\Delta}(\mathbb{T}) = \mathbb{Z}$. Of course, $\mathrm{H}_0^{\Delta}(\mathbb{T}) = \mathbb{Z}$. And in degree 1 we have:

$$B_1 = \mathbb{Z}(a+b-c) \quad \subset \quad \mathbb{Z}a \oplus \mathbb{Z}b \oplus \mathbb{Z}c = Z_1$$

so that $\mathrm{H}_1^\Delta(\mathbb{T}) = \mathbb{Z} \oplus \mathbb{Z}$.

Exercise 2.32. Compute the simplicial homology of the square, see Example 2.11. In fact, before you start, ask yourself what the simplicial homology *should be* and then verify that.

Commentary 2.33. The first problem sheet asks you to compute more examples of simplicial homology groups. At this point, it's probably good to observe that once you've written down the groups of chains and the boundary operators between them, computing the simplicial homology is an entirely mechanical process. If you don't have a preferred method already I recommend the following one:

1. First, determine the Smith Normal Form for each boundary operator in matrix form. (See the document with Preliminaries on the moodle page.)
2. Apply the following result: Given two matrices

$$\mathbb{Z}^l \xrightarrow{A} \mathbb{Z}^m \xrightarrow{B} \mathbb{Z}^n$$

with $BA = 0$, we have

$$\ker(B)/\mathrm{img}(A) \cong \left(\oplus_{i=1}^r \mathbb{Z}/a_i\right) \oplus \mathbb{Z}^{m-r-s}$$

where

- a_i are the invariant factors of A,
- $r = \mathrm{rk}(A)$, $s = \mathrm{rk}(B)$.

(This is Exercise 1.4 on sheet 1.)

We still need to justify Definition 2.23.

Lemma 2.34. *Let S be a Δ-set. Then $\partial_n \circ \partial_{n+1} = 0$. Equivalently, $B_n(S) \subseteq Z_n(S)$.*

Of course, we have set up things precisely so that this holds. For example, here is what happens when you apply twice the boundary operator to the standard 2-simplex $\Delta^2 = [v_0, v_1, v_2]$.

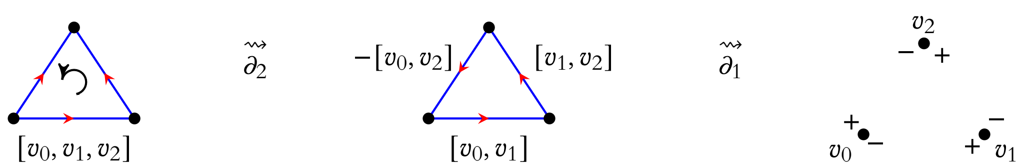

In the resulting formal sum every 0-simplex appears once with a plus and once with a minus sign, the two canceling each other out. This is what happens in general:

Proof. It is clear from the definitions of cycles and boundaries that the two statements are equivalent. So we only prove the first. For this, let $s \in S_{n+1}$. Then

$$\partial_n \partial_{n+1}(s) = \partial_n \left(\sum_{j=0}^{n+1} (-1)^j d_j^{n+1}(s) \right) = \sum_{i=0}^{n} \sum_{j=0}^{n+1} (-1)^{i+j} d_i^n d_j^{n+1}(s).$$

Fix a pair $0 \leq i < j \leq n+1$. By (2.15), the two summands $A := (-1)^{i+j} d_i^n d_j^{n+1}(s)$ and $B := (-1)^{i+j-1} d_{j-1}^n d_i^{n+1}$ cancel out. And every summand is of the form A or B (and not both) so that the sum vanishes.

If you like it more formally, we can do this by breaking up the sum in two:

$$\begin{aligned}
\sum_{i=0}^{n}\sum_{j=0}^{n+1}(-1)^{i+j} d_i d_j(s) &= \sum_{0\leq i<j\leq n+1} (-1)^{i+j} d_i d_j(s) \quad &+ \quad \sum_{0\leq j\leq i\leq n} (-1)^{i+j} d_i d_j(s) \\
&= \sum_{0\leq i<j\leq n+1} (-1)^{i+j} d_{j-1} d_i(s) \quad &+ \quad \sum_{0\leq j\leq i\leq n} (-1)^{i+j} d_i d_j(s) \\
&= \sum_{0\leq i\leq j\leq n} (-1)^{i+j-1} d_j d_i(s) \quad &+ \quad \sum_{0\leq j\leq i\leq n} (-1)^{i+j} d_i d_j(s) \\
&= 0
\end{aligned}$$

□

Commentary 2.35. While mathematicians had spoken of "cycles modulo boundaries" before, it seems like it was Poincaré who introduced formal sums of simplices and the boundary operator, leading to a precise definition of homology classes. However, he did not see the importance of studying the homology *groups* themselves as invariants of topological spaces. This was only noticed years later by Emmy Noether and, independently, Leopold Vietoris[6] and Walther Mayer. Mayer (1929) is also credited with introducing the term "chain complex" that we are about to discuss.

2.3 Chain complexes

Commentary 2.36. Just now we have been computing homology groups of Δ-sets in a two-step process, if you like:

$$S \quad \mapsto \quad (\Delta_\bullet(S), \partial_\bullet) \quad \mapsto \quad \mathrm{H}_\bullet(S)$$

It will be useful to discuss the second step of this process independently of the situation here since it will recur repeatedly throughout the course. We will see that the input should be viewed as a *chain complex* and the step can be described as taking the *homology* of a chain complex.

Definition 2.37. A *chain complex* $C_\bullet = (C_\bullet, \partial_\bullet)$ is a family of abelian groups C_n for $n \in \mathbb{Z}$ and maps (called *differentials*) $\partial_n : C_n \to C_{n-1}$ such that $\partial_n \circ \partial_{n+1} = 0$ for all n.

Remark 2.38. We often depict a chain complex as a sequence of abelian groups and maps between them, like so:

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

This should remind you of the pictures in, say, Examples 2.26 and 2.31. Of course, that's not a coincidence:

[6] Did you know that Vietoris was the oldest verified Austrian man ever? He lived in three different centuries and died at the age of 110.

Example 2.39. Let S be a Δ–set. We associated in Definition 2.20 abelian groups $\Delta_n(S) = \mathbb{Z}S_n$ and maps $\partial_n\colon \Delta_n(S) \to \Delta_{n-1}(S)$ that satisfy $\partial_n \circ \partial_{n+1} = 0$, by Lemma 2.34. At least we did this for $n \geq 0$. If we set $\Delta_n(S) = 0$ for $n < 0$ with zero maps between them we get a chain complex $\Delta_\bullet(S)$ that is called the *simplicial chain complex* associated with S.

Convention 2.40. In general if we only define C_n for n in some interval $[a, b]$ then it's understood that $C_n = 0$ for all $n \notin [a, b]$.

Example 2.41. Let $\iota : \mathbb{Z} \to \mathbb{R}$ be the inclusion of the integers inside the real numbers. We may view this as a chain complex placed in degrees, say, 1 and 0:

$$0 \to \mathbb{Z} \xrightarrow{\partial_1=\iota} \mathbb{R} \to 0$$

Definition 2.42. Let $C_\bullet$ be a chain complex.

- The *n-cycles* are $Z_n(C_\bullet) = \ker(\partial_n\colon C_n \to C_{n-1})$.
- The *n-boundaries* are $B_n(C_\bullet) = \operatorname{img}(\partial_{n+1}\colon C_{n+1} \to C_n)$.

Note that since $\partial_n \circ \partial_{n+1} = 0$ we have $B_n \subseteq Z_n$.

- The *nth homology group* is $\mathrm{H}_n(C_\bullet) := \frac{Z_n}{B_n}$.
- If $\mathrm{H}_n(C_\bullet) = 0$ one also says that the complex $C_\bullet$ is *exact in degree n*.

Example 2.43. For the simplicial chain complex $\Delta_\bullet(S)$ of a Δ–set S we recover the n–cycles, n–boundaries, and the nth homology group of S:

$$Z_n(\Delta_\bullet(S)) = Z_n(S), \qquad B_n(\Delta_\bullet(S)) = B_n(S), \qquad \mathrm{H}_n(\Delta_\bullet(S)) = \mathrm{H}_n(S)$$

Example 2.44. For the chain complex of Example 2.41 we get

$$\mathrm{H}_n = \begin{cases} \mathbb{R}/\mathbb{Z} & : n = 0 \\ 0 & : n \neq 0 \end{cases}$$

The group $\mathbb{R}/\mathbb{Z}$ can be identified with S^1 (viewed as a subgroup of the complex numbers with multiplication), via the map $\exp(2\pi i \cdot -)\colon \mathbb{R} \to S^1$.

Example 2.45. Let $C^\infty(\mathbb{R}^3)$ be the group of smooth functions $\mathbb{R}^3 \to \mathbb{R}$, or smooth scalar fields on $\mathbb{R}^3$. Then $V^\infty(\mathbb{R}^3) = C^\infty(\mathbb{R}^3)^3$ is the group of smooth vector fields on $\mathbb{R}^3$. The operators div, grad and curl can be viewed as homomorphisms

$$\begin{aligned} \mathsf{div}\colon V^\infty(\mathbb{R}^3) &\to C^\infty(\mathbb{R}^3), \\ \mathsf{grad}\colon C^\infty(\mathbb{R}^3) &\to V^\infty(\mathbb{R}^3), \\ \mathsf{curl}\colon V^\infty(\mathbb{R}^3) &\to V^\infty(\mathbb{R}^3). \end{aligned}$$

The facts that $\mathsf{div} \circ \mathsf{curl} = 0$ and $\mathsf{curl} \circ \mathsf{grad} = 0$ mean that we get a chain complex

$$0 \to C^\infty(\mathbb{R}^3) \xrightarrow{\mathsf{grad}} V^\infty(\mathbb{R}^3) \xrightarrow{\mathsf{curl}} V^\infty(\mathbb{R}^3) \xrightarrow{\mathsf{div}} C^\infty(\mathbb{R}^3) \to 0$$

which we think of as being concentrated in $[-3, 0]$. We have $\mathrm{H}_0 = \mathbb{R}$, the constant functions, and you might have learned in Vector Calculus that $\mathrm{H}_{-3} = \mathrm{H}_{-2} = \mathrm{H}_{-1} = 0$ (Poincaré Lemma).

Example 2.46. Let $C_\bullet$ be a chain complex with all differentials the zero maps. Then $Z_n = C_n$ and $B_n = 0$ so that $\mathrm{H}_n = C_n$.

Commentary 2.47. Whenever you define new objects it is a good practice (especially in algebraic topology) to ask: what are the morphisms between these objects? For example, soon after defining groups one also introduces group homomorphisms. After vector spaces linear transformations. After topological spaces continuous maps. And so on.

We do the same for chain complexes.

Definition 2.48. Let $(C_\bullet, \partial_\bullet)$ and $(C'_\bullet, \partial'_\bullet)$ be two chain complexes. A *chain map* $f_\bullet \colon C_\bullet \to C'_\bullet$ is a family of maps

$$f_n \colon C_n \to C'_n, \qquad \text{such that for all } n : \partial'_n \circ f_n = f_{n-1} \circ \partial_n.$$

In pictures we have 'commuting squares':

$$\begin{array}{ccccccccc} \cdots & \longrightarrow & C_{n+1} & \xrightarrow{\partial_{n+1}} & C_n & \xrightarrow{\partial_n} & C_{n-1} & \longrightarrow & \cdots \\ & & \downarrow{\scriptstyle f_{n+1}} & & \downarrow{\scriptstyle f_n} & & \downarrow{\scriptstyle f_{n-1}} & & \\ \cdots & \longrightarrow & C'_{n+1} & \xrightarrow{\partial'_{n+1}} & C'_n & \xrightarrow{\partial'_n} & C'_{n-1} & \longrightarrow & \cdots \end{array}$$

Lemma 2.49. *A chain map $f_\bullet \colon C_\bullet \to C'_\bullet$ restricts to maps*

- $f_n \colon Z_n(C_\bullet) \to Z_n(C'_\bullet)$*, as well as*
- $f_n \colon B_n(C_\bullet) \to B_n(C'_\bullet)$*,*

and hence descends to maps $f_n \colon \mathrm{H}_n(C_\bullet) \to \mathrm{H}_n(C'_\bullet)$.

Proof. See exercise sheet 1, problem 1.4. □

Definition 2.50. Let $S = (S_\bullet, d_\bullet)$ and $S' = (S'_\bullet, d'_\bullet)$ be two Δ-sets. A *map of Δ-sets* $f \colon S \to S'$ is a family of maps $f_n \colon S_n \to S'_n$ such that for all $0 \leq i \leq n$,

$$d'_i \circ f_n = f_{n-1} \circ d_i \;:\; S_n \to S'_{n-1}.$$

Lemma 2.51. *A map of Δ-sets $f_\bullet \colon S \to S'$ induces a chain map $f_\bullet \colon \Delta_\bullet(S) \to \Delta_\bullet(S')$.*

Proof. We define $f_n \colon \Delta_n(S) \to \Delta_n(S')$ on generators $s \in S_n$ by

$$\begin{aligned} \mathbb{Z}S_n &\longrightarrow \mathbb{Z}S'_n \\ s &\longmapsto f_n(s) \end{aligned}$$

We now need to verify that these maps commute with the boundary operators. For this we show:

$$\begin{aligned}
\partial'_n \circ f_n(s) &= \partial'_n(f_n(s)) \\
&= \sum_{i=0}^{n} (-1)^i d'_i(f_n(s)) && \text{definition of } \partial'_n \\
&= \sum_{i=0}^{n} (-1)^i f_{n-1}(d_i(s)) && \text{definition of map of } \Delta\text{-sets} \\
&= f_{n-1}\left(\sum_{i=0}^{n} (-1)^i d_i(s)\right) && \text{linearity of } f_{n-1} \\
&= f_{n-1} \circ \partial_n(s) && \text{definition of } \partial_n
\end{aligned}$$

□

Remark 2.52. Combining the last two lemmas we see that every map of Δ-sets induces a map in simplicial homology.

Exercise 2.53. What do you think is a/the appropriate notion of a morphism between Δ-complexes?

3 Singular homology ...

Commentary 3.1. In the previous section we defined the simplicial homology for Δ-complexes and computed it in some examples. At least two problems arise with this tool: First, topological spaces do not typically come with an obvious Δ-complex structure. In fact, some topological spaces admit *no* such structure at all. Secondly, even if a given space X admits such a structure, it might not be unique. And unfortunately we didn't prove that the simplicial homology is independent of the choice one then has to make.

In this section we are going to define a variant which avoids these difficulties (and eventually will allow us to prove the independence in the paragraph above). The idea of singular homology is to allow *all* (as always, continuous) maps $\sigma\colon \Delta^n \to X$ as simplices. Note the departure: X is not neatly built out of these simplices by identifying certain faces. For example, σ could be the constant map. To mark this departure we call such simplices *singular*, thus the name of this homology theory.

3.1 Definition

Definition 3.2. Let X be a topological space and $n \geq 0$. A *singular n-simplex* in X is a continuous map $\sigma\colon \Delta^n \to X$.

Example 3.3.

1. A singular 0-simplex $\Delta^0 = \{*\} \to X$ is nothing but a point $x \in X$. We will sometimes identify the two.
2. A singular 1-simplex $\gamma\colon \Delta^1 = [0,1] \to X$ is nothing but a path in X from $\gamma(0)$ to $\gamma(1)$.

3. If X is a Δ-complex then any n-simplex in X gives rise to a singular n-simplex in X.

Commentary 3.4. One good thing about the notion of singular simplices is that we already know how to define its 'oriented' boundary since we know it for the standard n-simplex. Namely, the 'oriented' boundary of $\sigma\colon \Delta^n \to X$ should be:

$$\sum_{i=0}^{n} (-1)^i \sigma|_{\partial_i \Delta^n}. \tag{3.5}$$

If we identify the ith face with the standard $(n-1)$-simplex, as in Definition 2.3, then this expression becomes a formal linear combination of singular $(n-1)$-simplices in X.

Definition 3.6. Let X be a topological space and $n \geq 0$.

1. The group of *singular n-chains* in X is the free abelian group on the singular n-simplices, denoted $C_n(X)$.
2. The *boundary operator* $\partial_n\colon C_n(X) \to C_{n-1}(X)$ is the homomorphism defined on a basis element $\sigma\colon \Delta^n \to X$ by the formula in (3.5).

The same proof as in Lemma 2.34 then gives:

Lemma 3.7. *Let X be a topological space. Then $\partial_n \circ \partial_{n+1}\colon C_{n+1}(X) \to C_{n-1}(X)$ is the zero map.*

In other words, $(C_\bullet(X), \partial_\bullet)$ is a chain complex and we can make the following definition.

Definition 3.8. Let X be a topological space. The *singular chain complex* of X is the complex $(C_\bullet(X), \partial_\bullet)$. The *singular homology groups* of X are

$$\mathrm{H}_n(X) := \mathrm{H}_n(C_\bullet(X)).$$

Example 3.9. Let $X = *$ be a point. For each $n \geq 0$ there is a unique singular n-simplex, the constant map $c_n\colon \Delta^n \to *$. Therefore the singular chain complex looks as follows:

$$\cdots \to \mathbb{Z} \xrightarrow{\partial_n} \mathbb{Z} \to \cdots \to \mathbb{Z} \xrightarrow{\partial_2} \mathbb{Z} \xrightarrow{\partial_1} \mathbb{Z} \to 0$$

Note that

$$\partial_n(c_n) = \sum_{i=0}^{n} (-1)^n c_n|_{\partial_i \Delta^n} = \sum_{i=0}^{n} (-1)^n c_{n-1} = \begin{cases} c_{n-1} & : n > 0 \text{ even} \\ 0 & : \text{else} \end{cases}$$

and hence

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

We conclude that

$$\mathrm{H}_n(*) = \begin{cases} \mathbb{Z} & : n = 0 \\ 0 & : n \neq 0 \end{cases}$$

Commentary 3.10. Here we can already glimpse one of the disadvantages of singular homology. Even for such a simple space the singular chain complex is non-trivial. For larger spaces, the number of singular n-simplices is so enormous that hands-on computations with the singular chain complex are rarely feasible. We will now give interpretations of H_0 and H_1 of a topological space but eventually we will need some high-powered machine that allows us to compute the singular homology groups without actually working with the singular chain complex directly. This we will do in Section 4.

Commentary 3.11. The singular (in contrast to the simplicial) version of homology as defined here is due to Samuel Eilenberg (1944) building on earlier work of Salomon Lefschetz (1933). The improvement of the former over the latter consists in the realization that the simplices should not be viewed as being oriented but rather as having an ordering on the vertices only.

3.2 Low-degree interpretation

Commentary 3.12. Let $\gamma\colon \Delta^1 = [0,1] \to X$ be a singular 1-simplex, that is a path, from $x = \gamma(0)$ to $y = \gamma(1)$, as in Example 3.3. Then we have

$$\partial_1(\gamma) = y - x \in C_0(X)$$

and we conclude that x and y are *homologous*, that is, their difference differs by a boundary. (We also write $x \sim y$.) What is not clear, but we will prove in a moment, is that the converse holds as well: Two points (viewed as singular 0-chains) are homologous if and only if there is a path connecting them.

Lemma 3.13. *Let X be non-empty and path-connected.*[7] *Then* $\mathrm{H}_0(X) = \mathbb{Z}$.

Proof. Define the degree homomorphism $\deg\colon C_0(X) \to \mathbb{Z}$ to send every basis element $x \in X$ (that is, a singular 0-simplex) to $1 \in \mathbb{Z}$. We now show the following points:

- *The map* deg *is surjective.* Indeed, since X is non-empty there exists $x \in X$ with image a generator $1 \in \mathbb{Z}$.
- *We have an inclusion* $B_0(X) \subseteq \ker(\deg)$. Indeed, if $\gamma\colon \Delta^1 \to X$ is a basis element in $C_1(X)$ then
$$\deg(\partial_1(\gamma)) = \deg(\gamma(1) - \gamma(0)) = 1 - 1 = 0.$$
- *Conversely,* $\ker(\deg) \subseteq B_0(X)$. Indeed, let $L = \sum_{x\in X} \lambda_x \cdot x \in C_0(X)$ be an n-chain whose degree vanishes ($\lambda_x \in \mathbb{Z}$). In other words,
$$\mathbb{Z} \ni 0 = \sum_x \lambda_x = \sum_{y,\lambda_y>0} \lambda_y - \sum_{z,\lambda_z<0} (-\lambda_z).$$

  Since both sums are equal we may rewrite L as
$$L = \sum (y - z)$$

[7] Often one takes 'path-connected' to imply 'non-empty'. We're just being extra careful.

for some choice of points $y, z \in X$, possibly repeating. Since X is path-connected there is a path γ from z to y so that $y - z = \partial_1(\gamma) \in B_0(X)$. We conclude that $L \in B_0(X)$ as we needed to show.

Altogether the first isomorphism theorem allows us to conclude:

$$\mathrm{H}_0(X) = \frac{Z_0(X)}{B_0(X)} = \frac{C_0(X)}{\ker(\deg)} \xrightarrow[\cong]{\deg} \mathbb{Z}$$

□

To deduce an interpretation of the 0th singular homology group for *any* space we need the following intuitive fact.

Proposition 3.14. *Let X be a topological space and $(X_\alpha)_\alpha$ its path-connected components. Then we have*

$$\mathrm{H}_n(X) = \oplus_\alpha \mathrm{H}_n(X_\alpha).$$

Proof. If $\sigma\colon \Delta^n \to X$ is a singular n-simplex then its image is path-connected and therefore lies entirely in one of the X_α. In other words we have $C_n(X) = \oplus_\alpha C_n(X_\alpha)$. Moreover, the 'oriented' boundary (3.5) of σ is a linear combination of $(n-1)$-simplices all of which also lie in X_α. That is, ∂_n is the sum of the boundary operators for each X_α. (Another way of saying this is that $C_\bullet(X) = \oplus_\alpha C_\bullet(X_\alpha)$ *as chain complexes*.) This decomposition therefore passes to cycles and boundaries, and eventually to homology. □

Corollary 3.15. *Let X be a topological space. Then* $\mathrm{H}_0(X) = \mathbb{Z}\pi_0(X)$.

Proof. This is a combination of the two previous results, Lemma 3.13 and Proposition 3.14. □

Commentary 3.16. Corollary 3.15 was a warm-up for what will occupy us in the rest of this section, namely an interpretation of $\mathrm{H}_1(X)$. Our first goal is to construct a morphism of groups

(3.17) $$h_1 : \pi_1(X, x) \to \mathrm{H}_1(X).$$

Subsequently we will show that this identifies $\mathrm{H}_1(X)$ with the abelianization of the fundamental group (at least when X is path-connected).

To define (3.17) on elements start with a loop $\gamma\colon \Delta^1 \to X$ based at x, that is, $\gamma(0) = \gamma(1) = x$. It follows that $\partial_1(\gamma) = \gamma(1) - \gamma(0) = x - x = 0 \in C_0(X)$ so that $\gamma \in Z_1(X)$ is a 1-cycle. We would like to define $h_1([\gamma]) = [\gamma] \in \mathrm{H}_1(X)$ as the homology class associated with γ. For this we need to know that if we change γ to a homotopic loop γ' then this homology class remains constant.

Lemma 3.18. *Let γ_1, γ_2 be two paths in X, homotopic relative to their endpoints. Then $\gamma_1 \sim \gamma_2$.*

Proof. We are given a homotopy H between γ_1 and γ_2 relative to their endpoints:

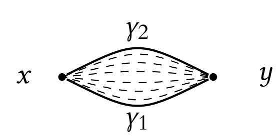

Thus $H\colon [0,1]^2 \to X$ looks as follows:

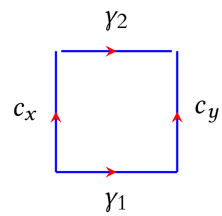

We can turn this into a 2-chain in X like so,

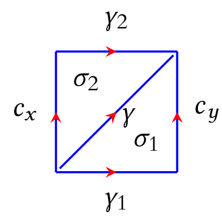

with $\gamma(t) = H(t,t)$ the 'diagonal'. Hence

$$\partial_2(\sigma_2 - \sigma_1) = \gamma_2 - \gamma + c_x - c_y + \gamma - \gamma_1 = (\gamma_2 - \gamma_1) + (c_x - c_y) \tag{3.19}$$

and it suffices to show that $c_x, c_y \in B_1(X)$. For this let $\sigma\colon \Delta^2 \to X$ be the constant map with value x. Then $\partial_2(\sigma) = c_x - c_x + c_x = c_x$ which shows what we want for x, and for y it's the same argument. □

At this point we have a well-defined map (3.17), $h_1([\gamma]) = [\gamma]$. (Note that the brackets on the left mean 'homotopy class' while on the right they mean 'homology class'.)

Lemma 3.20. *The map (3.17) so defined is a group homomorphism.*

Proof. We have already seen at the end of the last proof that the constant loop based at $x \in X$ is a boundary and therefore vanishes in homology. In other words, h_1 preserves the unit for the group structure.

Let γ_1, γ_2 be two loops based at x and let $\gamma_1\gamma_2$ be their composite, given for $t \in [0,1]$ by:

$$\gamma_1\gamma_2(t) = \begin{cases} \gamma_1(2t) & : t \leq 1/2 \\ \gamma_2(2t-1) & : t \geq 1/2 \end{cases}$$

In order to show that $\gamma_1\gamma_2 \sim \gamma_1 + \gamma_2$ we can produce a 2-simplex

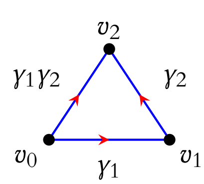

whose boundary is the difference $\gamma_1 + \gamma_2 - \gamma_1\gamma_2$. For this, project down orthogonally onto the face $[v_0, v_2]$ and then apply $\gamma_1\gamma_2$. □

Remark 3.21. Let γ_1, γ_2 be two paths in X with $\gamma_1(1) = \gamma_2(0)$ so that they can be concatenated. The preceding proof goes through and shows that in fact $\gamma_1\gamma_2 \sim \gamma_1 + \gamma_2$. In particular, for any path γ we have $\gamma^{-1} \sim -\gamma$. (Indeed, $\gamma + \gamma^{-1} \sim \gamma\gamma^{-1} \sim 0$, by Lemma 3.18.)

Commentary 3.22. We cannot expect (3.17) to be an isomorphism in general for $\mathrm{H}_1(X)$ is abelian while $\pi_1(X, x)$ is not necessarily. But this means that the map factors through the 'abelianization':

Definition 3.23. Let G be a group. Its *commutator* $[G, G]$ is the subgroup generated by the elements $ghg^{-1}h^{-1}$ for $g, h \in G$. (It is normal.) The *abelianization* of G, denoted G^{ab}, is the quotient $G/[G, G]$.

Example 3.24. If G is abelian then $G^{\mathrm{ab}} = G$.

Example 3.25. If $G = \mathbb{Z} * \mathbb{Z}$ then $G^{\mathrm{ab}} = \mathbb{Z} \oplus Z$. (See Exercise 2.2.)

For more examples, see the second exercise sheet.

Remark 3.26. By construction, G^{ab} is abelian, and it is in fact *universal* with respect to this property. This means that whenever $\phi\colon G \to A$ is a morphism to an abelian group A there exists a unique group morphism $\bar{\phi}\colon G^{\mathrm{ab}} \to A$ such that the composite $G \twoheadrightarrow G^{\mathrm{ab}} \to A$ is ϕ:

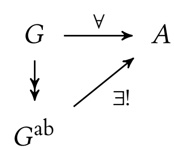

By this universal property of the abelianization and the fact that $\mathrm{H}_1(X)$ is abelian, the homomorphism (3.17) factors uniquely through a morphism between abelian groups,

$$\bar{h}_1\colon \pi_1(X, x)^{\mathrm{ab}} \to \mathrm{H}_1(X), \tag{3.27}$$

and we will now prove:

Theorem 3.28. *If X is non-empty and path-connected then the map (3.27) is an isomorphism.*

Proof. We will, in effect, define an inverse to this map.

Since X is path-connected we choose, for every given $y \in X$, once and for all, a path η_y from x to y. Given *any* path $\gamma\colon \Delta^1 \to X$ we associate a loop based at x, as the following concatenation:

$$g(\gamma) := \eta_{\gamma(0)}\gamma\eta_{\gamma(1)}^{-1} \tag{3.29}$$

Linearly extending, we obtain a homomorphism

$$g\colon Z_1(X) \subseteq C_1(X) \to \pi_1(X, x)^{\mathrm{ab}}.$$

What happens to the boundaries? If $\sigma\colon \Delta^2 \to X$ is a singular 2-simplex

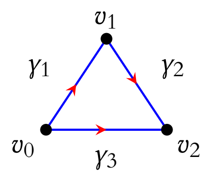

then $\gamma_1\gamma_2$ is homotopic to γ_3 relative to their endpoints. To see this, embed the 2-simplex in the square

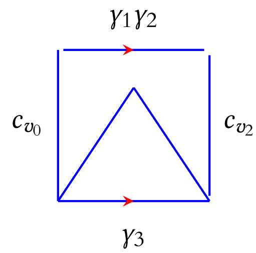

and project vertically onto the 2-simplex.

It follows that in $\pi_1(X,x)^{\mathrm{ab}}$ we have

$$\begin{aligned} g(\partial_2\sigma) &= g\left(\gamma_1+\gamma_2-\gamma_3\right) \\ &= g(\gamma_1)+g(\gamma_2)-g(\gamma_3) \\ &= [\eta_{v_0}\gamma_1\eta_{v_1}^{-1}] + [\eta_{v_1}\gamma_2\eta_{v_2}^{-1}] - [\eta_{v_0}\gamma_3\eta_{v_2}^{-1}] \\ &= [(\eta_{v_0}\gamma_1\eta_{v_1}^{-1})(\eta_{v_1}\gamma_2\eta_{v_2}^{-1})(\eta_{v_2}\gamma_3^{-1}\eta_{v_0}^{-1})] \\ &= [\eta_{v_0}\gamma_1\gamma_2\gamma_3^{-1}\eta_{v_0}^{-1}] \\ &= [\eta_{v_0}\eta_{v_0}^{-1}] \\ &= 0 \end{aligned}$$

so that (3.29) passes to $\bar{g}\colon \mathrm{H}_1(X) \to \pi_1(X,x)^{\mathrm{ab}}$.

Now we check that the two composites are the identity:

- Let γ be a loop based at x. Then $\bar{g}(\bar{h}_1([\gamma])) = \bar{g}([\gamma]) = [\eta_x\gamma\eta_x^{-1}] = [\eta_x]+[\gamma]-[\eta_x] = [\gamma] \in \pi_1(X,x)^{\mathrm{ab}}$ so $\bar{g}\circ\bar{h}_1 = \mathrm{id}$.
- Let $L = \sum \lambda_\gamma \gamma \in Z_1(X)$ be a 1-cycle, $\lambda_\gamma \in \mathbb{Z}$. Replacing $-\gamma$ by γ^{-1} if necessary we may assume that all $\lambda_\gamma > 0$ are positive, see Remark 3.21. Relabeling we can then write $L = \sum_{i=1}^k \gamma_i$ for paths γ_i in X, possibly with repetitions. If γ_1 is not a loop there must exist $i > 1$ such that $\gamma_1(1) = \gamma_i(0)$. Replacing $\gamma_1+\gamma_i$ by the concatenation $\gamma_1\gamma_i$ (again using Remark 3.21) and doing induction we finally reduce to showing the claim for $L = \gamma$ a single loop in X, based at y, say. In that case we have $\bar{h}_1(\bar{g}([\gamma])) = [\eta_y\gamma\eta_y^{-1}] = [\eta_y]+[\gamma]-[\eta_y] = [\gamma]$ and this completes the proof.

□

Corollary 3.30. *If X simply connected (hence non-empty and path-connected) then* $\mathrm{H}_1(X) = 0$.[8]

Corollary 3.31. *We have* $\mathrm{H}_1(S^1) = \mathbb{Z}$. *A generator is given by the (homology class of the) obvious surjective map* $\gamma_1 : \Delta^1 \to S^1$ *identifying the end points.*

Remark 3.32. If we view $\pi_1(X, x)$ as the set of homotopy classes of pointed maps $(S^1, *) \to (X, x)$ this Corollary provides another description of the map (3.17). As we will see in Section 4.1 below, any continuous map $f : S^1 \to X$ induces a morphism in homology $f_* : \mathrm{H}_1(S^1) \to \mathrm{H}_1(X)$ and (3.17) can also be described as:

$$h_1([f]) = f_*([\gamma_1]) = [f \circ \gamma_1].$$

This suggests an analogous map in arbitrary degrees which was first studied by Hurewicz. Once we prove that $\mathrm{H}_n(S^n) = \mathbb{Z}$,[9] generated by the homology class of some $\gamma_n \in Z_n(S^n)$, we can define

$$\begin{aligned} h_n : \pi_n(X, x) = [(S^n, *), (X, x)] &\longrightarrow \mathrm{H}_n(X) \\ [f : S^n \to X] &\longmapsto f_*([\gamma_n]) \end{aligned}$$

4 ... and its fundamental theorems

So far our understanding of the singular homology is restricted to degrees 0 and 1. For example, at this point we don't even know the higher homology groups of such basic spaces as the circle or higher-dimensional spheres. Since the singular chain complex is typically so large, hands-on computations become all but infeasible. The goal of this section, therefore, is to state and prove two fundamental theorems that allow us to compute the singular homology of topological spaces without ever having to deal with the singular chain complex. To show the power of these theorems we will also do some of these easy computations. More involved applications will be forthcoming in later sections.

4.1 Statement and immediate applications

To state the first theorem we observe the following.

Lemma 4.1. *Let $f : X \to Y$ be a continuous map. It induces a map in homology,*

$$f_* : \mathrm{H}_n(X) \to \mathrm{H}_n(Y),$$

satisfying:

1. $\mathrm{id}_* = \mathrm{id}$,
2. $(f \circ g)_* = f_* \circ g_*$.

[8] Of course, if X is empty then the conclusion holds as well.
[9] See Corollary 4.13.

Proof. If $\sigma\colon \Delta^n \to X$ is a singular n-simplex in X then $f \circ \sigma\colon \Delta^n \to Y$ is a singular n-simplex in Y, by continuity of f. Linearly extending we obtain a homomorphism $f_*\colon C_n(X) \to C_n(Y)$. Assembling these homomorphisms for all n yields a chain map $f_*\colon C_\bullet(X) \to C_\bullet(Y)$ and therefore a map in homology (Lemma 2.49). The two properties stated are immediate from the construction: These are true for the composition of maps. □

Theorem 4.2 (Homotopy invariance). *Suppose $f, g\colon X \to Y$ are homotopic maps. Then:*

$$f_* = g_*\colon \mathrm{H}_n(X) \to \mathrm{H}_n(Y)$$

Corollary 4.3. *If $X \simeq Y$ are homotopy equivalent then $\mathrm{H}_n(X) \cong \mathrm{H}_n(Y)$ are isomorphic.*

Proof. Let $f : X \to Y$ be a homotopy equivalence, with $g\colon Y \to X$ a homotopy inverse. That is, $f \circ g \simeq \mathrm{id}_Y$, $g \circ f \simeq \mathrm{id}_X$. Then:

$$f_* \circ g_* = (f \circ g)_* = (\mathrm{id}_Y)_* = \mathrm{id}, \qquad\qquad g_* \circ f_* = (g \circ f)_* = (\mathrm{id}_X)_* = \mathrm{id},$$

so that $g_* : \mathrm{H}_n(Y) \to \mathrm{H}_n(X)$ is an inverse to f_*. □

Corollary 4.4. *Let X be a (non-empty) contractible space. Then*

$$\mathrm{H}_n(X) = \begin{cases} \mathbb{Z} & : n = 0 \\ 0 & : n \neq 0 \end{cases}$$

Proof. This follows from Corollary 4.3 together with Example 3.9. □

Example 4.5. We therefore know the homology of real Euclidean space $\mathbb{R}^k$, of the unit ball D^k and unit cube $[0, 1]^k$, among others.

Commentary 4.6. By itself, Homotopy Invariance might not be as powerful. Keeping one of our motivating examples in mind, the sphere S^n is not contractible nor is it homotopy equivalent to some other space whose homology we *can* compute (at least when $n \geq 1$). However, we know how to build the sphere out of (contractible) n-balls, here for $n = 1$:

(4.7)

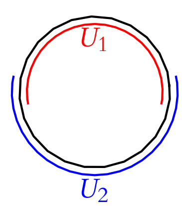

More generally, if X is a manifold we can cover it by open subspaces U_i, each of which is homeomorphic to $\mathbb{R}^n$, for some n, and therefore contractible. If we knew how to express the homology of X in terms of the homology of the U_i then we could apply Homotopy Invariance to conclude. Our second fundamental theorem, named after *Mayer* and *Vietoris*, allows us to do that, with the single provision you are familiar with from the Seifert-van Kampen Theorem: We also need to keep track of the intersection of U_1 and U_2.

Commentary 4.8. The Seifert–van Kampen Theorem is very pertinent for the discussion here so let us recall it in the following form. Let $X = U_1 \cup U_2$ be the union of two path-connected open subspaces $j_1 \colon U_1 \hookrightarrow X$, $j_2 \colon U_2 \hookrightarrow X$ such that $U_1 \cap U_2$ is also path-connected. For simplicity I will omit base-points. We then have a map

$$\pi_1(U_1) * \pi_1(U_2) \xrightarrow{(j_1)_* * (j_2)_*} \pi_1(X),$$

where:

- the map $(j_1)_* * (j_2)_*$ is surjective;
- its kernel is the normal subgroup generated by elements of the form $i(\gamma) = (i_1)_*(\gamma)(i_2)_*(\gamma)^{-1}$ and $i_\ell \colon U_1 \cap U_2 \hookrightarrow U_\ell$;

Passing to homology we should abelianize everything in sight (Theorem 3.28) and this simplifies the situation:

$$\mathrm{H}_1(U_1 \cap U_2) \xrightarrow{(i_1)_* - (i_2)_*} \mathrm{H}_1(U_1) \oplus \mathrm{H}_1(U_2) \xrightarrow{(j_1)_* + (j_2)_*} \mathrm{H}_1(X) \tag{4.9}$$

where now

- $(j_1)_* + (j_2)_*$ is surjective;
- its kernel is precisely the image of i.

We may express the last point as saying that the chain complex (4.9) is exact in the middle, see Definition 2.42. (If we extend the chain complex by zero further to the right, then the first point says precisely that the chain complex is also exact at $\mathrm{H}_1(X)$. However, as we will discuss now, this does not typically hold when $U_1 \cap U_2$ is not path-connected.)

Theorem 4.10 (Mayer–Vietoris long exact sequence)**.** *Let $X = U_1 \cup U_2$ be the union of two open subspaces $j_1 \colon U_1 \hookrightarrow X$, $j_2 \colon U_2 \hookrightarrow X$. There are 'connecting homomorphisms' $\partial \colon \mathrm{H}_n(X) \to \mathrm{H}_{n-1}(U_1 \cap U_2)$ such that*

$$\cdots \to \mathrm{H}_{n+1}(X) \xrightarrow{\partial} \mathrm{H}_n(U_1 \cap U_2) \xrightarrow{(i_1)_* - (i_2)_*} \mathrm{H}_n(U_1) \oplus \mathrm{H}_n(U_2) \xrightarrow{(j_1)_* + (j_2)_*} \mathrm{H}_n(X) \xrightarrow{\partial} \mathrm{H}_{n-1}(U_1 \cap U_2) \to \cdots$$

is an exact *chain complex.*

Remark 4.11. Recall that being a chain complex, the composite of any two maps in this sequence is zero. In other words, the image of each morphism is contained in the kernel of the next. And being exact means that the converse holds, that is, the image of each morphism is *precisely* the kernel of the next one. Equivalently, its homology vanishes in each degree. An exact chain complex also goes by the name of *long exact sequence* thus the name of the theorem.

Example 4.12. Let us go back to the circle (4.7). Since the intersection $U_1 \cap U_2 \simeq S^0$ consists of two points (up to homotopy equivalence) we have $\mathrm{H}_n(U_1 \cap U_2) = 0$ for $n > 0$, see Example 3.9 and Proposition 3.14. Since both U_ℓ are contractible, we also have $\mathrm{H}_n(U_\ell) = 0$ for $n > 0$. This means that most terms in the Mayer–Vietoris l.e.s. vanish:

$$0 \oplus 0 \xrightarrow{(j_1)_* + (j_2)_*} \mathrm{H}_n(S^1) \xrightarrow{\partial} 0$$

is exact for each $n > 1$ so that $\mathrm{H}_n(S^1) = 0$. With our known low-degree computations (Corollary 3.31 and Lemma 3.13) or analyzing the Mayer–Vietoris l.e.s. further we obtain

$$\mathrm{H}_n(S^1) = \begin{cases} \mathbb{Z} & : n = 0, 1 \\ 0 & : n \neq 0, 1 \end{cases}$$

This generalizes to spheres of higher dimensions:

Corollary 4.13. *Let $k \geq 1$. Then* $\mathrm{H}_n(S^k) = \begin{cases} \mathbb{Z} & : n = k, 0 \\ 0 & : n \neq k, 0 \end{cases}$

Proof. We may write S^k as the union of open neighborhoods U_1 (resp. U_2) of the lower (resp. upper) hemisphere:

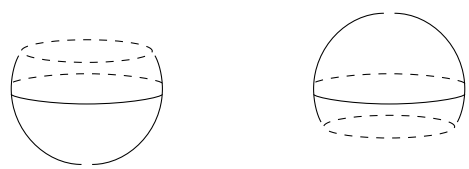

U_1 $\qquad$ U_2

Each of these pieces is contractible and $U_1 \cap U_2$ is homotopy equivalent to S^{k-1}. Mayer–Vietoris therefore yields a sequence

$$\cdots \to H_n(S^{k-1}) \longrightarrow 0 \longrightarrow H_n(S^k) \to$$
$$H_{n-1}(S^{k-1}) \longrightarrow 0 \longrightarrow H_{n-1}(S^k) \dashrightarrow$$
$$H_1(S^{k-1}) \longrightarrow 0 \longrightarrow H_1(S^k) \to$$
$$H_0(S^{k-1}) \longrightarrow H_0(U_1) \oplus H_0(U_2) \longrightarrow H_0(S^k) \longrightarrow 0$$

in which the kernel of each map is equal to the image of the preceding one. We do induction on k, with $k = 1$ having been established in Example 4.12. For $k > 1$, the last terms in the sequence are (by induction hypothesis)

$$0 \to \mathrm{H}_1(S^k) \xrightarrow{\partial} \mathbb{Z} \xrightarrow{(i_1)_* - (i_2)_*} \mathbb{Z} \oplus \mathbb{Z} \xrightarrow{(j_1)_* + (j_2)_*} \mathrm{H}_0(S^k) \to 0,$$

showing that $\mathrm{H}_1(S^k) = 0$ and $\mathrm{H}_0(S^k) = \mathbb{Z}$. (Of course, we already knew that from Lemma 3.13 and Theorem 3.28.) And the upper parts of the sequence give isomorphisms for $n > 1$:

$$\mathrm{H}_n(S^k) \cong \mathrm{H}_{n-1}(S^{k-1}) = \begin{cases} \mathbb{Z} & : n = k \\ 0 & : n \neq k \end{cases}$$

by induction hypothesis. □

Convention 4.14. There is a more convenient formulation of the previous result. For this let $\pi\colon X \to *$ be the unique morphism to the point. We then define *reduced homology*

$$\tilde{\mathrm{H}}_n(X) := \ker\left(\mathrm{H}_n(X) \xrightarrow{\pi_*} \mathrm{H}_n(*)\right).$$

Note that if $n > 0$ then $\tilde{\mathrm{H}}_n(X) = \mathrm{H}_n(X)$ so the only difference is in degree 0. Also, by convention, we only want to consider reduced homology of non-empty spaces. (Otherwise all sorts of pathologies can arise.)

Remark 4.15. The following reformulation of Corollary 4.13 also holds when $k = 0$:

$$\tilde{\mathrm{H}}_n(S^k) = \begin{cases} \mathbb{Z} & : n = k \\ 0 & : n \neq k \end{cases}$$

Remark 4.16. Mayer-Vietoris and Homotopy Invariance hold equally for reduced homology. This follows from their original formulation by keeping track of the induced map π_*. See Exercise 3.2.

Commentary 4.17. We end this section with two celebrated applications of our fundamental theorems. Both are due to Brouwer in the 1910's (although with different proofs).

To motivate the first one, 'the' fixed-point theorem (there are in fact many of these), consider a continuous map $f\colon [0,1] \to [0,1]$. It is quite intuitive and you probably already knew that f necessarily has a fixed point. (If you need a proof: use the Intermediate Value Theorem.)

What happens in higher dimensions? For example, imagine you go hiking in the Peak District and get lost along the way. Imagine also you're super old-fashioned and actually carry a(n infinitely precise) map with you that you lay flat on the earth in front of you. The Brouwer Fixed-point Theorem in dimension 2 asserts that there is a point on the map that represents exactly the point where it touches the earth.

Or, take a glass of red wine and slosh it around. (This makes you look like a connaisseur.) In dimension 3, the theorem asserts that there is a molecule in exactly the same spot as before the sloshing![10]

Corollary 4.18 (Brouwer Fixed-point Theorem)**.** *Every continuous map $f\colon D^k \to D^k$ has a fixed point.*

Proof. Suppose not. Then the ray starting at $f(x)$ in the direction of x meets S^{k-1} in exactly one point $g(x) \neq f(x)$. This $g\colon D^k \to S^{k-1}$ would define a retraction, contradicting Corollary 4.19. □

Corollary 4.19. *$S^{k-1} = \partial D^k$ is not a retract of D^k.*

[10]For another fun application of the theorem see: David Gale, *The Game of Hex and the Brouwer Fixed-Point Theorem*, The American Mathematical Monthly, vol. 86, 1979, pp. 818–827, weblink

Proof. Assume to the contrary that $i\colon S^{k-1} \hookrightarrow D^k$ admits a retraction $r\colon D^k \to S^{k-1}$. Then $\mathrm{id}_{S^k} = r \circ i$ so that

$$\mathbb{Z} = \tilde{\mathrm{H}}_{k-1}(S^{k-1}) \xrightarrow{i_*} \tilde{\mathrm{H}}_{k-1}(D^k) \xrightarrow{r_*} \tilde{\mathrm{H}}_{k-1}(S^{k-1}) = \mathbb{Z}$$

is the identity map, by Lemma 4.1. However, $\tilde{\mathrm{H}}_{k-1}(D^k) = 0$ since D^k is contractible. We have arrived at a contradiction. □

Corollary 4.20 (Invariance of domain)**.** *If $k \neq \ell$ then $\mathbb{R}^k \not\approx \mathbb{R}^\ell$.*

Proof. The argument is similar to Example 1.9. Namely, assume $f\colon \mathbb{R}^k \xrightarrow{\approx} \mathbb{R}^\ell$. (The statement is clear if one of k or ℓ is zero so we may assume $k, \ell > 0$.) Then also $f\colon \mathbb{R}^k\backslash\{0\} \xrightarrow{\approx} \mathbb{R}^\ell\backslash\{f(0)\}$. Since $\mathbb{R}^k\backslash\{0\} \simeq S^{k-1}$ and $\mathbb{R}^\ell\backslash\{f(0)\} \simeq S^{\ell-1}$ we obtain an isomorphism

$$\mathbb{Z} = \tilde{\mathrm{H}}_{k-1}(S^{k-1}) = \tilde{\mathrm{H}}_{k-1}(\mathbb{R}^k\backslash\{0\}) \underset{\cong}{\xrightarrow{f_*}} \tilde{\mathrm{H}}_{k-1}(\mathbb{R}^\ell\backslash\{f(0)\}) = \tilde{\mathrm{H}}_{k-1}(S^{\ell-1})$$

which implies $k = \ell$. □

4.2 Homotopy Invariance

Commentary 4.21. The goal of this section is to prove Homotopy Invariance. Recall (Theorem 4.2) that we start with a homotopy $H\colon X \times [0,1] \to Y$ between f and g and want to prove $f_* = g_*$ as maps in homology. The strategy is as follows:

1. Together with the *prism operator*, H produces a *chain homotopy* between these two chain maps f_* and g_*.
2. It is a very easy algebraic observation that chain homotopic chain maps induce the same map in homology.

In order to see what exactly we ought to produce in the first step we start by introducing chain homotopies.

Definition 4.22. Let $a_\bullet, b_\bullet\colon C_\bullet \to C'_\bullet$ be two chain maps. A *chain homotopy* from a to b is a collection of morphisms

$$\eta_n\colon C_n \to C'_{n+1}$$

such that

$$b_n - a_n = \partial'_{n+1}\eta_n + \eta_{n-1}\partial_n \tag{4.23}$$

for all $n \in \mathbb{Z}$. As usual, we say that a and b are *chain homotopic* if there exists a chain homotopy between them.

Remark 4.24. In pictures, a chain homotopy looks like (for better readability I'm omitting the subscripts that go with the maps)

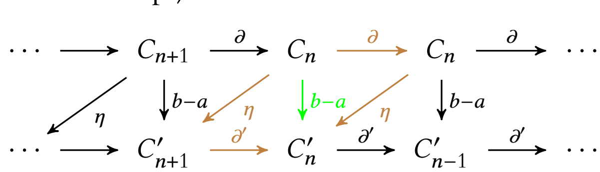

and where the defining identity (4.23) expresses the green path as the sum of the two brown paths.

Here is the easy observation alluded to above:

Lemma 4.25. *Let $a_\bullet$ and $b_\bullet$ be chain homotopic. Then their induced maps in homology are equal:*

$$a_n = b_n \colon \mathrm{H}_n(C_\bullet) \to \mathrm{H}_n(C'_\bullet)$$

Proof. Let $c \in Z_n(C_\bullet)$ be an n-cycle. Applying the defining equation (4.23) we see that

$$b_n(c) - a_n(c) = \partial'_{n+1}\eta_n(c) + \eta_{n-1}\partial_n(c) = \partial'_{n+1}(\eta_n(c))$$

is a boundary. In other words, $b_n(c)$ and $a_n(c)$ are homologous as was to be proven. □

More material on chain homotopies is on Exercise Sheets 3, 4.

Commentary 4.26. We now turn to the prism operator P and producing a chain homotopy between $a = f_*$ and $b = g_*$. The idea is as follows. Starting with a singular n-simplex $\sigma \colon \Delta^n \to X$ and the homotopy $H \colon X \times [0,1] \to Y$ we compose them to a continuous map

$$H \circ (\sigma \times \mathrm{id}) \colon \Delta^n \times [0,1] \to Y.$$

This provides a homotopy from $H \circ (\sigma \times \{0\}) = f \circ \sigma = f_*(\sigma)$ to $H \circ (\sigma \times \{1\}) = g \circ \sigma = g_*(\sigma)$.

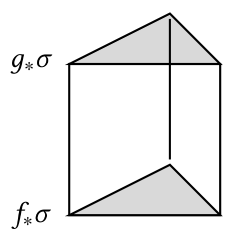

How could we produce an $(n+1)$-chain in Y out of this data? Well, we chop up the 'prism' $\Delta^n \times [0,1]$ into $(n+1)$-simplices. We first consider the situation for small n to get an idea what's going on. (To keep our notation straight let us label the vertices of $\Delta^n = [v_0, \dots, v_n]$. Then the vertices of $\Delta^n \times [0,1]$ are denoted by $v_{ij} = (v_i, j)$ where $0 \le i \le n$ and $0 \le j \le 1$.)

- For $n = 0$ we have that $\Delta^0 \times [0,1] = [v_{00}, v_{01}]$ is already a 1-simplex so we set

$$P(\Delta^0) := [v_{00}, v_{01}].$$

- For $n = 1$ we get a square $\Delta^1 \times [0,1]$. We can turn this into a 2-chain, using the exact same idea as in the proof of Lemma 3.18.

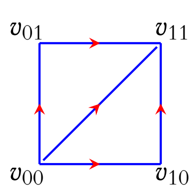

As in *loc. cit.* (or, trial and error shows that) a good 2-chain to consider is

$$P(\Delta^1) := [v_{00}, v_{01}, v_{11}] - [v_{00}, v_{10}, v_{11}]$$

because its boundary is

$$\begin{aligned}\partial_2(P(\Delta^1)) &= [v_{01}, v_{11}] - [v_{00}, v_{11}] + [v_{00}, v_{01}] - [v_{10}, v_{11}] + [v_{00}, v_{11}] - [v_{00}, v_{10}] \\ &= ([v_{01}, v_{11}] - [v_{00}, v_{10}]) - ([v_{10}, v_{11}] - [v_{00}, v_{01}]) \\ &= ([v_{01}, v_{11}] - [v_{00}, v_{10}]) - P(\partial\Delta^1),\end{aligned}$$

as required by (4.23). (Here, we think of P as linearly extended, so that $P(\partial\Delta^1) = P([v_1] - [v_0]) = P([v_1]) - P([v_0]) = [v_{10}, v_{11}] - [v_{00}, v_{01}]$.)

- For $n = 2$ we are dealing with $\Delta^2 \times [0,1]$ which we may subdivide into three 3-simplices as follows:

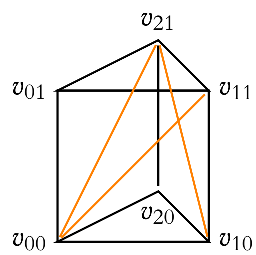

In order to satisfy the formula (4.23) we need to again take an alternating sum:

$$P(\Delta^2) = [v_{00}, v_{01}, v_{11}, v_{21}] - [v_{00}, v_{10}, v_{11}, v_{21}] + [v_{00}, v_{10}, v_{20}, v_{21}]$$

We are now ready to guess the general formula.

Construction 4.27. We define the *prism operator* on $\Delta^n = [v_0, \dots, v_n]$ as follows:

$$P(\Delta^n) = \sum_{i=0}^{n} (-1)^i [v_{00}, \dots, v_{i0}, v_{i1}, \dots, v_{n1}] \in C_{n+1}(\Delta^n \times [0,1])$$

One way to think about this is that the boundary of the prism $\Delta^n \times [0,1]$ is made of three pieces:

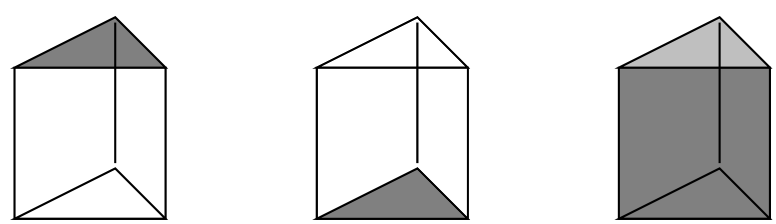

$\Delta^n \times \{1\}$ $\Delta^n \times \{0\}$ $\partial\Delta^n \times [0,1]$

The signs in the definition of $P(\Delta^n)$ ensure we can turn this into an equality of n-chains:

Lemma 4.28. *For every $n \geq 0$ we have in $C_n(\Delta^n \times [0,1])$:*

$$\partial P(\Delta^n) = [v_{01}, \dots, v_{n1}] - [v_{00}, \dots, v_{n0}] - P(\partial \Delta^n)$$

Remark 4.29. Just to avoid any confusion, in this formula, $\partial \Delta^n = \sum_{j=0}^{n} (-1)^j [v_0, \dots, \hat{v}_j, \dots, v_n]$ is the n-th boundary operator applied to the standard n-simplex. And we have linearly extended P so that the last term is

$$P(\partial \Delta^n) = \sum_{j=0}^{n} (-1)^j P([v_0, \dots, \hat{v}_j, \dots, v_n])$$

Before giving the proof of this lemma let us see how the prism operator can be used to finish the proof of Homotopy Invariance.

Proof of Theorem 4.2. Choose a homotopy $H\colon X \times [0,1] \to Y$ from f to g. Let $\sigma\colon \Delta^n \to X$ be a singular n-simplex in X and consider $H \circ (\sigma \times \mathrm{id})$, a map $\Delta^n \times [0,1] \to Y$. It then induces a map on $(n+1)$-chains

$$\left(H \circ (\sigma \times \mathrm{id})\right)_* : C_{n+1}(\Delta^n \times [0,1]) \to C_{n+1}(Y).$$

We define

$$\begin{aligned} \eta_n \colon C_n(X) &\longrightarrow C_{n+1}(Y) \\ \sigma &\longmapsto \left(H \circ (\sigma \times \mathrm{id})\right)_* (P(\Delta^n)) \end{aligned}$$

so that

$$\begin{aligned} \partial \eta_n(\sigma) &= \partial \left(H \circ (\sigma \times \mathrm{id})\right)_* (P(\Delta^n)) && \text{definition of } \eta_n \\ &= \left(H \circ (\sigma \times \mathrm{id})\right)_* (\partial P(\Delta^n)) && \text{chain map identity} \\ &= \left(H \circ (\sigma \times \mathrm{id})\right)_* ([v_{01}, \dots, v_{n1}] - [v_{00}, \dots, v_{n0}] - P(\partial \Delta^n)) && \text{Lemma 4.28} \\ &= g_*\sigma - f_*\sigma - \eta_{n-1}\partial(\sigma) && \text{unwinding definitions} \end{aligned}$$

This shows that η defines a chain homotopy from f_* to g_* so Lemma 4.25 kicks in. □

We are left to give the:

Proof of Lemma 4.28. This is a straightforward albeit tedious calculation.
We have

$$\partial_{n+1} P(\Delta^n) = \sum_{j \leq i} (-1)^{i+j} [v_{00}, \dots, \hat{v}_{j0}, \dots, v_{i0}, v_{i1}, \dots, v_{n1}] + \sum_{j \geq i} (-1)^{i+j+1} [v_{00}, \dots, v_{i0}, v_{i1}, \dots, \hat{v}_{j1}, \dots, v_{n1}] \tag{4.30}$$

- Consider first the terms with $i = j$. We get for these:

$$\sum_{i=0}^{n}[v_{00},\dots,v_{(i-1)0},v_{i1},\dots,v_{n1}] - \sum_{i=0}^{n}[v_{00},\dots,v_{i0},v_{(i+1)1},\dots,v_{n1}]$$

  all but two of which cancel each other out, leaving us with

$$[v_{01},\dots,v_{n1}] - [v_{00},\dots,v_{n0}].$$

- To account for the remaining terms ($i \neq j$) we apply the prism operator P to each face $[v_0,\dots,\hat{v}_j,\dots,v_n]$ of Δ^n,

$$P([v_0,\dots,\hat{v}_j,\dots,v_n]) = \sum_{i<j}(-1)^i[v_{00},\dots,v_{i0},v_{i1},\dots,\hat{v}_{j1},\dots,v_{n1}] + \sum_{j<i}(-1)^{i+1}[v_{00},\dots,\hat{v}_{j0},\dots,v_{i0},v_{i1},\dots,v_{n1}],$$

  hence, taking the alternating sum over all j, we get

$$P(\partial_n\Delta^n) = \sum_{i<j}(-1)^{i+j}[v_{00},\dots,v_{i0},v_{i1},\dots,\hat{v}_{j1},\dots,v_{n1}] + \sum_{j<i}(-1)^{i+j+1}[v_{00},\dots,\hat{v}_{j0},\dots,v_{i0},v_{i1},\dots,v_{n1}].$$

  We note that this is precisely the negative of the terms in (4.30) yet to be accounted for (that is, those with $i \neq j$). This concludes the proof.

□

4.3 Mayer–Vietoris: strategy

Commentary 4.31. The goal of this section is to prove (the existence of) the Mayer–Vietoris long exact sequence, Theorem 4.10. To re-familiarize ourselves with that result let us do one more application.

Example 4.32. Let us compute the homology $\mathrm{H}_\bullet(X)$ of the torus $X = \mathbb{T}$ and the Klein bottle $X = \mathbb{K}$. Recall that these can be obtained from the square by identifying opposite edges:

In each case we may cover them by two open subsets V_1 (in green) and V_2 (in gray):

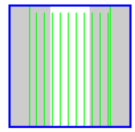

Note that in each case, $V_i \simeq S^1$ and $V_1 \cap V_2 \simeq S^1 \amalg S^1$ which we label L (for left) and R (for right). Here are the pictures for the torus on the left and the Klein bottle on the right:

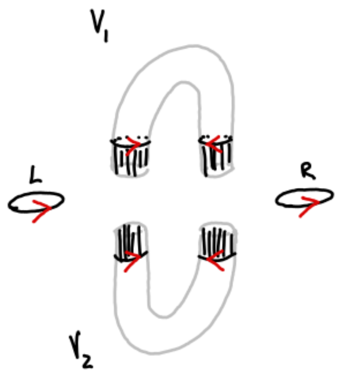

$\mathbb{T}$

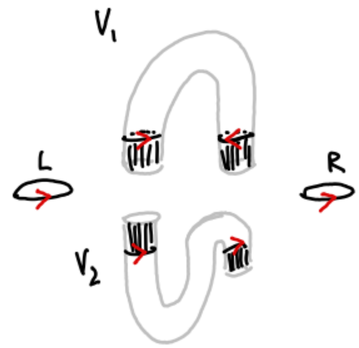

$\mathbb{K}$

The black part depicts the common intersection. And the red arrows amount to choosing generators of the first homology group (for $V_1, V_2, V_1 \cap V_2$), see Corollary 3.31. The MV long exact sequence vanishes in degrees > 2 and the only interesting bit is

$$0 \to \mathrm{H}_2(X) \xrightarrow{\partial} \mathrm{H}_1(L) \oplus \mathrm{H}_1(R) \xrightarrow{i} \mathrm{H}_1(V_1) \oplus \mathrm{H}_1(V_2) \xrightarrow{j} \mathrm{H}_1(X) \xrightarrow{\partial}$$
$$\mathrm{H}_0(L) \oplus \mathrm{H}_0(R) \to \mathrm{H}_0(V_1) \oplus \mathrm{H}_0(V_2) \to \mathrm{H}_0(X) \to 0$$

The bottom row is easy to analyze:

$$\mathbb{Z} \oplus \mathbb{Z} \xrightarrow{\begin{pmatrix} 1 & 1 \\ -1 & -1 \end{pmatrix}} \mathbb{Z} \oplus \mathbb{Z} \xrightarrow{(1\ 1)} \mathbb{Z} \to 0$$

so that the we get an exact sequence

$$0 \to \mathrm{H}_2(X) \xrightarrow{\partial} \mathbb{Z}^2 \xrightarrow{i} \mathbb{Z}^2 \xrightarrow{j} \mathrm{H}_1(X) \xrightarrow{\partial} \mathbb{Z} \to 0.$$

It follows from Exercise 2.6 that

$$\mathrm{H}_2(X) = \ker(i), \qquad \mathrm{H}_1(X) = \mathbb{Z} \oplus \operatorname{coker}(i).$$

In both cases, the inclusions $L \hookrightarrow V_1, V_2$ send generator to generator in H_1. Also in both cases, $R \hookrightarrow V_1$ sends generator to the negative of the generator. For $R \hookrightarrow V_2$, the situation for the torus and the Klein bottle are different. Putting everything together we find for i the matrix:

$$\mathbb{T}: \begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix} \qquad \mathbb{K}: \begin{pmatrix} 1 & -1 \\ -1 & -1 \end{pmatrix}$$

so that

$$\mathrm{H}_n(\mathbb{T}) = \begin{cases} \mathbb{Z} & : n = 0, 2 \\ \mathbb{Z}^2 & : n = 1 \\ 0 & : \text{else} \end{cases} \qquad \mathrm{H}_n(\mathbb{K}) = \begin{cases} \mathbb{Z} & : n = 0 \\ \mathbb{Z} \oplus \mathbb{Z}/2 & : n = 1 \\ 0 & : \text{else} \end{cases}$$

(Note that the result for the torus coincides with the simplicial homology groups computed in Example 2.31. This is an instance of the general comparison theorem between the two homology theories to be discussed in Section 8.1.)

Commentary 4.33. As for the proof of Homotopy Invariance (section 4.2), the proof of Mayer–Vietoris combines a topological step with a purely algebraic step:

1. The purely algebraic fact is that a 'short exact sequence of chain complexes' gives rise to a long exact sequence in homology. (We will define short exact sequences of chain complexes in a moment but you could try to guess already what it should be.)
2. The relevant short exact sequence of chain complexes associated with the cover $X = U_1 \cup U_2$ is

$$0 \to C_\bullet(U_1 \cap U_2) \to C_\bullet(U_1) \oplus C_\bullet(U_2) \to C_\bullet(U_1 + U_2) \to 0,$$

   where $C_\bullet(U_1 + U_2) \subseteq C_\bullet(X)$ is a subcomplex with the same homology groups. This last bit (showing that they have the same homology) is the hardest part. This involves the topological ingredient, a process called *barycentric subdivision*.

Commentary 4.34. Recall that a short exact sequence of abelian groups is an exact sequence of the form

$$0 \to A \xrightarrow{f} B \xrightarrow{g} C \to 0.$$

In other words,

- f is injective,
- g is surjective, and
- $\ker(g) = \mathrm{img}(f)$.

We now transfer this concept to the level of *chain complexes*.

Definition 4.35. A *short exact sequence of chain complexes*

$$0 \to A_\bullet \xrightarrow{f} B_\bullet \xrightarrow{g} C_\bullet \to 0$$

is the data of two chain maps, f and g, such that for each $n \in \mathbb{Z}$,

$$0 \to A_n \xrightarrow{f_n} B_n \xrightarrow{g_n} C_n \to 0$$

is a short exact sequence of abelian groups.

Example 4.36. Let $(B_\bullet, \partial_\bullet)$ be any chain complex. Define a new chain complex $A_\bullet$ by

$$A_n = \begin{cases} B_n & : n > 0 \\ \ker(\partial_0) & : n = 0 \\ 0 & : n < 0 \end{cases}$$

The differentials are the ones from $B_\bullet$ restricted to $A_\bullet$. We may then define $C_\bullet$ levelwise as $C_n = B_n/A_n$ with the induced differentials. The resulting short exact sequence $0 \to A_\bullet \to B_\bullet \to C_\bullet \to 0$ of chain complexes looks as follows:

$$\begin{array}{ccccccccc}
 & & \vdots & & \vdots & & \vdots & & \\
 & & \downarrow & & \downarrow & & \downarrow & & \\
0 & \longrightarrow & B_2 & = & B_2 & \longrightarrow & 0 & \longrightarrow & 0 \\
 & & \downarrow{\scriptstyle \partial_2} & & \downarrow{\scriptstyle \partial_2} & & \downarrow & & \\
0 & \longrightarrow & B_1 & = & B_1 & \longrightarrow & 0 & \longrightarrow & 0 \\
 & & \downarrow{\scriptstyle \partial_1} & & \downarrow{\scriptstyle \partial_1} & & \downarrow & & \\
0 & \longrightarrow & \ker(\partial_0) & \hookrightarrow & B_0 & \twoheadrightarrow & B_0/\ker(\partial_0) & \longrightarrow & 0 \\
 & & \downarrow & & \downarrow{\scriptstyle \partial_0} & & \downarrow{\scriptstyle \partial_0} & & \\
0 & \longrightarrow & 0 & \hookrightarrow & B_{-1} & = & B_{-1} & \longrightarrow & 0 \\
 & & \downarrow & & \downarrow{\scriptstyle \partial_{-1}} & & \downarrow{\scriptstyle \partial_{-1}} & & \\
0 & \longrightarrow & 0 & \hookrightarrow & B_{-2} & = & B_{-2} & \longrightarrow & 0 \\
 & & \downarrow & & \downarrow & & \downarrow & & \\
 & & \vdots & & \vdots & & \vdots & &
\end{array}$$

Proposition 4.37. *Let*

$$0 \to A_\bullet \xrightarrow{f} B_\bullet \xrightarrow{g} C_\bullet \to 0$$

be a short exact sequence of chain complexes. There are 'connecting homomorphisms' $\partial\colon \mathrm{H}_n(C_\bullet) \to \mathrm{H}_{n-1}(A_\bullet)$ *and a long exact sequence in homology:*

$$\cdots \to \mathrm{H}_n(A_\bullet) \xrightarrow{f_n} \mathrm{H}_n(B_\bullet) \xrightarrow{g_n} \mathrm{H}_n(C_\bullet) \xrightarrow{\partial} \mathrm{H}_{n-1}(A_\bullet) \xrightarrow{f_{n-1}} \mathrm{H}_{n-1}(B_\bullet) \xrightarrow{g_{n-1}} \mathrm{H}_{n-1}(C_\bullet) \xrightarrow{\partial} \mathrm{H}_{n-2}(A_\bullet) \to \cdots$$

Proof. This is a proof best done 'live' on the board or on paper/tablet etc. Give it a go yourself! (Or watch the lecture recording. Or, it's Hatcher, p. 116f.) In any case, here's what one has to do:

- To construct ∂ let $c \in Z_n(C_\bullet)$. Choose a lift $b \in B_n$ under g_n. Since $g_{n-1}\partial_n(b) = \partial_n g_n(b) = \partial_n(c) = 0$ there exists $a \in A_{n-1}$ such that $f_{n-1}(a) = \partial_n b$. One checks that $[c] \mapsto [a]$ is a well-defined homomorphism $\partial\colon \mathrm{H}_n(C_\bullet) \to \mathrm{H}_{n-1}(A_\bullet)$.
- One checks exactness at each of the spots in the sequence.

□

Commentary 4.38. We now return to the topological situation. Let $U_1, U_2 \subseteq X$ be two subspaces (not necessarily open). Denote by $C_n(U_1 + U_2) \subseteq C_n(X)$ the subgroup of n-chains that can be written as the sum of n-chains in U_1 and n-chains in U_2.

Remark 4.39. The boundary of an n-chain in U_ℓ is an $(n-1)$-chain in U_ℓ which implies that the differentials in $C_\bullet(X)$ restrict to $C_\bullet(U_1 + U_2)$. With these induced differentials we may view $C_\bullet(U_1 + U_2) \subseteq C_\bullet(X)$ as a sub-chain complex. In other words, the inclusion maps $C_n(U_1 + U_2) \hookrightarrow C_n(X)$ define a chain map.

Proposition 4.40. *Let $j_\ell\colon U_\ell \hookrightarrow X$ and $i_\ell\colon U_1 \cap U_2 \hookrightarrow U_\ell$ be the inclusions, for $\ell = 1, 2$. There is a short exact sequence of chain complexes,*

$$0 \to C_\bullet(U_1 \cap U_2) \xrightarrow{(i_1)_* - (i_2)_*} C_\bullet(U_1) \oplus C_\bullet(U_2) \xrightarrow{(j_1)_* + (j_2)_*} C_\bullet(U_1 + U_2) \to 0.$$

Proof. We know these are chain maps, so we need to show exactness at the three spots (in each degree n):

- The subgroup $C_n(U_1 + U_2)$ was defined precisely as the image of $(j_1)_* + (j_2)_*$ so the latter is surjective.
- The morphism $(i_1)_*\colon C_n(U_1 \cap U_2) \hookrightarrow C_n(U_1)$ is already injective hence so is $(i_1)_* - (i_2)_*$.
- The composite

  $$((j_1)_* + (j_2)_*) \circ ((i_1)_* - (i_2)_*) = (j_1)_*(i_1)_* - (j_2)_*(i_2)_* = (j_1 i_1)_* - (j_2 i_2)_* = 0$$

  since $j_1 i_1 = j_2 i_2$ is simply the inclusion $k\colon U_1 \cap U_2 \hookrightarrow X$.

  Conversely, if $(c_1, c_2) \in C_n(U_1) \oplus C_n(U_2)$ such that $(j_1)_*(c_1) + (j_2)_*(c_2) = 0$ then also $(j_1)_*(c_1) = (j_2)_*(-c_2)$ and we see that both sides must be supported on $U_1 \cap U_2$. That is, there exists $c \in C_n(U_1 \cap U_2)$ such that $k_*(c) = (j_1)_*(c_1) = (j_2)_*(-c_2)$. By injectivity of $(j_\ell)_*$ we deduce that already $(i_1)_*(c) = c_1$ and $-(i_2)_*(c) = c_2$.

This completes the proof. □

Commentary 4.41. We now state the other property about $C_\bullet(U_1 + U_2)$ that we need for the proof of MV. (The proof will be given shortly.)

Proposition 4.42. *Assume that the interiors of U_1 and U_2 jointly cover X. Then the inclusion $C_\bullet(U_1 + U_2) \hookrightarrow C_\bullet(X)$ induces isomorphisms in homology.*

Proof of Mayer-Vietoris (Theorem 4.10). By Proposition 4.37, the short exact sequence of Proposition 4.40 induces a long exact sequence in homology:

$$\cdots \to \mathrm{H}_{n+1}(C_\bullet(U_1{+}U_2)) \xrightarrow{\partial} \mathrm{H}_n(U_1{\cap}U_2) \xrightarrow{(i_1)_*-(i_2)_*} \mathrm{H}_n(U_1){\oplus}\mathrm{H}_n(U_2) \xrightarrow{(j_1)_*+(j_2)_*} \mathrm{H}_n(C_\bullet(U_1{+}U_2)) \xrightarrow{\partial} \cdots$$

By Proposition 4.42, we may replace the homology groups $\mathrm{H}_n(C_\bullet(U_1+U_2))$ by $\mathrm{H}_n(X)$ everywhere thus the claim. □

Remark 4.43. As we have just seen, the Mayer-Vietoris long exact sequence also exists when the subspaces $U_1, U_2 \subseteq X$ are not open. The important requirement is rather that their interiors jointly cover X.

(Of course, some condition is necessary since every space $X = V \cup V^c$ is the union of any subspace and its complement. These intersect trivially so that a Mayer-Vietoris long exact sequence would imply $\mathrm{H}_n(X) = \mathrm{H}_n(V) \oplus \mathrm{H}_n(V^c)$ for all n.)

Commentary 4.44. The remainder of this section is devoted to proving Proposition 4.42. Basically, the idea is that every singular n-simplex $\Delta^n \to X$ can be chopped up into n-simplices that lie completely either in U_1 or in U_2.

More precisely, we will construct 'subdivision' maps $S = S_n \colon C_n(X) \to C_n(X)$ such that

1. $S\colon C_\bullet(X) \to C_\bullet(X)$ is a chain map;
2. $S \simeq \mathrm{id}$ are chain homotopic.
3. For any $\sigma\colon \Delta^n \to X$ there exists $k \geq 1$ such that $S^k(\sigma) \in C_n(U_1 + U_2)$.

Since subdivision makes simplices only smaller, we have $S(C_n(U_1+U_2)) \subseteq C_n(U_1+U_2)$. And the chain homotopy T from id to S will also satisfy $T(C_n(U_1+U_2)) \subseteq C_{n+1}(U_1+U_2)$.

Assuming these we may complete the proof as follows.

Proof of Proposition 4.42. We show that the map $\iota : \mathrm{H}_n(C_\bullet(U_1+U_2)) \to \mathrm{H}_n(X)$ is injective and surjective:

- *Surjectivity:* Let $z \in Z_n(X)$ be an n-cycle in X. This can be written as a linear combination of finitely many n-simplices $\sigma_i\colon \Delta^n \to X$. Choose $k \gg 0$ such that $S^k(\sigma_i) \in C_n(U_1+U_2)$ for all i. Hence also $S^k(z) \in C_n(U_1+U_2)$ and this is a cycle. As shown in Exercise 3.1, we have $S^k \simeq \mathrm{id}$ via some chain homotopy η hence $z - S^k(z) = \partial\eta(z) + \eta\partial(z) = \partial\eta(z)$. In particular, $z \sim S^k z$ and in $\mathrm{H}_n(X)$ we have $[z] = \iota([S^k z])$.
- *Injectivity:* Let $w \in Z_n(C_\bullet(U_1+U_2))$ such that $\iota([w]) = 0$, that is, $w = \partial(z)$ for some $z \in C_{n+1}(X)$. As before there exists k and η such that $z - S^k(z) = \partial\eta(z) + \eta\partial(z)$. But then $\partial S^k(z) = \partial(z) - \partial^2\eta(z) - \partial\eta\partial(z) = w - \partial\eta w$ so that $[w] = 0$ as required.[11] □

[11]Note that in the last step we used that $\eta(w) \in C_{n+1}(U_1+U_2)$. Indeed, $\eta = -\sum_{i=0}^{k-1} TS^i$ sends $C_n(U_1+U_2)$ into $C_{n+1}(U_1+U_2)$.

4.4 Barycentric subdivision

Commentary 4.45. Start with a singular simplex $\sigma\colon \Delta^n \to X$. Since the interiors of U_1 and U_2 cover X, we have an open cover

$$\sigma^{-1}(\mathring{U}_1) \cup \sigma^{-1}(\mathring{U}_2)$$

of Δ^n. Let $A_i = \Delta^n \backslash \sigma^{-1}(\mathring{U}_i)$ be the closed complement, $i = 1, 2$, and define a function $f\colon \Delta^n \to \mathbb{R}$,

$$f(x) = \frac{d(x, A_1) + d(x, A_2)}{2},$$

the average distance of x to each of the two closed subsets. (If one of the A_i is empty then $\sigma(\Delta^n) \subseteq U_i$ and the discussion is moot.) By compactness of Δ^n, this function attains a minimum δ, which is necessarily > 0 (otherwise we wouldn't have a cover). One can then easily check that every simplex $[w_0, \dots, w_n] \subseteq \Delta^n$ of diameter $< \delta$ is entirely contained in one of the $\sigma^{-1}(\mathring{U}_i)$. Such a number δ is called a *Lebesgue number* for the given open cover of Δ^n.

The upshot of this is that as soon as we chop up Δ^n into simplices which are of diameter $< \delta$ then the restriction of σ to each of these lies in $C_n(U_1 + U_2)$. *Barycentric subdivison* is a process which systematically chops up linear simplices into smaller ones such that, iterating the process, the diameter tends to 0.

Convention 4.46. We have been using the following notion implicitly already, for example when discussing the prism operator. Given some euclidean space V and elements $v_0, \dots, v_n \in V$, recall that the *linear simplex* $[v_0, \dots, v_n]$ is the subspace

$$\{\sum_{i=0}^{n} x_i v_i \mid x_i \geq 0, \sum x_i = 1\} \subseteq V.$$

Remark 4.47. For example, the standard n-simplex in $V = \mathbb{R}^{n+1}$ is obtained by taking the standard basis vectors as the vertices v_i. For vectors v_i in 'general position', the linear simplex is at least homeomorphic to Δ^n (although not isometric). However, weird things can happen if we don't require the difference vectors $[v_0, v_1], \dots, [v_0, v_n]$ to be linearly independent. (For example, think of the n-simplex $[v_0, v_0, \dots, v_0]$.) So, typically one also requires the vectors to satisfy this condition.

Definition 4.48. The *barycenter* of a simplex $[v_0, \dots, v_n]$ is the point

$$b := \frac{1}{n+1} \sum_{i=0}^{n} v_i.$$

(Recall that the "barycenter" means center of mass.)

Example 4.49. Here are some examples in low dimensions:

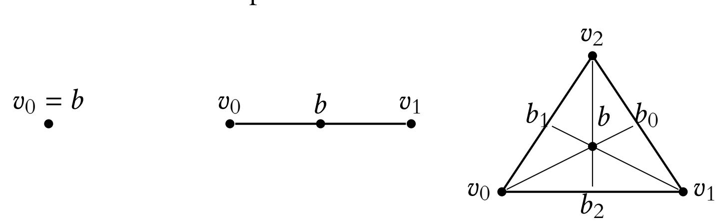

Commentary 4.50. In fact, one can find the barycenter recursively. Knowing the barycenter b_i on the ith face $[v_0, \dots, \hat{v}_i, \dots, v_n]$, let ℓ_i be the line connecting b_i and v_i. Then b is the intersection of all these lines ℓ_i.

We will not need this fact but it is helpful to have in mind since it will lead our thinking in subdividing $[v_0, \dots, v_n]$. The pieces will all be cones with apex b, and base given by one of the pieces we obtained, inductively, by subdividing one of the faces. (So there will be $(n+1)!$ many simplices.) The only thing to be careful about, as always, are the signs. Here is the construction.

Construction 4.51. Given a linear $(n-1)$-simplex $[w_1, \dots, w_n] \subseteq \Delta^n$ we define its *barycentric cone*

$$b[w_1, \dots, w_n] := [b, w_1, \dots, w_n] \subseteq \Delta^n.$$

We extend this to linear combinations of linear $(n-1)$-simplices.

The *barycentric subdivision* $S(\Delta^n) \in C_n(\Delta^n)$ of Δ^n is defined by induction on n:

- If $n = 0$, we set $S(\Delta^0) = \Delta^0$.
- If $n > 0$, we set

$$S(\Delta^n) = bS\partial\Delta^n := \sum_{i=0}^{n} (-1)^i bS(\partial_i \Delta^n).$$[12]

Example 4.52. Compare with the pictures in Example 4.49:

1. $S\Delta^1 = bS[v_1] - bS[v_0] = [b, v_1] - [b, v_0]$
2. $S\Delta^2 = [b, b_0, v_2] - [b, b_0, v_1] - [b, b_1, v_2] + [b, b_1, v_0] + [b, b_2, v_1] - [b, b_2, v_0]$.
3. In dimension 3 we offer the following picture instead of the lengthy formula (and ignoring the signs):

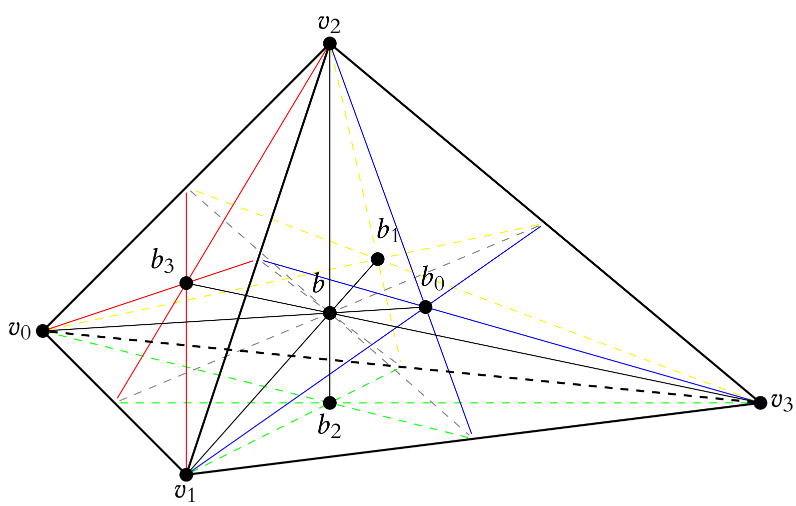

[12] Note that by induction, $S(\partial_i \Delta^n)$ is a linear combination of *linear* simplices hence the barycentric cone is defined.

Commentary 4.53. The signs are chosen so that the boundary of $S\Delta^n$ is the (subdivision of the) boundary of Δ^n, that is, all the internal boundaries cancel out. For example,

$$\partial S\Delta^1 = \partial([b, v_1] - [b, v_0]) = v_1 - b - v_0 + b = v_1 - v_0 = \partial\Delta^1.$$

We now prove this in general.

Lemma 4.54. *1. $\partial b(\sigma) + b\partial(\sigma) = \sigma$ for every linear simplex σ.*

2. We have $\partial S\Delta^n = S(\partial\Delta^n)$.

If we ignore the signs, then the first identity says something intuitively clear: the boundary of the cone consists of the base and the cones on its faces.

Proof. Let $\sigma = [w_1, \dots, w_n]$. Then

$$\begin{aligned}\partial b[w_1, \dots, w_n] &= \partial[b, w_1, \dots, w_n] \\ &= \sum_{i=0}^{n} (-1)^i \partial_i [b, w_1, \dots, w_n] \\ &= [w_1, \dots, w_n] - b\partial[w_1, \dots, w_n]\end{aligned}$$

and the first identity is established.

The second identity follows from the first, by induction. Namely, for $n = 0$ we have zero on both sides. And for $n > 0$ we have

$$\partial S\Delta^n = \partial bS\partial\Delta^n = (\mathrm{id} - b\partial)S\partial\Delta^n = S\partial\Delta^n - bS\partial^2\Delta^n = S\partial\Delta^n$$

where we used the first identity in the second equality, and induction in the third. □

Definition 4.55. Let X be a topological space and $\sigma\colon \Delta^n \to X$ a singular n-simplex. We define the *barycentric subdivision* of σ to be the n-chain

$$S(\sigma) := \sigma_* S\Delta^n \in C_n(X).$$

By linear extension we obtain a homomorphism

$$S\colon C_n(X) \to C_n(X)$$

Corollary 4.56. *$S\colon C_\bullet(X) \to C_\bullet(X)$ is a chain map.*

Proof. We have:

$$\begin{aligned}\partial S(\sigma) &= \partial\sigma_* S\Delta^n && \text{definition of } S \\ &= \sigma_*\partial S\Delta^n && \sigma_* \text{ chain map} \\ &= \sigma_* S(\partial\Delta^n) && \text{Lemma 4.54} \\ &= \sum_{i=0}^{n} (-1)^i \sigma_* S(\partial_i\Delta^n) && \\ &= \sum_{i=0}^{n} (-1)^i S(\partial_i\sigma) && \text{definition of } S \\ &= S(\partial\sigma)\end{aligned}$$

This establishes $\partial S = S\partial$, as required. □

Proposition 4.57. *The barycentric subdivision is chain homotopic to the identity:* $S \simeq \mathrm{id}\colon C_\bullet(X) \to C_\bullet(X)$.

Proof. That is, we need to define a chain homotopy $T\colon C_n(X) \to C_{n+1}(X)$ from S to id. The idea is to chop up the prism $\Delta^n \times [0,1]$ into $(n+1)$-simplices in such a way that the 'bottom face' stays intact (the 'identity') and the 'top face' is subdivided:

$\Delta^1 \times [0,1]$:

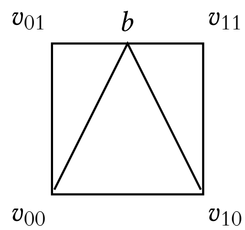

Let us write Δ^n_0 for the bottom face $\Delta^n \times \{0\}$, Δ^n_1 for the top face $\Delta^n \times \{1\}$, and b for the barycentric cone of Δ^n_1. Then we set, recursively,

$$T(\Delta^n) := b\Delta^n_0 - bT\partial\Delta^n_0 \in C_{n+1}(\Delta^n \times [0,1]).$$

- So, for example, for $n = 0$ we get $T(\Delta^0) = [b, v_{00}] = [v_{01}, v_{00}]$.
- And if $n = 1$ we get $T(\Delta^1) = [b, v_{00}, v_{10}] - [b, v_{11}, v_{10}] + [b, v_{01}, v_{00}]$.
- Here is a higher-dimensional example (with the appropriate signs omitted):

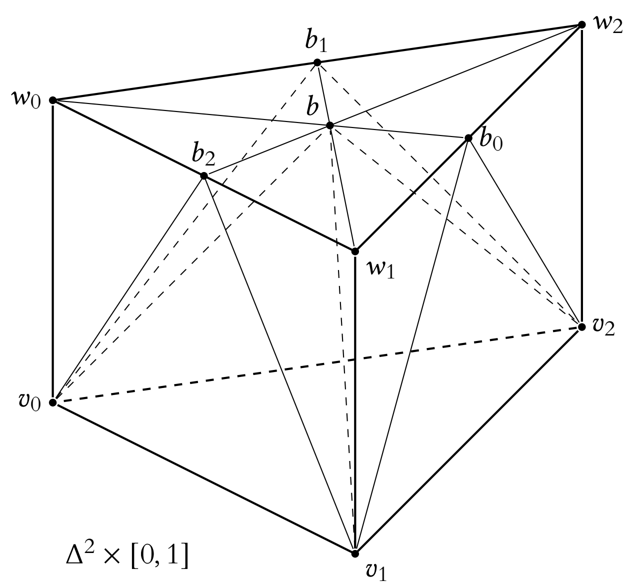

We now verify the following identity reminiscent of the chain homotopy condition (4.23):

(4.58) $$\partial T(\Delta^n) + T(\partial\Delta^n_0) = \Delta^n_0 - S\Delta^n_1$$

Indeed, for $n = 0$ both sides equal $[v_{00}] - [v_{01}]$, and for $n > 0$ we have

$$\begin{aligned}
\partial T(\Delta^n) &= \partial b\Delta_0^n - \partial bT\partial\Delta_0^n && \text{definition of } T(\Delta^n)\\
&= \Delta_0^n - b\partial\Delta_0^n - T\partial\Delta_0^n + b\partial T\partial\Delta_0^n && \text{Lemma 4.54}\\
&= \Delta_0^n - T\partial\Delta_0^n + b(\partial T\partial\Delta_0^n - \partial\Delta_0^n) &&\\
&= \Delta_0^n - T\partial\Delta_0^n - b(S\partial\Delta_1^n + T\partial^2\Delta_0^n) && \text{induction}\\
&= \Delta_0^n - T\partial\Delta_0^n - S\Delta_1^n && \partial^2 = 0, \text{definition of } S
\end{aligned}$$

Given a singular n-simplex $\sigma\colon \Delta^n \to X$ let $\sigma'\colon \Delta^n \times [0,1] \to \Delta^n \xrightarrow{\sigma} X$ be the composition of σ with the projection away from the second factor. If we define

$$\begin{aligned}
T\colon C_n(X) &\to C_{n+1}(X)\\
\sigma &\mapsto \sigma'_* T(\Delta^n)
\end{aligned}$$

then (4.58) implies that T defines a chain homotopy $S \simeq \mathrm{id}$ as claimed. □

Commentary 4.59. In view of Commentary 4.45, the last thing to establish in Commentary 4.44 is that the diameters of the simplices in $S^k\Delta^n$ tend to zero as $k \to \infty$, that is, as we repeatedly subdivide. This follows from the next lemma.

Lemma 4.60. *Let $[w_0, \dots, w_n]$ be a simplex in the barycentric subdivision of $[v_0, \dots, v_n]$. Then*

$$\operatorname{diam}([w_0, \dots, w_n]) \le \frac{n}{n+1} \operatorname{diam}([v_0, \dots, v_n]).$$

Exercise 4.61. Verify this by hand in low dimensions, see Example 4.49!

Proof. If $n = 0$ the claim is true since $[w_0] = [v_0]$ and both have diameter 0. So from now on we assume $n > 0$.

We start with the following observation about any linear simplex $[x_0, \dots, x_n]$ whatsoever:

For every x, its maximum distance to points in the simplex is attained at a vertex x_i.

Indeed, let $y \in [x_0, \dots, x_n]$ with $\|x - y\|$ maximal. So we may write $y = \sum t_i x_i$ with $\sum t_i = 1$ and $t_i \ge 0$. Then we have

$$\|x - y\| = \|x - \sum t_i x_i\| = \|\sum t_i(x - x_i)\| \le \sum t_i \|x - x_i\| \le \max \|x - x_i\|$$

with equality if y is one of those vertices x_i with $\|x - x_i\|$ maximal. This establishes the observation.

In particular, applying this observation twice we see that the diameter of $[w_0, \dots, w_n]$ is the length of the longest edge $[x, y]$. Let us distinguish two cases:

- If none of x, y is the barycenter b then they must be vertices of a simplex in the barycentric subdivision of one of the faces $[v_0, \dots, \hat{v}_i, \dots, v_n]$. By induction we then have

$$\operatorname{diam}[w_0, \dots, w_n] = \|x - y\| \le \frac{n-1}{n} \operatorname{diam}[v_0, \dots, \hat{v}_i, \dots, v_n] \le \frac{n}{n+1} \operatorname{diam}[v_0, \dots, v_n]$$

- If, say, $x = b$, then y lies on some face of $[v_0, \dots, v_n]$ and the observation above allows us to assume $y = v_i$ is one of the vertices of that face. Let b_i be the barycenter of $[v_0, \dots, \hat{v}_i, \dots, v_n]$, that is:

$$b_i = \frac{1}{n} \sum_{j \neq i} v_j, \qquad \text{hence} \qquad b = \frac{1}{n+1} \sum_j v_j = \frac{1}{n+1} v_i + \frac{n}{n+1} b_i.$$

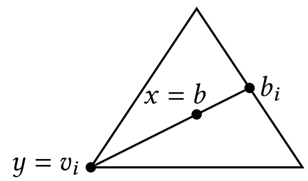

It follows that

$$\operatorname{diam}[w_0, \dots, w_n] = \|v_i - b\| \le \frac{n}{n+1} \|v_i - b_i\| \le \frac{n}{n+1} \operatorname{diam}[v_0, \dots, v_n]$$

□

5 Applications

5.1 Fundamental classes for spheres

Commentary 5.1. We saw in Corollary 4.13 that the homology of the sphere S^k (with $k \ge 1$) is a copy of $\mathbb{Z}$ in degrees $0, k$ (and vanishes in all other degrees). A generator of $H_k(S^k)$ is called a *fundamental class*. Our goal in this section is to describe explicitly fundamental classes (that is, cycles representing them) for all spheres.

More generally, orientable compact connected manifolds M of dimension k have $H_k(M) \cong \mathbb{Z}$, and a generator is called a fundamental class for M. It may be obtained by choosing a triangulation of the manifold and then choosing a compatible orientation of the simplices in the triangulation. (The fact that a triangulation exists is due to compactness. And that a compatible orientation exists is precisely the orientability condition. For disconnected manifolds, one can do this process on each connected component separately.)

Example 5.2. In Example 4.12 we already described a fundamental class for S^1. Viewing the circle as obtained from an interval by identifying the endpoints, the quotient map

$$\Delta^1 = [0, 1] \to [0, 1]/(0 \sim 1) \approx S^1$$

is a 1-cycle which generates $H_1(S^1)$.

Commentary 5.3. When we try to reproduce this construction for S^2 we run into a problem. The singular 2-simplex given by the quotient map

$$\sigma \colon \Delta^2 \to \Delta^2 / \partial\Delta^2 \approx S^2$$

has boundary

$$\partial\sigma = * - * + * = *$$

where $*\colon \Delta^1 \to S^2$ denotes the constant path at the point to which $\partial\Delta^2$ was contracted. In other words, σ is *not* a cycle.

In fact, this construction gives a fundamental class for S^k precisely when k is odd. Below, we will instead describe a construction that works in all dimensions.

Example 5.4. We gave two Δ-complex structures on S^1 in Examples 2.9 and 2.10. The first one is related to the cycle of Example 5.2. The second one suggests looking at the cycle

$$\sigma := \sigma_+ - \sigma_-$$

instead, where $\sigma_+\colon \Delta^1 \to S^1$ picks out the upper hemicircle, and $\sigma_-\colon \Delta^1 \to S^1$ picks out the lower hemicircle, running in the opposite direction.

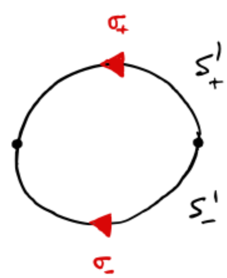

By Theorem 3.28 (or a direct verification), this cycle represents a fundamental class. This is the cycle that we will generalize to all dimensions.

Remark 5.5. Recall that $S^k \approx \partial\Delta^{k+1}$ are homeomorphic, see Exercise 2.7. The k-boundary $\partial(\Delta^{k+1}) \in C_k(\Delta^{k+1})$ actually lies in the subgroup $C_k(\partial\Delta^{k+1})$. Writing $\Delta^{k+1} = [v_0, \ldots, v_{k+1}]$ we therefore have explicitly

$$\partial\Delta^{k+1} = \sum_{i=0}^{k+1} (-1)^i [v_0, \ldots, \hat{v}_i, \ldots, v_{k+1}].$$

Proposition 5.6. *The cycle $\partial\Delta^{k+1} \in C_k(\partial\Delta^{k+1})$ represents a generator in homology.*

Proof. It is clear that this is a cycle, since it is a boundary in the chain complex $C_\bullet(\Delta^{k+1})$. We now prove the statement by induction on k, using Mayer-Vietoris.

Choose open subspaces $U_1, U_2 \subset \Delta^{k+1}$ as follows (see also the picture).

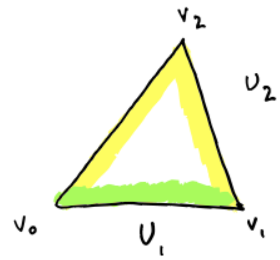

U_1 is an open neighborhood of the last face $\partial_{k+1}\Delta^{k+1}$ which deformation retracts onto the latter. U_2 is an open neighborhood of the remaining faces $\cup_{0\le i\le k}\partial_i\Delta^{k+1}$ which deformation retracts onto these. We choose them in such a way that $U_1 \cap U_2$ deformation retracts onto $\partial(\partial_{k+1}\Delta^{k+1}) = \partial[v_0,\dots,v_k]$ and $U_1 \cup U_2$ deformation retracts onto $\partial[v_0,\dots,v_{k+1}]$. By induction hypothesis, we know that $\mathrm{H}_{k-1}(U_1 \cap U_2) \cong \mathbb{Z}$ is generated by

$$\partial[v_0,\dots,v_k] = \sum_{i=0}^{k}(-1)^i[v_0,\dots,\hat{v}_i,\dots,v_k]. \tag{5.7}$$

In the Mayer-Vietoris l.e.s. for the pair U_1, U_2 covering $U_1 \cup U_2 \simeq \partial\Delta^{k+1}$, the connecting homomorphism

$$\mathrm{H}_k(U_1 \cup U_2) \to \mathrm{H}_{k-1}(U_1 \cap U_2)$$

is an isomorphism and we only need to show that $\partial\Delta^{k+1}$ is sent to the generator (5.7) or its negative.

We shall now use the explicit construction of the connecting homomorphism, see Proposition 4.37. The cycle $\partial\Delta^{k+1}$ has an obvious lift to $C_{k+1}(U_1) \oplus C_{k+1}(U_2)$, namely

$$\left((-1)^{k+1}[v_0,\dots,v_k], \sum_{i=0}^{k}(-1)^i[v_0,\dots,\hat{v}_i,\dots,v_{k+1}]\right).$$

It follows that its image under the connecting homomorphism is the unique $(k-1)$-cycle σ in $U_1 \cap U_2$ which satisfies

$$\Big((i_1)_*(\sigma)\ ,\ -(i_2)_*(\sigma) \Big) = \left((-1)^{k+1}\partial[v_0,\dots,v_k]\ ,\ \sum_{i=0}^{k}(-1)^i\partial[v_0,\dots,\hat{v}_i,\dots,v_{k+1}] \right).$$

It is clear that $\sigma = (-1)^{k+1}\partial[v_0,\dots,v_k]$ does the trick. □

Remark 5.8. Let S^k_+ (resp. S^k_-) be the upper (resp. lower) hemisphere of S^k. Choose homeomorphisms

$$\sigma_+ \colon \Delta^k \xrightarrow{\approx} S^k_+, \qquad \sigma_- \colon \Delta^k \xrightarrow{\approx} S^k_-$$

such that

- both σ_+ and σ_- map the boundary $\partial\Delta^k$ homeomorphically onto the equator $S^k_+ \cap S^k_-$; and
- the composition $\partial\Delta^k \xrightarrow{\sigma_+} S^k_+ \cap S^k_- \xrightarrow{(\sigma_-)^{-1}} \partial\Delta^k$ is the identity.

Corollary 5.9. *The cycle $\sigma_+ - \sigma_- \in C_k(S^k)$ represents a fundamental class for S^k.*

Proof. For $k = 1$ we have already seen this in Example 5.4. We may therefore assume $k > 1$. Our two assumptions on σ_+ and σ_- imply that $\sigma_+ - \sigma_-$ is indeed a cycle. Choose open neighborhoods U_+, U_- of the two hemispheres, respectively, which deformation retract on to them and whose intersection deformation retracts onto the equator. The connecting homomorphism

$$\mathrm{H}_k(S^k) \to \mathrm{H}_{k-1}(U_+ \cap U_-)$$

in the Mayer-Vietoris sequence is an isomorphism and sends $\sigma_+ - \sigma_-$ to $\partial\sigma_+ = \partial\sigma_-$. We saw in Proposition 5.6 that $\partial\Delta^k$ is a generator of $\mathrm{H}_{k-1}(\partial\Delta^k)$ so we win. □

5.2 The Jordan curve theorem

Recall that a *Jordan curve* is a simple closed curve in $\mathbb{R}^2$. The Jordan curve theorem, due to Jordan (1887), states:

> *Every Jordan curve splits the plane into two regions.*

One of the two regions is necessarily bounded and is thus interpreted as the *interior*, while the other region is necessarily unbounded, thus interpreted to be the *exterior*. The Jordan curve is then the boundary of each of these regions.

The theorem is intuitively clear if one imagines a 'nice' (say smooth) curve. However, imagine for example fractal curves or an Osgood curve. In these cases, the intuition breaks down. Luckily, homology comes to our rescue! (Exercise 3.6 treats another class of 'nice' curves.)

Proposition 5.10 (Jordan Curve Theorem). *Let $\gamma\colon S^1 \hookrightarrow \mathbb{R}^2$ be an injective continuous map with image C. Then:*

$$\mathrm{H}_n(\mathbb{R}^2 \backslash C) = \begin{cases} \mathbb{Z}^2 & : n = 0 \\ \mathbb{Z} & : n = 1 \\ 0 & : n > 1 \end{cases}$$

Remark 5.11. According to Proposition 3.14 (and the fact that $\mathbb{R}^2 \backslash C$ is locally path-connected), the part $\mathrm{H}_0(\mathbb{R}^2 \backslash C) = \mathbb{Z}^2$ of the proposition says precisely that the complement of C has two connected components.

We will translate the problem as follows. Let $\mathbb{R}^2 \hookrightarrow \mathbb{R}^2 \cup \{\infty\} \approx S^2$ be the one-point compactification. We are going to prove the following:

$$\mathrm{H}_n(S^2 \backslash C) = \begin{cases} \mathbb{Z}^2 & : n = 0 \\ 0 & : n > 0 \end{cases} \tag{5.12}$$

Exercise 5.13. Let $X = S^2 \backslash C$ and $U = \mathbb{R}^2 \backslash C$, and let V be an open disk around ∞ which does not meet C. Apply Mayer-Vietoris to the cover $X = U \cup V$ to deduce Proposition 5.10 from (5.12).

We start by proving the following auxiliary result (which we are going to apply to parts of the curve C). This says in particular that a non-intersecting path in $\mathbb{R}^2$ cannot separate it.

Lemma 5.14. *Let $\kappa\colon [0,1] \to S^2$ be an injective continuous map with image $D = \kappa([0,1])$. Then*

$$\mathrm{H}_n(S^2 \backslash D) = \begin{cases} \mathbb{Z} & : n = 0 \\ 0 & : n > 0 \end{cases}$$

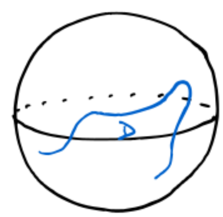

Proof. For an interval $I \subseteq [0,1]$ we set $D_I := \kappa(I)$ so that $D = D_{[0,1]}$. Let $U = S^2 \backslash D_{[0,1/2]}$ and $V = S^2 \backslash D_{[1/2,1]}$. Note that

$$U \cap V = S^2 \backslash D, \qquad\qquad U \cup V = S^2 \backslash D_{[1/2,1/2]} \approx \mathbb{R}^2 \simeq *.$$

The Mayer-Vietoris l.e.s. gives isomorphisms for $n \geq 1$,

$$\mathrm{H}_n(S^2 \backslash D) \cong \mathrm{H}_n(S^2 \backslash D_{[0,1/2]}) \oplus \mathrm{H}_n(S^2 \backslash D_{[1/2,1]}),$$

and a short exact sequence

$$0 \to \mathrm{H}_0(S^2 \backslash D) \to \mathrm{H}_0(S^2 \backslash D_{[0,1/2]}) \oplus \mathrm{H}_0(S^2 \backslash D_{[1/2,1]}) \to \mathbb{Z} \to 0.$$

Assume that for some $n \geq 1$ there exists $\sigma \in \mathrm{H}_n(S^2 \backslash D)$ non-zero. By the isomorphism just exhibited it remains non-zero in at least one of the two groups $\mathrm{H}_n(S^2 \backslash D_{I_1})$, I_1 being the first or second half of the interval. Repeating this argument we find a nested sequence of intervals

$$[0,1] = I_0 \supset I_1 \supset I_2 \supset I_3 \supset \cdots$$

with $\bigcap_{\ell \geq 0} I_\ell = \{p\}$ and such that $\sigma \neq 0$ in all $\mathrm{H}_n(S^2 \backslash D_{I_\ell})$.

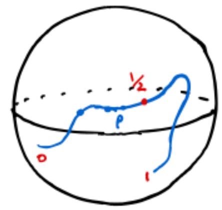

But note that $S^2 \backslash D_{[p,p]} = S^2 \backslash \{\kappa(p)\} \approx \mathbb{R}^2$ is contractible hence σ vanishes in $\mathrm{H}_n(S^2 \backslash \{\kappa(p)\})$. Let $\tau \in C_{n+1}(S^2 \backslash D_{[p,p]})$ such that $\partial \tau = \sigma$. Write τ as linear combination of singular simplices. Each of these has compact image in $S^2 \backslash D_{[p,p]}$. The union of these images is covered by the open subsets $(S^2 \backslash D_{I_\ell})_\ell$ so by compactness, there exists ℓ such that $\tau \in C_{n+1}(S^2 \backslash D_{I_\ell})$. But then $\sigma = 0 \in \mathrm{H}_n(S^2 \backslash D_{I_\ell})$, contradicting our assumption. We deduce that $\mathrm{H}_n(S^2 \backslash D) = 0$ for all $n > 0$.

We argue similarly in degree $n = 0$. Assume $x, y \in S^2 \backslash D$ are two points in different path-connected components. Then we find a nested sequence of intervals I_ℓ as before such that x, y are in different path-connected components of $S^2 \backslash D_{I_\ell}$ for all ℓ. Since $S^2 \backslash \{\kappa(p)\}$ is contractible, it must contain a path connecting x with y. By compactness again, this path misses D_{I_ℓ} for $\ell \gg 0$ and hence x and y lie in the same path-connected component of $S^2 \backslash D_{I_\ell}$ for $\ell \gg 0$. A contradiction. We deduce that $\mathrm{H}_0(S^2 \backslash D) = \mathbb{Z}$. □

Proof of Proposition 5.10. By Remark 5.11, we should compute the homology of $S^2 \backslash C$. Choose S^1_+ and S^1_- to be the upper and lower hemicircles so that $S^1_+ \cap S^1_- = S^0$. We apply MV with

- $X = S^2 \backslash \gamma(S^0)$,
- $U_+ = S^2 \backslash \gamma(S^1_+)$,
- $U_- = S^2 \backslash \gamma(S^1_-)$,
- $U_+ \cap U_- = S^2 \backslash C$.

As $S^1_\pm \approx [0,1]$ we know the homology groups of $U_\pm$ from Lemma 5.14. And as $S^2 \backslash \gamma(S^0) \simeq S^1$ we also know *its* homology groups. Thus Mayer-Vietoris gives the result.

In more detail, for $n > 0$ we have an exact sequence

$$\mathrm{H}_{n+1}(S^2 \backslash \gamma(S^0)) \to \mathrm{H}_n(S^2 \backslash C) \to \mathrm{H}_n(U_+) \oplus \mathrm{H}_n(U_-),$$

in which both outer terms vanish. Hence so does $\mathrm{H}_n(S^2 \backslash C)$.

And for $n = 0$ we have an exact sequence

$$0 \to \mathrm{H}_1(S^2 \backslash \gamma(S^0)) \to \mathrm{H}_0(S^2 \backslash C) \to \mathrm{H}_0(U_+) \oplus \mathrm{H}_0(U_-) \to \mathrm{H}_0(S^2 \backslash \gamma(S^0)) \to 0$$

which is of the form

$$0 \to \mathbb{Z} \to \mathrm{H}_0(S^2 \backslash C) \to \mathbb{Z}^2 \xrightarrow{(1\ 1)} \mathbb{Z} \to 0.$$

It follows that $\mathrm{H}_0(S^2 \backslash C) = \mathbb{Z}^2$, as claimed. □

5.3 Relative homology

Given a subspace $A \subseteq X$, we are interested in the relation between $\mathrm{H}_n(A)$ and $\mathrm{H}_n(X)$. As you observed in Exercise 2.4, the induced map $\mathrm{H}_n(A) \to \mathrm{H}_n(X)$ is not injective in general (nor surjective, of course). The goal of this section is to give a precise measure of the failure of injectivity and surjectivity. This will be expressed in terms of the *relative homology groups* $\mathrm{H}_\bullet(X, A)$ and we will prove *excision*, a powerful tool in computing these groups.

Commentary 5.15. Before passing to homology, that is, at the level of singular chains,

$$C_n(A) \subseteq C_n(X)$$

is a subgroup. Namely, tautologically, the free abelian group on singular n-simplices in $A \subseteq X$. Of course, this actually defines a sub-chain complex $C_\bullet(A) \subseteq C_\bullet(X)$.

Definition 5.16. We define $C_n(X, A)$ to be the quotient $C_n(X)/C_n(A)$ with the differential $C_n(X, A) \to C_{n-1}(X, A)$ induced from $\partial \colon C_n(X) \to C_{n-1}(X)$. We define the *relative homology groups* as the homology groups of the relative singular chain complex $C_\bullet(X, A)$:

$$\mathrm{H}_n(X, A) := \mathrm{H}_n(C_\bullet(X, A))$$

Commentary 5.17. Let us take a moment to break this construction down a bit.

- Let us call an n-chain $z \in C_n(X)$ a *relative n-cycle* if $\partial(z) \in C_{n-1}(A)$. For example, for a singular n-simplex $\sigma \colon \Delta^n \to X$ this just means that image of the boundary $\partial\Delta^n$ is contained in the subspace A.
- Let us also call an n-chain $z \in C_n(X)$ a *relative n-boundary* if it is homologous to an n-chain in A, that is if there exist $w \in C_{n+1}(X)$ and $a \in C_n(A)$ such that $z = a + \partial w$. Note that every relative n-boundary is a relative n-cycle since $\partial z = \partial a$.

By construction,

$$\mathrm{H}_n(X, A) \cong \frac{\text{relative } n\text{-cycles}}{\text{relative } n\text{-boundaries}}.$$

Therefore, the intuition is that $\mathrm{H}_n(X, A)$ measures the homology of X 'discarding A'. We will see later that in good cases this intuition can be made precise (Theorem 5.23, Proposition 5.28).

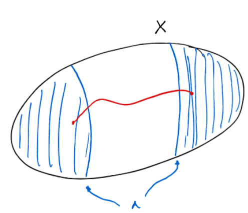

A relative 1-cycle that is not a cycle in X

We now come back to the question raised at the start of this section. The *long exact sequence for the pair* (X, A) appearing in the following result provides a precise way of measuring the failure of the maps $\mathrm{H}_n(A) \to \mathrm{H}_n(X)$ to be isomorphisms.

Corollary 5.18. *There is an exact sequence*

$$\cdots \longrightarrow \mathrm{H}_n(A) \longrightarrow \mathrm{H}_n(X) \longrightarrow \mathrm{H}_n(X,A) \to \mathrm{H}_{n-1}(A) \to \mathrm{H}_{n-1}(X) \to \mathrm{H}_{n-1}(X,A) \to \cdots$$

Proof. By construction, we have a short exact sequence of chain complexes

$$0 \to C_\bullet(A) \to C_\bullet(X) \to C_\bullet(X,A) \to 0.$$

Therefore the claim follows from Proposition 4.37. □

Exercise 5.19. Show that the connecting homomorphism takes a homology class $[z] \in \mathrm{H}_n(X,A)$ represented by a relative cycle $z \in C_n(X)$ (cf. Commentary 5.17) to $[\partial z] \in \mathrm{H}_{n-1}(A)$.

Exercise 5.20. Let $x \in X$ be a point and set $A = \{x\}$. We write (X,x) for the pair $(X,\{x\})$. Show that there is an isomorphism

$$\mathrm{H}_n(X,x) \cong \tilde{\mathrm{H}}_n(X)$$

for all n.

Remark 5.21. Let $k \geq 1$. Recall (Commentary 5.3) that the canonical singular k-simplex

$$\sigma\colon \Delta^k \to \Delta^k/\partial\Delta^k \approx S^k,$$

given by the quotient map, described a fundamental class for S^k precisely when k was odd. The problem for even k may seem silly: It's that σ is not a cycle. We can rectify that using relative homology (and reduced homology, given Exercise 5.20). Indeed, σ is a perfectly acceptable *relative* k-cycle for the pair

$$(\Delta^k/\partial\Delta^k, \partial\Delta^k/\partial\Delta^k) \approx (S^k, *).$$

Moreover, its class in homology generates the group

$$\mathrm{H}_k(S^k, *).$$

Note that $\mathrm{H}_k(S^k) \xrightarrow{\sim} \mathrm{H}_k(S^k,*)$ is an isomorphism so, in some sense, σ *does* define a fundamental class in all dimensions. Can you explain the apparent contradiction with Commentary 5.3?

Commentary 5.22. We now come to excision which is essentially equivalent to Mayer-Vietoris (Theorem 4.10). The way to think about it is as follows. If it is indeed true that $\mathrm{H}_n(X,A)$ 'ignores' what's going on inside of A then, surely, this group won't change if I remove (=excise) a piece of A?

Theorem 5.23 (Excision). *Let $Z \subseteq A \subseteq X$ be subspaces and assume that $\bar{Z} \subseteq \mathring{A}$. Then*

$$\mathrm{H}_n(X,A) \cong \mathrm{H}_n(X\backslash Z, A\backslash Z).$$

Proof. We are going to employ the following trick. Set $B = X\backslash Z$. Note that:

- $A \cap B = A\backslash Z$,
- $\mathring{A} \cup \mathring{B} = \mathring{A} \cup (X\backslash \bar{Z}) = X$ by our assumption, and
- $C_n(X\backslash Z, A\backslash Z) = C_n(B, A \cap B) = \frac{C_n(B)}{C_n(A\cap B)} \cong \frac{C_n(A+B)}{C_n(A)}$ by the second isomorphism theorem.

We now have a 'morphism of short exact sequences of chain complexes' (meaning simply, a commutative diagram albeit in three dimensions):

$$\begin{array}{ccccccccc} 0 & \longrightarrow & C_\bullet(A) & \longrightarrow & C_\bullet(A+B) & \longrightarrow & C_\bullet(X\backslash Z, A\backslash Z) & \longrightarrow & 0 \\ & & \downarrow{=} & & \downarrow & & \downarrow & & \\ 0 & \longrightarrow & C_\bullet(A) & \longrightarrow & C_\bullet(X) & \longrightarrow & C_\bullet(X, A) & \longrightarrow & 0 \end{array}$$

Passing to homology we get a morphism of chain complexes

$$\begin{array}{ccccccccc} \mathrm{H}_n(A) & \longrightarrow & \mathrm{H}_n(A+B) & \longrightarrow & \mathrm{H}_n(X\backslash Z, A\backslash Z) & \longrightarrow & \mathrm{H}_{n-1}(A) & \longrightarrow & \mathrm{H}_{n-1}(A+B) \\ \downarrow{=} & & \downarrow & & \downarrow & & \downarrow{=} & & \downarrow \\ \mathrm{H}_n(A) & \longrightarrow & \mathrm{H}_n(X) & \longrightarrow & \mathrm{H}_n(X, A) & \longrightarrow & \mathrm{H}_{n-1}(A) & \longrightarrow & \mathrm{H}_{n-1}(X) \end{array}$$

in which both rows are exact. By barycentric subdivision, the second and the fifth vertical arrow are isomorphisms, see Proposition 4.42. By the (aptly-named) five lemma (Exercise 4.1), the middle vertical arrow is an isomorphism as well. This shows the claim. □

Commentary 5.24. Recall that a *topological manifold of dimension k* is a Hausdorff space so that every point has an open neighborhood homeomorphic to $\mathbb{R}^k$. Of course, every smooth manifold is a topological manifold.

Corollary 5.25. *Let M be a k-dimensional topological manifold and $x \in M$ a point. Then*

$$\mathrm{H}_n(M, M\backslash x) \cong \mathrm{H}_n(\mathbb{R}^k, \mathbb{R}^k\backslash *) \cong \begin{cases} \mathbb{Z} & : n = k \\ 0 & : n \neq k \end{cases}$$

Proof. Let $x \in U$ be an open neighborhood such that $U \approx \mathbb{R}^k$. Then Excision gives the first isomorphism (with $X = M$, $A = M\backslash x$, and $Z = M\backslash U$). For the second isomorphism we consider the long exact sequence of the pair $(\mathbb{R}^k, \mathbb{R}^k\backslash *)$. (Note that the claim is clear if $k = 0$ so we will assume $k \geq 1$ from now on.) Using that $\mathbb{R}^k\backslash * \simeq S^{k-1}$ and that $\mathbb{R}^k$ is contractible we get $\mathrm{H}_n(\mathbb{R}^k, \mathbb{R}^k\backslash *) = 0$ for $n > k$, and for $n < k$, while for $n = k$ we find an exact sequence

$$0 \to \mathrm{H}_k(\mathbb{R}^k, \mathbb{R}^k\backslash *) \to \mathrm{H}_{k-1}(S^{k-1}) \to \mathrm{H}_{k-1}(\mathbb{R}^k).$$

If $k > 1$ the last term vanishes, giving the required isomorphism. If $k = 1$, the last map identifies with $(\,1\;1\,) : \mathbb{Z}^2 \to \mathbb{Z}$ giving, again, the required isomorphism. □

The next result generalizes Corollary 4.20.

Corollary 5.26 (Invariance of domain, revisited). *Let $U \subseteq \mathbb{R}^k$, $V \subseteq \mathbb{R}^\ell$ be non-empty open subsets. If $U \approx V$ then $k = \ell$.*

We conclude this section with another version of the intuition, according to which relative homology $\mathrm{H}_n(X, A)$ 'discards A'.

Definition 5.27. A pair (X, A) is called *good* if

1. $A \subseteq X$ is a closed subset, and
2. there exists an open neighborhood $A \subseteq V$ which deformation retracts onto A.

Proposition 5.28. *Let (X, A) be a good pair. Then the quotient map $X \to X/A$ induces isomorphisms*

$$\mathrm{H}_n(X, A) \cong \mathrm{H}_n(X/A, A/A) \cong \tilde{\mathrm{H}}_n(X/A).$$

Proof. The long exact sequence for the pairs (X, A) and (X, V), respectively, and the five lemma give an isomorphism

$$\mathrm{H}_n(X, A) \xrightarrow{\sim} \mathrm{H}_n(X, V).$$

Excision gives a further isomorphism $\mathrm{H}_n(X, V) \cong \mathrm{H}_n(X \backslash A, V \backslash A)$. Combining the two we have

$$\mathrm{H}_n(X, A) \cong \mathrm{H}_n(X \backslash A, V \backslash A). \tag{5.29}$$

Now, repeat this argument with $(X/A, A/A)$ instead of (X, A), yielding

$$\mathrm{H}_n(X/A, A/A) \cong \mathrm{H}_n((X/A) \backslash (A/A), (V/A) \backslash (A/A)) = \mathrm{H}_n(X \backslash A, V \backslash A) \tag{5.30}$$

Finally, (5.29) and (5.30) together prove the Proposition. □

Example 5.31. Clearly, $(\Delta^k, \partial\Delta^k)$ is a good pair. We deduce that

$$\mathrm{H}_n(\Delta^k, \partial\Delta^k) \cong \tilde{\mathrm{H}}_n(S^k) \cong \begin{cases} \mathbb{Z} & : n = k \\ 0 & : n \neq k \end{cases}$$

6 Degrees

The degree of a map $S^k \to S^k$, despite its simplicity, is a powerful concept. We will illustrate this in this section, after studying some of its properties.

6.1 Basic properties and examples

Definition 6.1. Let $f : S^k \to S^k$ be a (as always, continuous) map. Then the induced homomorphism in homology,

$$f_* : \tilde{\mathrm{H}}_k(S^k) \to \tilde{\mathrm{H}}_k(S^k),$$

is an endomorphism of an infinite cyclic group, and hence given by multiplication by a unique integer $\deg(f) \in \mathbb{Z}$. This is called the *degree* of f.

Let us record some basic properties around this construction.

Lemma 6.2. *Let $f, g\colon S^k \to S^k$. Then:*

1. $\deg(\mathrm{id}) = 1$
2. $\deg(g \circ f) = \deg(g) \cdot \deg(f)$
3. *if $f \simeq g$ then* $\deg(f) = \deg(g)$
4. *if f is a homotopy equivalence then* $\deg(f) = \pm 1$
5. *if f is not surjective (for example, constant) then* $\deg(f) = 0$

Proof. The first two properties are clear, and the third follows from Homotopy Invariance. The next follows from the first three, cf. the proof of Corollary 4.3. For the last one, let $x \in S^k$ not in the image of f. Then f factors as

$$\begin{array}{ccc} S^k & \xrightarrow{\quad f \quad} & S^k \\ & \searrow \quad \nearrow & \\ & S^k \backslash x & \end{array}$$

Upon passing to reduced homology we see that f_* factors through $\tilde{\mathrm{H}}_k(S^k \backslash x) = 0$ since $S^k \backslash x$ is contractible:

$$\begin{array}{ccc} \mathbb{Z} & \xrightarrow{\quad f_* \quad} & \mathbb{Z} \\ & \searrow \quad \nearrow & \\ & 0 & \end{array}$$

We deduce that $\deg(f) = 0$. □

Commentary 6.3. Note that these basic properties do not produce a map $f\colon S^k \to S^k$ whose degree is different from $0, 1$. However, for $k = 1$ we know how to produce maps of arbitrary degree (see Example 6.5 below). For $k > 1$ we will deduce the existence of such maps from the case $k = 1$, see Proposition 6.6.

Exercise 6.4. Let $k = 0$. What are the possible degrees? Give an example of each.

Example 6.5. Let n be an integer, and let $f\colon z \mapsto z^n$ be the n-power map of $S^1 \subseteq \mathbb{C}^\times$ viewed as the complex numbers of norm 1. We saw in Example 5.2 that the loop

$$\begin{aligned} \sigma\colon [0,1] &\to S^1 \\ t &\mapsto e^{2\pi i t} \end{aligned}$$

represents a generator in $\mathrm{H}_1(S^1) \cong \tilde{\mathrm{H}}_1(S^1)$. Then $f_*([\sigma]) = [f \circ \sigma]$ is represented by the loop $t \mapsto e^{2\pi i n t}$ which is homologous to $n\sigma$. (For example, you can use Theorem 3.28 to see this most easily.) We deduce that f has degree n.

Proposition 6.6. *Let $k \geq 1$. For every integer $n \in \mathbb{Z}$ there exists a map $f\colon S^k \to S^k$ of degree n.*

Proof. Recall (Exercise 2.9) the suspension SX of a space X, which is obtained from $X \times [-1,1]$ by collapsing each of $X \times \{1\}$ and $X \times \{-1\}$ to a point. The two opens

$$C_+X = X \times (-\epsilon, 1]/(X \times \{1\}), \quad C_-X = X \times [-1, \epsilon)/(X \times \{-1\})$$

are contractible, and Mayer-Vietoris gives an isomorphism

$$\partial\colon \mathrm{H}_{k+1}(SX) \xrightarrow{\sim} \mathrm{H}_k(X).$$

A map $f\colon X \to X$ induces a map $Sf\colon SX \to SX$ in an obvious way, and one checks (using the definition of the connecting homomorphism) that the square

$$\begin{array}{ccc} \mathrm{H}_{k+1}(SX) & \xrightarrow{Sf_*} & \mathrm{H}_{k+1}(SX) \\ \downarrow \partial & & \downarrow \partial \\ \mathrm{H}_k(X) & \xrightarrow{f_*} & \mathrm{H}_k(X) \end{array}$$

commutes.

Applying this with $X = S^{k-1}$ and noticing that $S(S^{k-1}) \approx S^k$, we get $\deg(Sf) = \deg(f)$ and thereby reduce, inductively, to $k = 1$. This we saw in Example 6.5. □

Commentary 6.7. Given a continuous map $f\colon S^k \to S^k$, how would one go about computing its degree? This is certainly not obvious just from the definitions. We have described in Corollary 5.9 a rather explicit fundamental class for S^k, represented by a cycle $\sigma_+ - \sigma_-$. Therefore we need to find the integer n such that the k-chains

(6.8) $$n\sigma_+ - n\sigma_- \qquad \text{and} \qquad f \circ \sigma_+ - f \circ \sigma_-$$

are homologous. Unfortunately, we don't seem to have a good way of solving this problem. Since the isomorphism $\mathrm{H}_k(S^k) \cong \mathbb{Z}$ was obtained by induction on k, through the connecting homomorphism in the Mayer-Vietoris sequence $\mathrm{H}_k(S^k) \xrightarrow[\sim]{\partial} \mathrm{H}_{k-1}(S^{k-1})$, the natural thing to try would be to apply ∂ to both k-chains. Eventually, we would reduce to the case $k = 1$ where things might be sufficiently explicit. However, you can see how these reductions become quite rapidly impractical in most cases. (Think about how the connecting homomorphsm was defined!)

We will actually see examples (okay: *one* example) where the integer n in (6.8) can be determined explicitly. However, besides that we will mostly develop general tools for computing degrees. And at the same time, and arguably more importantly, we will learn some interesting things about spheres.

6.2 Antipodes

Lemma 6.9. *Let $S^k \subset \mathbb{R}^{k+1}$ be the unit circle. Let $f\colon S^k \to S^k$ be the reflection in a hyperplane through 0 in $\mathbb{R}^{k+1}$. Then* $\deg(f) = -1$.

Proof. Let $H \subseteq \mathbb{R}^{k+1}$ be the fixed hyperplane. It splits the sphere S^k into two hemispheres S^k_+ and S^k_- as in Remark 5.8. If we pick some homeomorphism $\sigma_+ \colon \Delta^k \xrightarrow{\approx} S^k_+$ and set $\sigma_- = f \circ \sigma_+$ then these satisfy the assumptions of Remark 5.8. By Corollary 5.9, $s = [\sigma_+ - \sigma_-]$ generates $\tilde{\mathrm{H}}_k(S^k)$, and

$$f_*(s) = [f \circ \sigma_+] - [f \circ \sigma_-] = [\sigma_-] - [\sigma_+] = -f_*(s),$$

so that $\deg(f) = -1$. □

Example 6.10. Let $T \colon \mathbb{R}^{k+1} \to \mathbb{R}^{k+1}$ be a linear orthogonal transformation. It restricts to a homeomorphism

$$f \colon S^k \xrightarrow{\approx} S^k.$$

We therefore already know $\deg(f) = \pm 1$. Let us show that in fact $\deg(f) = \det(T)$.

Choosing a suitable orthonormal basis, T can be represented by a block sum matrix where each block is either a 2×2-rotation matrix, or a 1×1-matrix with entry ± 1, with at most one -1. One can homotope any rotation matrix to the identity, via $t \mapsto \begin{pmatrix} \cos(\alpha t) & -\sin(\alpha t) \\ \sin(\alpha t) & \cos(\alpha t) \end{pmatrix}$, so we further reduce to the case where T is the identity, or a reflection at a hyperplane through 0. In both cases we have computed the degree and it coincides with the determinant of T.

Corollary 6.11. *Let $f \colon S^k \to S^k$ be the antipodal map $x \mapsto -x$. Then* $\deg(f) = (-1)^{k+1}$.

Proof. The map $x \mapsto -x \colon \mathbb{R}^{k+1} \to \mathbb{R}^{k+1}$ is the composition of $k+1$ reflections through the coordinate hyperplanes. We conclude with Lemma 6.9. □

Corollary 6.12. *Suppose $f \colon S^k \to S^k$ has no fixed points. Then* $\deg(f) = (-1)^{k+1}$.

Proof. It is sufficient to prove that f is homotopic to the antipodal map, by Corollary 6.11. The line through $f(x)$ and $-x$ goes through the origin if and only if $f(x) = x$.

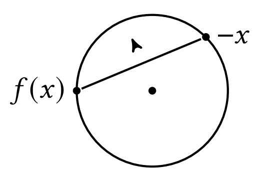

Since we are excluding this possibility, the following expression defines a homotopy from the antipodal map to f:

$$(t, x) \mapsto \frac{tf(x) + (1-t)(-x)}{\|tf(x) + (1-t)(-x)\|}$$

□

Commentary 6.13. Recall that a vector field on S^k is a continuous map $\mathrm{v} \colon S^k \to \mathbb{R}^{k+1}$. It is a *tangent vector field* if $\mathrm{v}(x) \perp x$ for all $x \in S^k$.

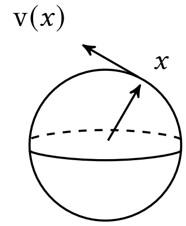

Example 6.14. Obviously, the constant map $v(x) = 0$ is a tangent vector field on every sphere. We are interested in tangent vector fields that vanish *nowhere*.

Example 6.15. The map

$$x = (x_1, \dots, x_{2m}) \mapsto v(x) = (-x_2, x_1, \dots, -x_{2m}, x_{2m-1})$$

defines a non-vanishing tangent vector field on S^{2m-1}.

Obviously this particular map can only be defined on odd-dimensional spheres. Is there another, maybe more clever, construction for even-dimensional ones? The answer is: No!

Corollary 6.16 (Hairy Ball Theorem)**.** *Every tangent vector field on an even-dimensional sphere vanishes at some point.*

Proof. Assume to the contrary that $v\colon S^k \to \mathbb{R}^{k+1}$ is a non-vanishing tangent vector field with k even. Consider then the map

$$S^k \times [0, 1] \ni (x, t) \mapsto \cos(\pi t)x + \sin(\pi t)v(x) \in \mathbb{R}^{k+1}.$$

Since $x \perp v(x)$ and $v(x) \neq 0$ the two vectors $x, v(x)$ are linearly independent. It follows that the map lands in $\mathbb{R}^{k+1} \backslash \{0\}$ and we can divide by the norm to get a homotopy $S^k \times [0, 1] \to S^k$. At time $t = 0$ (resp. $t = 1$) it is the identity (resp. antipodal) map. But this is impossible since $\deg(\mathrm{id}) = 1$ while $\deg(-\mathrm{id}) = (-1)^{k+1} = -1$, by Corollary 6.11. □

Commentary 6.17. Take an (ordinary, 3-dimensional) ball and imagine it comes with hairs on its surface, for some weird reason. (For example, like a coconut comes with hairs. Or you may think of a hedgehog that is curling up into a ball instead.) The theorem says that you cannot comb all these hairs flat at the same time. Here is a picture (from wikipedia) illustrating a failed attempt:

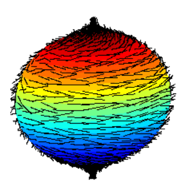

You can see how there are tufts at each of the two poles.[13]

[13] It is not a coincidence that there are two tufts here either (at least if the vector field is also supposed smooth). If counted with the correct 'multiplicities' (=indices) the number is always equal to 2 which is the Euler characteristic of the 2-sphere, as we will discuss later. This statement is known as the Poincaré-Hopf Theorem.

The theorem is due to Brouwer (1912). For a (maybe slightly) more practical implication see Exercise 5.8.

6.3 Local degrees

Commentary 6.18. We will now redo the definition of the degree, but this time as a *local* invariant around a point. The main result is Proposition 6.21 that expresses the 'global' degree in terms of local ones. This is a very useful tool for computing the degree. See Example 6.22 for a nice application.

Remark 6.19. Let $k \geq 1$, let $f\colon S^k \to S^k$ be a continuous map and assume $f^{-1}(y) = \{x_1, \ldots, x_n\}$ is finite. We may choose disjoint open balls $U_i \subseteq S^k$ around x_i.

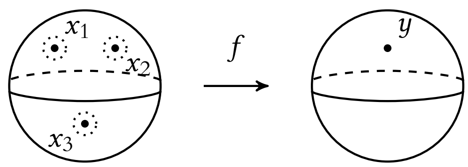

So f induces a map of pairs $(U_i, U_i - x_i) \to (S^k, S^k - y)$ and hence by excision, a morphism in homology,

$$\mathrm{H}_k(S^k, S^k - x_i) \cong \mathrm{H}_k(U_i, U_i - x_i) \xrightarrow{f_*} \mathrm{H}_k(S^k, S^k - y). \tag{6.20}$$

We saw in Corollary 5.25 that both of these groups are infinite cyclic. In fact, the map in the long exact sequence for the pair $\mathrm{H}_k(S^k) \xrightarrow{\sim} \mathrm{H}_k(S^k, S^k - x_i)$ is an isomorphism, and similarly $\mathrm{H}_k(S^k) \xrightarrow{\sim} \mathrm{H}_k(S^k, S^k - y)$. Therefore we may view (6.20) as an endomorphism

$$f|_{x_i}\colon \mathrm{H}_k(S^k) \to \mathrm{H}_k(S^k)$$

of an infinite cyclic group, which is multiplication by some integer, called the *local degree* of f at x. We denote it by $\deg(f|_{x_i})$.

Proposition 6.21. *In the situation above we have*

$$\deg(f) = \sum_{i=1}^{n} \deg(f|_{x_i})$$

Example 6.22. Let $p(z) = p_n z^n + \cdots + p_0 \in \mathbb{C}[z]$ be a complex polynomial viewed as a map $\mathbb{C} \to \mathbb{C}$. It extends to a continuous map on the one-point compactification

$$f\colon S^2 \approx \mathbb{C} \cup \{\infty\} \to \mathbb{C} \cup \{\infty\} \approx S^2.$$

Let $w \in \mathbb{C}$ be such that $p'(z) \neq 0$ for all $z \in f^{-1}(w)$. This means that f is invertible around z and therefore $\deg(f|_z) = \pm 1$. In fact, since polynomials are orientation preserving, we must have $\deg(f|_z) = 1$. It follows from Proposition 6.21 that

$$\deg(f) = \#f^{-1}(w) = n = \deg(p).$$

Proof of Proposition 6.21. Consider the relative homology group $\mathrm{H}_k(S^k, S^k - f^{-1}(y))$ which by excision may be identified with

$$\mathrm{H}_k(\amalg U_i, \amalg(U_i - x_i)) \cong \oplus \mathrm{H}_k(U_i, U_i - x_i) \cong \mathrm{H}_k(S^k, S^k - x_i).$$

The composite with the obvious maps in the long exact sequences for the various pairs

$$\mathrm{H}_k(S^k) \xrightarrow{\Sigma} \oplus \mathrm{H}_k(S^k, S^k - x_i) \cong \mathrm{H}_k(S^k, S^k - f^{-1}(y)) \xrightarrow{f_*} \mathrm{H}_k(S^k, S^k - y) \xleftarrow{\sim} \mathrm{H}_k(S^k)$$

is, by definition, multiplication by $\sum \deg(f|_{x_i})$. But the composition of the first two arrows $\mathrm{H}_k(S^k) \to \mathrm{H}_k(S^k, S^k - f^{-1}(y))$ is simply the canonical map in the long exact sequence of the pair, and hence fits into the following commutative square

$$\begin{array}{ccc} \mathrm{H}_k(S^k) & \longrightarrow & \mathrm{H}_k(S^k, S^k - f^{-1}(y)) \\ \downarrow f_* & & \downarrow f_* \\ \mathrm{H}_k(S^k) & \xrightarrow{\sim} & \mathrm{H}_k(S^k, S^k - y) \end{array}$$

Unwinding the definitions, this amounts to precisely the claim. □

7 Manifolds

Many of the topological spaces one typically encounters "in nature" are manifolds. In this section we collect some topics which are specific to these nice objects. This includes foremost the notion of *orientation*. We will also see how homology is a fine enough invariant to classify low-dimensional manifolds.

7.1 Examples

We already recalled the notion of a (topological) manifold in Commentary 5.24: A Hausdorff space so that every point has an open neighborhood homeomorphic to $\mathbb{R}^k$, for some fixed k, which is called the *dimension* of the manifold.

Example 7.1.

1. Of course, $\mathbb{R}^k$ itself is a manifold of dimension k.
2. Let P, Q be the north and south poles of S^k. Stereographic projection gives homeomorphisms:

$$\pi_P : S^k \setminus \{P\} \xrightarrow{\cong} \mathbb{R}^k$$
$$\pi_Q : S^k \setminus \{Q\} \xrightarrow{\cong} \mathbb{R}^k$$

Which shows that S^k is a k-manifold.

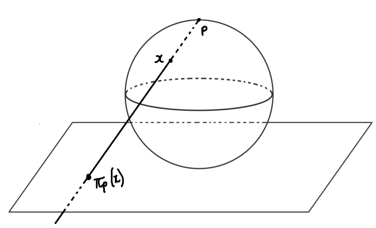

3. An open subspace of a k-manifold is itself a k-manifold.

Example 7.2. Besides $\mathbb{R}^2$ and S^2, all of $\mathbb{T}$, $\mathbb{K}$ and $\mathbb{RP}^2$ are 2-manifolds:

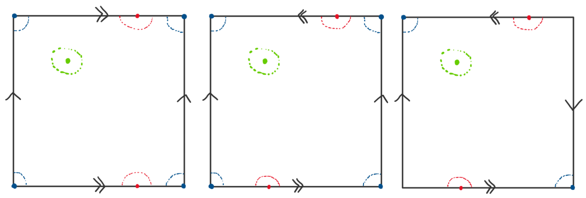

We will see many more examples of 2-manifolds in Section 7.3.

Remark 7.3. The interval $[0, 1]$ is not a manifold. Indeed, no open neighborhood of 0 (or 1) is homeomorphic to $\mathbb{R}^k$ for any k. (Instead, it is what is called a manifold with boundary. We will not discuss these in this course.)

Example 7.4. A 0-dimensional manifold is simply a discrete set.

Commentary 7.5. How does a 1-dimensional manifold M look like? Start at a point $x \in M$. Since locally around x, M looks like the real line, we may pick one of the two directions and follow it. If we ever come back to x we have found $M \approx S^1$. Otherwise we could follow the other direction and clearly could not come back to x either. For example, this could occur if $M \approx \mathbb{R}$.[14] The following exercise makes this argument more precise.

Exercise 7.6. Show that the only compact, connected 1-manifold is (homeomorphic to) S^1.

Sketch. Let M be a compact, connected 1-manifold. By assumption there are open subsets $U_1, \ldots, U_n \subseteq M$ such that $U_i \approx \mathbb{R}^1$. We may choose n to be minimal. Note that $n > 1$. (Why?) Since M is connected we may assume, after relabeling, that $U_1 \cap U_2 \neq \emptyset$. If $\pi_0(U_1 \cap U_2) = *$ then $U_1 \cup U_2 \approx \mathbb{R}^1$, contradicting minimality. The hardest bit is to show that the only other possibility is $\pi_0(U_1 \cap U_2) = * \coprod *$. In that case $M = U_1 \cup U_2 \approx S^1$. (In particular, $n = 2$.) □

7.2 Orientations

Let V be a non-zero finite-dimensional real vector space. Recall that two bases define the same orientation if the change of basis transformation from one to the other has positive determinant. This defines an equivalence relation whose two equivalence classes are the two possible orientations of V. Reversing the logic, we can think of invertible transformations with positive (resp. negative) determinant as orientation-preserving (resp. orientation-reversing). For example, one of the basic observations is that a reflection is orientation-reversing.

Since manifolds locally look like finite-dimensional vector spaces we expect that orientations can be generalized to manifolds. What could play the role of bases and change of

[14]These are in fact the only two possibilities if one assumes M to be connected and not too wild, for example second-countable. Look up the long line for a wild 1-manifold.

basis transformations in this context? The result (Lemma 6.9) that reflections on spheres have degree -1 suggests the following dictionary:

- basis $\leftrightsquigarrow$ generator of the infinite-cyclic group $\mathrm{H}_k(S^k)$
- linear transformation $\leftrightsquigarrow$ endomorphism of $\mathrm{H}_k(S^k)$
- determinant $\leftrightsquigarrow$ degree

The connection with manifolds is given by Corollary 5.25: Namely, the kth *local homology group* of a k-dimensional manifold M is always infinite-cyclic and can be identified with

$$\mathrm{H}_k(M, M\backslash x) \cong \mathrm{H}_{k-1}(S^{k-1}).$$

Definition 7.7. A *local orientation* of M at x is a choice of generator of $\mathrm{H}_k(M, M\backslash x)$.

Commentary 7.8. Here is one way to think about this. To fix our ideas let M be a 2-manifold and choose an open neighborhood $U \approx \mathbb{R}^2$ around x. The long exact sequence for the pair $(U, U\backslash x)$ induces an isomorphism

$$\mathrm{H}_2(U, U\backslash x) \xrightarrow{\sim} \mathrm{H}_1(U\backslash x) \cong \mathrm{H}_1(S^1)$$

since $U\backslash x$ deformation retracts onto a small circle around x. Choosing a local orientation ω_x at x therefore amounts to choosing in which direction to loop around this circle.

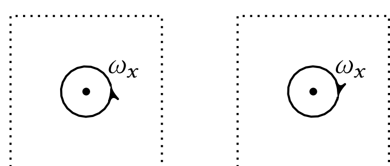

Commentary 7.9. It is clear how choosing a local orientation at $0 \in \mathbb{R}^2$ determines a local orientation at every other point $x \in \mathbb{R}^2$. We can *globally* choose the clockwise or counter-clockwise orientation. (Similarly, we can globally choose the right-handed or left-handed orientation of $\mathbb{R}^3$, etc.) The same is not true on the (open) Möbius strip, however. If we choose a local orientation at x and try to 'transport' it along a loop around the strip we end up in x with the *opposite* orientation.

We will now make a series of definitions which precisely distinguishes between these two behaviours. We will then express this by saying that the plane is *orientable*, while the open Möbius strip is not. Compared to the theory for vector spaces discussed above, non-orientability is a genuinely new concept.

Convention 7.10. Let $B \subseteq M$ be a subset of a k-manifold. We say that B is a *small open* (resp. *closed*) *ball* if it has an open neighborhood $B \subset U \approx \mathbb{R}^k$ in which it identifies with an open (resp. closed) ball of finite radius.

The point of this convention is that by excision (and where B identifies with the open ball $\mathring{B}(x, r)$ of radius r around $x \in \mathbb{R}^k$)

$$\mathrm{H}_k(M, M\backslash B) \cong \mathrm{H}_k(U, U\backslash B) \cong \mathrm{H}_k(\mathbb{R}^k, \mathbb{R}^k\backslash \mathring{B}(x, r)) \cong \mathrm{H}_{k-1}(\partial B(x, r)),$$

which is infinite-cyclic. So we may think of a generator here as an orientation of the boundary sphere of B.

Note that for every $y \in B$ we thereby get an induced local orientation through the canonical map

$$\mathrm{H}_k(M, M\backslash B) \xrightarrow{\sim} \mathrm{H}_k(M, M\backslash y).$$

We say that a family of local orientations $(\omega_y, y \in B)$ is *consistent* if there is a generator $\omega_B \in \mathrm{H}_k(M, M\backslash B)$ mapping to each ω_y.

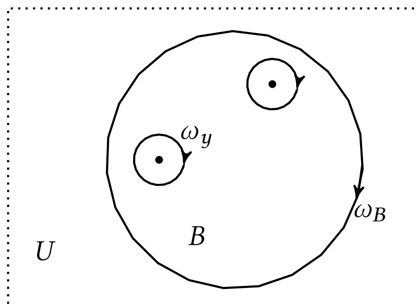

Definition 7.11.

- An *orientation* of a k-manifold M is a function $x \mapsto \omega_x \in \mathrm{H}_k(M, M\backslash x)$ assigning to $x \in M$ a local orientation, which is locally consistent in the sense above. That is, every $x \in M$ sits inside a small open ball B such that the local orientations $(\omega_y, y \in B)$ are consistent.
- M is called *orientable* if it admits an orientation. Otherwise it is called *non-orientable*.

Example 7.12. The k-sphere is orientable. Indeed, after choosing a fundamental class in $\mathrm{H}_k(S^k)$, the canonical map

$$\mathrm{H}_k(S^k) \to \mathrm{H}_k(S^k, S^k\backslash x) \tag{7.13}$$

induces local orientations at each point $x \in S^k$. And these are locally consistent since (7.13) factors through $\mathrm{H}_k(S^k, S^k\backslash B)$ for any small open ball B around x.

Construction 7.14. We now construct the *orientation bundle* $\widetilde{M}$ associated with M. Define $\widetilde{M}$ as a set to be the pairs (x, ω_x) where $x \in M$ and ω_x is a local orientation at x. It comes with an obvious map $\pi\colon \widetilde{M} \to M$ sending such a pair to x.

For the topology, if $B \subseteq M$ is a small open ball then we saw that there are exactly two collections of consistent local orientations (y, ω_y) where y ranges over the points in B. In other words, $\pi^{-1}(B) = B \amalg B$, and this gives $\widetilde{M}$ the structure of a manifold itself so that $\pi\colon \widetilde{M} \to M$ is a 2-sheeted cover.

Exercise 7.15. Show that for any manifold M, its orientation bundle $\widetilde{M}$ is orientable. In fact, it has a *canonical* orientation. Moreover, the deck transformation $(x, \omega_x) \mapsto (x, -\omega_x)$ reverses this orientation.

Example 7.16. Let M be the open Möbius strip

$$M := [0,1] \times (0,1)/\sim$$

where

$$(0,y) \sim (1, 1-y)$$

for all $y \in (0,1)$.

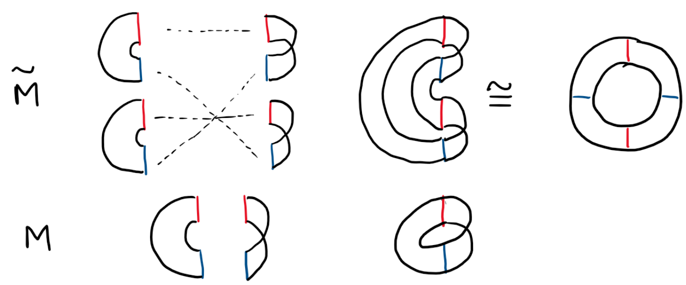

So $\widetilde{M} \cong S^1 \times (0,1)$.

Lemma 7.17. *Giving an orientation of M is equivalent to giving a continuous section to π.*

Proof. Giving a section $\omega\colon M \to \widetilde{M}$ (not necessarily continuous) amounts to choosing, for each $x \in M$, a local orientation ω_x at x. The map ω is continuous if and only if for each small open ball $B \subseteq M$, and $\pi^{-1}(B) = B \amalg B =: B_+ \amalg B_-$, the preimages $\omega^{-1}(B_+)$ and $\omega^{-1}(B_-)$ are both open in B. Since these two preimages are disjoint and jointly cover B this condition is equivalent to $\omega(B) = B_+$ or $\omega(B) = B_-$. Which means precisely that the local orientations $(\omega_y, y \in B)$ are consistent. □

Remark 7.18. Let $x \in M$ and choose a local orientation ω_x. A path in M from x to y has a unique lift to $\widetilde{M}$ starting at (x, ω_x) and ending at (y, ω_y) for some ω_y. In other words, the path determines a unique local orientation at y.

Corollary 7.19. *Assume M is a connected manifold. Then:*

- *either, $\widetilde{M} \to M$ is a non-trivial 2-sheeted cover and M is non-orientable;*
- *or, $\widetilde{M} = M \amalg M$ and M admits precisely two orientations.*

Proof. Assume $\pi\colon \widetilde{M} \to M$ is a non-trivial 2-sheeted cover, and $\omega\colon M \to \widetilde{M}$ a continuous section to π. Let $x \in M$ and $\pi^{-1}(x) = \{(x, \omega_x), (x, -\omega_x)\}$. By assumption there is a path γ in $\widetilde{M}$ from (x, ω_x) to $(x, -\omega_x)$. Then

$$\gamma, \qquad \omega\pi\gamma$$

are two paths in $\widetilde{M}$ lifting $\pi\gamma = \pi\omega\pi\gamma$ and starting at (x, ω_x). But the first path ends at $(x, -\omega_x)$ while the second one ends at (x, ω_x). This is a contradiction to uniqueness of lifting paths, see Remark 7.18. We deduce that M is non-orientable.

The second case is clear. □

Example 7.20. We deduce from Example 7.16 that the open Möbius strip is not orientable.

Corollary 7.21. *Any simply-connected manifold is orientable.*

Example 7.22. In particular, euclidean space $\mathbb{R}^k$ is orientable.

Example 7.23. Let $k \geq 1$. The quotient map $S^k \to \mathbb{RP}^k$ is the unique non-trivial two-sheeted cover of real projective space. Moreover, the non-trivial desk transformation is given by the antipodal map $x \mapsto -x$ which has degree $(-1)^{k+1}$. Hence:

- If k is even so that this degree is -1 we deduce that the desk transformation is orientation-reversing and $S^k = \widetilde{\mathbb{RP}^k}$. As S^k is connected, $\mathbb{RP}^k$ must be non-orientable in this case.
- If k is odd so that this degree is 1 we deduce that the desk transformation is orientation-preserving and $S^k \neq \widetilde{\mathbb{RP}^k}$. The latter must then be the trivial 2-sheeted cover and $\mathbb{RP}^k$ is orientable.

Remark 7.24. Suppose M is a manifold and $M \approx N$ a homeomorphism with another topological space. Then N is also a manifold (of the same dimension). Moreover, M is orientable if and only if N is orientable.

On the other hand, if M is a manifold, with $M \simeq N$ for some topological space N then N is not necessarily a manifold. And even if it is, (non-)orientability is not preserved. For example, the open Möbius strip is homotopy equivalent to S^1 but only the latter is orientable.

7.3 Surfaces: topology

In the remainder of this section we discuss a particular class of manifolds, namely surfaces. A *surface* (in this course) is a compact, connected 2-manifold. (In particular, a surface is non-empty.) After recalling the classification of surfaces[15] we will (in section 7.4) apply the 'algebraic topology' we have learnt so far to these objects. This is a particularly instructive illustration of some of the theory we've developed. Of course, the mathematics here is very old even if in the, say, 19th century it wasn't expressed in this exact language.

Example 7.25. We have seen several examples of surfaces already: S^2, $\mathbb{T}$, $\mathbb{K}$, $\mathbb{RP}^2$. We will now describe a general procedure for constructing new examples from old ones.

Construction 7.26. Let S_1 and S_2 be two surfaces and let $D_i \subseteq S_i$ be two small closed disks. We can then glue $S_1 \backslash \mathring{D}_1$ and $S_2 \backslash \mathring{D}_2$ along $\partial D_1 \approx \partial D_2$. We call the resulting space the *connected sum* $S_1 \# S_2$.

[15] You may have seen this classification in other courses before. In any case, this isn't a central part of this course which is why we will leave some results without proof.

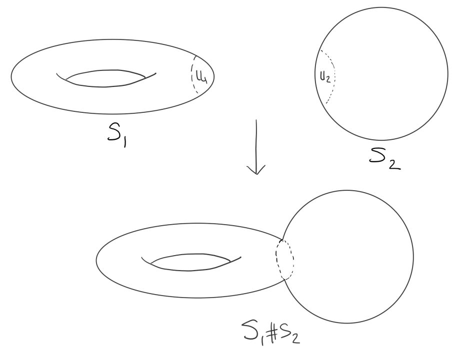

Connected sum of $S_1 = \mathbb{T}$ and $S_2 = S^2$ with $U_i := \mathring{D}_i$

Remark 7.27. It is not too hard to show that the connected sum of two surfaces is a surface. It is non-trivial but also true that the homeomorphism type of $S_1 \# S_2$ is independent of the choice of the disks D_i and the homeomorphism between their boundaries. In fact, the #-operation becomes associative, commutative and unital on the set of homeomorphism types of surfaces. As the picture in the construction above suggest, the unit for this operation is S^2.

Example 7.28. One can obtain the g-holed torus as the connected sum of g tori, $\Sigma_g = \mathbb{T} \# \cdots \# \mathbb{T}$:

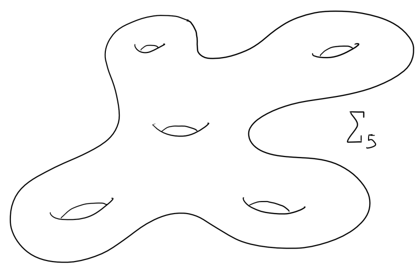

Alternatively, Σ_g is what one gets by identifying the boundary edges in a $4g$-gon according to the word

$$W_g = a_1 b_1 a_1^{-1} b_1^{-1} \cdots a_g b_g a_g^{-1} b_g^{-1} = [a_1, b_1] \cdots [a_g, b_g].$$

This follows from the next lemma.

Remark 7.29. More generally, whenever W is a word in n letters, with each letter occurring twice, one can consider the regular $2n$-gon $P(W)$ with edges labeled according to the word W. And then $M(W) = P(W)/\sim$ is the quotient obtained by identifying the corresponding oriented edges. For example, here is the $4g$-gon giving rise to Σ_g:

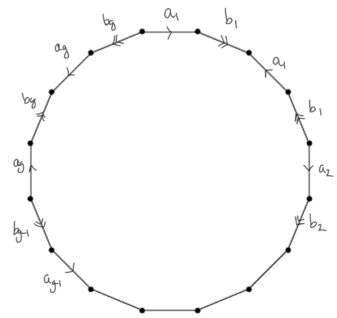

Lemma 7.30. *For words W and W' as above (with disjoint alphabets) we have $M(W)\#M(W') \approx M(WW')$ where WW' is the concatenation of the two words.*

Proof. We will not prove this but the idea is to 'lift' the operation of connected sum to the level of $2n$-gons. □

Example 7.31. The surface $N_1 := M(z^2)$ is homeomorphic to $\mathbb{RP}^2$. By the lemma, $N_h := M(z_1^2 \cdots z_h^2) \approx \mathbb{RP}^2\# \cdots \#\mathbb{RP}^2$. In Exercise 5.7 you show that $N_2 \approx \mathbb{RP}^2\#\mathbb{RP}^2 \approx \mathbb{K}$ is a familiar surface, namely the Klein bottle.

Example 7.32. The surface associated with the empty word is the 2-sphere. We can think of it as $\Sigma_0 = M(\emptyset) = M(xx^{-1})$.

Remark 7.33 (Classification of surfaces)**.** We now recall the result that each surface is homeomorphic to exactly one surface of the following list:

1. Σ_g, $g \geq 0$;[16]
2. N_h, $h > 0$.[17]

Remark 7.34. It follows that every surface is homeomorphic to a quotient of a regular polygon.

Remark 7.35. Given all we said so far, you should be wondering: what happens if I take the connected sum of an N_h and a Σ_g? After all I should end up with another surface of this list. Which one is it? It turns out that

$$\mathbb{T}\#\mathbb{RP}^2 \approx N_3,$$

which is due to Walther von Dyck (1888). So the commutative monoid of homeomorphism types of surfaces is isomorphic to

$$\langle t, r \mid t + r = 3r \rangle.$$[18]

[16] The integer g is of course the *genus* of the surface.

[17] The integer h is also called the *non-orientable genus* of the surface. We will see below that the N_h are indeed precisely the non-orientable surfaces in this list.

[18] Beware that the monoid is not cancellative. That is, from $t + r = 3r$ we cannot deduce $t = 2r$ (which is false).

7.4 Surfaces: Homology and orientation

Commentary 7.36. Recall that a compact 0-manifold is just a finite discrete set. Hence the 0th homology group classifies them (is a 'complete invariant'). Recall also that a compact 1-manifold is just a finite disjoint union of circles.[19] Hence, again, H_0 classifies them. Moreover, taking into account H_1 as well we can distinguish 1-manifolds from 0-manifolds.

In the remainder of this section, we will prove that the story continues in dimension 2. Namely, the classification of Remark 7.33 is reflected in the homology groups H_0, H_1, H_2 of the surfaces. Moreover, they are also distinguished from 1- and 0-manifolds. In summary: $H_0 \oplus H_1 \oplus H_2$ is a complete invariant for compact manifolds (up to homeomorphism) of dimension at most 2.

Commentary 7.37. One thing that is new in the 2-dimensional story is the occurrence of non-orientable manifolds. This doesn't exist in lower dimensions. On the other hand, in dimension 3, classification questions like this become much harder. And homology is not a complete invariant anymore. For example, there are many non-homeomorphic compact 3-manifolds with the homology of S^3. (As an aside, if you also assume that the manifold is simply-connected then S^3 is the only one. This is the celebrated Poincaré Conjecture (cf. Commentary 1.34) proved by Grigori Perelman in the early 2000's.)

In preparation of the homology computations below we recall from Exercise 5.10 the following notion.

Definition 7.38. Let (X_α, x_α) be a family of pointed spaces. Their *wedge sum* is the quotient space

$$\bigvee_\alpha (X_\alpha, x_\alpha) := \frac{\coprod_\alpha X_\alpha}{\coprod_\alpha \{x_\alpha\}}.$$

It is naturally a pointed space.

Example 7.39. $(S_1, *) \vee (S_1, *)$ is the figure eight: two circles touching at a point.

Lemma 7.40. *Assume that each (X_α, x_α) is a good pair. Then $\tilde{H}_n(\bigvee(X_\alpha, x_\alpha)) \cong \oplus \tilde{H}_n(X_\alpha)$.*

Proof. As in Exercise 5.10, the statement follows from Proposition 5.28 together with Proposition 3.14. □

Proposition 7.41. *The homology of the g-holed torus is*

$$H_n(\Sigma_g) = \begin{cases} \mathbb{Z} & : n = 0, 2 \\ \mathbb{Z}^{2g} & : n = 1 \\ 0 & : n > 2 \end{cases}$$

Proof. We compute its homology using the following open cover:

[19] In this discussion, we always mean: up to homeomorphism.

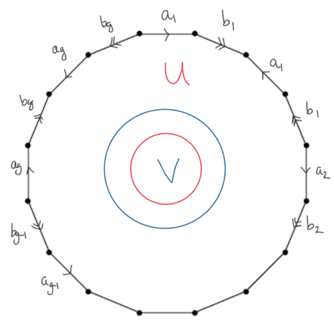

The subspace V is contractible, $U \cap V$ is homotopy equivalent to a loop γ,[20] and U deformation retracts onto the boundary, which is a wedge sum of $2g$ loops $a_1, b_1, \ldots, a_g, b_g$. By Lemma 7.40, $\tilde{\mathrm{H}}_n(U)$ is then the direct sum of the reduced homologies of the loops (which is concentrated in degree 1). We deduce that the only non-vanishing bit in the reduced MV exact sequence is

$$0 \to \tilde{\mathrm{H}}_2(\Sigma_g) \to \tilde{\mathrm{H}}_1(U \cap V) \to \tilde{\mathrm{H}}_1(U) \to \tilde{\mathrm{H}}_1(\Sigma_g) \to 0$$

where the middle map $\mathbb{Z}\gamma \to \mathbb{Z}^{2g} = \langle a_1, b_1, \ldots, a_g, b_g \rangle$ sends the generator g to

$$a_1 + b_1 - a_1 - b_1 + \cdots + a_g + b_g - a_g - b_g = 0.$$

The claim follows. □

Corollary 7.42. *The surfaces Σ_g, $g \geq 0$, are orientable.*

Proof. We use the computation of the homology in Proposition 7.41. According to that, we have $\mathrm{H}_2(\Sigma_g) = \mathbb{Z}$. We now claim that, as in Example 7.12 (for $g = 0$), the homomorphism $\mathrm{H}_2(\Sigma_g) \to \mathrm{H}_2(\Sigma_g, \Sigma_g - x)$ is an isomorphism for every point $x \in \Sigma_g$. As there, this would show that the surface Σ_g is orientable. To establish the isomorphism we look at the long exact sequence for the pair:[21]

$$\mathrm{H}_2(\Sigma_g - x) \to \mathrm{H}_2(\Sigma_g) \to \mathrm{H}_2(\Sigma_g, \Sigma_g - x) \to \mathrm{H}_1(\Sigma_g - x) \to \mathrm{H}_1(\Sigma_g) \tag{7.43}$$

Letting U be as in the proof of Proposition 7.41 the inclusion $U \hookrightarrow \Sigma_g - x$ is a homotopy equivalence. In particular, $\Sigma_g - x$ is a $2g$-fold wedge of circles and $\mathrm{H}_2(\Sigma_g - x) = 0$. Moreover, we saw in that proof that $\mathrm{H}_1(U) \xrightarrow{\sim} \mathrm{H}_1(\Sigma_g)$ hence also $\mathrm{H}_1(\Sigma_g - x) \xrightarrow{\sim} \mathrm{H}_1(\Sigma_g)$, that is, the last map in (7.43) is invertible. This shows that the third map in (7.43) is zero. Altogether the second map in (7.43) is an isomorphism as required. □

Remark 7.44. Using this proposition one can describe quite explicitly a fundamental class for Σ_g,[22] and observe the close connection between fundamental classes and orientations. (See

[20] What we mean by this is that $U \cap V \simeq S^1$ and we choose a generator $\gamma \in \mathrm{H}_1(S^1)$. Say in clockwise direction.

[21] In Example 7.12, the terms just before and after the map in question vanished since $S^k - x$ was contractible. This is not true if $g > 0$ so we need to be more careful.

[22] An alternative is to wait for Theorem 8.1 below that shows $\mathrm{H}_2(\Sigma_g) = \mathrm{H}_2^{\Delta}(\Sigma_g)$ and describe a fundamental class via a Δ-structure. We did this for the torus in Example 2.31.

the lectures.) While the argument in the Corollary (and also in Example 7.12) does not apply to other manifolds, it is still an indication of a general phenomenon. Namely, a compact, connected k-manifold M is orientable if and only if $\mathrm{H}_k(M) \cong \mathbb{Z}$. In that case, the canonical map $\mathrm{H}_k(M) \to \mathrm{H}_k(M, M-x)$ is an isomorphism for every $x \in M$, and a fundamental class for M induces an orientation. We will not prove this here but see Theorem 3.26 of Hatcher's book.

Next we turn to the other class of surfaces N_h, $h > 0$.

Proposition 7.45. *The homology of N_h is*

$$\mathrm{H}_n(N_h) = \begin{cases} \mathbb{Z} & : n = 0 \\ \mathbb{Z}^{h-1} \oplus \mathbb{Z}/2\mathbb{Z} & : n = 1 \\ 0 & : n \geq 2 \end{cases}$$

Proof. We take the same open cover as in the previous proof. Now, the boundary is a wedge of h circles and the interesting 'middle' map $\mathbb{Z}\gamma \to \mathbb{Z}^h = \langle x_1, \ldots, x_h \rangle$ in the MV exact sequence sends the generator γ to

$$2(x_1 + \cdots x_h),$$

thus the claim. □

Proposition 7.46. *The surfaces N_h, $h > 0$, are non-orientable.*

This follows from the computation in Proposition 7.45, in particular the fact that $\mathrm{H}_2(N_h) = 0$, together with the characterization of orientable compact connected manifolds mentioned in Remark 7.44. However, we can also give a direct proof as follows.

Proof. We use the observation that removing a small closed disk from $\mathbb{RP}^2$ yields a space homeomorphic to the open Möbius strip. (Presumably you show this in the course of doing Exercise 5.7. If you haven't done this yet: Think of the projective plane as obtained from the square in the usual way. Remove a disk which is cut in half by the top and bottom edges. Identify the corresponding quotient with the open Möbius strip.) It follows that for $h > 0$, the space $N_h = \mathbb{RP}^2 \# \cdots \# \mathbb{RP}^2$ contains the open Möbius strip as an open subspace. Since the latter is non-orientable (Example 7.20), we deduce that the former isn't either. □

Remark 7.47. This also gives an arguably more geometric criterion for the orientability of surfaces. Namely, it shows that a surface is non-orientable if and only if it contains an open Möbius strip.

The rest of this section is non-examinable but you might find it instructive.

We now sketch an alternative proof orientability of surfaces, based on the classification of surfaces and the following intuitive fact.

Lemma 7.48. *Let S_1, S_2 be two surfaces. Then $S_1 \# S_2$ is orientable if and only if both S_1 and S_2 are.*

Commentary 7.49. We won't give the proof of this statement here but if you would like to familiarize yourself better with both, connected sums and orientations, this would provide a good opportunity. See Exercise 6.1.

Alternative proof of Corollary 7.42 and Proposition 7.46. We know that $\mathbb{RP}^2$ is non-orientable, see Example 7.23. It follows from the lemma that every N_h is non-orientable.

Assume $\mathbb{T}$ was non-orientable. Then, by the lemma, all Σ_g would be non-orientable. And with the observation in the first paragraph we would conclude that *all* surfaces apart from the sphere are non-orientable. But consider the orientation bundle of the torus, $\widetilde{\mathbb{T}}$. As a two-sheeted cover of $\mathbb{T}$ it is itself a surface and orientable by Exercise 7.15. It follows that $\widetilde{\mathbb{T}} = S^2$. Since the latter is simply-connected it must be the universal cover, and we deduce that $\pi_1(\mathbb{T}) = \mathbb{Z}/2$. This is absurd since we know that $\pi_1(\mathbb{T}) = \mathbb{Z}^2$. □

Remark 7.50. How does one figure out the orientation bundle of the non-orientable surfaces, $\widetilde{N_h}$? You are asked to do this for the Klein bottle $\mathbb{K} \approx N_2$ in Exercise 5.7 and we will develop a very efficient tool, the *Euler characteristic*, to answer this question later in this course. But here is a picture to guide you.

Place the surface Σ_g inside $\mathbb{R}^3$ in such a way that every reflection at a coordinate hyperplane sends the surface to itself. You'll easily convince yoursef that this is indeed possible. It follows that the antipodal map $-\mathrm{id}\colon \mathbb{R}^3 \to \mathbb{R}^3$ also takes Σ_g to itself. The induced map in homology $-\mathrm{id}\colon \mathrm{H}_2(\Sigma_g) \to \mathrm{H}_2(\Sigma_g)$ is multiplication by -1. (You can see this from the identification of $\mathrm{H}_2(\Sigma_g)$ with $\mathrm{H}_1(U \cap V)$ in the proof of Proposition 7.41 and the fact that reflection of a circle ($\simeq U \cap V$) has degree -1.) If you chose the embedding $\Sigma_g \hookrightarrow \mathbb{R}^3$ in the way I'm thinking of, the antipodal map also has no fixed-points so that the quotient $\Sigma_g/(x \sim -x)$ is again a surface S. It follows that $\tilde{S} = \Sigma_g$ and hence that S is non-orientable. In fact, we will see later that $S = N_{g+1}$.

Of course, this recovers the known fact that $S^2/(x \sim -x) = \mathbb{RP}^2$.

8 Comparison

Recall that our definition of simplicial homology for Δ-complexes is still lacking justification: We haven't proven at this point that $\mathrm{H}_n^\Delta(X)$ is independent of the Δ-complex structure on X. Our first goal in this section is to show that, in fact, whatever Δ-complex structure one puts on the space X, the resulting simplicial homology is isomorphic to the singular homology of X.

Some of the ideas in the proof will recur when we introduce yet another homology theory: *cellular homology* for CW complexes. And again, we will show that cellular homology and singular homology agree. So, this section is all about comparing different homology theories.

8.1 Simplicial = singular

Let X be a topological space endowed with a Δ-complex structure $(T, |T| \approx X)$. Every n-simplex $s \in T$ gives rise to a canonical continuous map (abusively still denoted) $s\colon \Delta^n \to X$. This extends to a homomorphism $\Delta_n(T) \to C_n(X)$ and, in fact, to a chain map $\Delta_\bullet(T) \to C_\bullet(X)$ since the boundary operator is defined in the same way on both sides. Recall that we 'defined' $\mathrm{H}_n^\Delta(X) = \mathrm{H}_n(T)$ $(= \mathrm{H}_n(\Delta_\bullet(T)))$.

Theorem 8.1. *The induced map* $\mathrm{H}_n^\Delta(X) \xrightarrow{\sim} \mathrm{H}_n(X)$ *is an isomorphism.*

Proof. In other words, our goal is to show that $\mathrm{H}_n(T) \xrightarrow{\sim} \mathrm{H}_n(|T|)$ is an isomorphism.

For sub-Δ-sets $T'' \subseteq T'$ we define $\Delta_\bullet(T', T'') = \Delta_\bullet(T')/\Delta_\bullet(T'')$ and accordingly, $\mathrm{H}_n(T', T'') = \mathrm{H}_n(\Delta_\bullet(T', T''))$. By Proposition 4.37, it sits inside a familiar long exact sequence involving the simplicial homology of T' and T''. We may apply this to $T' = T^k$ and $T'' = T^{k-1}$, the Δ-sets of simplices in T of dimension at most k and $k-1$, respectively. This gives rise to a morphism of exact sequences, like so (and continuing in both directions):

$$\begin{array}{ccccccccc}
\mathrm{H}_{n+1}(T^k, T^{k-1}) & \longrightarrow & \mathrm{H}_n(T^{k-1}) & \longrightarrow & \mathrm{H}_n(T^k) & \longrightarrow & \mathrm{H}_n(T^k, T^{k-1}) & \longrightarrow & \mathrm{H}_{n-1}(T^{k-1}) \\
\downarrow & & \downarrow & & \downarrow & & \downarrow & & \downarrow \\
\mathrm{H}_{n+1}(|T^k|, |T^{k-1}|) & \to & \mathrm{H}_n(|T^{k-1}|) & \to & \mathrm{H}_n(|T^k|) & \to & \mathrm{H}_n(|T^k|, |T^{k-1}|) & \to & \mathrm{H}_{n-1}(|T^{k-1}|)
\end{array}$$

When $k = 0$, $|T^0|$ is a discrete topological space on the set T_0, and the map $\mathrm{H}_n(T^0) \to \mathrm{H}_n(|T^0|)$ is clearly an isomorphism. By induction and the five lemma, the middle vertical arrow in the diagram above is an isomorphism if the first and the fourth are. We will show this in Lemma 8.2 below. Assuming this for now, let us complete the proof as follows. (When $T = T^k$ for some k then we are already done. Only when T has simplices of arbitrarily large dimension we need to supply an additional argument.)

Let $z \in Z_n(T)$ be an n-cycle whose image in $\mathrm{H}_n(|T|)$ vanishes. In other words, there exists $\tau \in C_{n+1}(|T|)$ with $\partial\tau = z$. It is a general fact about the geometric realization of a Δ-set that every compact subspace of $|T|$ must lie in some $|T^k|$. Hence $\tau \in C_{n+1}(|T^k|)$ for some $k > n$ and we deduce that z maps to zero in $\mathrm{H}_n(|T^k|)$. By the previous argument then $z = 0 \in \mathrm{H}_n(T^k) = \mathrm{H}_n(T)$. This shows injectivity of $\mathrm{H}_n(T) \to \mathrm{H}_n(|T|)$.

Similarly, let $\sigma \in Z_n(|T|)$ be an n-cycle. As we observed before $\sigma \in Z_n(|T^k|)$ for some $k > n$ and by the argument above, $[\sigma]$ comes from $\mathrm{H}_n(T^k) = \mathrm{H}_n(T)$. This shows surjectivity of $\mathrm{H}_n(T) \to \mathrm{H}_n(|T|)$ and we are done. □

Lemma 8.2. *The canonical homomorphism* $\mathrm{H}_n(T^k, T^{k-1}) \xrightarrow{\sim} \mathrm{H}_n(|T^k|, |T^{k-1}|)$ *is an isomorphism.*

Proof. We are going to identify both sides independently with the free abelian group on the set T_k of k-simplices, and then observe that generators are sent to generators in the obvious way.

Note that $\Delta_\bullet(T^{k-1})$ is a chain complex in degrees $k-1, k-2, \ldots, 0$. And $\Delta_\bullet(T^k)$ is the same chain complex with an additional term $\mathbb{Z}T_k$ in degree k. It follows that $\Delta_\bullet(T^k, T^{k-1})$ is the chain complex with $\mathbb{Z}T_k$ concentrated in degree k. In particular, we have:

$$\mathrm{H}_n(T^k, T^{k-1}) = \begin{cases} \mathbb{Z}T_k & : n = k \\ 0 & : n \neq k \end{cases} \tag{8.3}$$

By construction of the geometric realization, we have homeomorphisms

$$\frac{|T^k|}{|T^{k-1}|} \approx \frac{\Delta^k \times T_k}{\partial\Delta^k \times T_k} \approx \bigvee_{T_k} \frac{\Delta^k}{\partial\Delta^k}.$$

We deduce from Lemma 7.40 and Example 5.31 that[23]

$$\mathrm{H}_n(|T^k|, |T^{k-1}|) \cong \tilde{\mathrm{H}}_n(|T^k|/|T^{k-1}|) \cong \begin{cases} \mathbb{Z}T_k & : n = k \\ 0 & : n \neq k \end{cases} \tag{8.4}$$

with generators corresponding to elements $s \in T_k$ via the relative cycles $s \colon \Delta^k \to |T^k|$. It follows that the map in the statement identifies the groups in (8.3) and (8.4). □

Corollary 8.5. *The simplicial homology $\mathrm{H}^{\Delta}_{\bullet}(X)$ depends on X only (and not on the Δ-complex structure).*

Corollary 8.6. *Suppose X has a Δ-complex structure with simplices in dimension $\leq k$ only. Then $\mathrm{H}_n(X) = 0$ for all $n > k$.*

Example 8.7. Recall (Corollary 5.9) our construction of a fundamental class for the sphere S^k. We now describe an alternative approach, using Theorem 8.1.

Namely, S^k has a Δ-complex structure obtained by glueing two copies of Δ^k along the boundary. Thus

$$\Delta_k(S^k) = \mathbb{Z}^2 = \langle \sigma_+, \sigma_- \rangle, \qquad \Delta_{k-1}(S^k) = \mathbb{Z}^{k+1} = \langle \tau_0, \dots, \tau_k \rangle$$

with $\partial(\sigma_\pm) = \sum (-1)^i \tau_i$. It follows that

$$\mathrm{H}_k(S^k) \cong \mathrm{H}^{\Delta}_k(S^k) = \mathbb{Z} = \langle \sigma_+ - \sigma_- \rangle.$$

Remark 8.8. Similarly, one could construct fundamental classes for the orientable surfaces Σ_g using a Δ-complex structure, cf. Remark 7.44.

8.2 CW complexes

Recall that a *CW complex* is a topological space X obtained as follows:

1. Start with a discrete space X^0, the 0-*cells*.
2. Inductively, the *n-skeleton* X^n is obtained from X^{n-1} by attaching *n-cells* D^n_α along maps $\phi_\alpha \colon \partial D^n_\alpha = S^{n-1}_\alpha \to X^{n-1}$.
3. If $X = X^k$ and k is minimal with this property then X is *of dimension* k. More generally we can have $X = \cup_n X^n$, in which case a subspace $U \subseteq X$ is open iff $U \cap X^n \subseteq X^n$ is open for all n.

The ϕ_α are called *attaching maps*. Since $X^n = X^{n-1} \cup_{\amalg \phi_\alpha} (\amalg_\alpha D^n_\alpha)$, the attaching map extends to a *characteristic map* $\Phi_\alpha \colon D^n_\alpha \to X^n \to X$.

Remark 8.9. As for Δ-complexes, this terminology is abusive. More accurately, we described what a CW complex structure on a space X consists of. There can be many such structures.

[23] We use Exercise 6.3 to argue that these are good pairs.

Example 8.10. The sphere S^k (for $k > 0$) admits a CW complex structure with a single 0-cell $*$ and a single k-cell. The attaching map $\partial D^k = S^{k-1} \to *$ is the unique map.

Commentary 8.11. Every Δ-complex is a CW complex. The main difference between the two is that the attaching maps $\phi_\alpha\colon \partial D^n_\alpha \to X^{n-1}$ for CW complexes are allowed to be any continuous map while for Δ-complexes, they restrict to the inclusion of an $(n-1)$-cell for each face of the n-simplex.

Example 8.10 shows how the greater flexibility allows for more parsimonious cell structures.

Example 8.12. A 1-dimensional CW complex is the same thing as a 1-dimensional Δ-complex. Both of these can be identified with graphs.

Example 8.13. The torus $\mathbb{T}$ admits a CW complex structure with a single 0-cell to which one attaches two loops (1-cells) a, b, and finally one 2-cell via the attaching map described by the loop $[a, b] = aba^{-1}b^{-1}$.

Example 8.14. We may think of real projective k-space $\mathbb{RP}^k$ as the quotient $S^k/(x \sim -x)$ of the sphere under the antipodal map. Alternatively, it is the quotient of one of the two hemispheres D^k with antipodal points on the boundary $\partial D^k = S^{k-1}$ identified. In other words, $\mathbb{RP}^k$ is obtained by attaching a k-cell to $\mathbb{RP}^{k-1}$ along the canonical quotient map $S^{k-1} \to \mathbb{RP}^{k-1}$. Inductively, it follows that $\mathbb{RP}^k$ has a CW structure with one cell in each dimension $0, 1, \ldots, k$.

Example 8.15. We don't need to stop the process in the previous example. Continuing we get the infinite union $\mathbb{RP}^\infty = \cup_k \mathbb{RP}^k$ as a CW complex with a single cell in each dimension.

More examples of CW complexes can be found on the exercise sheet 6.

8.3 Cellular homology

Commentary 8.16. Our next goal is to define a homology theory for CW complexes, similarly to how we defined simplicial homology in Section 2. Of course, we would like to take as n-chains $C_n^{\mathrm{CW}}(X)$ the free abelian group on the n-cells. But how to define the boundary operator? The greater flexibility in attaching maps also means that there isn't an obvious way to say what the 'oriented boundary' of an n-cell is.

We will proceed somewhat differently. In this section, we define the *cellular chain complex* of a CW complex in terms of relative singular homology. The definition is quite abstract but:

- Lemma 8.17 shows that its degree-n term is the free abelian group on the n-cells;
- Corollary 8.25 shows that the cellular chain complex computes singular homology.
- In the next section will we identify the boundary operator and thereby make cellular homology an eminently computable theory.

Lemma 8.17. *Let X be a CW complex with n-cells $\{D^n_\alpha\}$, $n > 0$. Then*

$$\mathrm{H}_m(X^n, X^{n-1}) \cong \begin{cases} \oplus_\alpha \mathbb{Z} \cdot [D^n_\alpha] & : m = n \\ 0 & : m \neq n \end{cases}$$

Proof. This is exactly as in the proof of Lemma 8.2:

$$\mathrm{H}_m(X^n, X^{n-1}) \cong \tilde{\mathrm{H}}_m(X^n/X^{n-1}) \cong \tilde{\mathrm{H}}_m(\vee_\alpha D^n_\alpha/\partial D^n_\alpha) \cong \oplus_\alpha \tilde{\mathrm{H}}_m(D^n_\alpha/\partial D^n_\alpha)$$

□

Remark 8.18. We may describe the isomorphism in the lemma explicitly as follows. Choose a homeomorphism $\Delta^n \approx D^n$. We then get a continuous map

$$\Delta^n \approx D^n \xrightarrow{\Phi_\alpha} X^n$$

which is in fact a relative cycle for the pair (X^n, X^{n-1}). Its relative homology class generates the copy of $\mathbb{Z}$ corresponding to D^n_α.

The following is the main technical lemma we need to define the cellular chain complex and identify its homology with singular homology. You could compare it with simplicial homology for Δ-complexes where the corresponding statements are obviously true.

Lemma 8.19. *Let X be a CW complex.*

1. *If X is of dimension k then $\mathrm{H}_n(X) = 0$ for all $n > k$.*
2. *The map $\mathrm{H}_n(X^m) \to \mathrm{H}_n(X)$ induced by the inclusion $X^m \to X$ is an isomorphism if $n < m$ and surjective if $n = m$.*

Proof. 1. We do induction on k, with $k = 0$ being clear (Example 3.9). Assuming the statement for $k - 1$ let us consider the long exact sequence for the pair (X^k, X^{k-1}):

$$\cdots \to \mathrm{H}_{n+1}(X^k, X^{k-1}) \to \mathrm{H}_n(X^{k-1}) \to \mathrm{H}_n(X^k) \to \mathrm{H}_n(X^k, X^{k-1}) \to \cdots \tag{8.20}$$

By induction, the second term vanishes for $n > k > k - 1$. By Lemma 8.17, the last term vanishes for $n > k$. Therefore the third term vanishes for $n > k$.

2. Consider again the long exact sequence (8.20) for the pair above. When

   - $n < k - 1$, both outer terms vanish and the middle arrow $\mathrm{H}_n(X^{k-1}) \xrightarrow{\sim} \mathrm{H}_n(X^k)$ is an isomorphism;
   - $n < k$, the last term vanishes and the middle arrow $\mathrm{H}_n(X^{k-1}) \twoheadrightarrow \mathrm{H}_n(X^k)$ is surjective.

   In particular, for $n < m$ we have isomorphisms

$$\mathrm{H}_n(X^m) \xrightarrow{\sim} \mathrm{H}_n(X^{m+1}) \xrightarrow{\sim} \mathrm{H}_n(X^{m+2}) \xrightarrow{\sim} \cdots \tag{8.21}$$

   A compactness argument as in the proof of Theorem 8.1 shows that these groups are necessarily $\mathrm{H}_n(X)$.

   On the other hand, if $n = m$, then the first arrow in (8.21) is (in general) only surjective while the remaining ones are isomorphisms. This completes the proof.

□

We will now contemplate the following diagram that is 'spliced together' from the long exact sequences for the various pairs and which define the horizontal arrows in red.

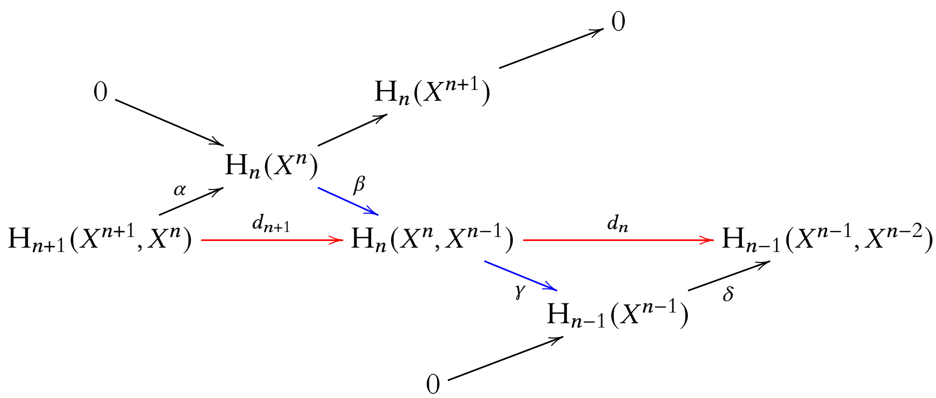

Note that the various zeros come from Lemma 8.19.

Lemma 8.22. *1. The composition $d_n \circ d_{n+1}$ is zero.*

2. The quotient group $\ker(d_n)/\mathrm{img}(d_{n+1})$ is canonically isomorphic to $H_n(X)$.

Proof. 1. The composition $d_n \circ d_{n+1}$ factors through the two blue arrows in the diagram. These are part of a long exact sequence hence their composite is zero.

2. Since δ is injective, we have

$$\ker(d_n) = \ker(\gamma) = \mathrm{img}(\beta). \tag{8.23}$$

We also know by Lemma 8.19 that $H_n(X^{n+1}) \cong H_n(X)$ canonically. That is,

$$H_n(X) \cong \mathrm{coker}(\alpha) \cong \mathrm{img}(\beta)/\mathrm{img}(\beta \circ \alpha) \overset{(8.23)}{=} \ker(d_n)/\mathrm{img}(d_{n+1}),$$

where the second isomorphism comes from the fact that β is injective.

□

Definition 8.24. Let X be a CW complex. We define the *cellular chain complex* $C_\bullet^{CW}(X)$ so that

$$\cdots \to C_{n+1}^{CW}(X) = H_{n+1}(X^{n+1}, X^n) \xrightarrow{d_{n+1}} C_n^{CW}(X) = H_n(X^n, X^{n-1}) \xrightarrow{d_n} \cdots$$

The *cellular homology* groups are the homology of this chain complex:

$$H_n^{CW}(X) = H_n(C_\bullet^{CW}(X))$$

Corollary 8.25. *For any CW complex X there are canonical isomorphisms $H_n^{CW}(X) \cong H_n(X)$.*

8.4 Computing cellular homology

By Lemma 8.22, we know that cellular homology theoretically computes singular homology for CW complexes. Our goal now is to make this a practical tool.

Remark 8.26. Let X be a CW complex so that for $n > 0$

$$X^n = X^{n-1} \cup_{\phi_\alpha} (\amalg_\alpha D^n_\alpha)$$

for the attaching maps $\phi_\alpha \colon \partial D^n_\alpha \to X^{n-1}$. Note that we have[24]

$$X^n/X^{n-1} \approx \vee_\alpha D^n_\alpha/\partial D^n_\alpha = \vee_\alpha S^n_\alpha, \tag{8.27}$$

and in particular there are canonical quotient maps, obtained by collapsing all but one sphere to a point:

$$\pi_\alpha \colon X^n \twoheadrightarrow X^n/X^{n-1} \twoheadrightarrow S^n_\alpha$$

Remark 8.28. By Lemma 8.17, the cellular chain complex in degree n is the free abelian group on the n-cells. So, it remains to identify the boundary operator in terms of a chosen basis. Recall that a basis $([D^n_\alpha])_\alpha$ for $C^{\mathrm{CW}}_n(X)$ may be obtained by choosing a generator $[D^n] \in \mathrm{H}_n(D^n, \partial D^n)$:

$$[D^n_\alpha] = (\Phi_\alpha)_*([D^n]) \in \mathrm{H}_n(X^n, X^{n-1}) \tag{8.29}$$

Proposition 8.30. *There are generators $[D^n]$ such that in terms of the basis (8.29) for the groups $C^{\mathrm{CW}}_\bullet(X)$, the boundary operator is given by the following formula:*

- *in degree $n = 1$:*
$$d_1([D^1_\alpha]) = [\phi_\alpha(1)] - [\phi_\alpha(-1)]$$
- *in degrees $n > 1$:*
$$d_n([D^n_\alpha]) = \sum_\beta d_{\alpha\beta}[D^{n-1}_\beta],$$
where
$$d_{\alpha\beta} = \deg\left(\Delta_{\alpha\beta} \colon S^{n-1}_\alpha \xrightarrow{\phi_\alpha} X^{n-1} \xrightarrow{\pi_\beta} S^{n-1}_\beta\right)$$

Remark 8.31. If one changes the generator $[D^n]$ (for all α simultaneously!) then the formula for the boundary operator changes by a sign. Of course, this does not affect the homology of the complex so in practice this is not so important. In the sequel we will not dwell on this point.

Before giving the proof let us do some examples.

Example 8.32. Whenever there is a single 0-cell, the first boundary operator $d_1 \colon C^{\mathrm{CW}}_1(X) \to C^{\mathrm{CW}}_0(X)$ is the zero map.

Example 8.33. Let us consider the 'minimal' CW complex structure on the sphere S^k, $k > 0$, as in Example 8.10. The cellular chain complex has a copy of the integers in degrees 0 and k. If $k = 1$, the boundary operator $d_1 = 0$ by the previous example. If $k > 1$, then all boundary operators are clearly zero as well. We deduce that $\mathrm{H}^{\mathrm{CW}}_\bullet(S^k) = C^{\mathrm{CW}}_\bullet(S^k)$ has two copies of the integers, in degrees 0 and k. (Of course, we knew that already.)

[24] To avoid any ambiguities in the sequel we fix once and for all homeomorphisms $D^n/\partial D^n \approx S^n$ for all $n > 0$.

Example 8.34. Let $X = \Sigma_g$ be the orientable surface of genus $g \geq 0$. It admits a CW complex structure with a single 0-cell, $2g$ 1-cells $a_1, b_1, \ldots, a_g, b_g$, and a single 2-cell with attaching map $\gamma\colon \partial D^2 = S^1 \to X^1$ given by the loop $[a_1, b_1] \cdots [a_g, b_g]$. The map $S^1 \xrightarrow{\gamma} X^1 \xrightarrow{\pi_\beta} S^1_\beta$ is homotopic to the constant map since each a_i and b_i appears together with its inverse. It follows that the cellular chain complex has the shape:

$$0 \to \mathbb{Z} \xrightarrow{0} \mathbb{Z}^{2g} \xrightarrow{0} \mathbb{Z} \to 0$$

We deduce that the cellular homology groups $\mathrm{H}^{\mathrm{CW}}_\bullet(\Sigma_g)$ coincide with those found in Proposition 7.41.

Exercise 8.35. Repeat this computation for the non-orientable surface N_h, cf. Exercise 6.4.

Example 8.36. Let us compute the cellular homology of $\mathbb{RP}^k$, with the CW complex structure of Example 8.14. Thus it has a single cell in each dimension $n = 0, \ldots, k$. To determine the boundary operator d_n we consider—as we should—the map

$$f\colon S^{n-1} \xrightarrow{2:1} \mathbb{RP}^{n-1} \twoheadrightarrow \mathbb{RP}^{n-1}/\mathbb{RP}^{n-2} \approx S^{n-1},$$

and use the local degree formula (Proposition 6.21) to compute its degree. At the 'North pole' N (importantly, away from the equator), the map is locally a homeomorphism hence has degree ± 1.[25] At the 'South pole' S, the map is also locally a homeomorphism, namely the antipodal map followed by the one at N. Since the antipodal map has degree $(-1)^n$, it follows that

$$\deg(f) = \deg(f|_N) + (-1)^n \deg(f|_N) = \begin{cases} \pm 2 & : n \text{ even} \\ 0 & : n \text{ odd} \end{cases}$$

We conclude that the cellular chain complex looks as follows:

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

with $d_k = \pm 2$ if k is even, and $d_k = 0$ if k is odd. We conclude that

$$\mathrm{H}_n(\mathbb{RP}^k) = \begin{cases} \mathbb{Z} & : n = 0, \text{ and } n = k \text{ if odd} \\ \mathbb{Z}/2\mathbb{Z} & : 0 < n < k \text{ odd} \\ 0 & : \text{else} \end{cases}$$

Example 8.37. It follows from the computations in the previous example that

$$\mathrm{H}_n(\mathbb{RP}^\infty) = \begin{cases} \mathbb{Z} & : n = 0 \\ \mathbb{Z}/2 & : 0 < n \text{ odd} \\ 0 & : \text{else} \end{cases}$$

[25]Pedant remark (which you should probably ignore): The sign of the degree actually depends on the chosen homeomorphism $D^{n-1}/\partial D^{n-1} \approx S^{n-1}$ which we have fixed throughout the discussion.

Proof is non-examinable

Proof of Proposition 8.30. We choose the 'obvious' generator $[D^0]$ of $\mathrm{H}_0(D^0, \emptyset) = \mathbb{Z}$ given by the unique simplex $\Delta^0 \to D^0 = *$. Next we choose a homeomorphism $\Delta^1 \approx D^1$ which gives a path *from* -1 *to* 1. As in Remark 8.18, this gives a generator $[D^1] \in \mathrm{H}_1(D^1, \partial D^1)$. Inductively, we define $[D^n]$ (for $n > 1$) as the generator corresponding to $[D^{n-1}] \in \mathrm{H}_{n-1}(D^{n-1}, \partial D^{n-1})$ under the sequence of isomorphisms

$$\mathrm{H}_n(D^n, S^{n-1}) \xrightarrow{\partial} \mathrm{H}_{n-1}(S^{n-1}) \cong \mathrm{H}_{n-1}(D^{n-1}/S^{n-2}) \cong \mathrm{H}_{n-1}(D^{n-1}, S^{n-2}). \tag{8.38}$$

Here, the first isomorphism is the connecting homomorphism in the long exact sequence for the pair, the second isomorphism is induced by our chosen homeomorphism $S^{n-1} \approx D^{n-1}/\partial D^{n-1}$, and the last isomorphism is an instance of Proposition 5.28.

With this chosen basis in each dimension we now verify the formula for the boundary operator in the statement:

In degree $n = 1$, the boundary operator is given by the connecting homomorphism $\mathrm{H}_1(X^1, X^0) \to \mathrm{H}_0(X^0)$. Start with a 1-cell D^1_α. The 'characteristic map'

$$\tilde{\Phi}_\alpha \colon \Delta^1 \approx D^1_\alpha \xrightarrow{\Phi_\alpha} X^1$$

is a relative cycle and, by Exercise 5.19, the connecting homomorphism takes this to the class of $\partial\tilde{\Phi}_\alpha = \phi_\alpha(1) - \phi_\alpha(-1)$, by our choice of the homeomorphism $\Delta^1 \approx D^1$. We deduce that

$$d_1([D^n_\alpha]) \overset{(8.29)}{=} d_1\Big((\tilde{\Phi}_\alpha)_*([\Delta^1])\Big) = d_1([\tilde{\Phi}_\alpha]) = [\phi_\alpha(1)] - [\phi_\alpha(-1)].$$

In degree $n > 1$, we consider the following diagram:

$$\begin{array}{ccccccccc}
\mathrm{H}_n(D^n_\alpha, \partial D^n_\alpha) & \xrightarrow[\sim]{\partial} & \mathrm{H}_{n-1}(S^{n-1}_\alpha) & \xrightarrow{(\Delta_{\alpha\beta})_*} & \mathrm{H}_{n-1}(S^{n-1}_\beta) & \overset{\sim}{\text{——}} & \mathrm{H}_{n-1}(D^{n-1}_\beta/\partial D^{n-1}_\beta) & \xleftarrow{\sim} & \mathrm{H}_{n-1}(D^{n-1}_\beta, \partial D^{n-1}_\beta) \\
\downarrow{\scriptstyle (\Phi_\alpha)_*} & & \downarrow{\scriptstyle (\phi_\alpha)_*} & \nearrow{\scriptstyle (\pi_\beta)_*} & & & \uparrow & & \\
\mathrm{H}_n(X^n, X^{n-1}) & \xrightarrow{\partial} & \mathrm{H}_{n-1}(X^{n-1}) & \rightarrow & \mathrm{H}_{n-1}(X^{n-1}, X^{n-2}) & \xrightarrow{\sim} & \mathrm{H}_{n-1}(X^{n-1}/X^{n-2}) & &
\end{array}$$

It is clearly commutative. Note that the composite of the two red arrows is the boundary operator d_n. Now, start with $[D^n]$ in the top left. Following the path down-right-up-right all the way to the top right, we get a multiple of $[D^{n-1}_\beta]$, the coefficient of $d_n([D^n_\alpha]$ for the basis element $[D^{n-1}_\beta]$. On the other hand, following the path straight to the right we get $d_{\alpha\beta} \cdot [D^{n-1}_\beta]$, by (8.38). □

9 Thr Euler characteristic

We often measure mathematical objects by their 'size': sets by their cardinality, vector spaces by their dimension, abelian groups by their rank etc. When you try to define an analogous notion of size in algebraic topology, you end up with the Euler characteristic. Even more than the homology groups it is a way of breaking down the complexity of a space, in this case to a single number (okay, integer). Given this, it is quite surprising how powerful this notion turns out to be.

9.1 Definition

Euler didn't speak of 'Euler characteristics', obviously, but he was one of the first to systematically apply this notion. Let us recall the formula he found (Proposition 9.4).

Definition 9.1. A *plane graph* is a finite 1-dimensional CW complex embedded in the real plane $\mathbb{R}^2$. Equivalently, it is a finite graph in the plane in which the edges don't cross.

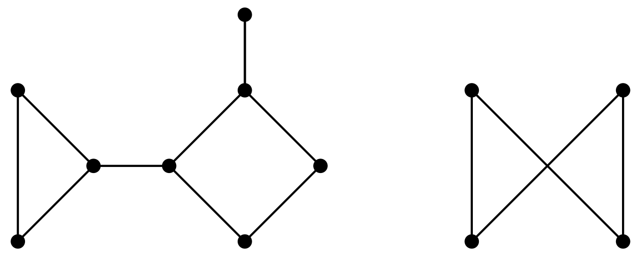

a plane graph   not a plane graph

Commentary 9.2. There are finite graphs that admit no such an embedding in $\mathbb{R}^2$, and these are called *nonplanar*. (Note that the example on the right above *can* be embedded in $\mathbb{R}^2$, as a square, hence it is *planar*.) An example of a nonplanar graph is the complete graph K_5 on five vertices. Another one is the complete bipartite graph $K_{3,3}$:

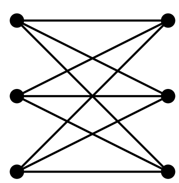

There are theoretical criteria as well as practical algorithms for deciding whether a graph is planar.

Remark 9.3. A *face* of a plane graph $\mathcal{G}$ is a connected component of $\mathbb{R}^2 \backslash \mathcal{G}$:

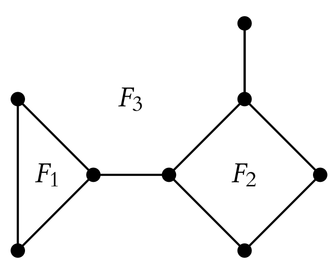

Recall now:

Proposition 9.4 (Euler's formula)**.** *For a planar graph with v vertices, e edges and f faces we have*

$$v - e + f = 2.$$

Example 9.5. Recall the complete bipartite graph $K_{3,3}$. We have $v = 6$ and $e = 9$. Moreover, every face has an even number of boundary edges, that is, at least 4. So $4f \leq 2e = 18$, or $f \leq 4$. Hence

$$v - e + f \leq -3 + 4 = 1.$$

Euler's formula shows that $K_{3,3}$ cannot be embedded in $\mathbb{R}^2$, that is, $K_{3,3}$ is nonplanar.

Commentary 9.6. The argument for Proposition 9.4 we are going to give will link the situation with homology and introduce several important ideas. That's not to say that Euler thought of it in this way.

Remark 9.7. By adding a point at ∞, a planar graph yields a CW complex structure on the 2-sphere, with v many 0-cells, e 1-cells and f 2-cells. Right?

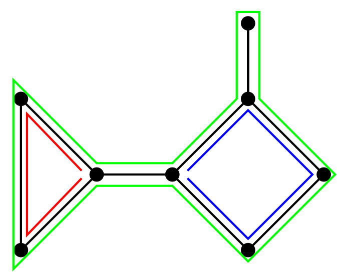

The discussion around the Jordan curve theorem should have made you wary of such a claim. And rightly so: What if the edges aren't as nice? Is it still true that they form the boundaries of 2-cells?

The answer is yes, but this is (a special case of) a non-trivial result called the *Schoenflies' Theorem*. (Which we are going to assume without proof. Alternatively, we give a proof of Euler's formula for those plane graphs that underlie a CW structure on S^2.) So, Euler's formula will follow from the following statement since the minimal CW complex structure on S^2 has a single 0-cell and a single 2-cell (Example 8.10).

Proposition 9.8. *Let X be a topological space that admits the structure of a finite CW complex. Then the alternating sum*

$$\chi(X) = \sum_{n \geq 0} (-1)^n \#\{n\text{-cells}\}$$

is independent of the choice of CW complex structure.

Remark 9.9. The alternating sum $\chi(X)$ is called the *Euler characteristic* of X. We will define it in greater generality below.

Definition 9.10. Let A be a finitely generated abelian group. It decomposes as $A = T \oplus F$, the direct sum of the torsionfree part $T \cong \mathbb{Z}^r$ (which is necessarily free abelian) and a finite abelian group F. The integer r is well-defined and called the *rank* of A.

Alternatively, $\mathrm{rk}(A)$ is the dimension of the $\mathbb{Q}$-vector space $A \otimes_{\mathbb{Z}} \mathbb{Q}$. This follows from the fact that $\mathbb{Z}^r \otimes \mathbb{Q} = \mathbb{Q}^r$ while $F \otimes \mathbb{Q} = 0$.

Exercise 9.11. Here is a warm-up exercise for you:

1. Let $0 \to V_1 \to V_2 \to V_3 \to 0$ be a short exact sequence, with V_i finite dimensional $\mathbb{Q}$-vector spaces. Then $\dim(V_0) - \dim(V_2) + \dim(V_3) = 0$.
2. Let $\cdots \to C_{n+1} \xrightarrow{d_{n+1}} C_n \xrightarrow{d_n} C_{n-1} \to \cdots$ be an exact chain complex. Then its $\mathbb{Q}$-linearization $\cdots \to C_{n+1} \otimes \mathbb{Q} \xrightarrow{d_{n+1}} C_n \otimes \mathbb{Q} \xrightarrow{d_n} C_{n-1} \otimes \mathbb{Q} \to \cdots$ is still exact.
3. Conclude that for a short exact sequence $0 \to A_1 \to A_2 \to A_3 \to 0$ of finitely generated abelian groups we have
$$\mathrm{rk}(A_1) - \mathrm{rk}(A_2) + \mathrm{rk}(A_3) = 0.$$

Corollary 9.12. *Let $C_\bullet$ be a chain complex with only finitely many non-zero terms, all of which are finitely generated abelian groups. Then*

$$\sum_n (-1)^n \,\mathrm{rk}(C_n) = \sum_n (-1)^n \,\mathrm{rk}(\mathrm{H}_n(C_\bullet))$$

Proof. We have two short exact sequences, for all n:

$$0 \to Z_n \to C_n \to B_{n-1} \to 0$$
$$0 \to B_n \to Z_n \to \mathrm{H}_n \to 0$$

We therefore get

$$\begin{aligned}\sum_n (-1)^n \,\mathrm{rk}(C_n) &= \sum_n (-1)^n [\mathrm{rk}(Z_n) + \mathrm{rk}(B_{n-1})] && \text{Exercise 9.11}\\ &= \sum_n (-1)^n [\mathrm{rk}(Z_n) - \mathrm{rk}(B_n)] \\ &= \sum_n (-1)^n \,\mathrm{rk}(\mathrm{H}_n) && \text{Exercise 9.11}\end{aligned}$$

□

Proof of Proposition 9.8. By the Corollary, $\chi(X) = \sum_n (-1)^n \,\mathrm{rk}(\mathrm{H}_n(X))$ and the right-hand side is independent of the cell structure. □

The proof suggests the following definition.

Definition 9.13. Let X be a space with only finitely many non-zero homology groups, all of which are finitely generated abelian groups. Then its *Euler characteristic* is defined as

$$\chi(X) := \sum_{n \geq 0} (-1)^n \,\mathrm{rk}(\mathrm{H}_n(X)).$$

Remark 9.14. To emphasize, if X is a finite CW complex then Corollary 9.12 shows that

$$\chi(X) = \sum_n (-1)^n \,\mathrm{rk}(\mathrm{H}_n) = \sum_n (-1)^n \#\{n\text{-cells}\},$$

so this is compatible with the notation earlier in Proposition 9.8. As we will see below, going back and forth between these two expressions can be a powerful tool.

Example 9.15. As we observed above, $\chi(S^2) = 2$. More generally, we have

$$\chi(S^k) = \begin{cases} 2 & : k \text{ even} \\ 0 & : k \text{ odd} \end{cases}$$

This can be seen

- either from the Definition 9.13 and our homology computation in Corollary 4.13,
- or from the 'minimal' cell structure Example 8.10.

Note that although there are many CW complex structures on S^k none of it has an odd number of cells!

Remark 9.16. Let $C_\bullet$ be a chain complex that has only finitely many non-zero terms, all of which are finitely generated abelian groups. If $C_\bullet$ is exact then $\sum_n (-1)^n \chi(C_n) = 0$. This follows from Corollary 9.12.

Example 9.17. Consider a CW structure on S^2 made of v vertices, e edges, and 2-cells made of p pentagons and h hexagons. For example, consider a football:

We have

- $2e = 6h + 5p$,
- $3v = 2e$,
- $2 = \chi(S^2) = v - e + p + h$.

Putting these together we get

$$6p = 12 + 6e - 6v - 6h = 12 + 2e - 6h = 12 + 6h + 5p - 6h = 12 + 5p$$

that is, $p = 12$.

Note that for the football pictured above one has $h = 20$. However, the regular dodecahedron has $h = 0$, for example.

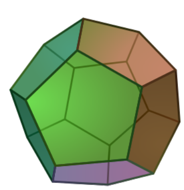

Example 9.18. Let us come back to the complete bipartite graph $K_{3,3}$. As we showed in Example 9.5, any drawing would need $f \leq 4$, and the graph is therefore nonplanar. However, it *can* be drawn on $\mathbb{T}$, $\mathbb{K}$ with $f = 3$, and on $\mathbb{RP}^2$ with $f = 4$:

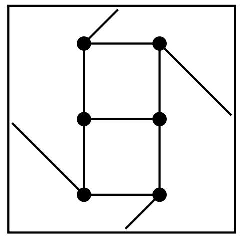

So

$$v - e + f = \begin{cases} 0 & \text{on } \mathbb{T} \text{ or } \mathbb{K} \\ 1 & \text{on } \mathbb{RP}^2 \end{cases}$$

These are the Euler characteristic of the three spaces as we know from either their cell structure (for example, as Δ-complexes) or the homology we computed earlier.

See also Exercise 7.4.

9.2 Properties

In this section we prove many of the properties one would expect a measure of size to satisfy. To get an intuition for what these are note that any set X can be viewed as a topological space with the discrete topology. If X is finite then $\chi(X) = \mathrm{rk}(\mathrm{H}_0(X)) = |X|$ is its cardinality. In the case of finite discrete spaces, some of the properties we will prove reduce to the following trivialities:

- $|X \cup Y| = |X| + |Y| - |X \cap Y|$
- $|X \times Y| = |X| \cdot |Y|$
- If $f : Y \to X$ is a $d : 1$-map then $|Y| = d \cdot |X|$.

Proposition 9.19. *Let $X = U \cup V$ and assume that*

1. *either, X is a CW complex and U, V are subcomplexes,*
2. *or $U, V \subseteq X$ are open subsets.*

If $\chi(U)$, $\chi(V)$ and $\chi(U \cap V)$ are all defined then so is $\chi(X)$ and we have

$$\chi(X) = \chi(U) + \chi(V) - \chi(U \cap V).$$

Proof. In the first case, if U and V are finite CW complexes then so is X. Looking at the number c_n of cells in each dimension n we have

$$c_n(X) = c_n(U) + c_n(V) - c_n(U \cap V)$$

which implies the identity in the statement.

For the second case look at the Mayer-Vietoris long exact sequence and use that the alternating sum of the ranks is 0, by Remark 9.16. □

Example 9.20. Recall the connected sum $S_1\#S_2$ of surfaces (section 7.3). We have for a small disk $D \subseteq S_i$,

$$\chi(S_i) = \chi(D) + \chi(S^2\backslash\mathring{D}) - \chi(S^1)$$

so that $\chi(S^2\backslash\mathring{D}) = \chi(S_i) - 1$. It follows that

$$\begin{aligned}\chi(S_1\#S_2) &= \chi(S_1\backslash\mathring{D}) + \chi(S_2\backslash\mathring{D}) - \chi(S^1)\\ &= \chi(S_1) + \chi(S_2) - 2.\end{aligned}$$

(To be sure, we are applying Proposition 9.19 not directly to the (sub)spaces $S_i\backslash\mathring{D}$, D, or S^1 but rather to suitable open neighborhoods of these.)

Corollary 9.21. *We have:*

$$\begin{aligned}\chi(\Sigma_g) &= 2 - 2g\\ \chi(N_h) &= 2 - h\end{aligned}$$

Remark 9.22. Alternatively, this follows directly from the homology computations in Propositions 7.41 and 7.45.

Remark 9.23. The Euler characteristic of a surface can therefore attain any integer ≤ 2. The sphere is the only surface with $\chi = 2$ and N_h the only one with $\chi = 2 - h$ as long as h is odd. Every even nonpositive integer is the Euler characteristic of precisely two surfaces, one orientable and one non-orientable.

In particular: Surfaces are completely classified by

- whether they are orientable or not, and
- their Euler characteristic.

Proposition 9.24. *Let X and Y be finite CW complexes. Then so is $X \times Y$ and*

$$\chi(X \times Y) = \chi(X) \cdot \chi(Y).$$

Proof. This follows from Exercise 7.6 in which you show that

$$c_n(X \times Y) = \sum_{a+b=n} c_a(X) \cdot c_b(Y). \qquad \square$$

Remark 9.25. In fact, the statement remains true for arbitrary topological spaces whenever $\chi(X)$ and $\chi(Y)$ are defined. Obviously this imposes restrictions on when a space can be written as a product of certain other spaces. For example, if a space has non-zero Euler characteristic it isn't homeomorphic to $S^3 \times Y$, for any Y.

We will not prove nor use that result, however.

Proposition 9.26. *Let $p\colon Y \to X$ be a d-sheeted cover and assume that X is a finite CW complex. Then so is Y and we have*

$$\chi(Y) = d \cdot \chi(X).$$

Proof. Let $\Phi_\alpha\colon D^n_\alpha \to X$ be one of the characteristic maps. By the lifting property (and since D^n is simply-connected) there are precisely d lifts $\Phi_{\alpha,i}\colon D^n_{\alpha,i} \to Y$, and these, I claim, are the characteristic maps for a CW complex structure. So, the main point is that Y admits a CW structure such that above each cell D^n_α in X there are precisely d cells $D^n_{\alpha,1}, \ldots, D^n_{\alpha,d}$. This is a good thing to know. And of course it immediately yields the identity in the proposition.

In the remainder I will prove the claim above. It relies on basic properties of covers and is not terribly important in the sequel. We use Exercise 6.2 where you show (or read in Hatcher's book) that the collection $(\Phi_{\alpha,i})_{n,\alpha,i}$ defines a CW complex structure on Y if and only if the following two conditions hold:[26]

- Each $\Phi_{\alpha,i}$ restricts to a homeomorphism from the interior $\mathring{D}^n_{\alpha,i} \xrightarrow{\approx} e^n_{\alpha,i}$ to its image and $Y = \coprod_{n,\alpha,i} e^n_{\alpha,i}$.
- The image $\Phi_{\alpha,i}(\partial D^n_{\alpha,i})$ is contained in cells of dimension less than n.

For the first condition, since $\Phi_\alpha = p \circ \Phi_{\alpha,i}$ is injective on the interior so is $\Phi_{\alpha,i}$. Let $U \subseteq \mathring{D}^n_{\alpha,i}$ be a (non-empty) open subset. To show that $\Phi_{\alpha,i}(U) \subseteq e^n_{\alpha,i}$ is open pick a point y in the image, with $x = p(y) \in X$. There exists a small open neighborhood $x \in V \subseteq e^n_\alpha$ such that $p^{-1}(V) = \coprod_{j=1}^d V_j$ with each $V_j = V$. It follows that $V_i \cap p^{-1}(\Phi_\alpha(U) \cap V)$ is an open neighborhood of y inside $e^n_{\alpha,i}$.

If $y \in Y$ is an arbitrary point with image $x = p(y)$ downstairs, there exist unique n, α such that $x \in e^n_\alpha$. By the unique lifting property, $y \in e^n_{\alpha,i}$ for exactly one i. And y cannot belong to some $e^m_{\beta,j}$ with $(n,\alpha) \neq (m,\beta)$ either as otherwise $x \in e^m_\beta$. At this point we have established the first condition.

For the second condition, if $y = \Phi_{\alpha,i}(z)$ for some point $z \in \partial D^n_{\alpha,i}$ on the boundary, then $p(y) = x \in e^m_\beta$ for some $m < n$ and some β. Therefore $y \in e^m_{\beta,j}$ for some (unique) j. □

Remark 9.27. Again, under suitable finiteness assumptions this statement is true for spaces without CW complex structures.

Example 9.28. We know that all surfaces N_h for $h > 0$ have a non-trivial orientation bundle $\tilde{N}_h$. Necessarily it is one of the orientable surfaces Σ_g. In Remark 7.50, we described how to visualize the covering map $\Sigma_g \to N_h$. As an application of Proposition 9.26 we can determine the genus g very easily.

Indeed, it gives us the identity

$$\begin{aligned} \chi(\Sigma_g) &= 2 \cdot \chi(N_h) && \Leftrightarrow \\ 2 - 2g &= 2 \cdot (2 - h) && \Leftrightarrow \\ g &= h - 1 \end{aligned}$$

Example 9.29. For any $k \geq 0$, $S^k \to \mathbb{RP}^k$ is a 2-sheeted cover. It follows that

$$\chi(\mathbb{RP}^k) = \frac{1}{2}\chi(S^k) = \begin{cases} 1 & : k \text{ even} \\ 0 & : k \text{ odd} \end{cases}$$

[26] The third condition there is superfluous for *finitely many* cells.

If we endow $\mathbb{RP}^k$ with the cell structure of Example 8.14 which has one cell in each dimension $\leq k$ then the cell structure on S^k produced in the proof of Proposition 9.26 has two cells in each dimension $\leq k$. In fact, it can be described very easily as follows.

Start with ± 1 as 0-cells. Attach two 1-cells as the upper and lower hemicircles. Think of the result as the equator in S^2. Attach then two 2-cells as the upper and lower hemisphere. Continue like that to get S^k. The cover map $S^k \to \mathbb{RP}^k$ precisely identifies the two cells in each dimension.

One advantage of this cell structure on S^k (in contrast to the minimal one of Example 8.10) is that the n-skeleton of S^k identifies with S^n as long as $n \leq k$. So, the union $S^\infty := \cup_{k\geq 0} S^k$ makes sense and has an obvious structure of CW complex.

Example 9.30 (Unimportant quiz)**.** What should the Euler characteristic $\chi(\mathbb{RP}^\infty)$ of infinite real projective space be? To be sure, this is not something we have defined since this space has infinitely many nonvanishing homology groups, as we saw in Example 8.37. So, it is more of an idle question.

- However, you might argue, while the homology groups are non-zero in all odd degrees, these are cyclic of order 2 hence their rank is zero. They should not contribute to the Euler characteristic. Which is then just

$$\chi(\mathbb{RP}^\infty) = \mathrm{rk}(\mathrm{H}_0(\mathbb{RP}^\infty)) = 1.$$

  More generally, this may suggest that we could define the Euler characteristic for spaces X with finitely generated abelian homology groups and such that $\mathrm{rk}(\mathrm{H}_n(X)) = 0$ for $n \gg 0$. But: read on!

- Of course, when your CW complex has infinitely many cells you cannot expect to be able to compute the Euler characteristic as an alternating sum of numbers of cells in each dimension. Just to drive this point home, for $\mathbb{RP}^\infty$ we have one cell in each dimension hence we should get

$$\chi(\mathbb{RP}^\infty) = \sum_{n\geq 0} (-1)^n,$$

  which is a divergent series. Another way of saying this is that since $\mathbb{RP}^\infty = \cup_{k\geq 0} \mathbb{RP}^k$ we would expect

$$\chi(\mathbb{RP}^\infty) = \varinjlim_{k\to\infty} \chi(\mathbb{RP}^k) = \varinjlim_{k\to\infty} \frac{1+(-1)^k}{2},$$

  which does not converge.

- The CW complex S^∞ considered in Example 9.29 comes with a canonical 2:1-cover map $S^\infty \to \mathbb{RP}^\infty$. Now the homology of S^∞ is very easy, by Lemma 8.19. Since $\mathrm{H}_n(S^k)$ vanishes whenever $0 < n < k$ we have

$$\mathrm{H}_n(S^\infty) = \begin{cases} \mathbb{Z} & : n = 0 \\ 0 & : n > 0 \end{cases}$$

and therefore $\chi(S^\infty) = 1$. (Note that this *is* defined!) By the formula for covers in Proposition 9.26 we would therefore expect

$$\chi(\mathbb{RP}^\infty) = \frac{1}{2}.$$

So, the point is that in extending the notion of Euler characteristic, one needs to be careful in extending the basic properties we established.[27]

10 Homology theories

Let's take stock. We defined no less than three homology theories,

- simplicial,
- singular, and
- cellular,

each with its own advantages and disadvantages. Importantly, we also proved that they are all the 'same'. The goal of this last section is to explain why this is not a coincidence: that everything 'behaving like a homology theory' *is* in fact the same.

Of course, the important bit here is what is in scare quotes. The insight of Eilenberg and Steenrod (1945) was that a few axioms are sufficient for the characterization of a homology theory. It won't surprise you that some of the most important theorems we proved in the course of the past weeks will show up again as axioms.

The whole section 10 is non-examinable.

Commentary 10.1. The topics we have discussed up till this point were all developed in what one might call the 'early period of algebraic topology' (with the already mentioned Analysis Situs papers by Poincaré around 1900 seen as its starting point). One could say that it took a couple of decades for the subject of homology to become so stable that an axiomatization was made possible and universally accepted (as far as I know). This axiomatization, then, seems like a good place to end our foray into algebraic topology. Needless to say, homology theory is only a small part of algebraic topology, with cohomology and homotopy theory two very natural directions to head to thereafter, see Section 10.3.

[27] There is actually an invariant you can associate to a space X that has infinitely many non-zero homology groups all of which are finitely generated abelian. This is the *Poincaré series*

$$P(X,t) := \sum_{n\geq 0} \mathrm{rk}(\mathrm{H}_n(X))t^n \in \mathbb{Z}[[t]]$$

which is a formal series. If the Euler characteristic is defined then $\chi(X) = P(X,-1)$. And it satisfies similar properties as the Euler characteristic with respect to unions and products.

Also, there *is* an invariant that spits out $1/2$ for $\mathbb{RP}^\infty$. It is called the *homotopy cardinality*. However, it behaves more multiplicatively than additively.

10.1 Categories and functors

It's certainly no accident that Eilenberg is one of the authors (the other being Mac Lane) who introduced around the same time the notions of categories, functors, and natural transformations. These are used in the axiomatization of homology theory.

Commentary 10.2. So, what is a category?

A category is simply a formalization of the good practice stipulated in Commentary 2.47. It is a language for speaking about objects and morphisms without reference to the subject matter at hand. It is therefore necessarily abstract. Here, we will only introduce the very basics.

Slightly more precisely, a *category* $\mathscr{C}$ is a collection of *objects* $\mathrm{Ob}(\mathscr{C})$ and for any pair of objects $c, c' \in \mathrm{Ob}(\mathscr{C})$ a set of *morphisms* denoted $\mathrm{Hom}_{\mathscr{C}}(c, c')$, which can be composed.

Before unpacking the last part of the definition let us look at some examples.

Example 10.3.

1. The category of topological spaces Top has as objects topological spaces and morphisms are continuous maps.
2. The category of abelian groups Ab has as objects abelian groups and morphisms are homomorphisms of abelian groups.
3. The category of chain complexes Cpx has as objects chain complexes and as morphisms chain maps.

Commentary 10.4. Let $f : X \to Y$ and $g : Y \to Z$ be continuous maps between topological spaces. We can compose the two maps to obtain another continuous map

$$g \circ f : X \to Z, \qquad x \mapsto g(f(x)).$$

This composition rule satisfies the following:

Associativity $h \circ (g \circ f) = (h \circ g) \circ f$,

Identity the identity map $\mathrm{id}_Y : Y \to Y$ is a neutral element for composition: $g \circ \mathrm{id}_Y = g$ and $\mathrm{id}_Y \circ f = f$.

We require a category to come with a composition rule that satisfies the analogous properties:

Definition 10.5. A *category* $\mathscr{C}$ is a collection of objects and sets of morphisms as in Commentary 10.2 together with composition rules (for every triple of objects),

$$\circ\colon \mathrm{Hom}_{\mathscr{C}}(c', c'') \times \mathrm{Hom}_{\mathscr{C}}(c, c') \longrightarrow \mathrm{Hom}_{\mathscr{C}}(c, c''),$$

that is associative and admits an identity element.

We write $f\colon c \to c'$ if $f \in \mathrm{Hom}_{\mathscr{C}}(c, c')$.

Commentary 10.6. This notion is rather abstract and therefore general. In particular, it does not stipulate that the composition rule must be actual 'composition of maps'. The following two examples are meant to illustrate that. However, they won't play a role in the sequel.

Example 10.7. Let G be a group. Define a category BG with a single object $*$ and morphism set $\mathrm{Hom}_{BG}(*,*) = G$ where the composition is multiplication of group elements:

$$g \circ g' := g \cdot g'$$

Associativity follows from the associativity of group multiplication. The identity is given by the identity element of the group. (We did not use inverses. BG remains a category if G is a monoid only.)

Example 10.8. Let $(P, \leq)$ be a partially ordered set. Define a category $\tilde{P}$ with objects the elements of P and morphism sets

$$\mathrm{Hom}_{\tilde{P}}(x,y) := \begin{cases} * & : x \leq y \\ \emptyset & : \text{else} \end{cases}$$

I leave it as an exercise to check that there is a unique composition rule making $\tilde{P}$ into a category.

Definition 10.9. A morphism $f\colon c \to d$ in a category $\mathscr{C}$ is an *isomorphism* if there exists $g\colon d \to c$ such that $f \circ g = \mathrm{id}_d$ and $g \circ f = \mathrm{id}_c$. The morphism g is unique if it exists and is called the *inverse* of f.

Example 10.10. This recovers the following notions in the examples considered above:

1. in Top: homeomorphisms
2. in Ab: isomorphisms (that is, bijective homomorphisms)
3. in Cpx: isomorphisms in each degree
4. in BG: every homomorphism is an isomorphism if G is a group
5. in $\tilde{P}$: (if you require the partially ordered set to be antisymmetric then) the only isomorphisms are the identities

Definition 10.11. Let $\mathscr{C}$ and $\mathscr{D}$ be categories. A *functor* $F\colon \mathscr{C} \to \mathscr{D}$ associates to each object $c \in \mathrm{Ob}(\mathscr{C})$ an object $F(c) \in \mathrm{Ob}(\mathscr{D})$, and comes with maps

$$F\colon\ \mathrm{Hom}_{\mathscr{C}}(c,c') \to \mathrm{Hom}_{\mathscr{D}}(F(c),F(c'))$$

that are compatible with composition and identities. In short: $F(\mathrm{id}) = \mathrm{id}$ and $F(f \circ g) = F(f) \circ F(g)$.

Here is an immediate consequence.

Lemma 10.12. *If f is an isomorphism in $\mathscr{C}$ then $F(f)$ is an isomorphism in $\mathscr{D}$.*

Proof. Let g be an inverse to f. Then

$$F(f) \circ F(g) = F(f \circ g) = F(\mathrm{id}) = \mathrm{id}$$

and

$$F(g) \circ F(f) = F(g \circ f) = F(\mathrm{id}) = \mathrm{id}$$

so that $F(g)$ is an inverse to $F(f)$. □

Example 10.13. Here are some examples of functors we have seen during the course of the past weeks:

1. $\mathrm{H}_n \colon \mathsf{Top} \to \mathsf{Ab}$,
2. $\mathrm{H}_n \colon \mathsf{Cpx} \to \mathsf{Ab}$,
3. $C_\bullet \colon \mathsf{Top} \to \mathsf{Cpx}$,[28]
4. $(-)^{\mathrm{ab}} \colon \mathsf{Grp} \to \mathsf{Ab}$

Example 10.14. Let $f \colon G \to G'$ be a group homomorphism. There is then an induced functor

$$Bf \colon BG \to BG'$$

which (necessarily) sends the unique object of BG to the unique object of BG' and on morphism sets is given by

$$f \colon G = \mathrm{Hom}_{BG}(*, *) \to \mathrm{Hom}_{BG'}(*, *) = G'.$$

In fact it is clear that group homomorphisms $G \to G'$ and functors $BG \to BG'$ are in bijection.

10.2 Axioms

To state the axioms for homology theory we need to introduce a variant of the category of topological spaces.

Definition 10.15. We denote by CW_2 the category of *CW pairs*. Its objects are pairs of CW complexes (X, Y) where $Y \subseteq X$ is a subcomplex. A morphism $(X, Y) \to (X', Y')$ is a continuous map $f \colon X \to X'$ such that $f(Y) \subseteq Y'$. These are composed in the obvious way.

We identify CW complexes X with pairs $(X, \emptyset)$.

Definition 10.16. A *homology theory* is a collection of functors (for $n \geq 0$),

$$h_n \colon \mathsf{CW}_2 \to \mathsf{Ab},$$

and *natural* maps

$$\partial_n \colon h_n(X, Y) \to h_{n-1}(Y) := h_{n-1}(Y, \emptyset)$$

satisfying the following axioms:

[28] In fact, one can compose functors in a pretty obvious way. (What does that suggest? Right, there is a category of (small) categories with functors as morphisms.) And then the following triangle commutes by definition:

$$\begin{array}{ccc} \mathsf{Top} & \xrightarrow{C_\bullet} & \mathsf{Cpx} \\ & \searrow^{\mathrm{H}_n} & \downarrow{\scriptstyle \mathrm{H}_n} \\ & & \mathsf{Ab} \end{array}$$

Homotopy invariance If $f \simeq g \colon X \to Y$ then $h_n(f) = h_n(g) \colon h_n(X) \to h_n(Y)$.

Collapse [29] The map $h_n(X, Y) \to h_n(X/Y, Y/Y)$ is an isomorphism.

Exactness The sequence

$$\cdots \to h_{n+1}(X, Y) \xrightarrow{\partial_{n+1}} h_n(Y) \to h_n(X) \to h_n(X, Y) \xrightarrow{\partial_n} h_{n-1}(Y) \to \cdots$$

is exact.

Additivity For any collection (X_α, Y_α), the map $\oplus_\alpha h_n(X_\alpha, Y_\alpha) \to h_n(\amalg_\alpha (X_\alpha, Y_\alpha))$ is an isomorphism.

Dimension $h_n(*) = \delta_{n0}\mathbb{Z}$

Proposition 10.17. *Singular homology defines a homology theory.*

Proof. **Homotopy invariance** is Theorem 4.2, **Exactness** is Corollary 5.18, **Collapse** is Proposition 5.28, **Additivity** is Proposition 3.14, and **Dimension** is Example 3.9. □

Remark 10.18. Here, naturality of the connecting homomorphism ∂_n means that for every $f \colon (X, Y) \to (X', Y')$ of CW pairs the following square commutes:

$$\begin{array}{ccc} h_n(X, Y) & \xrightarrow{\partial_n} & h_{n-1}(Y) \\ \downarrow{\scriptstyle h_n(f)} & & \downarrow{\scriptstyle h_{n-1}(f)} \\ h_n(X', Y') & \xrightarrow{\partial_n} & h_{n-1}(Y') \end{array}$$

For singular homology this follows from the construction of the connecting homomorphism in Proposition 4.37.

In categorical language, this says that ∂_n is a *natural transformation* (of functors $\mathsf{CW}_2 \to \mathsf{Ab}$):

Definition 10.19. A *natural transformation* $\phi \colon F \to G$ between two functors $F, G \colon \mathscr{C} \to \mathscr{D}$ is a collection of maps $\phi_c \colon F(c) \to G(c)$ in $\mathscr{D}$ such that the following square commutes for every $f \colon c \to c'$:

$$\begin{array}{ccc} F(c) & \xrightarrow{\phi_c} & G(c) \\ \downarrow{\scriptstyle F(f)} & & \downarrow{\scriptstyle G(f)} \\ F(c') & \xrightarrow{\phi_{c'}} & G(c') \end{array}$$

For the rest of this section we fix a homology theory h_n. Let us see what we can compute just from the axioms. We define, as before, the reduced homology group as $\tilde{h}_n(X) := \ker(h_n(X) \to h_n(*))$.

[29] This is non-standard terminology. Presumably, others would call this **Excision** instead. The two results (Theorem 5.23 and Proposition 5.28) are certainly related but I would like to distinguish them.

Example 10.20. 1. For each triple $X \supseteq Y \supseteq Z$, the sequence[30]

$$\cdots \to h_{n+1}(X,Y) \xrightarrow{\partial_{n+1}} h_n(Y,Z) \to h_n(X,Z) \to h_n(X,Y) \xrightarrow{\partial_n} h_{n-1}(Y,Z) \to \cdots$$

is exact. This is a diagram chase, starting with **Exactness**. Cf. Exercise 4.5.

2. If X is a contractible space then $\tilde{h}_n(X) = 0$ for all n. Indeed, the map $h_n(X) \to h_n(*)$ is an isomorphism, by **Homotopy Invariance**, cf. Corollary 4.3.
3. Consider the exact sequence above for the triple $(D^k, \partial D^k, *)$:

$$\cdots \to \tilde{h}_n(D^k) \to h_n(D^k, S^{k-1}) \to \tilde{h}_{n-1}(S^{k-1}) \to \tilde{h}_{n-1}(D^k) \to \cdots$$

As the two outer terms vanish we have an isomorphism in the middle. And by **Collapse**, **Additivity** and **Dimension** we get

$$\tilde{h}_n(S^k) \xleftarrow{\sim} h_n(D^k, S^{k-1}) \xrightarrow{\sim} \tilde{h}_{n-1}(S^{k-1}) \xleftarrow{\sim} \cdots \xrightarrow{\sim} \tilde{h}_{n-k}(S^0) = \delta_{nk}\mathbb{Z}.$$

4. We have by **Collapse**

$$h_n(X^k, X^{k-1}) \xrightarrow{\sim} \tilde{h}_n(X^k/X^{k-1}) \xrightarrow{\sim} \tilde{h}_n(\vee_\alpha S^k_\alpha) \xleftarrow{\sim} \oplus_\alpha \tilde{h}_n(S^k_\alpha) = \oplus_\alpha \delta_{nk}\mathbb{Z}$$

where the wrong-way isomorphism is deduced from **Additivity** and **Collapse**, as in Lemma 7.40. Here, the α range over the k-cells in X.

In fact, *everything* we computed (in terms of homology) in this course can be computed just from the axioms:

Theorem 10.21. *If (h_n, ∂_n) is a homology theory then there are natural isomorphisms*

$$h_n(X,Y) \cong H_n(X,Y).$$

Sketch of proof. Example 10.20 allows us to define a 'cellular chain complex'

$$C_n^{\mathrm{CW},h}(X) := h_n(X^n, X^{n-1}) \cong \oplus_\alpha \mathbb{Z}[D^n_\alpha]$$

with differentials exactly as for singular homology. To deduce that the homology of this chain complex is our old cellular homology, we need to identify the differentials. If you recall the cellular boundary formula Proposition 8.30, you see that it suffices to show that $\deg^h(f) = \deg(f)$ for maps $f\colon S^k \to S^k$, $k > 0$. You can prove the local degree formula Proposition 6.21 for the theory h and deduce that this is true at least for all f we constructed in Proposition 6.6. The result for general f then follows from the following non-trivial fact (see Hatcher, Corollary 4.25, and also section 10.3 below): The degree map induces an isomorphism

$$\deg\colon [S^k, S^k] \xrightarrow{\sim} \mathbb{Z}$$

[30]The connecting homomorphisms are the composites $h_{n+1}(X,Y) \to h_n(Y,\emptyset) \to h_n(Y,Z)$.

between homotopy classes of maps $S^k \to S^k$ and the integers.[31] That is, up to homotopy, we have already constructed all such self-maps in Proposition 6.6.

At this point we have proved $h_n(X) \cong h_n^{\mathrm{CW}}(X) \cong \mathrm{H}_n^{\mathrm{CW}}(X) \cong \mathrm{H}_n(X)$, by Corollary 8.25. We deduce that $h_n(X,Y) \cong \tilde{h}_n(X/Y) \cong \tilde{H}_n(X/Y) \cong \mathrm{H}_n(X,Y)$ which completes the sketch of the proof. □

Remark 10.22. One can make the natural isomorphism in the theorem commute with the 'connecting homomorphisms' ∂_n. It follows that in a suitable category of homology theories, all objects are isomorphic.

Example 10.23. If h_n is a homology theory (for example, singular homology $h_n = H_n$) then so is h'_n where $h'_n(X,Y) = h_n(X,Y)$ and $\partial'_n = -\partial_n$. By the previous remark, these two homology theories are isomorphic nonetheless. An explicit isomorphism is given by $(-1)^n \,\mathrm{id}\colon h_n(X,Y) \to h'_n(X,Y)$.

10.3 Generalized (co)homology theories

Perhaps you find Theorem 10.21 satisfying (I do), or perhaps you find it rather disappointing. Indeed, wasn't the goal of algebraic topology to distinguish spaces, as discussed in Section 1.1? And if two spaces can't be distinguished using singular homology then—Theorem 10.21 tells us—they can't be distinguished by *any* homology theory.

But not so fast! Looking at the history of 'homology theory' since Theorem 10.21 was established, what seems to have happened is the following. Rather than giving up on the subject, mathematicians started weakening and modifying the axioms and studying theories that satisfied these new sets of axioms. I will touch upon 2 ½ modifications below, very briefly. These can also be read as suggestions of where to go next if you would like to study more algebraic topology.

Coefficients Actually even before the axiomatization, mathematicians had considered homology that spit out abelian groups other than the integers on a point. Here is a precise definition.

Definition 10.24. Let A be an abelian group. A *homology theory with coefficients in A* is a collection of functors $h_n\colon \mathsf{CW}_2 \to \mathsf{Ab}$ together with natural maps ∂_n satisfying the same axioms as in Definition 10.16 except that the dimension axiom is replaced by $h_n(*) = \delta_{n0}A$.

Remark 10.25. To obtain such a theory replace the singular chain complex $C_\bullet(X)$ by the tensor product $C_\bullet(X) \otimes_{\mathbb{Z}} A$. That is, in degree n you consider the group

$$C_n(X;A) := \bigoplus_{\sigma\colon \Delta^n \to X} A[\sigma]$$

[31]The addition on the left-hand side takes two (pointed) maps $f,g\colon S^k \to S^k$ to the composite

$$S^k \twoheadrightarrow S^k \vee S^k \xrightarrow{f\vee g} S^k$$

where the first map collapses the equator (containing the base point). See Section 10.3 below for more details.

with the 'same' differential as before. Then we set $\mathrm{H}_n(X;A) := \mathrm{H}_n(C_\bullet(X;A))$ and similarly for relative homology. This *singular homology with coefficients in A* satisfies the axioms as in the previous definition.

A particularly useful group to consider is $A = \mathbb{Z}/2\mathbb{Z}$.

Example 10.26. The cellular chain complex for $\mathbb{RP}^k$ with $\mathbb{Z}/2$-coefficients looks as follows (compare with Example 8.36):

$$0 \to \mathbb{Z}/2 \xrightarrow{0} \mathbb{Z}/2 \to \cdots \to \mathbb{Z}/2 \xrightarrow{0} \mathbb{Z}/2 \xrightarrow{0} \mathbb{Z}/2 \to 0,$$

so that

$$\mathrm{H}_n(\mathbb{RP}^k; \mathbb{Z}/2) = \begin{cases} \mathbb{Z}/2 & : 0 \leq n \leq k \\ 0 & : n > k \end{cases}$$

and

$$\mathrm{H}_n(\mathbb{RP}^\infty; \mathbb{Z}/2) = \mathbb{Z}/2$$

for all $n \geq 0$.

Here is an application related to Exercise 5.8.

Proposition 10.27. *Let $f\colon S^k \to S^k$ be an odd map. Then $f_*\colon \mathbb{Z}/2 = \mathrm{H}_k(S^k;\mathbb{Z}/2) \xrightarrow{\sim} \mathrm{H}_k(S^k;\mathbb{Z}/2) = \mathbb{Z}/2$ is an isomorphism.*

Proof. See Hatcher, Proposition 2B.6, which really uses $\mathbb{Z}/2$-coefficients in an essential way. □

Corollary 10.28 (Borsuk-Ulam). *For any map $f\colon S^k \to \mathbb{R}^k$ there is $x \in S^k$ such that $f(-x) = f(x)$.*

Proof. See Exercise 5.8. □

Commentary 10.29. This implies that at any point in time there are two antipodal points on earth with the exact same temperature and wind speed.

On the other hand, in terms of distinguishing more spaces, this modification of homology theory is of not much help:

Proposition 10.30. *If $\mathrm{H}_n(X) \cong \mathrm{H}_n(Y)$ for all n then $h_n(X) \cong h_n(Y)$ for all homology theories with coefficients in A.*

Proof. There is a similar uniqueness theorem for homology theories with coefficients in A as Theorem 10.21, see Hatcher, Theorem 4.59. That is, every such homology theory is actually isomorphic to $\mathrm{H}_\bullet(-;A)$, the one constructed just above. And for this theory one has the *universal coefficient theorem* (Hatcher, Corollary 3A.4) which expresses it in terms of the ordinary singular homology:

$$\mathrm{H}_n(X;A) \cong (\mathrm{H}_n(X) \otimes A) \oplus \mathrm{Tor}(\mathrm{H}_{n-1}(X), A)$$

where the Tor-groups are some purely algebraic invariant of abelian groups. □

Dimension axiom Here is a more drastic departure from Definition 10.16.

Definition 10.31. A *generalized homology theory* is a collection of functors $h_n\colon \mathsf{CW}_2 \to \mathsf{Ab}$ together with natural maps ∂_n satisfying the same axioms as in Definition 10.16 except for (possibly) the dimension axiom.

Commentary 10.32. Several theories that satisfied only these weaker axioms were found in the following decades, arguably most prominently *bordism* and *stable homotopy*. A *lot* of work in algebraic topology since then has gone into investigating these and other examples, but many questions remain open!

However, one thing that is clear is that no uniqueness theorem in the form of Theorem 10.21 holds for generalized homology theories. The two examples just mentioned (and many others) are genuinely distinct.

My goal is to define stable homotopy for you, and give you an example of a question about them that is still open.

Definition 10.33. The homotopy groups for $n > 0$ of a pointed space are homotopy classes of pointed maps,

$$\pi_n(X,x) = [S^n, X].$$

The group law is defined by the following diagram.

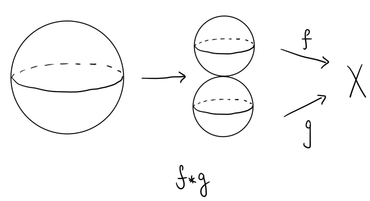

Here, the first map is the 'pinching map', sending the equator to the common base point. (The base point of the sphere on the left is assumed to lie on the equator so this pinching map becomes a map of pointed spaces.)

Remark 10.34. The group $\pi_n(X,x)$ turns out to be abelian if $n \geq 2$. To see this, imagine thickening the equator in the sphere on the left, thereby pinching a whole strip in the middle. Then you can swap the remaining caps around the north and south pole continuously without them meeting the base point (which, remember we placed somewhere on the equator). Finally you shrink the strip back to the equator. This shows that $f \cdot g \simeq g \cdot f$.

Exercise 10.35. Why does the previous argument not show that $\pi_1(X,x)$ is abelian? In other words, what goes wrong on the 1-sphere?

Commentary 10.36. Of course, we know that homotopy equivalences or homeomorphisms induce isomorphisms in homology, and we saw very few instances (surfaces, mainly) where the converse holds. The following gives an idea of how powerful homotopy groups are. Note that it applies to any CW complex whatsoever.

Theorem 10.37 (Whitehead's Theorem)**.** *Let $f\colon X \to Y$ be a continuous map between CW complexes. If it induces a bijection on connected components and*

$$f_*\colon \pi_n(X,x) \xrightarrow{\sim} \pi_n(Y,f(x))$$

for all $x \in X$ and all n then f is a homotopy equivalence.

Proof. See Hatcher, Theorem 4.5. □

Given this power, it is not surprising that computing these homotopy groups is very difficult in general. As mentioned in Example 1.10, the homotopy groups of spheres are mostly unknown.

Example 10.38 (Hopf fibration)**.** Consider the map $S^3 \to \mathbb{C} \cup \{\infty\}$,

$$(x_1,x_2,x_3,x_4) \mapsto \frac{x_1+ix_2}{x_3+ix_4}$$

which maps to ∞ when the denominator vanishes. We may identify the one-point compactification $\mathbb{C} \cup \{\infty\}$ with the 2-sphere, in which case this map $\eta\colon S^3 \to S^2$ is known as the *Hopf fibration* (due to Heinz Hopf (1931)). Over each point in S^2 the fiber of η consists of a circle S^1, and for any pair of distinct points in S^2 these circles pass through each other precisely once! (See the wikipedia page for some illustrations, or even better, watch this video.)

What is the relation with homotopy groups? One might think that just as $\pi_n(S^k) = 0$ for $n < k$, one also has $\pi_n(S^k) = 0$ for $n > k$. (On the 'diagonal' $\pi_k(S^k) = \mathbb{Z}$ one always has a copy of the integers, see below.) Hopf put that idea to rest. He used the fact that the linking number between any two distinct fibers is 1 to prove that the Hopf fibration is not null-homotopic. In fact, $\pi_3(S^2) \cong \mathbb{Z}[\eta]$ generated by the class of the Hopf fibration.

However, there is a sort of stability phenomenon happening that was discovered by Hans Freudenthal (1936).

Theorem 10.39 (Freudenthal Suspension Theorem)**.** *Let X be a pointed CW complex. The suspension map*

$$\pi_{n+k}(\Sigma^k X) = [S^{n+k}, \Sigma^k X] \to [\Sigma S^{n+k}, \Sigma^{k+1} X] = [S^{n+k+1}, \Sigma^{k+1} X] = \pi_{n+k+1}(\Sigma^{k+1} X)$$

is an isomorphism for $k \gg n$.

Proof. See Hatcher, Corollary 4.24. □

Example 10.40. In fact, there are precise bounds on k, n (depending on X) for this map to become an isomorphism. For example it implies rather easily that all maps in the sequence

$$\mathbb{Z} \cong \pi_1(S^1) \xrightarrow{\sim} \pi_2(S^2) \xrightarrow{\sim} \pi_3(S^3) \to \cdots,$$

are isomorphisms.

Example 10.41. Again by the Freudenthal suspension theorem, there is a sequence

$$\mathbb{Z}[\eta] = \pi_3(S^2) \twoheadrightarrow \pi_4(S^3) \xrightarrow{\sim} \pi_5(S^4) \xrightarrow{\sim} \cdots$$

It turns out that the first map is not an isomorphism. Instead it identifies $\pi_4(S^3)$ with $\mathbb{Z}[\eta]/2[\eta]$.

The theorem says that the group $\pi_{n+k}(\Sigma^k X)$ is independent of k for large k. Therefore we can define:

Definition 10.42. Let X be a pointed CW complex. The *stable homotopy groups* are

$$\pi_n^s(X) = \varinjlim_{k\to\infty} \pi_{n+k}(\Sigma^k X).$$

Note that these make sense even for $n < 0$, and by Remark 10.34, they are all abelian groups.

So, these groups record only *stable* phenomena at the level of homotopy groups. For this reason, they are easier to compute than the original 'unstable' homotopy groups, although still (too) difficult.

Example 10.43.

1. From Example 10.40 we deduce that $\pi_0^s(*) = \mathbb{Z}$.
2. It is not too hard to see that $\pi_n^s(*) = 0$ if $n < 0$.
3. It is also known (Serre (1953)) that $\pi_n^s(*)$ is a finite group for all $n > 0$. For example, $\pi_1^s(*) = \mathbb{Z}/2$ generated by the class of η, as seen in Example 10.41.
4. Except for small-ish values of $n > 0$ (in the region of two to three digits), the groups $\pi_n^s(*)$ are largely unknown. Nevertheless it is known (at least since Adams (1965)) that infinitely many are non-zero. You can find some intriguing information about these in Hatcher, p. 384ff.

Commentary 10.44. The stable homotopy groups of spheres are linked, among many other things, with the geometry of higher-dimensional manifolds. Because the latter is hard, it is not too surprising that the former is hard as well. Just as one example, computations of $\pi_n^s(*)$ allowed mathematicians to determine the number of smooth structures on spheres.

Proposition 10.45. *Stable homotopy groups π_n^s define a generalized homology theory.*

Proof. Hatcher, Proposition 4F.1. □

Remark 10.46. Note that this theory fails the dimension axiom in an extreme way: The values $\pi_n^s(*)$ are non-zero for infinitely many n and they are not even known!

Exercise 10.47. You could say: Okay, the groups $\pi_n^s(*)$ are hard to compute but assuming these, would it then be possible to compute $\pi_n^s(X)$ for any CW complex X? An obvious idea would be to try to repeat the proof of Theorem 10.21, and to construct a cellular chain complex with terms given by stable homotopy groups by spheres. What goes wrong?

Hint: Observe that $\pi_\bullet^s(S^n) = \pi_{\bullet-n}^s(*)$.

Cohomology Another modification to the axioms of a homology theory one can make is to 'dualize' everything. This leads to the notion of (ordinary and generalized) *co*homology theory.

Thus, a cohomology theory is a collection of *contravariant* functors $h^n : \mathsf{CW}_2 \to \mathsf{Ab}$ and natural transformations $h^{n-1}(Y) \to h^n(X, Y)$ satisfying axioms dual to the ones for a homology theory.[32] We will not make this precise here but see Hatcher, p. 202. Instead, it turns out that the 'ordinary' cohomology theories are again uniquely determined by their value on a point, so there is essentially one example which we will discuss now.

Definition 10.48. The *singular cochain complex* of a space X has terms the abelian groups

$$C_n^*(X) := \mathrm{Hom}(C_{-n}(X), \mathbb{Z})$$

and differential given by precomposition with the differential in the singular chain complex. Note that it has non-zero terms in non-positive degrees only.

The *singular cohomology* of X is the collection of groups

$$\mathrm{H}^n(X) := \mathrm{H}_{-n}(C_\bullet^*(X))$$

Remark 10.49. This dualization process actually has some positive side-effects, making cohomology far from a trivial modification of homology. The starting point is the observation that there are maps

$$\cup : \mathrm{H}^m(X) \times \mathrm{H}^n(X) \to \mathrm{H}^{m+n}(X \times X) \xrightarrow{\Delta^*} \mathrm{H}^{m+n}(X),$$

where the first map can be thought of as taking an 'm-cell e_m and an n-cell e_n' in X and producing the '$(m+n)$-cell $e_m \times e_n$ in $X \times X$. This can be made precise using the cellular cochain complex (for CW complexes). The second map is induced by the diagonal embedding $\Delta : X \to X \times X$ (recall that maps in cohomology go in the *opposite* direction!).

This is called the *cup product* and it makes $\mathrm{H}^*(X) := \oplus_n \mathrm{H}^n(X)$ into a (graded) ring, a structure that is not visible at the level of homology. The cup product admits a more geometric interpretation in good cases. For example, on surfaces S, the product $\mathrm{H}^1(S) \times \mathrm{H}^1(S) \to \mathrm{H}^2(S)$ can be determined by computing intersections (after deformation) of their fundamental loops (cf. Section 7.3).

Example 10.50. For $k > 0$, the cohomology ring for the sphere is $\mathrm{H}^*(S^k) = \mathbb{Z}[t]/t^2$ (with t in degree k).

Example 10.51. One also has cohomology with coefficients (by taking $\mathrm{Hom}(C_{-n}(X), A)$ instead of $\mathrm{Hom}(C_{-n}, \mathbb{Z})$). As rings,

$$\mathrm{H}^*(\mathbb{RP}^\infty; \mathbb{Z}/2) = \mathbb{Z}/2[t]$$

(with t in degree 1).

[32] A functor $F : \mathscr{C} \to \mathscr{D}$ is contravariant if it induces morphisms in the opposite direction, that is, comes with maps

$$\mathrm{Hom}_{\mathscr{C}}(c, c') \to \mathrm{Hom}_{\mathscr{D}}(F(c'), F(c)),$$

still assumed compatible with composition and identities.

Example 10.52. Recall from Exercise 5.10 the space $X = S^1 \vee S^1 \vee S^2$ whose homology groups are isomorphic to the ones of the torus $\mathbb{T}$. It is 'clear' that these spaces are not homeomorphic (nor homotopy equivalent), and you proved this in that exercise. Here is a way to see this using cohomology. In both cases, one has isomorphic cohomology groups and they coincide with the homology groups:

$$\mathrm{H}^0 = \mathbb{Z}, \quad \mathrm{H}^1 = \mathbb{Z}^2, \quad \mathrm{H}^2 = \mathbb{Z},$$

with all other groups zero. The cup product of the two generators α, β in degree 1, however, is non-zero for the torus, and zero for X. Hence $\mathbb{X} \not\simeq \mathbb{T}$.

That $\alpha \cup \beta \neq 0$ in $\mathbb{T}$ can be explained by the fact that however you deform the two circles

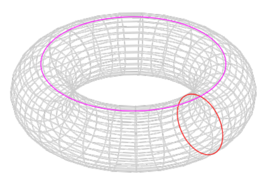

they will continue to intersect.

That $\alpha \cup \beta = 0$ in X can be explained by the fact that the cohomology ring of a wedge sum is a subring of the product:

$$\mathrm{H}^*(S^1 \vee S^1 \vee S^2) \subseteq \mathrm{H}^*(S^1 \vee S^1) \times \mathrm{H}^*(S^2)$$

and since $\alpha \cup \beta \in \mathrm{H}^2(S^1 \vee S^1) = 0$.

Remark 10.53. There is also a notion of *generalized* cohomology theory, obtained by dropping the dimension axiom. A famous such theory is *topological K-theory*. It was used by Adams and Atiyah (1966) to give an elegant proof of the Hopf invariant one problem to which we now turn.

Definition 10.54. A *real division algebra* is a finite-dimensional real vector space $\mathbb{A}$ together with an operation $\cdot : \mathbb{A}^2 \to \mathbb{A}$ such that for all $a \in \mathbb{A}$, the maps

$$\mathbb{A} \to \mathbb{A}$$
$$x \mapsto a \cdot x$$
$$x \mapsto x \cdot a$$

are $\mathbb{R}$-linear, and invertible if $a \neq 0$. (No assumptions are made on associativity, identity, or commutativity. Nor need these be normed algebras.)

Example 10.55. The classical examples are $\mathbb{R}$ itself, the complex numbers $\mathbb{C}$, the quaternions $\mathbb{H}$, and the octonions $\mathbb{O}$. The first two are commutative, associative and have an identity. The third is associative with identity but is not commutative. The last is made of pairs of quaternions with multiplication given by

$$(a, b)(c, d) = (ac - \bar{d}b, da + b\bar{c}).$$

It is not even associative (but has an identity).

Another example is the complex numbers with the non-standard multiplication $zw = \overline{zw}$. It has no identity.

Note that these division algebras have real dimension 1,2,4,8.

Remark 10.56 (Hopf invariant)**.** Let $\phi \colon S^{2n-1} \to S^n$ be a continuous map for $n > 1$. We may use this to construct a CW complex C_ϕ with cells e_0, e_n, e_{2n}, with ϕ the attaching map for the $2n$-cell. By cellular cohomology, we have as only non-trivial cohomology groups

$$\mathrm{H}^0 = \mathbb{Z}[e_0], \quad \mathrm{H}^n = \mathbb{Z}[e_n], \quad \mathrm{H}^{2n} = \mathbb{Z}[e_{2n}].$$

The cohomology ring is completely determined by the cup product of a generator $[e_n] \in \mathrm{H}^1$ with itself. We have

$$[e_n] \cup [e_n] = h(\phi) \cdot [e_{2n}]$$

for some integer $h(\phi) \in \mathbb{Z}$, defined up to sign. This is called the *Hopf invariant* of ϕ. (With more care, it can be defined without the sign indeterminacy.)

Example 10.57. For $n = 2$, we have $h(\eta) = 1$ which is related to the fact that the knots in distinct fibers have linking number 1. There are analogues of the Hopf fibration in higher dimensions,

$$\eta_4 \colon S^7 \to S^4 \text{ (with fibers } S^3), \qquad \eta_8 \colon S^{15} \to S^8 \text{ (with fibers } S^7),$$

also with Hopf invariant one.

In fact, $X_\eta = \mathbb{CP}^2$, $X_{\eta_4} = \mathbb{HP}^2$, $X_{\eta_8} = \mathbb{OP}^2$ and in each case the cohomology ring is $\mathbb{Z}[t]/t^3$ with t in degree $2, 4, 8$. (For $n = 1$ one could take $\mathbb{Z}/2$-coefficients. The canonical 2:1-cover $\eta_1 \colon S^1 \to S^1$ produces $X_{\eta_1} = \mathbb{RP}^2$ and $\mathrm{H}^*(\mathbb{RP}^2; \mathbb{Z}/2) = \mathbb{Z}/2[t]/t^3$.)

Remark 10.58 (Hopf invariant one problem)**.** The Hopf invariant problem asks: are there maps other than $\eta = \eta_2$, η_4, and η_8 with Hopf invariant one?

Theorem 10.59 (Adams (1960); Adams–Atiyah (1966))**.** *Maps with Hopf invariant one only occur for $n = 2, 4, 8$ (and $n = 1$ if you wish).*

The first proof which is quite long uses operations in cohomology (an ordinary cohomology theory), the second is shorter and uses K-theory (a generalized cohomology theory).

Corollary 10.60. *There are no real division algebras except in dimensions $1, 2, 4, 8$.*

Remark 10.61. The connection between Hopf fibrations and division algebras is indicated by the following observation. Let $\mathbb{A}$ be one of the four classical real division algebras of Example 10.55 of dimension $n = 1, 2, 4, 8$. Then there is a map

$$\mathbb{S}(\mathbb{A}^2) \to \mathbb{AP}^1$$

from the elements in $\mathbb{A}^2$ of unit norm to the projective line over $\mathbb{A}$. It sends an element a of unit norm to the line through the origin and a.

Now, the left-hand side identifies with the sphere S^{2n-1}, while the right-hand side identifies with S^n so this is a continuous map $S^{2n-1} \to S^n$. Guess which one? Right, it's one of the Hopf fibrations from Example 10.57.