Topology I and II, 2023–2024, HU Berlin

Chris Wendl

Contents

First semester (Topologie I)

1. Introduction and motivation (April 18, 2023)

To start with, let us discuss what kinds of problems are studied in topology. This lecture is only intended as a sketch of ideas, so nothing in it is intended to be precise—we'll introduce precise definitions in the next lecture.

(1) Classification of spaces. Let's assume for the moment that we understand what the word "space" means. We'll be more precise about it next week, but in this course, a "space" X is a set with some extra structure on it such that we have well-defined notions of things like *open* subsets (*offene Teilmengen*) $\mathcal{U} \subset X$ and *continuous maps/mappings* (*stetige Abbildungen*) $f : X \to Y$ (where Y is another space). It is then natural to consider two spaces X and Y equivalent if there is a **homeomorphism** (*Homöomorphismus*) between them: this means a continuous bijection $f : X \to Y$ whose inverse $f^{-1} : Y \to X$ is also continuous. We say in this case that X and Y are **homeomorphic** (*homöomorph*).

So for instance, one can try to classify all *surfaces* (*Flächen*) up to homeomorphism:

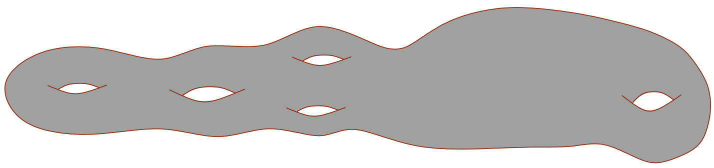

The space in this picture is known as a "closed orientable surface of genus (*Geschlecht*) five". The genus is a nonnegative integer that, roughly speaking, counts the number of "handles" you would need to attach to a sphere in order to construct the surface. The notation Σ_g is often used for a surface of genus $g \geqslant 0$.

There are also closed surfaces that cannot be embedded in $\mathbb{R}^3$, though they are harder to visualize. Here are two examples.

Example 1.1. Here is a picture of the **Klein bottle** (*Kleinsche Flasche*), a surface that can be "immersed" (with self-intersections) in $\mathbb{R}^3$, but not embedded:

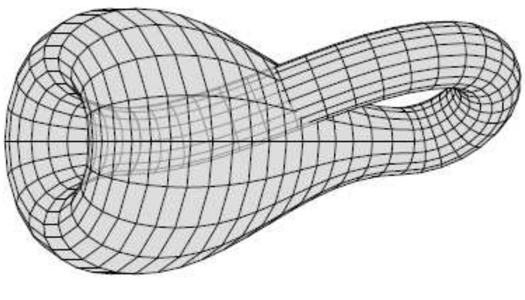

We'll give a more precise definition of the Klein bottle as a topological space later.

EXAMPLE 1.2. The **real projective plane** (*reelle projektive Ebene*) $\mathbb{RP}^2$ is a space that can be described in various equivalent ways:

(1) $\mathbb{RP}^2 := S^2/\sim$, i.e. the set of equivalence classes of elements in the unit sphere $S^2 := \{\mathbf{x} \in \mathbb{R}^3 \mid |\mathbf{x}| = 1\}$, with the equivalence relation defined by $\mathbf{x} \sim -\mathbf{x}$ for each $\mathbf{x} \in S^2$. In other words, every element of $\mathbb{RP}^2$ is a set of two elements $\{\mathbf{x}, -\mathbf{x}\}$, with both belonging to the unit sphere. (See Remark 1.3 below on notation for defining equivalence relations.)
(2) $\mathbb{RP}^2 := \mathbb{D}^2/\sim$, where $\mathbb{D}^2 := \{\mathbf{x} \in \mathbb{R}^2 \mid |\mathbf{x}| \leqslant 1\}$ and the equivalence relation is defined by $z \sim -z$ for every point z on the boundary of the disk. One obtains this from the first description of $\mathbb{RP}^2$ by restricting attention to only one hemisphere of S^2; no information is lost since the other hemisphere is identified with it, but along the equator between them, there is still an identification of antipodal points.
(3) $\mathbb{RP}^2$ is the space of all lines through 0 in $\mathbb{R}^3$. This is equivalent to the first description since every line through the origin in $\mathbb{R}^3$ hits S^2 at exactly two points, which are antipodal to each other.
(4) $\mathbb{RP}^2$ is the space constructed by gluing a disk $\mathbb{D}^2$ to a **Möbius strip** (*Möbiusband*)
$$\mathbb{M} := \left\{(\theta, t\cos(\pi\theta), t\sin(\pi\theta)) \in \mathbb{R}/\mathbb{Z} \times \mathbb{R}^2 \;\middle|\; \theta \in \mathbb{R},\ t \in [-1, 1]\right\}.$$
To see this, draw a picture of the unit sphere S^2 and think of $\mathbb{RP}^2$ as $S^2/\sim$. After identifying antipodal points of the sphere in this way, a neighborhood of the equator looks like a Möbius strip, and everything else is a disk (it looks like two disks in the picture, but the two are identified with each other).

More generally, for each integer $n \geqslant 0$ one can define the **n-sphere**
$$S^n = \left\{\mathbf{x} \in \mathbb{R}^{n+1} \;\middle|\; |\mathbf{x}| = 1\right\}$$
and the **real projective n-space**
$$\mathbb{RP}^n = S^n/\{\mathbf{x} \sim -\mathbf{x}\} = \left\{\text{lines through 0 in } \mathbb{R}^{n+1}\right\}.$$

REMARK 1.3. In topology, we often specify an equivalence relation $\sim$ on a set X with words such as "the equivalence relation defined by $x \sim f(x)$ for all $x \in A$" where $A \subset X$ is a subset and $f : A \to X$ a map. This should always be interpreted to mean that $\sim$ is the *smallest* equivalence relation for which the stated property is true, i.e. since every equivalence relation must also be reflexive and symmetric, it is implied that $x \sim x$ for all $x \in X$ and $f(x) \sim x$ for all $x \in A$, even if we do not say so explicitly. Transitivity may then imply further equivalences that are not explicitly specified: for an extreme example, "the equivalence relation on $\mathbb{Z}$ such that $n \sim n+1$ for all $n \in \mathbb{Z}$" makes every integer equivalent to every other integer, i.e. there is only one equivalence class.

Here is a result we will be able to prove later in the course:

THEOREM 1.4. *A closed orientable surface Σ_g of genus g is homeomorphic to a closed orientable surface Σ_h of genus h if and only if $g = h$.*

The hard part is showing that if $g \neq h$, then there cannot exist any continuous bijective map $f : \Sigma_g \to \Sigma_h$ with a continuous inverse. This requires techniques from the subject known as *algebraic* topology. The main idea will be that we can associate to each topological space X an algebraic object (e.g. a group) $H(X)$ such that any continuous map $f : X \to Y$ induces a homomorphism $f_* : H(X) \to H(Y)$, and such that compositions of continuous maps satisfy
$$(f \circ g)_* = f_* \circ g_*$$
and the identity map $\mathrm{Id} : X \to X$ gives rise to the identity map $H(X) \to H(X)$. These properties imply that whenever $f : X \to Y$ is a homeomorphism, $f_* : H(X) \to H(Y)$ must be an

isomorphism. Thus it suffices to compute the algebraic objects $H(\Sigma_g)$ and $H(\Sigma_h)$ and show that they are not isomorphic. (Recognizing non-isomorphic groups is often easier than recognizing non-homeomorphic spaces.)

The full classification of closed orientable surfaces up to homeomorphism is completed by the following result:

THEOREM 1.5. *Every closed connected and orientable surface is homeomorphic to Σ_g for some $g \geqslant 0$.*

The previous theorem implies of course that for any given surface, the value of g in this result is unique. For the moment, you can understand the word "orientable" to mean "embeddable in $\mathbb{R}^3$". There is a similar result for the non-orientable surfaces: notice that by the fourth definition we gave above for $\mathbb{RP}^2$, one can understand $\mathbb{RP}^2$ as the result of taking S^2, cutting out a hole (e.g. removing the southern hemisphere, thus leaving the northern hemisphere, which is also a disk $\mathbb{D}^2$) and then gluing in a Möbius strip. That is the first example of the following more general construction:

THEOREM 1.6. *Every closed connected and non-orientable surface is homeomorphic to a surface obtained from S^2 by cutting out finitely many holes and gluing in Möbius strips.*

Surfaces are the simplest interesting examples of more general topological spaces called **manifolds** (*Mannigfaltigkeiten*): a surface is a 2-dimensional manifold, while a smooth curve such as the circle S^1 is a 1-dimensional manifold. In general, one can consider n-dimensional manifolds (abbreviated as "n-manifolds") for any integer $n \geqslant 0$; obvious examples include $\mathbb{R}^n$, S^n and $\mathbb{RP}^n$. The classification problem becomes much harder when $n \geqslant 3$, e.g. the following difficult problem was open for almost exactly 100 years:

POINCARÉ CONJECTURE (solved by G. Perelman, c. 2004). *Suppose X is a closed and connected 3-manifold that is "simply connected" (i.e. every continuous map $f : S^1 \to X$ can be extended continuously to $\mathbb{D}^2 \to X$). Then X is homeomorphic to S^3.*

One of the more surprising developments in topology in the 20th century was that the analogue of this problem in dimensions greater than three turns out to be easier. We'll introduce the notion of "homotopy equvalence" (*Homotopieäquivalenz*) in a few weeks; it turns out that for closed 3-manifolds, the condition of being *simply connected* is equivalent to being homotopy equivalent to S^3. Thus the following two results are higher-dimensional versions of the Poincaré conjecture, but they were proved much earlier:

THEOREM 1.7 (S. Smale, c. 1960). *For every $n \geqslant 5$, every closed connected n-manifold homotopy equivalent to S^n is also homeomorphic to S^n.*

THEOREM 1.8 (M. Freedman, c. 1980). *Every closed connected 4-manifold homotopy equivalent to S^4 is also homeomorphic to S^4.*

(2) Differential topology. Though we will not have much time to talk about it in this semester, the neighboring field of "differential" topology modifies the classification problem by studying the following stronger notion of equivalence between spaces: X and Y are **diffeomorphic** (*diffeomorph*) if there exists a homeomorphism $f : X \to Y$ such that both f and f^{-1} are infinitely differentiable, i.e. C^∞, and f is in this case called a **diffeomorphism** (*Diffeomorphismus*). From your analysis courses, you at least know what this means if X and Y are open subsets of Euclidean spaces—defining "differentiability" on spaces more general than that requires some notions from the subject of *differential geometry.* In a nutshell, it requires X and Y to be spaces on which any map $X \to Y$ can at least *locally* (i.e. in a sufficiently small neighborhood of any point) be identified with a map between open subsets of Euclidean spaces, for which we know how to define derivatives.

Identifying a small neighborhood in X with an open subset of $\mathbb{R}^n$ is another way of saying that we can choose a set of n independent "coordinates" to describe the points in that neighborhood, and this is the fundamental property that defines X as an n-dimensional manifold. So talking about smooth maps and diffeomorphisms doesn't make sense for arbitrary topological spaces, but it does make sense for at least some class of manifolds, and these are the main objects of study in differential topology.

It turns out that up to dimension three, classification up to diffeomorphism is equivalent to classification up to homeomorphism:

THEOREM 1.9. *For $n \leqslant 3$, two n-manifolds X and Y are diffeomorphic if and only if they are homeomorphic.*

For $n = 1$ and $n = 2$, this theorem can be explained by the fact that both versions of the classification problem for n-manifolds are not that hard to solve explicitly (this was already understood in the 19th century), and the answer for both versions turns out to be the same. The story of $n = 3$ is much more complicated, as a complete classification of 3-manifolds is not known, but this theorem was proved in the first half of the 20th century by using the more combinatorial notion of "piecewise linear" manifolds as an intermediary notion between "smooth" and "topological" manifolds.

From dimension four upwards, all hell breaks loose. For example, there are "exotic" $\mathbb{R}^4$'s:

THEOREM 1.10. *There exist 4-manifolds that are homeomorphic but not diffeomorphic to $\mathbb{R}^4$.*

And from dimension seven upwards, there also tend to exist "exotic spheres":

THEOREM 1.11 (Kervaire and Milnor, 1963). *There exist exactly 28 distinct manifolds that are homeomorphic to S^7 but not diffeomorphic to each other.*

As you might guess, there is an algebraic phenomenon behind the appearance of the number 28 in this theorem: it is the order of a group. In every dimension n, one can define a group structure on the set of all smooth manifolds up to diffeomorphism that are homeomorphic to S^n. Milnor and Kervaire proved that when $n = 7$, this group has order 28. In the mean time, this group is quite well understood in most cases: it is sometimes trivial (e.g. for $n = 1, 2, 3, 5, 6$) and often nontrivial, but always finite. The only case for which almost nothing is known is $n = 4$; dimension four turns out to be the hardest case in differential topology, because it is on the borderline between "low dimensional" and "high dimensional" methods, where often neither set of methods applies. If you can solve the following open problem, you deserve an instant Ph.D. (and also a permanent job as a research mathematician, and possibly a Fields medal):

CONJECTURE 1.12 ("smooth Poincaré conjecture"). *Every manifold homeomorphic to S^4 is also diffeomorphic to S^4.*

It is difficult to say whether this conjecture is generally believed to be true or false.

(3) Fixed point problems. Here is a simpler class of problems on which we'll actually be able to prove something in this semester. Suppose $f : X \to X$ is a continuous map. We say $x \in X$ is a **fixed point** (*Fixpunkt*) of f if $f(x) = x$. The question is: under what assumptions on X is f guaranteed to have a fixed point? Note that this is fundamentally different from the fixed point results you've probably seen in analysis, e.g. the Banach fixed point theorem (also known as the *contraction mapping principle*) is a result about a special class of maps satisfying analytical conditions, it does not just apply to *every* continuous map on a certain space.

The simplest fixed point theorem in topology is a statement about maps on the n-dimensional disk $\mathbb{D}^n := \{\mathbf{x} \in \mathbb{R}^n \mid |\mathbf{x}| \leqslant 1\}$.

THEOREM 1.13 (Brouwer's fixed point theorem). *For every integer $n \geqslant 1$, every continuous map $f : \mathbb{D}^n \to \mathbb{D}^n$ has a fixed point.*

The case $n = 1$ is an easy consequence of the intermediate value theorem, but for $n \geqslant 2$, we need some techniques from algebraic topology. Here is a sketch of the argument; we will fill in the gaps over the course of the semester.

We argue by contradiction, so suppose there exists a continuous map $f : \mathbb{D}^n \to \mathbb{D}^n$ such that $f(x) \neq x$ for every $x \in \mathbb{D}^n$. Then there is a unique line in $\mathbb{R}^n$ connecting $f(x)$ to x for each $x \in \mathbb{D}^n$. Let $g(x) \in S^{n-1}$ denote the point on the boundary of $\mathbb{D}^n$ obtained by following the unique line from $f(x)$ through x until that line reaches the boundary of the disk. Note that if x is already on the boundary, then by this definition $g(x) = x$. It is not hard to convince yourself that what we've just defined is a continuous map

$$g : \mathbb{D}^n \to S^{n-1},$$

and if $i : S^{n-1} \hookrightarrow \mathbb{D}^n$ denotes the natural inclusion map for the subset $S^{n-1} \subset \mathbb{D}^n$, then g satisfies

$$g \circ i = \mathrm{Id}_{S^{n-1}} . \tag{1.1}$$

We claim that, actually, no such map can exist. The proof of this requires an algebraic invariant, whose complete construction will require some time and effort, but for now I'll just tell you the result: one can associate to each space X an abelian group $H_{n-1}(X)$ called the **singular homology** (*singuläre Homologie*) of X in dimension $n-1$, which satisfies the usual desirable properties that continuous maps $f : X \to Y$ induce group homomorphisms $f_* : H_{n-1}(X) \to H_{n-1}(Y)$ satisfying $(f \circ g)_* = f_* \circ g_*$ and $\mathrm{Id}_* = \mathbb{1}$. Crucially, one can also compute this invariant for both $\mathbb{D}^n$ and S^{n-1}, and the answers are

$$H_{n-1}(\mathbb{D}^n) = \{0\}, \qquad H_{n-1}(S^{n-1}) \cong \mathbb{Z}.$$

Now the relation (1.1) implies that $g_* \circ i_*$ is the identity map on $H_{n-1}(S^{n-1}) \cong \mathbb{Z}$, so in particular it is an isomorphism. But $g_* \circ i_*$ also factors through the trivial group $H_{n-1}(\mathbb{D}^n) \cong \{0\}$, and therefore can only be the trivial homomorphism. This is a contradiction, thus proving Brouwer's theorem.

We will discuss the construction of singular homology and carry out the required computations for the above argument in the last few weeks of this semester; homology and the closely related subject of **cohomology** (*Kohomologie*) will then be the main topic of Topology 2 next semester. But before all that, we will also spend considerable time on other invariants in algebraic topology, notably the fundamental group, which underlies the notion of "simply connected" spaces appearing in the Poincaré conjecture.

2. Metric spaces (April 20, 2023)

We now begin in earnest with point-set topology, which will be the main topic for the next three or four weeks. This subject is important but a little dry, so we will cover only the portions of it that seem absolutely necessary as groundwork for studying the more geometrically motivated questions discussed in the previous lecture.

The subject begins with metric spaces, because these are the most familiar examples of topological spaces. For most students, this material will be a review of things you've seen before in analysis courses. Almost everything in this lecture will be generalized to a wider and slightly more abstract context when we introduce topologies and topological spaces next week.

DEFINITION 2.1. A **metric space** (*metrischer Raum*) is a set X endowed with a function $d : X \times X \to \mathbb{R}$ that satisfies the following conditions for all $x, y, z \in X$:

(i) $d(x, y) \geqslant 0$;

(ii) $d(x, x) = 0$;
(iii) $d(x, y) = d(y, x)$, i.e. "symmetry";
(iv) $d(x, z) \leqslant d(x, y) + d(y, z)$, i.e. the "triangle inequality" (*Dreiecksungleichung*);
(v) $d(x, y) > 0$ whenever $x \neq y$.

The function d is then called a **metric** (*Metrik*). If d satisfies the first four conditions but not necessarily the fifth, then it is called a **pseudometric** (*Pseudometrik*).

Much of the theory of metric spaces makes sense for pseudometrics just as well as metrics, but we will see that some desirable and intuitively "obvious" facts become false when the positivity condition is dropped.

In any metric space (X, d), one can define the **open ball** (*offene Kugel*) of radius $r > 0$ about a given point $x \in X$ as

$$B_r(x) := \left\{ y \in X \;\middle|\; d(x, y) < r \right\}.$$

An arbitrary subset $\mathcal{U} \subset X$ is then called **open** (*offen*) if for every $x \in \mathcal{U}$, the ball $B_\epsilon(x)$ is contained in $\mathcal{U}$ for all $\epsilon > 0$ sufficiently small. (Of course it only needs to be true for one particular $\epsilon > 0$, since then it is true for all smaller ϵ as well.) Given a subset $A \subset X$, another subset $\mathcal{U} \subset X$ is called a **neighborhood** (*Umgebung*) of A in X if $\mathcal{U}$ contains some open subset of X that also contains A. Some books require the neighborhood itself to be open, but we will not require this; it makes very little difference in practice, but this bit of extra freedom in our definition will allow us to make certain other definitions and proofs a few words shorter now and then.

A subset $A \subset X$ is **closed** (*abgeschlossen*) if its complement $X \backslash A$ is open. Achtung: this is not the same thing as saying that A is not open. It is a common trap for beginners to think that every subset must be either open or closed, but in reality, most are neither—and some (e.g. X itself) are both.[1]

Whenever you encounter a set of axioms, you should ask yourself why we are studying these axioms in particular—why not a slightly different set of axioms? In the case of metrics, it's fairly obvious why we would want any notion of "distance" to satisfy conditions (i)–(iii) and (v), but perhaps the triangle inequality seems slightly less obvious. So, let us point out two obviously desirable properties that follow mainly from the triangle inequality:

- The "open ball" $B_r(x) \subset X$ is also an open subset in the sense of the definition given above. Indeed, for any $y \in B_r(x)$, we have $B_\epsilon(y) \subset B_r(x)$ for every $\epsilon < r - d(x, y)$ since every $z \in B_\epsilon(y)$ then satisfies
$$d(x, z) \leqslant d(x, y) + d(y, z) < d(x, y) + \epsilon < d(x, y) + r - d(x, y) = r.$$
- The function $d : X \times X \to [0, \infty)$ is *continuous* (see below for a review of the definition of continuity), since one can use the triangle inequality to show that for every $x, y, x', y' \in X$,
$$|d(x, y) - d(x', y')| \leqslant d(x, x') + d(y, y').$$

Also, while I'm sure you already accept without question that the distance between two distinct points should always be positive rather than zero, let us point out one "obvious" fact that would cease to be true if condition (v) were removed:

- For every $x \in X$, the subset $\{x\} \subset X$ is closed. Indeed, $X \backslash \{x\}$ is an open subset of X because for every $y \in X \backslash \{x\}$, the ball $B_\epsilon(y)$ is contained in $X \backslash \{x\}$ for all $\epsilon < d(x, y)$. (This of course presupposes that $d(x, y) > 0$.)

You're probably not used to thinking about pseudometric spaces much, so here is an example.

[1] Yes, the empty set $\varnothing \subset X$ is always open. Reread the definition carefully until you are convinced that this is true.

EXAMPLE 2.2. Let $X = (\mathbb{R} \times \{0,1\})/\sim$ for an equivalence relation defined by $(x,0) \sim (x,1)$ for every $x \neq 0$. We can think of this intuitively as a "real line with two zeroes" because it mostly looks just the same as $\mathbb{R}$ (each number $x \neq 0$ corresponding to the equivalence class of $(x,0)$ and $(x,1)$), but $x = 0$ is an exception, where there really are *two* distinct points $[(0,0)]$ and $[(0,1)]$ in X. We can then define $d : X \times X \to \mathbb{R}$ by

$$d([(x,i)],[(y,j)]) := |x - y| \quad \text{for } i,j \in \{0,1\},\ x,y \in \mathbb{R}.$$

This satisfies conditions (i)–(iv) for all the same reasons that the usual metric on $\mathbb{R}$ does, but condition (v) fails because

$$d([(0,0)],[(0,1)]) = 0$$

even though $[(0,0)] \neq [(0,1)]$.

EXERCISE 2.3. Show that for the pseudometric space X in Example 2.2, $\{[(0,0)]\} \subset X$ is not a closed subset.

DEFINITION 2.4. In a metric space (X,d), a sequence (*Folge*) $x_n \in X$ indexed by $n \in \mathbb{N}$ **converges to** (*konvergiert gegen*) a point $x \in X$ if for every $\epsilon > 0$, we have $x_n \in B_\epsilon(x)$ for all n sufficiently large. Equivalently, this means that for every neighborhood $\mathcal{U} \subset X$ of x, $x_n \in \mathcal{U}$ for all n sufficiently large. We use the notation

$$x_n \to x \quad \text{or} \quad \lim x_n = x$$

to indicate that x_n converges to x.

Note that in the second formulation of this definition, involving arbitrary neighborhoods instead of the open ball $B_\epsilon(x)$, one can understand the definition without knowing what the metric is—one only has to know what a "neighborhood" is, which means knowing which subsets are open and which are not. This will be the formulation that we need when we generalize sequences and convergence to arbitrary topological spaces.

Here is a similarly standard definition from analysis, for which we give three equivalent formulations.

DEFINITION 2.5. For two metric spaces (X,d_X) and (Y,d_Y), a map (*Abbildung*) $f : X \to Y$ is called **continuous** (*stetig*) if it satisfies any of the following equivalent conditions:

(a) For every $x_0 \in X$ and $\epsilon > 0$, there exists a number $\delta > 0$ such that $d_Y(f(x), f(x_0)) < \epsilon$ whenever $d_X(x,x_0) < \delta$, i.e. $f(B_\delta(x_0)) \subset B_\epsilon(f(x_0))$.

(b) For every open subset $\mathcal{U} \subset Y$, the preimage

$$f^{-1}(\mathcal{U}) := \{x \in X \mid f(x) \in \mathcal{U}\}$$

is an open subset of X.

(c) For every convergent sequence $x_n \in X$, $x_n \to x$ implies $f(x_n) \to f(x)$.

The equivalence of (a) and (b) is pretty easy to see: if (a) holds and $\mathcal{U} \subset Y$ is open, then for every $x_0 \in f^{-1}(\mathcal{U})$, the openness of $\mathcal{U}$ guarantees an $\epsilon > 0$ such that $f(x_0) \in B_\epsilon(f(x_0)) \subset \mathcal{U}$. But then condition (a) gives a $\delta > 0$ such that $f(B_\delta(x_0)) \subset B_\epsilon(f(x_0)) \subset \mathcal{U}$, implying $B_\delta(x_0) \subset f^{-1}(\mathcal{U})$, hence $\mathcal{U}$ is open and (b) therefore holds. Conversely, if (b) holds, then (a) holds because $B_\epsilon(f(x_0))$ is open and thus so is $f^{-1}(B_\epsilon(f(x_0)))$, which contains x_0 and therefore also (by openness) contains $B_\delta(x_0)$ for some $\delta > 0$.

Notice that conditions (b) and (c) do not require specific knowledge of the metric, but again only require knowing what an open subset is. Condition (b) is the one we will later use to define continuity in general topological spaces. It may be instructive to review why (b) and (c) are equivalent—especially because this is something that will turn out to be *false* in general for topological spaces, at least without some extra assumption.

Proof that (b) ⇔ (c). To show that (b) ⇒ (c), suppose $x_n \to x$ and $\mathcal{U} \subset Y$ is a neighborhood of $f(x)$. Then $\mathcal{U}$ contains an open set $\mathcal{V}$ containing $f(x)$, hence $f^{-1}(\mathcal{U})$ contains $f^{-1}(\mathcal{V})$ which contains x, and by condition (b), $f^{-1}(\mathcal{V})$ is also open, implying $f^{-1}(\mathcal{U})$ is a neighborhood of x. Convergence then implies that $x_n \in f^{-1}(\mathcal{U})$ and thus $f(x_n) \in \mathcal{U}$ for all n sufficiently large, which proves $f(x_n) \to f(x)$ since the neighborhood $\mathcal{U}$ was arbitrary.

For the other direction, we shall prove the contrapositive, i.e. we show that if (b) is false then so is (c). So assume there is an open subset $\mathcal{U} \subset Y$ such that $f^{-1}(\mathcal{U}) \subset X$ is not open. Being not open means that for some $x \in f^{-1}(\mathcal{U})$, no open ball about x is contained in $f^{-1}(\mathcal{U})$. As a consequence, for every $n \in \mathbb{N}$, we can find a point

$$x_n \in B_{1/n}(x) \quad \text{such that} \quad x_n \notin f^{-1}(\mathcal{U}),$$

meaning $f(x_n) \notin \mathcal{U}$. The sequence x_n then converges to x, since every neighborhood of x contains $B_{1/n}(x)$ for n sufficiently large, implying that x_n belongs to the given neighborhood for all large n. But $f(x_n)$ cannot converge to $f(x)$ since it never belongs to $\mathcal{U}$, which is a neighborhood of $f(x)$. □

I want to point out two things about the above proof. First, the proof that (b) ⇒ (c) never mentioned the metric, it only talked about neighborhoods and open sets—as a consequence, that implication will remain true when we reconsider all these notions in general topological spaces. But the proof that (c) ⇒ (b) did refer to the metric, because it used the precise definition of openness in terms of open balls. We will see that this implication does not actually hold in arbitrary topological spaces, though a mild modification of it does.

Definition 2.6. A map $f : X \to Y$ is a **homeomorphism** (*Homöomorphismus*) if it is continuous and bijective and its inverse $f^{-1} : Y \to X$ is also continuous.

Example 2.7. Consider $\mathbb{R}^n$ with the **standard Euclidean metric**

$$d_E(\mathbf{x}, \mathbf{y}) := |\mathbf{x} - \mathbf{y}| = \sqrt{\sum_{j=1}^{n} (x_j - y_j)^2}$$

for vectors $\mathbf{x} = (x_1, \ldots, x_n)$ and $\mathbf{y} = (y_1, \ldots, y_n)$ in $\mathbb{R}^n$. We claim that for any $\mathbf{x} \in \mathbb{R}^n$ and $r > 0$, $(B_r(\mathbf{x}), d_E)$ is homeomorphic to $(\mathbb{R}^n, d_E)$. (It follows of course that all open balls in $\mathbb{R}^n$ are also homeomorphic to each other, though it is perhaps easier to prove the latter directly.) To construct a homeomorphism, choose any continuous, increasing, bijective function $f : [0, r) \to [0, \infty)$ and define $F : B_r(\mathbf{x}) \to \mathbb{R}^n$ by

$$F(\mathbf{x}) = \mathbf{x} \quad \text{and} \quad F(\mathbf{x} + \mathbf{y}) = \mathbf{x} + f(|\mathbf{y}|)\frac{\mathbf{y}}{|\mathbf{y}|} \text{ for all } \mathbf{y} \in B_r(0)\backslash\{0\} \subset \mathbb{R}^n.$$

It is easy to check that both F and F^{-1} are then continuous.

One conclusion to draw from the above example is that the notion of "boundedness," which is very important in analysis, is not going to make much sense in topology. Indeed, we would like to consider two spaces as "equivalent" whenever they are homeomorphic, so topologically it would be meaningless to call a space bounded if another space homeomorphic to it is not. What plays this role instead is the somewhat stricter notion of *compactness*. To write down the correct definition, we need to have the notion of an **open covering** (*offene Überdeckung*): assume I is any set (the so-called "index set") and $\{\mathcal{U}_\alpha\}_{\alpha \in I}$ is a collection of open subsets $\mathcal{U}_\alpha \subset X$ labeled by elements $\alpha \in I$. We call $\{\mathcal{U}_\alpha\}_{\alpha \in I}$ an open covering/cover of a subset $A \subset X$ if

$$A \subset \bigcup_{\alpha \in I} \mathcal{U}_\alpha.$$

Definition 2.8. A subset K in a metric space (X, d) is **compact** (*kompakt*) if either of the following equivalent conditions is satisfied:

(a) Every open cover $\{\mathcal{U}_\alpha\}_{\alpha\in I}$ of K has a finite subcover (*eine endliche Teilüberdeckung*), i.e. there is a finite subset $\{\alpha_1, \dots, \alpha_N\} \subset I$ such that

$$K \subset \bigcup_{i=1}^{N} \mathcal{U}_{\alpha_i}.$$

(b) Every sequence $x_n \in K$ has a convergent subsequence with limit in K.

We call (X, d) itself a **compact space** if X is a compact subset of itself.

Compactness is probably the least intuitive definition in this course so far, and at this stage we can only justify it by saying that it has stood the test of time: many beautiful and useful theorems have turned out to be true for compact spaces and *only* compact spaces. The first of these is the following, which explains why, unlike boundedness, compactness really is a topologically invariant notion, i.e. if X is compact, then so is every space that is homeomorphic to it.

Theorem 2.9. *If $f : X \to Y$ is continuous and $K \subset X$ is compact, then so is $f(K) \subset Y$.*

Proof. If $\{\mathcal{U}_\alpha\}_{\alpha\in I}$ is an open cover of $f(K)$, then the sets $f^{-1}(\mathcal{U}_\alpha)$ are all open in X and thus form an open cover of K, which is compact, so there is a finite subset $\{\alpha_1, \dots, \alpha_N\} \subset I$ such that

$$K \subset \bigcup_{i=1}^{N} f^{-1}(\mathcal{U}_{\alpha_i}),$$

implying $f(K) \subset \bigcup_{i=1}^{N} \mathcal{U}_{\alpha_i}$, hence we have found a finite subcover of our given open cover of $f(K)$. □

One more remark about compactness: the equivalence of conditions (a) and (b) in Definition 2.8 is not so obvious, but is a fairly deep theorem called the *Bolzano-Weierstrass* theorem which you've probably seen proved in your analysis classes. We will prove an analogue of that theorem for topological spaces in Lecture 5, but it does not say that these two definitions are always equivalent—as with continuity, characterizing compactness via sequences becomes a slightly subtler issue in topological spaces, though the equivalence does hold for most of the spaces we actually care about.

Let's see some more examples now.

Example 2.10. For any metric space (X, d) and an arbitrary subset $A \subset X$, (A, d) is also a metric space. So for instance, we can use the Euclidean metric d_E on $\mathbb{R}^{n+1}$ to define a metric on the subset

$$S^n = \left\{ \mathbf{x} \in \mathbb{R}^{n+1} \;\middle|\; |\mathbf{x}| = 1 \right\},$$

the n-dimensional sphere.

Example 2.11. Any set X can be assigned the **discrete metric** (*diskrete Metrik*), defined by

$$d_D(x, y) = \begin{cases} 0 & \text{if } x = y, \\ 1 & \text{otherwise.} \end{cases}$$

This metric keeps every point at a measured distance away from every other point. So for instance, we can assign the discrete metric to $\mathbb{R}^n$ and compare it with the Euclidean metric d_E. We claim that the identity map on $\mathbb{R}^n$ defines a continuous map from $(\mathbb{R}^n, d_D)$ to $(\mathbb{R}^n, d_E)$, but it is not a homeomorphism, i.e. its inverse is not continuous. This follows immediately from the next exercise.

EXERCISE 2.12. Show that on any set X with the discrete metric d_D, every subset is open. In particular this includes the set $\{x\} \subset X$ for every $x \in X$. Conclude that a sequence x_n converges to x if and only if $x_n = x$ for all n sufficiently large, i.e. the sequence is "eventually constant". Then use this to prove the following statements:

(a) All maps from (X, d_D) to any other metric space are continuous.
(b) All continuous maps from $(\mathbb{R}^n, d_E)$ to (X, d_D) are constant.

EXAMPLE 2.13. Given two metric spaces (X, d_X) and (Y, d_Y), one can define a **product metric** on $X \times Y$ by

$$d_{X\times Y}((x,y),(x',y')) := \sqrt{d_X(x,x')^2 + d_Y(y,y')^2}.$$

This is the obvious generalization of the Euclidean metric, e.g. if X and Y are both $\mathbb{R}$ with its standard Euclidean metric, then $d_{X\times Y}$ becomes d_E on $\mathbb{R}^2$. But this is not the only reasonable choice of metric on $X \times Y$: for instance, one can also define a metric by

$$d'_{X\times Y}((x,y),(x',y')) := \max\left\{d_X(x,x'), d_Y(y,y')\right\}.$$

This metric is indeed different: for instance, if we again take X and Y to be the Euclidean $\mathbb{R}$, then an open ball with respect to $d'_{X\times Y}$ in $\mathbb{R}^2$ does not look circular, it looks rather like a square. On the other hand, this does not have a huge impact on the notion of open sets: it is not hard to show that the identity map from $(X \times Y, d_{X\times Y})$ to $(X \times Y, d'_{X\times Y})$ is always a homeomorphism.

DEFINITION 2.14. Two metrics d and d' on the same set X are called (topologically) **equivalent** if the identity map from (X, d) to (X, d') is a homeomorphism.

In light of the various ways we now have for defining what "continuous" means, equivalence of metrics can also be understood as follows:

- d and d' are equivalent if they both define the same notion of open subsets in X;
- d and d' are equivalent if they both define the same notion of convergence of sequences in X.

The characterization in terms of sequences is the subject of the next exercise.

EXERCISE 2.15. Suppose d_1 and d_2 are two metrics on the same set X. Show that the identity map defines a homeomorphism $(X, d_1) \to (X, d_2)$ if and only if the following condition is satisfied: for every sequence $x_n \in X$ and $x \in X$,

$$x_n \to x \text{ in } (X, d_1) \quad \Longleftrightarrow \quad x_n \to x \text{ in } (X, d_2).$$

EXAMPLE 2.16. In functional analysis, one often studies metric spaces whose elements are functions, and the exact choice of metric on such a space needs to be handled rather carefully. Consider for instance the set

$$X = C^0[-1,1] := \{\text{continuous functions } f : [-1,1] \to \mathbb{R}\}.$$

If we think of this as an infinite-dimensional vector space whose elements $f \in X$ are described by the (infinitely many) "coordinates" $f(t) \in \mathbb{R}$ for $t \in [-1,1]$, then the natural generalization of the Euclidean metric to such a space is

$$d_2(f,g) := \sqrt{\int_{-1}^{1} |f(t) - g(t)|^2 \, dt}.$$

This is the metric corresponding to the so-called "L^2-norm" on the space of functions $[-1,1] \to \mathbb{R}$. On the other hand, our alternative product metric discussed in Example 2.13 above generalizes to this space in the form

$$d_\infty(f,g) := \max_{t\in[-1,1]} |f(t) - g(t)|,$$

which is well defined since continuous functions on compact intervals always attain maxima. It is not hard to see that the identity map from (X, d_∞) to (X, d_2) is continuous, but is not a homeomorphism. Indeed, if $f_n \to f$ in (X, d_∞), then

$$d_2(f_n, f)^2 = \int_{-1}^{1} |f_n(t) - f(t)|^2 \, dt \leqslant \int_{-1}^{1} \max_t |f_n(t) - f(t)|^2 \, dt \leqslant 2 d_\infty(f_n, f)^2 \to 0,$$

proving that $f_n \to f$ also in (X, d_2). On the other hand, there exist sequences $f_n \in X$ such that $f_n \to 0$ with respect to d_2 but $d_\infty(f_n, 0) = 1$ for all n: just take a sequence of "bump" functions $f_n : [-1, 1] \to [0, 1]$ that all satisfy $f_n(0) = 1$ but vanish outside of progressively smaller neighborhoods of 0. These will satisfy $d_2(f_n, 0)^2 = \int_{-1}^{1} |f_n(t)|^2 \, dt \to 0$, but $d_\infty(f_n, 0) = \max_t |f_n(t)| = 1$ for all n, preventing convergence to 0 with respect to d_∞.

EXERCISE 2.17. Suppose (X, d_X) is a metric space and $\sim$ is an equivalence relation on X, with the resulting set of equivalence classes denoted by $X/\sim$. For equivalence classes $[x], [y] \in X/\sim$, define

$$d([x], [y]) := \inf \left\{ d_X(x, y) \;\middle|\; x \in [x],\ y \in [y] \right\}. \tag{2.1}$$

(a) Show that d is a metric on $X/\sim$ if the following assumption is added: for every triple $[x], [y], [z] \in X/\sim$, there exist representatives $x \in [x]$, $y \in [y]$ and $z \in [z]$ such that

$$d_X(x, y) = d([x], [y]) \quad \text{and} \quad d_X(y, z) = d([y], [z]).$$

Comment: The hard part is proving the triangle inequality.

(b) Consider the real projective n-space

$$\mathbb{RP}^n := S^n/\sim,$$

where $S^n := \{\mathbf{x} \in \mathbb{R}^{n+1} \mid |\mathbf{x}| = 1\}$ and the equivalence relation identifies antipodal points, i.e. $\mathbf{x} \sim -\mathbf{x}$. If d_X is the metric on S^n induced by the standard Euclidean metric on $\mathbb{R}^{n+1}$, show that the extra assumption in part (a) is satisfied, so that (2.1) defines a metric on $\mathbb{RP}^n$.

(c) For the metric defined on $\mathbb{RP}^n$ in part (b), show that the natural quotient projection $\pi : S^n \to \mathbb{RP}^n$ sending each $\mathbf{x} \in S^n$ to its equivalence class $[\mathbf{x}] \in \mathbb{RP}^n$ is continuous, and a subset $\mathcal{U} \subset \mathbb{RP}^n$ is open if and only if $\pi^{-1}(\mathcal{U}) \subset S^n$ is open (with respect to the metric d_X).

(d) Here is a very different example of a quotient space. Define

$$X = (-1, 1)^2 \backslash \{(0, 0)\} \subset \mathbb{R}^2$$

with the metric d_X induced by the Euclidean metric on $\mathbb{R}^2$. Now fix the function $f : X \to \mathbb{R} : (x, y) \mapsto xy$ and define the relation $p_0 \sim p_1$ for $p_0, p_1 \in X$ to mean that there exists a continuous curve $\gamma : [0, 1] \to X$ with $\gamma(0) = p_0$ and $\gamma(1) = p_1$ such that $f \circ \gamma$ is constant. Show that for this equivalence relation, the extra assumption of part (a) is not satisfied, and the distance function defined in (2.1) does not satisfy the triangle inequality.

(e) Despite our failure to define $X/\sim$ as a metric space in part (d), it is natural to consider the following notion: define a subset $\mathcal{U} \subset X/\sim$ to be *open* if and only if $\pi^{-1}(\mathcal{U})$ is an open subset of (X, d_X), where $\pi : X \to X/\sim$ denotes the natural quotient projection. We can then define a sequence $[x_n] \in X/\sim$ to be *convergent* to an element $[x] \in X/\sim$ if for every open subset $\mathcal{U} \subset X/\sim$ containing $[x]$, $[x_n] \in \mathcal{U}$ for all n sufficiently large. Find a sequence $[x_n] \in X/\sim$ and two elements $[x], [y] \in X/\sim$ such that

$$[x_n] \to [x] \quad \text{and} \quad [x_n] \to [y], \quad \text{but} \quad [x] \neq [y].$$

This could not happen if we'd defined convergence on $X/\sim$ in terms of a metric. (Why not?)

Exercise 2.18.

(a) Show that for any metric space (X, d),
$$d'(x, y) := \min\{1, d(x, y)\}$$
defines another metric on X which is equivalent to d. In particular, this means that every metric is equivalent to one that is bounded.

(b) Suppose (X, d_X) and (Y, d_Y) are metric spaces satisfying
$$d_X(x, x') \leqslant 1 \text{ for all } x, x' \in X, \qquad d_Y(y, y') \leqslant 1 \text{ for all } y, y' \in Y.$$
Now let $Z = X \cup Y$, and for $z, z' \in Z$ define
$$d_Z(z, z') = \begin{cases} d_X(z, z') & \text{if } z, z' \in X, \\ d_Y(z, z') & \text{if } z, z' \in Y, \\ 2 & \text{if } (z, z') \text{ is in } X \times Y \text{ or } Y \times X. \end{cases}$$
Show that d_Z is a metric on Z with the following property: a subset $\mathcal{U} \subset Z$ is open in (Z, d_Z) if and only if it is the union of two (possibly empty) open subsets of (X, d_X) and (Y, d_Y). In particular, X and Y are each both open and closed subsets of Z. (Recall that subsets of metric spaces are closed if and only if their complements are open.)

(c) Suppose (Z, d) is a metric space containing two disjoint subsets $X, Y \subset Z$ that are each both open and closed. Show that there exists no continuous map $\gamma : [0, 1] \to Z$ with $\gamma(0) \in X$ and $\gamma(1) \in Y$.

(d) Show that if (X, d) is a metric space with the discrete metric, then for every point $x \in X$, the subset $\{x\} \subset X$ is both open and closed.

3. Topological spaces (April 25, 2023)

We saw in the last lecture that most of the notions we want to consider in topology (continuous maps, homeomorphisms, convergence of sequences…) can be defined on metric spaces without specific reference to the metric, but using only our knowledge of which subsets are *open*. Moreover, one can define distinct but "equivalent" metrics on the same space for which the open sets match and therefore all these notions are the same. This suggests that we should view the notion of open sets as something more fundamental than a metric. The starting point of topology is to endow a set with the extra structure of a distinguished collection of subsets that we will call "open". The first question to answer is: what properties should we require this collection of subsets to have?

To motivate the axioms, let's revisit metric spaces for a moment and recall two important definitions. Both will also make sense in the context of topological spaces once we have fixed a definition for the latter.

Definition 3.1. Suppose X is a metric (or topological) space.

(a) The **interior** (*offener Kern* or *Inneres*) of a subset $A \subset X$ is the set
$$\mathring{A} = \{x \in A \mid \text{some neighborhood of } x \text{ in } X \text{ is contained in } A\}.$$
Points in this set are called **interior points** (*innere Punkte*) of A.

(b) The **closure** (*abgeschlossene Hülle* or *Abschluss*) of a subset $A \subset X$ is the set
$$\bar{A} = \{x \in X \mid \text{every neighborhood of } x \text{ in } X \text{ intersects } A\}.$$
Points in this set are called **cluster points** (*Berührpunkte*) of A.

The following exercise is easy, but it's worth thinking through why it is true.

Exercise 3.2. Show that for any subset $A \subset X$, the interior $\mathring{A}$ is the largest open subset of X that is contained in A, and the closure $\bar{A}$ is the smallest closed subset of X that contains A, i.e.

$$\mathring{A} = \bigcup_{\mathcal{U} \subset X \text{ open}, \, \mathcal{U} \subset A} \mathcal{U} \quad \text{and} \quad \bar{A} = \bigcap_{\mathcal{U} \subset X \text{ closed}, \, A \subset \mathcal{U}} \mathcal{U}.$$

I worded this exercise in a slightly sneaky way by calling the union of all the open sets inside A the "largest open subset of X that is contained in A": how do we actually know that this union of subsets is also open? This is the point: we know it because in a metric space, *arbitrary unions* of open subsets are also open. This follows almost immediately from the definitions in the previous lecture. It also implies (by taking complements) that arbitrary intersections of closed subsets are also closed, hence writing $\bar{A}$ as an intersection as in the exercise reveals that $\bar{A}$ is also a closed subset. These are properties you'd expect any reasonable notion of "open" or "closed" sets to have, so we will want to keep them.

What about intersections of open sets? Well, in metric spaces, arbitrary intersections of open sets need not be open, e.g. the intervals $(-1/n, 1/n) \subset \mathbb{R}$ are open for all $n \in \mathbb{N}$, but

$$\bigcap_{n \in \mathbb{N}} \left(-\frac{1}{n}, \frac{1}{n} \right) = \{0\}$$

is not an open subset of $\mathbb{R}$. Something slightly weaker is true, however: the intersection of any *two* open sets is open, and by an easy inductive argument, it follows that any *finite* intersection of open sets is open. Indeed, if $\mathcal{U}, \mathcal{V} \subset X$ are both open and $x \in \mathcal{U} \cap \mathcal{V}$, we know that $\mathcal{U}$ and $\mathcal{V}$ each contain balls about x for sufficiently small radii, so it suffices to take any radius small enough to fit inside both of them. (Why doesn't this necessarily work for an infinite intersection of open sets? Look at the example of the intervals $(-1/n, 1/n)$ above if you're not sure.) Taking complements, we also deduce from this discussion that arbitrary unions of closed subsets are not always closed, but *finite* unions are.

One last remark before we proceed: in any metric space X, the empty set $\varnothing$ and X itself are both open (and therefore also closed) subsets. With these observations as motivation, here is the definition on which everything else in this course will be based.

Definition 3.3. A **topology** (*Topologie*) on a set X is a collection[2] $\mathcal{T}$ of subsets of X satisfying the following axioms:

(i) $\varnothing \in \mathcal{T}$ and $X \in \mathcal{T}$;
(ii) For every subcollection $I \subset \mathcal{T}$, $\bigcup_{\mathcal{U} \in I} \mathcal{U} \in \mathcal{T}$;
(iii) For every pair $\mathcal{U}_1, \mathcal{U}_2 \in \mathcal{T}$, $\mathcal{U}_1 \cap \mathcal{U}_2 \in \mathcal{T}$.

The pair $(X, \mathcal{T})$ is then called a **topological space** (*topologischer Raum*), and we call the sets $\mathcal{U} \in \mathcal{T}$ the **open** subsets (*offene Teilmengen*) in $(X, \mathcal{T})$.

We can now repeat several definitions from the previous lecture in our newly generalized context.

Definitions 3.4. Assume $(X, \mathcal{T}_X)$ and $(Y, \mathcal{T}_Y)$ are topological spaces.

(1) A subset $A \subset X$ is **closed** (*abgeschlossen*) if $X \backslash A \in \mathcal{T}_X$.

[2] I am calling $\mathcal{T}$ a "collection" instead of a "set" in an attempt to minimize the inevitable confusion caused by $\mathcal{T}$ being a set whose elements are also sets. Strictly speaking, there is nothing wrong with saying "$\mathcal{T}$ is a subset of 2^X satisfying the following axioms. . . ," where 2^X is the set-theoretician's fancy notation for the set consisting of all subsets of X. But if you found that sentence confusing, my recommendation is to call $\mathcal{T}$ a "collection" instead of a "set".

(2) A map $f : X \to Y$ is **continuous** (*stetig*) if for all $\mathcal{U} \in \mathcal{T}_Y$, $f^{-1}(\mathcal{U}) \in \mathcal{T}_X$. Note that if we prefer to describe the topology in terms of closed rather than open subsets, then it is equivalent to say that for all $\mathcal{U} \subset Y$ closed, $f^{-1}(\mathcal{U}) \subset X$ is also closed.
(3) A **neighborhood** (*Umgebung*) of a subset $A \subset X$ is any subset $\mathcal{U} \subset X$ such that $A \subset \mathcal{V} \subset \mathcal{U}$ for some $\mathcal{V} \in \mathcal{T}_X$.
(4) A sequence (*Folge*) $x_n \in X$ **converges to** (*konvergiert gegen*) $x \in X$ (written "$x_n \to x$") if for every neighborhood $\mathcal{U} \subset X$ of x, $x_n \in \mathcal{U}$ holds for all $n \in \mathbb{N}$ sufficiently large.

Remark 3.5. One can equivalently define a topology $\mathcal{T}$ on a set X by specifying the *closed* sets $\mathcal{T}' := \{X \backslash \mathcal{U} \mid \mathcal{U} \in \mathcal{T}\}$. Then condition (ii) in Definition 3.3 is equivalent to

$$\bigcap_{A \in I} A \in \mathcal{T}' \quad \text{for all subcollections } I \subset \mathcal{T}',$$

and condition (iii) is equivalent to

$$A_1 \cup A_2 \in \mathcal{T}' \quad \text{for all } A_1, A_2 \in \mathcal{T}'.$$

For many topologies that one encounters in practice, it is not so easy to say what *all* the open sets look like, but much easier to describe a smaller subcollection that "generates" them.

Definition 3.6. Suppose $(X, \mathcal{T})$ is a topological space and $\mathcal{B} \subset \mathcal{T}$ is a subcollection of the open sets.

- We call $\mathcal{B}$ a **base** or **basis** (*Basis*)[3] for $\mathcal{T}$ if every set $\mathcal{U} \in \mathcal{T}$ is a union of sets in $\mathcal{B}$, i.e.
$$\mathcal{U} = \bigcup_{\mathcal{V} \in I} \mathcal{V} \quad \text{for some subcollection } I \subset \mathcal{B}.$$
- We call $\mathcal{B}$ a **subbase** or **subbasis** (*Subbasis*) for $\mathcal{T}$ if every set $\mathcal{U} \in \mathcal{T}$ is a union of finite intersections of sets in $\mathcal{B}$, i.e.
$$\mathcal{U} = \bigcup_{\alpha \in I} \mathcal{U}_\alpha$$
for some collection of subsets $\mathcal{U}_\alpha \subset X$ indexed by a (possibly empty) set I, such that for each $\alpha \in I$,
$$\mathcal{U}_\alpha = \mathcal{U}_\alpha^1 \cap \ldots \cap \mathcal{U}_\alpha^{N_\alpha}$$
for some $N_\alpha \in \mathbb{N}$ and $\mathcal{U}_\alpha^1, \ldots, \mathcal{U}_\alpha^{N_\alpha} \in \mathcal{B}$.

Every base is obviously also a subbase, though we'll see in a moment that the converse is not true. You should take a moment to convince yourself that given any collection $\mathcal{B}$ of subsets of X that cover all of X (meaning $X = \bigcup_{\mathcal{U} \in \mathcal{B}} \mathcal{U}$), $\mathcal{B}$ is a subbase of a unique topology on X, namely the smallest topology that contains $\mathcal{B}$. It consists of all unions of finite intersections of sets from $\mathcal{B}$, and we say in this case that the topology $\mathcal{T}$ is **generated by** the collection $\mathcal{B}$.

Example 3.7. The **standard topology** on $\mathbb{R}$ has the collection of all open intervals $\{(a, b) \subset \mathbb{R} \mid -\infty \leqslant a < b \leqslant \infty\}$ as a base. The smaller subcollection of half-infinite open intervals $\{(-\infty, a) \mid a \in \mathbb{R}\} \cup \{(a, \infty) \mid a \in \mathbb{R}\}$ is also a subbase, though not a base. (Why not?)

[3] Things got slightly confusing in Tuesday's lecture because when I stated the definition of a base, I neglected at first to require $\mathcal{B} \subset \mathcal{T}$, i.e. not only is every open set a union of sets from $\mathcal{B}$, but the sets in $\mathcal{B}$ are themselves also open, and as a result, *every* union of sets from $\mathcal{B}$ is also an open set. If one did not require the latter, then some stupid examples would be possible, e.g. the collection of one-point subsets would be a base for every topology. With the correct definition, however, $\mathcal{B}$ determines $\mathcal{T}$ uniquely, so taking $\mathcal{B}$ to consist of all one-point subsets automatically makes $\mathcal{T}$ the discrete topology.

Example 3.8. If (X, d) is any metric (or pseudometric) space, the natural topology on X induced by the metric is defined via the base

$$\mathcal{B} = \{B_r(x) \subset X \mid x \in X,\ r > 0\}.$$

Note that if d and d' are equivalent metrics as in Definition 2.14, then they induce the same topology on X: indeed, if the identity map $(X, d) \to (X, d')$ is a homeomorphism then it maps open sets to open sets. A topology that arises in this way from a metric is called **metrizable** (*metrisierbar*).

Example 3.9. On any set X, the **discrete topology** is the collection $\mathcal{T}$ consisting of *all* subsets of X. Take a moment to convince yourself that this is a topology, and moreover, it is metrizable—it can be defined via the discrete metric, see Definition 2.11. (Can you think of another metric on X that defines the same topology?) As a base for $\mathcal{T}$, we can take $\mathcal{B} = \{\{x\} \subset X \mid x \in X\}$. Note that since all subsets are open, all subsets are also closed! Moreover:

- Every map $f : X \to \mathbb{R}$ is continuous.
- A map $f : \mathbb{R} \to X$ is continuous if and only if it is constant. Here is a quick proof: for every $x \in X$, $\{x\} \subset X$ is both open and closed, so continuity requires $f^{-1}(x) \subset \mathbb{R}$ also to be both open and closed, but the only subsets of $\mathbb{R}$ with this property are $\mathbb{R}$ itself and the empty set.
- A sequence $x_n \in X$ converges to $x \in X$ if and only if $x_n = x$ for all $n \in \mathbb{N}$ sufficiently large.

Example 3.10. Also on any set X, one can define the **trivial** (also sometimes called the "indiscrete") topology $\mathcal{T} = \{\varnothing, X\}$. This topology has the distinguishing feature that every point $x \in X$ has only one neighborhood, namely the whole set. We then have:

- A map $f : X \to \mathbb{R}$ is continuous if and only if it is constant. Proof: Suppose f is continuous, $x_0 \in X$ and $f(x_0) = t \in \mathbb{R}$. Then for every $\epsilon > 0$, $f^{-1}(t - \epsilon, t + \epsilon)$ is an open subset of X containing x_0, so it is not $\varnothing$ and is therefore X. This proves
$$f(X) \subset \bigcap_{\epsilon > 0} (t - \epsilon, t + \epsilon) = \{t\}.$$
- All maps $f : \mathbb{R} \to X$ are continuous.
- $x_n \to x$ holds *always*, i.e. all sequences in X converge to all points! This proves that $(X, \mathcal{T})$ is not metrizable, as the limit of a convergent sequence in a metric space is always unique. (Prove it!)

Example 3.11. The **cofinite** topology on a set X is defined such that a proper subset $A \subset X$ is closed if and only if it is finite. Take a moment to convince yourself that this really defines a topology—see Remark 3.5. (Note that X itself is automatically closed but does not need to be finite, since it is not a *proper* subset of itself.) The neighborhoods of a point $x \in X$ are then all of the form $X \backslash \{x_1, \dots, x_N\}$ for arbitrary finite subsets $x_1, \dots, x_N \in X$ that do not include x.

Suppose $\mathcal{T}_1$ and $\mathcal{T}_2$ are two topologies on the same set X such that

$$\mathcal{T}_1 \subset \mathcal{T}_2,$$

meaning every open set in $(X, \mathcal{T}_1)$ is also an open set in $(X, \mathcal{T}_2)$. In this case we say that $\mathcal{T}_2$ is **stronger/finer/larger than** (*stärker/feiner als*) $\mathcal{T}_1$, and $\mathcal{T}_1$ is **weaker/coarser/smaller than** (*schwächer/gröber als*) $\mathcal{T}_2$. For example, since the open sets $\mathbb{R} \backslash \{x_1, \dots, x_N\}$ for the cofinite topology on $\mathbb{R}$ are also open with respect to its standard topology, we can say that the standard topology of $\mathbb{R}$ is stronger than the cofinite topology. On any set, the discrete topology is the strongest, and the trivial topology is the weakest. In general, having a stronger topology means that fewer sequences converge, fewer maps into X from other spaces are continuous, but more functions defined on X are continuous. In various situations, it is common and natural to specify a topology on a set as being the "strongest" or "weakest" possible topology subject to the condition that some given collection of maps are all continuous. We will see some examples of this below.

There are several natural ways in which a given topology on one or more spaces can induce a topology on some related space.

DEFINITION 3.12. $(X, \mathcal{T})$ determines on any subset $A \subset X$ the so-called **subspace topology** (*Unterraumtopologie*)

$$\mathcal{T}_A := \{\mathcal{U} \cap A \mid \mathcal{U} \in \mathcal{T}\}.$$

This is the weakest topology on A such that the natural inclusion $A \hookrightarrow X$ is a continuous map. (Prove it!)

EXAMPLE 3.13. The standard topology on $\mathbb{R}^{n+1}$ is the one defined via the Euclidean metric. We then assign the subspace topology to the set of unit vectors $S^n \subset \mathbb{R}^{n+1}$, meaning a subset $\mathcal{V} \subset S^n$ will be considered open in S^n if and only if $\mathcal{V} = S^n \cap \mathcal{U}$ for some open subset $\mathcal{U} \subset \mathbb{R}^{n+1}$. As you might expect, this is the same as the topology induced by the metric on S^n defined by restricting the Euclidean metric, but for a given open set $\mathcal{V} \subset S^n$, it is not always so easy to see an open set $\mathcal{U} \subset \mathbb{R}^{n+1}$ such that $\mathcal{V} = \mathcal{U} \cap S^n$. Such a set can be constructed as follows: for each $\mathbf{x} \in \mathcal{V}$, choose $\epsilon_{\mathbf{x}} > 0$ such that every $\mathbf{y} \in S^n$ satisfying $|\mathbf{y} - \mathbf{x}| < \epsilon_{\mathbf{x}}$ is also in $\mathcal{V}$. Then the set

$$\mathcal{U} := \bigcup_{\mathbf{x} \in \mathcal{V}} \{\mathbf{y} \in \mathbb{R}^{n+1} \mid |\mathbf{y} - \mathbf{x}| < \epsilon_{\mathbf{x}}\}$$

is a union of open balls and is thus open in $\mathbb{R}^{n+1}$, and satisfies $\mathcal{U} \cap S^n = \mathcal{V}$.

EXERCISE 3.14. Convince yourself that for any metric space (X, d) and subset $A \subset X$, the natural metrizable topology on (A, d) is precisely the subspace topology with respect to the topology on X induced by d.

DEFINITION 3.15. Given a collection of topological spaces $\{(X_\alpha, \mathcal{T}_\alpha)\}_{\alpha \in I}$ indexed by a set I such that $X_\alpha \cap X_\beta = \varnothing$ for all $\alpha \neq \beta$, the **disjoint union** (*disjunkte Vereinigung*) is the set $X := \bigcup_{\alpha \in I} X_\alpha$ with the topology

$$\mathcal{T} := \left\{ \bigcup_{\alpha \in I} \mathcal{U}_\alpha \;\middle|\; \mathcal{U}_\alpha \in \mathcal{T}_\alpha \text{ for all } \alpha \in I \right\}.$$

We typically denote the topological space $(X, \mathcal{T})$ defined in this way by

$$\coprod_{\alpha \in I} X_\alpha,$$

or for finite collections $I = \{1, \ldots, N\}$, $X_1 \amalg \ldots \amalg X_N$. The topology on this space is called the **disjoint union topology**.

EXERCISE 3.16. Show that the disjoint union topology $\mathcal{T}$ on $X = \coprod_\alpha X_\alpha$ is the strongest topology on this set such that for every $\alpha \in I$, the inclusion $X_\alpha \hookrightarrow X$ is continuous.

REMARK 3.17. A key feature of the disjoint union topology is that for every individual $\alpha \in I$, the subset $X_\alpha \subset X$ is both open and closed. It follows that there is no continuous path $\gamma : [0, 1] \to X$ with $\gamma(0) \in X_\alpha$ and $\gamma(1) \in X_\beta$ for $\alpha \neq \beta$, cf. Exercise 2.18(c).

REMARK 3.18. It is also often useful to be able to discuss disjoint unions $\coprod_\alpha X_\alpha$ in which the sets X_α and X_β need not be disjoint for $\alpha \neq \beta$, e.g. a common situation is where all X_α are taken to be the same fixed set Y. In this case we still want to treat X_α and X_β as disjoint "copies" of the same subset when $\alpha \neq \beta$, so that no element in the union can belong to more than one of them. One way to do this is by redefining the set $X = \coprod_\alpha X_\alpha$ as

$$X := \left\{ (\alpha, x) \;\middle|\; \alpha \in I,\; x \in X_\alpha \right\},$$

so that the disjoint union topology now literally becomes the collection of all subsets in X of the form

$$\bigcup_{\alpha \in I} \{\alpha\} \times \mathcal{U}_\alpha$$

with $\mathcal{U}_\alpha \subset X_\alpha$ open for every α, and in analogy with Exercise 3.16, this is the strongest topology on X for which the injective maps $X_\alpha \to X : x \mapsto (\alpha, x)$ are continuous for all $\alpha \in I$. We will usually not bother with this cumbersome notation when examples arise: just remember that whenever X_1 and X_2 are two sets, disjoint or otherwise, the set $X_1 \amalg X_2$ is defined so that its subsets $X_1 \subset X_1 \amalg X_2$ and $X_2 \subset X_1 \amalg X_2$ are disjoint.

Exercise 3.19. Let $I = \mathbb{R}$ and define X_α for each $\alpha \in \mathbb{R}$ to be the same space consisting of only one element; for concreteness, say $X_\alpha := \{0\} \subset \mathbb{R}$. According to the definition described above, this sets up an obvious bijection

$$\coprod_{\alpha \in \mathbb{R}} \{0\} := \{(\alpha, 0) \in \mathbb{R} \times \{0\}\} \to \mathbb{R},$$
$$(\alpha, 0) \mapsto \alpha.$$

Show that this bijection is a homeomorphism if we assign the discrete topology to $\mathbb{R}$ on the right hand side.

4. Products, sequential continuity and nets (April 27, 2023)

From now on, we'll adopt the following convention of terminology: if I say that X is a "**space**", then I mean X is a *topological* space unless I specifically say otherwise or the context clearly indicates that I mean something different (e.g. that X is a vector space). Similarly, if X and Y are spaces in the above sense and I refer to $f : X \to Y$ as a "**map**", then I typically mean that f is a *continuous* map unless the context indicates otherwise. We will sometimes have occasion to speak of maps $f : I \to X$ where X is a space but I is only a **set**, on which no topology has been specified: in this case no continuity is assumed since that notion is not well defined, but I will often try to be extra clear about it by calling f a "(not necessarily continuous) function" or something to that effect. I do not promise to be completely consistent about this, but hopefully my intended meaning will never be in doubt.

The previous lecture introduced two ways of inducing new topologies from old ones, namely on subspaces and on disjoint unions. It remains to discuss the natural topologies defined on products and quotients. We'll deal with the former in this lecture, and then use it to construct a surprising example illustrating the distinction between continuity and sequential continuity.

Definition 4.1. Given two spaces $(X_1, \mathcal{T}_1)$ and $(X_2, \mathcal{T}_2)$, the **product topology** $\mathcal{T}$ on $X_1 \times X_2$ is generated by the base

$$\mathcal{B} := \left\{ \mathcal{U}_1 \times \mathcal{U}_2 \subset X_1 \times X_2 \;\middle|\; \mathcal{U}_1 \in \mathcal{T}_1,\; \mathcal{U}_2 \in \mathcal{T}_2 \right\}.$$

Notice that if $X_1 \times X_2$ is endowed with the product topology, then both of the projection maps

$$\pi_1 : X_1 \times X_2 \to X_1 : (x_1, x_2) \mapsto x_1$$
$$\pi_2 : X_1 \times X_2 \to X_2 : (x_1, x_2) \mapsto x_2$$

are continuous. Indeed, for any open set $\mathcal{U}_1 \subset X_1$, $\pi_1^{-1}(\mathcal{U}_1) = \mathcal{U}_1 \times X_2$ is the product of two open sets and is therefore open in $X_1 \times X_2$; similarly, $\pi_2^{-1}(\mathcal{U}_2) = X_1 \times \mathcal{U}_2$ is open if $\mathcal{U}_2 \subset X_2$ is open.

Notice moreover that the intersection of these two sets is $\mathcal{U}_1 \times \mathcal{U}_2$, so one can form all open sets in the product topology as unions of sets that are finite intersections of the form $\pi_1^{-1}(\mathcal{U}_1) \cap \pi_2^{-1}(\mathcal{U}_2)$. In other words, the subcollection

$$\left\{\pi_1^{-1}(\mathcal{U}) \;\middle|\; \mathcal{U} \in \mathcal{T}_1\right\} \cup \left\{\pi_2^{-1}(\mathcal{U}) \;\middle|\; \mathcal{U} \in \mathcal{T}_2\right\}$$

forms a subbase for the product topology $\mathcal{T}$. This makes $\mathcal{T}$ the weakest (i.e. smallest) topology for which the projection maps π_1 and π_2 are both continuous.

That last observation leads us to the natural generalization of this discussion to infinite products, but the outcome turns out to be slightly different from what you probably would have expected.

Suppose $\{(X_\alpha, \mathcal{T}_\alpha)\}_{\alpha \in I}$ is a collection of spaces, indexed by an arbitrary (possibly infinite) set I. Their product can be defined as the set

$$\prod_{\alpha \in I} X_\alpha := \left\{ \text{functions } f : I \to \bigcup_{\alpha \in I} X_\alpha : \alpha \mapsto x_\alpha \text{ such that } x_\alpha \in X_\alpha \text{ for all } \alpha \in I \right\}.$$

Note that since I in this discussion is only a set with no topology, there is no assumption of continuity for the functions $\alpha \mapsto x_\alpha$. Whether the set I is infinite or finite, we can denote elements of the product space by

$$\{x_\alpha\}_{\alpha \in I} \in \prod_{\alpha \in I} X_\alpha,$$

so we think of each of the individual elements $x_\alpha \in X_\alpha$ as "coordinates" on the product.

Definition 4.2. The **product topology** (*Produkttopologie*) on $\prod_{\alpha \in I} X_\alpha$ is the weakest topology such that all of the projection maps

$$\pi_\alpha : \prod_{\beta \in I} X_\beta \to X_\alpha : \{x_\beta\}_{\beta \in I} \mapsto x_\alpha$$

for $\alpha \in I$ are continuous.

In particular, the product topology must contain $\pi_\alpha^{-1}(\mathcal{U}_\alpha)$ for every $\alpha \in I$ and $\mathcal{U}_\alpha \in \mathcal{T}_\alpha$, and it is the smallest topology that contains them, which means the sets $\pi_\alpha^{-1}(\mathcal{U}_\alpha)$ form a subbase. It is important to spell out precisely what this means. We have

$$\pi_\alpha^{-1}(\mathcal{U}_\alpha) = \left\{ \{x_\beta\}_{\beta \in I} \in \prod_{\beta \in I} X_\beta \;\middle|\; x_\alpha \in \mathcal{U}_\alpha \right\},$$

so in each of these sets, only a single coordinate is constrained. It follows that in a finite intersection of sets of this form, only *finitely many* of the coordinates will be constrained, while the rest remain completely free. This implies:

Proposition 4.3. *A base for the product topology on $\prod_{\alpha \in I} X_\alpha$ is formed by the collection of all subsets of the form $\prod_{\alpha \in I} \mathcal{U}_\alpha$ where $\mathcal{U}_\alpha \subset X_\alpha$ is open for every $\alpha \in I$ and $\mathcal{U}_\alpha \neq X_\alpha$ is satisfied for at most finitely many $\alpha \in I$.* □

The last part of the above statement makes no difference when the product is finite, but for infinite products, it means that arbitrary subsets of the form $\prod_{\alpha \in I} \mathcal{U}_\alpha \subset \prod_{\alpha \in I} X_\alpha$ are not open just because $\mathcal{U}_\alpha \subset X_\alpha$ is open for every α. Dropping the "at most finitely many" condition would produce a much stronger topology with very different properties (see Exercise 4.6 below).

Exercise 4.4. Show that a sequence $\{x_\alpha^n\}_{\alpha \in I} \in \prod_{\alpha \in I} X_\alpha$ for $n \in \mathbb{N}$ converges as $n \to \infty$ to $\{x_\alpha\}_{\alpha \in I} \in \prod_{\alpha \in I} X_\alpha$ in the product topology if and only if for all $\alpha \in I$, the individual sequences x_α^n converge in X_α to x_α.

Exercise 4.5. Show that for any other space Y, a map $f : Y \to \prod_{\alpha \in I} X_\alpha$ is continuous if and only if $\pi_\alpha \circ f : Y \to X_\alpha$ is continuous for every $\alpha \in I$.

There is a special notation for the product set in the case where all the X_α are taken to be the same fixed space X: the product $\prod_{\alpha \in I} X$ has an obvious identification with the set of all (not necessarily continuous) functions $I \to X$, and we write

$$X^I := \prod_{\alpha \in I} X = \{\text{(not necessarily continuous) functions } f : I \to X\}.$$

For example we could now write $\mathbb{R}^n = \mathbb{R}^{\{1,\dots,n\}}$ if we preferred. The notation is motivated in part by the combinatorial observation that if X and I are both finite sets with a and b elements respectively, then X^I has a^b elements. The case $X = \{0, 1\}$ is popular in abstract set theory since $\{0, 1\}^I = \{f : I \to \{0, 1\}\}$ has a straightforward interpretation as the set of all subsets of I, which is often abbreviated as $2^I := \{0, 1\}^I$. But this example is not very interesting for topology since $\{0, 1\}$ is not a very interesting topological space (no matter which topology you put on it—there are only four choices). When X is a more interesting space, the most important thing to understand about X^I comes from Exercise 4.4: a sequence of functions $f_n \in X^I$ converges to $f \in X^I$ if and only if it converges **pointwise**, i.e.

$$f_n(\alpha) \to f(\alpha) \qquad \text{for every } \alpha \in I.$$

The product topology on X^I is therefore also sometimes called the **topology of pointwise convergence** (*punktweise Konvergenz*).

Exercise 4.6. Assume I is an infinite set and $\{(X_\alpha, \mathcal{T}_\alpha)\}_{\alpha \in I}$ is a collection of topological spaces. In addition to the usual product topology on $\prod_\alpha X_\alpha$, one can define the so-called *box topology*, which has a base of the form

$$\left\{ \prod_{\alpha \in I} \mathcal{U}_\alpha \;\middle|\; \mathcal{U}_\alpha \in \mathcal{T}_\alpha \text{ for all } \alpha \in I \right\}.$$

(a) Compared with the usual product topology, is the box topology stronger, weaker, or neither?
(b) What does it mean for a sequence in $\prod_\alpha X_\alpha$ to converge in the box topology? In particular, consider the case where all the X_α are a fixed space X and $\prod_\alpha X$ is identified with the space of all functions $X^I = \{f : I \to X\}$; what does it mean for a sequence of functions $f_n : I \to X$ to converge in the box topology to a function $f : I \to X$?

With examples like these at our disposal, we can now address the following important question in full generality:

Question 4.7. *To what extent are the following conditions for maps $f : X \to Y$ between topological spaces equivalent?*

- *$f^{-1}(\mathcal{U}) \subset X$ is open for every open set $\mathcal{U} \subset Y$;*
- *For every convergent sequence $x_n \to x$ in X, $f(x_n) \to f(x)$ in Y.*

The first condition is ordinary continuity, while the second is called **sequential continuity** (*Folgenstetigkeit*). We proved in Lecture 2 that these two conditions are equivalent for maps between *metric spaces*, and if you look again at the proof that (b)⇒(c) in the discussion following Definition 2.5, you'll see that it still makes sense in arbitrary topological spaces, proving:

Theorem 4.8. *For arbitrary topological spaces X and Y, all continuous maps $X \to Y$ are sequentially continuous.* □

The converse is trickier. Look again at the proof in Lecture 2 that (c)⇒(b) for Definition 2.5. That proof specifically referred to open balls about a point, so it is not so clear how to make sense of it in topological spaces where there is no metric. We can see however that the argument still works if we can remove all mention of open balls and replace it with the following lemma:

"LEMMA" 4.9. *In any topological space X, a subset $A \subset X$ is not open if and only if there exists a point $x \in A$ and a sequence $x_n \in X \backslash A$ such that $x_n \to x$.*

I've put the word "lemma" in quotation marks here for a very good reason: as written, the statement is *false*, and so is the converse of Theorem 4.8! Sequential continuity does not always imply continuity. Here is a counterexample.

EXAMPLE 4.10 (cf. [**Jän05**, §6.3]). Let $X = C^0([0,1],[-1,1]) \subset [-1,1]^{[0,1]}$, i.e. X is the set of all continuous functions $f : [0,1] \to [-1,1]$, and we assign to it the subspace topology as a subset of the space $[-1,1]^{[0,1]}$ of *all* functions $f : [0,1] \to [-1,1]$. In other words, X carries the topology of pointwise convergence. Next, define Y to be the same set, but with the topology induced by the L^2-metric

$$d_2(f,g) = \sqrt{\int_0^1 |f(t) - g(t)|^2 \, dt}.$$

Now consider the identity map from X to Y:

$$\Phi : X \to Y : f \mapsto f.$$

If $f_n \to f$ is a convergent sequence in X, then the functions converge pointwise, so $|f_n - f|^2$ converges pointwise to 0, and we claim that this implies $\int_0^1 |f_n(t) - f(t)|^2 \, dt \to 0$. This requires a fundamental result from measure theory, Lebesgue's *dominated convergence theorem* (see e.g. [**LL01**, §1.8] or [**Rud87**, Theorem 1.34]): it states that if g_n is a sequence of measurable functions that converge almost everywhere to g and all satisfy $|g_n| \leqslant G$ for some Lebesgue integrable function G, then $\int g_n$ converges to $\int g$. In the present case, the hypotheses are satisfied since the functions f_n take values in the bounded domain $[-1,1]$, which bounds $|f_n - f|$ uniformly below the constant (and thus integrable) function 2. We conclude that $d_2(f_n, f) \to 0$, hence Φ is sequentially continuous.

To show however that Φ is continuous, we would need to find for every $\epsilon > 0$ a neighborhood $\mathcal{U} \subset X$ of 0 such that $\Phi(\mathcal{U}) \subset B_\epsilon(0) \subset Y$. The trouble here is that neighborhoods in X (with the product topology) are somewhat peculiar objects: if $\mathcal{U}$ is one, then it contains some open set containing 0, which means it contains at least one of the sets $\prod_{\alpha \in [0,1]} \mathcal{U}_\alpha$ in our base for the product topology, where the $\mathcal{U}_\alpha$ are all open neighborhoods of 0 in $[-1,1]$ but there is at most a finite subset $I \subset [0,1]$ consisting of $\alpha \in [0,1]$ for which $\mathcal{U}_\alpha \neq [-1,1]$. Now choose a continuous function $f : [0,1] \to [0,1]$ that vanishes on the finite subset I but equals 1 on a "large" subset of $[0,1] \backslash I$. Depending how many points are in I, you may have to make this function oscillate very rapidly back and forth between 0 and 1, but since I is only finite, you can still do this such that the measure of the domain on which $f = 1$ is as close to 1 as you like, which makes $d_2(f, 0)$ also only slightly less than 1. In particular, f belongs to the neighborhood $\mathcal{U}$ in X but not to $B_\epsilon(0) \subset Y$ if ϵ is sufficiently small.

We deduce from the above example that "Lemma" 4.9 is not always true, since it would imply that continuity and sequential continuity are equivalent. We are led to ask: what extra hypotheses could be added so that the lemma holds?

DEFINITION 4.11. Given a point x in a space X, a **neighborhood base** (*Umgebungsbasis*) for x is a collection $\mathcal{B}$ of neighborhoods of x such that every neighborhood of x contains some $\mathcal{U} \in \mathcal{B}$.

Recall that a set I is **countable** (*abzählbar*) if it admits an injection into the natural numbers $\mathbb{N}$. This definition allows I to be either finite or infinite; if it is "countably infinite" then we can equivalently say that I admits a bijection with $\mathbb{N}$. This is also equivalent to saying that there exists a sequence $\{x_n \in I\}_{n \in \mathbb{N}}$ that includes every point of I. For example, it is easy to show that the set $\mathbb{Q}$ of rational numbers is countable, but Cantor's famous "diagonal" argument shows that $\mathbb{R}$ is not.

DEFINITION 4.12 (the countability axioms). A space X is called **first countable** ("*X erfüllt das erste Abzählbarkeitsaxiom*") if every point in x has a countable neighborhood base. We call X **second countable** ("*X erfüllt das zweite Abzählbarkeitsaxiom*") if its topology has a countable base.

It is easy to see that every second countable space is also first countable: if X has a countable base $\mathcal{B}$, then for each $x \in X$, the collection of sets in $\mathcal{B}$ that contain x is a countable neighborhood base for x. The next example shows that the converse is false.

EXAMPLE 4.13. If X has the discrete topology, then it is first countable because for each $x \in X$, one can form a neighborhood base out of the single open set $\{x\} \subset X$. But X is second countable if and only if X itself is a countable set (prove it!), so e.g. $\mathbb{R}$ with the discrete topology is first but not second countable.

EXAMPLE 4.14. All metric spaces are first countable. Indeed, for every $x \in X$, the collection of open balls $B_{1/n}(x) \subset X$ for $n \in \mathbb{N}$ forms a countable neighborhood base. (Note that Example 4.13 is a special case of this, so not all metric spaces are second countable.)

We can now prove a corrected version of "Lemma" 4.9. Let us first make a useful general observation that follows directly from the axioms of a topology.

LEMMA 4.15. *In any space X, a subset $A \subset X$ is open if and only if every point $x \in A$ has a neighborhood $\mathcal{V} \subset X$ that is contained in A.*

PROOF. If the latter condition holds, then A is the union of open sets contained in such neighborhoods and is therefore open. Conversely, if A is open, then A itself can be taken as the desired neighborhood of every $x \in A$. □

LEMMA 4.16. *In any first countable topological space X, a subset $A \subset X$ is not open if and only if there exists a point $x \in A$ and a sequence $x_n \in X \backslash A$ such that $x_n \to x$.*

PROOF. If $A \subset X$ is open, then for every $x \in A$ and sequence $x_n \in X$ converging to x, we cannot have $x_n \in X \backslash A$ for all n since A is a neighborhood of x. This is true so far for *all* topological spaces, with or without the first countability axiom, but the latter will be needed in order to prove the converse. So, suppose now that $A \subset X$ is not open, which by Lemma 4.15, means there exists a point $x \in A$ such that no neighborhood $\mathcal{V} \subset X$ of x is contained in A. Fix a countable neighborhood base $\mathcal{U}_1, \mathcal{U}_2, \mathcal{U}_3, \dots$ for x.

It will make our lives slightly easier if the neighborhood base is a nested sequence, meaning

$$X \supset \mathcal{U}_1 \supset \mathcal{U}_2 \supset \mathcal{U}_3 \supset \dots \ni x,$$

and we claim that this can be assumed without loss of generality. Indeed, set $\mathcal{U}_1' := \mathcal{U}_1$, and if $\mathcal{U}_2$ is not contained in $\mathcal{U}_1'$, consider instead the set $\mathcal{U}_2 \cap \mathcal{U}_1'$, which is also a neighborhood of x and therefore (by the definition of a neighborhood base) contains $\mathcal{U}_n$ for some $n \in \mathbb{N}$. Since $\mathcal{U}_n$ is contained in $\mathcal{U}_1'$, we then set $\mathcal{U}_2' := \mathcal{U}_n$. Now continue this process by setting $\mathcal{U}_3' := \mathcal{U}_m$ such that $\mathcal{U}_m \subset \mathcal{U}_2' \cap \mathcal{U}_3$ and so forth. This algorithm produces a nested sequence $\mathcal{U}_1' \supset \mathcal{U}_2' \supset \mathcal{U}_3' \supset \dots$ such that $\mathcal{U}_n' \subset \mathcal{U}_n$ for every n, hence the new neighborhoods also form a neighborhood base for x. Let us replace our original sequence with the nested sequence and continue to call it $\{\mathcal{U}_n\}_{n \in \mathbb{N}}$.

With this new assumption in place, observe that since none of the neighborhoods $\mathcal{U}_n$ can be contained in A, there exists a sequence of points

$$x_n \in \mathcal{U}_n \quad \text{such that} \quad x_n \notin A.$$

This sequence converges to x since every neighborhood $\mathcal{V} \subset X$ of x contains one of the $\mathcal{U}_N$, implying that for all $n \geqslant N$,

$$x_n \in \mathcal{U}_n \subset \mathcal{U}_N \subset \mathcal{V}.$$

□

Combining this lemma with our proof in Lecture 2 that sequential continuity implies continuity in metric spaces yields:

COROLLARY 4.17. *For any spaces X and Y such that X is first countable, every sequentially continuous map $X \to Y$ is also continuous.* □

It is possible to generalize this result beyond first countable spaces, but it requires expanding our notion of what a "sequence" can be. If you think of a sequence in X as a map from the (ordered) set of natural numbers $\mathbb{N}$ to X, then one possible way to generalize is to consider more general partially ordered sets as domains. Recall that a binary relation $\prec$ defined on some subset of all pairs of elements in a set I is called a **partial order** (*Halbordnung* or *Teilordnung*) if it satisfies (i) $x \prec x$ for all x, (ii) $x \prec y$ and $y \prec x$ implies $x = y$, and (iii) $x \prec y$ and $y \prec z$ implies $x \prec z$. We write "$x \succ y$" as a synonym for "$y \prec x$", and the set I together with its partial order $\prec$ is called a **partially ordered set** (*partiell geordnete Menge*). One obvious example is $(\mathbb{N}, \leqslant)$, though unlike this example (which is *totally* ordered), it is not generally required in a partially ordered set $(I, \prec)$ that every pair of elements $x, y \in I$ satisfy either $x \prec y$ or $y \prec x$. We will see more exotic examples below.

DEFINITION 4.18. A **directed set** (*gerichtete Menge*) $(I, \prec)$ consists of a set I with a partial order $\prec$ such that for every pair $\alpha, \beta \in I$, there exists an element $\gamma \in I$ with $\gamma \succ \alpha$ and $\gamma \succ \beta$.

The natural numbers $(\mathbb{N}, \leqslant)$ clearly form a directed set, but in topology, one also encounters many interesting examples of directed sets that need not be totally ordered or countable.

EXAMPLE 4.19. If X is a space and $x \in X$, one can define a directed set $(I, \prec)$ where I is the set of all neighborhoods of x in X, and $\mathcal{U} \prec \mathcal{V}$ for $\mathcal{U}, \mathcal{V} \in I$ means $\mathcal{V} \subset \mathcal{U}$. This is a directed set because given any pair of neighborhoods $\mathcal{U}, \mathcal{V} \subset X$ of x, the intersection $\mathcal{U} \cap \mathcal{V}$ is also a neighborhood of x and thus defines an element of I with $\mathcal{U} \cap \mathcal{V} \subset \mathcal{U}$ and $\mathcal{U} \cap \mathcal{V} \subset \mathcal{V}$. Note that neither of $\mathcal{U}$ and $\mathcal{V}$ need be contained in the other, so they might not satisfy either $\mathcal{U} \prec \mathcal{V}$ or $\mathcal{V} \prec \mathcal{U}$.

DEFINITION 4.20. Given a space X, a **net** (*Netz*) $\{x_\alpha\}_{\alpha \in I}$ in X is a function $I \to X : \alpha \mapsto x_\alpha$, where $(I, \prec)$ is a directed set.

DEFINITION 4.21. We say that a net $\{x_\alpha\}_{\alpha \in I}$ in X **converges** to $x \in X$ if for every neighborhood $\mathcal{U} \subset X$ of x, there exists an element $\alpha_0 \in I$ such that $x_\alpha \in \mathcal{U}$ for every $\alpha \succ \alpha_0$.

Convergence of nets is also sometimes referred to in the literature as *Moore-Smith convergence*, see e.g. [**Kel75**]. Note that a net $\{x_\alpha\}_{\alpha \in I}$ whose underlying directed set is $(I, \prec) = (\mathbb{N}, \leqslant)$ is simply a sequence, and the above definition then reduces to the usual notion of convergence for a sequence. We can now prove the most general corrected version of "Lemma" 4.9.

LEMMA 4.22. *In any space X, a subset $A \subset X$ is not open if and only if there exists a point $x \in A$ and a net $\{x_\alpha\}_{\alpha \in I}$ in X that converges to x but satisfies $x_\alpha \notin A$ for every $\alpha \in I$.*

PROOF. If $A \subset X$ is open then it is a neighborhood of every $x \in A$, so the nonexistence of such a net is an immediate consequence of Definition 4.21. Conversely, if A is not open, then Lemma 4.15 provides a point $x \in A$ such that for every neighborhood $\mathcal{V} \subset X$ of x, there exists a point

$$x_{\mathcal{V}} \in \mathcal{V} \quad \text{such that} \quad x_{\mathcal{V}} \notin A.$$

Taking $(I, \prec)$ to be the directed set of all neighborhoods of x, ordered by inclusion as in Example 4.19, the collection of points $\{x_{\mathcal{V}}\}_{\mathcal{V} \in I}$ is now a net which converges to x since for every neighborhood $\mathcal{U} \subset X$ of x,

$$\mathcal{V} \succ \mathcal{U} \quad \Rightarrow \quad x_{\mathcal{V}} \in \mathcal{V} \subset \mathcal{U}.$$

□

Putting all this together leads to the following statement equating continuity with a generalized notion of sequential continuity. The proof is just a repeat of arguments we've already worked through, but we'll spell it out for the sake of completeness.

THEOREM 4.23. *For any spaces X and Y, a map $f : X \to Y$ is continuous if and only if for every net $\{x_\alpha\}_{\alpha \in I}$ in X converging to a point $x \in X$, the net $\{f(x_\alpha)\}_{\alpha \in I}$ in Y converges to $f(x)$.*

PROOF. Suppose f is continuous and $\{x_\alpha\}_{\alpha \in I}$ is a net in X converging to $x \in X$. Then for any neighborhood $\mathcal{U} \subset Y$ of $f(x)$, $f^{-1}(\mathcal{U}) \subset X$ is a neighborhood of x, hence there exists $\alpha_0 \in I$ such that $\alpha \succ \alpha_0$ implies $x_\alpha \in f^{-1}(\mathcal{U})$, or equivalently, $f(x_\alpha) \in \mathcal{U}$. This proves that $\{f(x_\alpha)\}_{\alpha \in I}$ converges in the sense of Definition 4.21 to $f(x)$.

To prove the converse, let us suppose that $f : X \to Y$ is not continuous, so there exists an open set $\mathcal{U} \subset Y$ for which $f^{-1}(\mathcal{U}) \subset X$ is not open. Then by Lemma 4.22, there exists a point $x \in f^{-1}(\mathcal{U})$ and a net $\{x_\alpha\}_{\alpha \in I}$ in X that converges to x but satisfies $x_\alpha \notin f^{-1}(\mathcal{U})$ for every $\alpha \in I$. Now $\{f(x_\alpha)\}_{\alpha \in I}$ is a net in Y that does not converge to $f(x)$, since $\mathcal{U}$ is an open neighborhood of $f(x)$ but $f(x_\alpha)$ is never in $\mathcal{U}$. □

Nets take a bit of getting used to in comparison with sequences. The following addendum to Example 4.10 may help in this regard, but it may also make you feel deeply unsettled.

EXAMPLE 4.24. For the identity map $\Phi : X \to Y$ in Example 4.10, one could extract from the above proof an example of a net $\{x_\alpha\}_{\alpha \in I}$ in X that converges to 0 without $\{\Phi(x_\alpha)\}_{\alpha \in I}$ converging to 0 in Y, but here is perhaps a slightly simpler example. Define I as the set of all finite subsets of $[0, 1]$, with the partial order $A \prec B$ for $A, B \subset [0, 1]$ defined to mean $A \subset B$. Note that $(I, \prec)$ is a directed set since for any two finite subsets $A, B \subset [0, 1]$, $A \cup B$ is also a finite subset and thus an element of I. Now choose for each $A \in I$ a continuous function

$$f_A : [0, 1] \to [0, 1]$$

such that $f_A|_A = 0$ but $\int_0^1 |f_A(t)|^2\, dt > 1/4$. The net $\{\Phi(f_A)\}_{A \in I}$ in Y clearly does not converge to 0 since none of these functions belong to the ball $B_{1/2}(0)$ in Y. But $\{f_A\}_{A \in I}$ does converge to 0 in X: indeed, since X has the product topology, any neighborhood $\mathcal{U} \subset X$ of 0 contains some open neighborhood of 0 that is of the form $\prod_{\alpha \in [0,1]} \mathcal{U}_\alpha$ for open neighborhoods $\mathcal{U}_\alpha \subset [-1, 1]$ of 0 such that $\mathcal{U}_\alpha = [-1, 1]$ for all α outside of some finite subset $A_0 \subset [0, 1]$. It follows that for all $A \in I$ with $A \succ A_0 \in I$,

$$f_A(\alpha) = 0 \in \mathcal{U}_\alpha \text{ for all } \alpha \in A_0,$$

implying $f_A \in \mathcal{U}$.

5. Compactness (May 2, 2023)

We saw in our discussion of metric spaces (Lecture 2) that boundedness is not a meaningful notion in topology, i.e. even if we have data such as a metric with which to define what a "bounded" set is, it may still be homeomorphic to sets that are not bounded. Instead, we consider *compact* sets, a notion that is topologically invariant. The main definition carries over from Lecture 2 with no change.

DEFINITION 5.1. Given a space X and subset $A \subset X$, an **open cover/covering** (*offene Überdeckung*) of A is a collection of open subsets $\{\mathcal{U}_\alpha \subset X\}_{\alpha \in I}$ such that $A \subset \bigcup_{\alpha \in I} \mathcal{U}_\alpha$.

We will also occasionally use the notation

$$A \subset \bigcup_{\mathcal{U} \in \mathcal{O}} \mathcal{U}$$

to indicate an open covering of A, where $\mathcal{O}$ is a collection of open subsets of X, i.e. $\mathcal{O} \subset \mathcal{T}$, where $\mathcal{T}$ is the topology of X.

DEFINITION 5.2. A subset $A \subset X$ is **compact** (*kompakt*) if every open cover of A has a finite subcover (*eine endliche Teilüberdeckung*), i.e. given an arbitrary open cover $\{\mathcal{U}_\alpha\}_{\alpha \in I}$ of A, one can always find a finite subset $\{\alpha_1, \ldots, \alpha_N\} \subset I$ such that $A \subset \mathcal{U}_{\alpha_1} \cup \ldots \cup \mathcal{U}_{\alpha_N}$. We say that X itself is a **compact space** if X is a compact subset of itself.

EXERCISE 5.3. Show that a subset $A \subset X$ is compact if and only if A with the subspace topology is a compact space.

EXAMPLE 5.4. For any space X with the discrete topology, a subset $A \subset X$ is compact if and only if A is finite. Indeed, the collection of subsets $\{\{x\} \subset X\}_{x \in A}$ forms an open covering of A in the discrete topology, and it has a finite subcovering if and only if A is finite, hence compactness implies finiteness. The converse follows from the next example.

EXAMPLE 5.5. In any space X, every finite subset $A \subset X$ is compact. Indeed, for $A = \{a_1, \ldots, a_N\}$ with an open covering $\{\mathcal{U}_\alpha\}_{\alpha \in I}$, pick any $\alpha_i \in I$ with $a_i \in \mathcal{U}_{\alpha_i}$ for $i = 1, \ldots, N$, then the sets $\mathcal{U}_{\alpha_1}, \ldots, \mathcal{U}_{\alpha_N}$ form an open subcover.

EXAMPLE 5.6. A subset $A \subset \mathbb{R}^n$ in Euclidean space with its standard topology is compact if and only if it is closed and bounded. This is known as the *Heine-Borel theorem*, and in one direction it is easy to prove; see Exercise 5.7 below. For the other direction, you have probably seen a proof in your analysis classes of the *Bolzano-Weierstrass theorem*, stating that if A is closed and bounded then every sequence in A has a convergent subsequence with limit in A; we say in this case that A is *sequentially compact*. We will prove in the following that compactness and sequential compactness are equivalent for second countable spaces, and every subset of $\mathbb{R}^n$ is second countable (see Exercise 5.9 below). A frequently occurring concrete example is the sphere

$$S^n \subset \mathbb{R}^{n+1},$$

which is a closed and bounded subset of $\mathbb{R}^{n+1}$ and is therefore compact.

EXERCISE 5.7. Show that in any metric space, compact subsets must be both closed and bounded.
Hint: For closedness, you may want to assume the theorem proved below that compact first countable spaces are also sequentially compact—recall that all metric spaces are first countable.

REMARK 5.8. Note that the converse of Exercise 5.7 is generally false: being closed and bounded is not enough for compactness in arbitrary metric spaces. Here is an important class of examples from functional analysis: a vector space $\mathcal{H}$ with an inner product $\langle\ ,\ \rangle$ is called a **Hilbert**

space (*Hilbertraum*) if it is complete (meaning all Cauchy sequences converge) with respect to the metric $d(x,y) = \sqrt{\langle x-y, x-y\rangle}$. The closed unit ball $\bar{B}_1(0) = \{x \in \mathcal{H} \mid \langle x,x\rangle \leqslant 1\}$ is clearly both closed and bounded in $\mathcal{H}$, and it is compact if $\mathcal{H}$ is finite dimensional since, in this case, $\mathcal{H}$ is both linearly isomorphic and homeomorphic to $\mathbb{R}^n$ (or $\mathbb{C}^n$ in the complex case) with its standard inner product. But if $\mathcal{H}$ is infinite dimensional, then $\bar{B}_1(0)$ contains an infinite orthonormal set $e_1, e_2, e_3, \ldots$, i.e. satisfying

$$\langle e_i, e_i\rangle = 1 \text{ for all } i, \qquad \langle e_i, e_j\rangle = 0 \text{ if } i \neq j.$$

It then follows by a standard argument of Euclidean geometry that $d(e_i, e_j) = \sqrt{2}$ whenever $i \neq j$, so for any $r < \sqrt{2}/2$, no ball of radius r in $\mathcal{H}$ can contain more than one of these vectors. It follows that $\{B_r(x) \mid x \in \mathcal{H}\}$ is an open cover of $\bar{B}_1(0)$ that has no finite subcover. This way of characterizing the distinction between finite- and infinite-dimensional Hilbert spaces in terms of the compactness of the unit ball has useful applications, e.g. in the theory of elliptic PDEs. The latter has many quite deep applications in geometry and topology, for instance the index theory of Atiyah-Singer (see [**Boo77**, **BB85**]), gauge-theoretic invariants of smooth manifolds [**DK90**], and the theory of pseudoholomorphic curves in symplectic topology [**MS12**, **Wen18**].

Exercise 5.9. A space X is called **separable** (*separabel*) if it contains a countable subset $A \subset X$ that is also **dense** (*dicht*), meaning the closure[4] of A is X.

(a) Show that if X is a metric space and $A \subset X$ is a dense subset, then the collection of open balls $\{B_{1/n}(x) \subset X \mid n \in \mathbb{N},\ x \in A\}$ forms a base for the topology of X.
(b) Deduce that every separable and metrizable space is second countable.
(c) Show that $\mathbb{R}^n$ with its standard topology is separable.
(d) Show that if X is any second countable space, then every subset $A \subset X$ with the subspace topology is also second countable.

Example 5.10. A union of finitely many compact subsets in a space X is also compact. (This is an easy exercise.)

The next result implies that closed subsets in compact spaces are also compact.

Proposition 5.11. *For any compact subset $K \subset X$, if $A \subset X$ is closed and also is contained in K, then A is compact.*

Proof. Suppose $\{\mathcal{U}_\alpha\}_{\alpha\in I}$ is an open cover of A. Since A is closed, $X \backslash A$ is open, so that supplementing the collection $\{\mathcal{U}_\alpha\}_{\alpha\in I}$ with $X\backslash A$ defines an open cover of X, and therefore also an open cover of K. Since K is compact, there is then a finite subset $\{\alpha_1, \ldots, \alpha_N\} \subset I$ such that

$$K \subset \mathcal{U}_{\alpha_1} \cup \ldots \cup \mathcal{U}_{\alpha_N} \cup (X\backslash A).$$

But $A \subset K$ is disjoint from $X\backslash A$, so this means $A \subset \mathcal{U}_{\alpha_1} \cup \ldots \cup \mathcal{U}_{\alpha_N}$, and we have found the desired finite subcover for A. □

The following theorem is just a repeat of Theorem 2.9, but in the more general context of topological rather than metric spaces. The proof carries over word for word.

Theorem 5.12. *If $f : X \to Y$ is continuous and $K \subset X$ is compact, then so is $f(K) \subset Y$.* □

Now would be a good moment to introduce the quotient topology, since it provides a large class of new examples of compact spaces.

[4] We gave the definition of the term *closure* in Lecture 3 (see Definition 3.1), originally in the context of metric spaces, but the same definition carries over to general topological spaces without change.

Definition 5.13. Suppose X is a space and $\sim$ is an equivalence relation on X, with the set of equivalence classes denoted by $X/\sim$. The **quotient topology** on $X/\sim$ is the strongest topology for which the natural projection map $\pi : X \to X/\sim$ sending each point $x \in X$ to its equivalence class $[x] \in X/\sim$ is continuous. Equivalently, a subset $\mathcal{U} \subset X/\sim$ is open in the quotient topology if and only if $\pi^{-1}(\mathcal{U})$ is an open subset of X.

I suggest you pause for a moment to make sure you understand why the two descriptions of the quotient topology in that definition are equivalent. Applying Theorem 5.12 to the continuous projection $\pi : X \to X/\sim$, we now have:

Corollary 5.14. *For any compact space X with an equivalence relation $\sim$, $X/\sim$ with the quotient topology is also compact.* □

Example 5.15. Since S^n is compact, so is $\mathbb{RP}^n = S^n/\{\mathbf{x} \sim -\mathbf{x}\}$ if we assign it the quotient topology. (Note that by Exercise 2.17(c), the quotient topology on $\mathbb{RP}^n$ is metrizable, and can be defined in terms of a natural metric induced on the quotient from the Euclidean metric restricted to S^n.)

Exercise 5.16. The space S^1, known as the **circle**, is normally defined as the unit circle in $\mathbb{R}^2$ and endowed with the subspace topology (induced by the Euclidean metric on $\mathbb{R}^2$). Show that the following spaces with their natural quotient topologies are both homeomorphic to S^1:

(a) $\mathbb{R}/\mathbb{Z}$, meaning the set of equivalence classes of real numbers where $x \sim y$ means $x - y \in \mathbb{Z}$.
(b) $[0, 1]/\sim$, where $0 \sim 1$.

For the next example, we introduce a convenient piece of standard notation. The quotient of a space X by a subset $A \subset X$ is defined as

$$X/A := X/\sim$$

with the quotient topology, where the equivalence relation is defined such that $x \sim y$ for every $x, y \in A$ and otherwise $x \sim x$ for all $x \in X$. In other words, X/A is the result of modifying X by "collapsing A to a point".

(c) Convince yourself that for every $n \in \mathbb{N}$, S^n is homeomorphic to $\mathbb{D}^n/S^{n-1}$, where

$$\mathbb{D}^n := \{\mathbf{x} \in \mathbb{R}^n \mid |\mathbf{x}| \leqslant 1\}.$$

Remark: Part (b) becomes a special case of part (c) if we replace $[0, 1]$ *by* $\mathbb{D}^1 = [-1, 1]$.

The remainder of this lecture will be concerned with the extent to which compactness is equivalent to the notion of **sequential compactness** (*Folgenkompaktheit*), defined as follows:

Definition 5.17. A subset $A \subset X$ is **sequentially compact** if every sequence in A has a subsequence that converges to a point in A.

As you might guess from our discussion of sequential continuity in the previous lecture, compactness and sequential compactness are not generally equivalent without some extra condition. But as with continuity, one obtains a result free of extra conditions by replacing sequences with *nets.*

Definition 5.18. Suppose $(I, \prec)$ is a directed set and $\{x_\alpha\}_{\alpha \in I}$ is a net in a space X. A point $x \in X$ is called a **cluster point** (*Häufungspunkt*) of $\{x_\alpha\}_{\alpha \in I}$ if for every neighborhood $\mathcal{U} \subset X$ of x and every $\alpha_0 \in I$, there exists $\alpha \succ \alpha_0$ such that $x_\alpha \in \mathcal{U}$.

Notice that the above definition is almost identical to that of *convergence* of $\{x_\alpha\}_{\alpha \in I}$ to x (see Definition 4.21), only the roles of "for every" and "there exist" have been reversed at the end. Informally, x being a cluster point does not require x_α to be arbitrarily close to x for *all* sufficiently large α, but only that one should be able to find *some* α arbitrarily large for which x_α is arbitrarily close. You should take a moment to think about what this definition means in the special case $(I, \prec) = (\mathbb{N}, \leqslant)$, where the net becomes a sequence, so the notion should be already familiar.

Definition 5.19. Given two directed sets $(I, \prec)$ and $(J, \prec)$, and nets $\{x_\alpha\}_{\alpha \in I}$ and $\{y_\beta\}_{\beta \in J}$ in a space X, we call $\{y_\beta\}_{\beta \in J}$ a **subnet** (*Teilnetz*) of $\{x_\alpha\}_{\alpha \in I}$ if $y_\beta = x_{\phi(\beta)}$ for all $\beta \in J$ and some function $\phi : J \to I$ with the property that for every $\alpha_0 \in I$, there exists $\beta_0 \in J$ for which $\beta \succ \beta_0$ implies $\phi(\beta) \succ \alpha_0$.

If $(I, \prec)$ and $(J, \prec)$ in the above definition are both $(\mathbb{N}, \leqslant)$ so that $\{x_\alpha\}_{\alpha \in I}$ and $\{y_\beta\}_{\beta \in I}$ become sequences x_n and y_k respectively, then y_k will be a subnet of x_n if it is of the form $y_k = x_{n_k}$ for some sequence $n_k \in \mathbb{N}$ satisfying $\lim_{k \to \infty} n_k = \infty$. This agrees with at least one of the standard definitions of the term **subsequence** (*Teilfolge*); a slightly stricter definition would require the sequence n_k to be monotone, but this difference is harmless. One should however be careful not to fall into the trap of thinking that a subnet of a sequence is always a subsequence—even if $(I, \prec) = (\mathbb{N}, \leqslant)$, Definition 5.19 allows much more general choices for the directed set $(J, \prec)$ and the function $\phi : J \to \mathbb{N}$ underlying a subnet of a sequence. In particular, the following lemma cannot be used to find convergent subsequences without imposing further conditions (cf. Lemma 5.22 below).

Lemma 5.20. *A net $\{x_\alpha\}_{\alpha \in I}$ in X has a cluster point at $x \in X$ if and only if it has a subnet convergent to x.*

Proof. Let us prove that a convergent subnet can always be derived from a cluster point x. Let $\mathcal{N}_x$ denote the set of all neighborhoods of x in X, and define $J = I \times \mathcal{N}_x$ with a partial order $\prec$ defined by

$$(\alpha, \mathcal{U}) \succ (\beta, \mathcal{V}) \quad \Leftrightarrow \quad \alpha \succ \beta \text{ and } \mathcal{U} \subset \mathcal{V}.$$

This makes $(J, \prec)$ a directed set since $(I, \prec)$ is already a directed set and the intersection of two neighborhoods is a neighborhood contained in both. Now since x is a cluster point of the net $\{x_\alpha\}_{\alpha \in I}$, there exists a function $\phi : J \to I$ such that for all $(\beta, \mathcal{U}) \in J$, $\phi(\beta, \mathcal{U}) =: \alpha$ satisfies $\alpha \succ \beta$ and $x_\alpha \in \mathcal{U}$. It is then straightforward to check that $\{x_{\phi(\beta,\mathcal{U})}\}_{(\beta,\mathcal{U}) \in J}$ is a subnet convergent to x.

The converse is easier, so I will leave it as an exercise. □

Here is the most general result relating compactness to nets.

Theorem 5.21. *A space X is compact if and only if every net in X has a convergent subnet.*

Proof. We prove first that if X is compact, then every net $\{x_\alpha\}_{\alpha \in I}$ has a cluster point (and therefore by Lemma 5.20 a convergent subnet). Arguing by contradiction, suppose no $x \in X$ is a cluster point of $\{x_\alpha\}_{\alpha \in I}$. Then one can associate to every $x \in X$ a neighborhood $\mathcal{U}_x$ and an element $\alpha_x \in I$ such that for every $\alpha \succ \alpha_x$, $x_\alpha \notin \mathcal{U}_x$. Without loss of generality let us suppose the neighborhoods $\mathcal{U}_x$ are all open. Then the collection of sets $\{\mathcal{U}_x\}_{x \in X}$ forms an open cover of X, and therefore has a finite subcover since X is compact. This means there is a finite set of points $x_1, \dots, x_N \in X$ such that $X = \mathcal{U}_{x_1} \cup \dots \cup \mathcal{U}_{x_N}$. Now since $(I, \prec)$ is a directed set, we can find an element $\beta \in I$ satisfying

$$\beta \succ \alpha_{x_i} \text{ for all } i = 1, \dots, N,$$

hence $x_\beta \notin \mathcal{U}_{x_i}$ for every $i = 1, \dots, N$. But the latter sets cover X, so this is impossible, and we have found a contradiction.

For the converse, we shall prove that if X is not compact then there exists a net with no cluster point. Being noncompact means one can find a collection $\mathcal{O}$ of open subsets such that $X = \bigcup_{\mathcal{U} \in \mathcal{O}} \mathcal{U}$ but no finite subcollection of them has union equal to X. Define I to be the set of all finite subcollections of the sets in $\mathcal{O}$, so by assumption, one can associate to every $\mathcal{A} \in I$ a point $x_{\mathcal{A}} \in X$ satisfying

$$x_{\mathcal{A}} \notin \bigcup_{\mathcal{U} \in \mathcal{A}} \mathcal{U}. \tag{5.1}$$

Define a partial order $\prec$ on I by

$$\mathcal{A} \prec \mathcal{B} \quad \Leftrightarrow \quad \mathcal{A} \subset \mathcal{B},$$

and notice that $(I, \prec)$ is now a directed set since the union of any two finite subcollections is another finite subcollection that contains both. This makes $\{x_{\mathcal{A}}\}_{\mathcal{A} \in I}$ a net in X, and we claim that it has no cluster point. Indeed, if $x \in X$ is a cluster point of $\{x_{\mathcal{A}}\}_{\mathcal{A} \in I}$, then since the sets in $\mathcal{O}$ cover X, there is a set $\mathcal{V} \in \mathcal{O}$ that is a neighborhood of x, and it follows that there must exist some $\mathcal{A} \succ \{\mathcal{V}\}$ in I for which

$$x_{\mathcal{A}} \in \mathcal{V} \subset \bigcup_{\mathcal{U} \in \mathcal{A}} \mathcal{U}.$$

This contradicts (5.1) and thus proves the claim that there is no cluster point. □

The next step is to impose countability axioms so that Theorem 5.21 gives us corollaries about sequential compactness.

Lemma 5.22. *If $x_n \in X$ is a sequence with a cluster point at $x \in X$ and x has a countable neighborhood base, then x_n has a subsequence converging to x.*

Proof. As in the proof of Lemma 4.16, we can assume without loss of generality that our countable neighborhood base has the form of a nested sequence of neighborhoods

$$X \supset \mathcal{U}_1 \supset \mathcal{U}_2 \supset \ldots \ni x.$$

Since x is a cluster point, we can choose $k_1 \in \mathbb{N}$ so that $x_{k_1} \in \mathcal{U}_1$, and then inductively for each $n \in \mathbb{N}$, choose $k_n \in \mathbb{N}$ such that $x_{k_n} \in \mathcal{U}_n$ and $k_n > k_{n-1}$. Then x_{k_n} is a subsequence of x_n and it converges to x, since for all neighborhoods $\mathcal{V} \subset X$ of x, we have $\mathcal{V} \supset \mathcal{U}_N$ for some $N \in \mathbb{N}$, implying

$$n \geqslant N \quad \Rightarrow \quad x_{k_n} \in \mathcal{U}_n \subset \mathcal{U}_N \subset \mathcal{V}.$$

□

Corollary 5.23. *If X is compact and first countable, then it is also sequentially compact.* □

Example 5.24. Though it is not so easy to see this, the space $[0, 1]^{\mathbb{R}}$ of (not necessarily continuous) functions $\mathbb{R} \to [0, 1]$ with the topology of pointwise convergence is compact, but not sequentially compact. Compactness follows directly from a deep result known as Tychonoff's theorem, which we will discuss in the next lecture. For the construction of a sequence in $[0, 1]^{\mathbb{R}}$ with no convergent subsequence, see Exercise 6.5.

To prove compactness from sequential compactness, it turns out that we will need to invoke the second countability axiom. In practice, almost all of the spaces that topologists spend their time thinking about are second countable, resulting from the fact that most of them are separable and metrizable (see Exercise 5.9). One useful property shared by all second countable (but not necessarily compact) spaces is the following.

Lemma 5.25. *If X is second countable, then every open cover of X has a countable subcover.*

Proof. Assume $\{\mathcal{U}_\alpha\}_{\alpha \in I}$ is an open cover of X and $\mathcal{B}$ is a countable base. Then each $\mathcal{U}_\alpha$ is a union of sets in $\mathcal{B}$, and the collection of all sets in $\mathcal{B}$ that are contained in some $\mathcal{U}_\alpha$ is a countable subcollection $\mathcal{B}' \subset \mathcal{B}$ that also covers X. Let us denote $\mathcal{B}' = \{\mathcal{V}_1, \mathcal{V}_2, \mathcal{V}_3, \ldots\}$. We can now choose for each $\mathcal{V}_n \in \mathcal{B}'$ an element $\alpha_n \in I$ such that $\mathcal{V}_n \subset \mathcal{U}_{\alpha_n}$, and $\{\mathcal{U}_{\alpha_n}\}_{n \in \mathbb{N}}$ is then a countable subcover of $\{\mathcal{U}_\alpha\}_{\alpha \in I}$. □

If you now take the second half of the proof of Theorem 5.21 and redo it with the focus on sequences instead of nets, and with Lemma 5.25 in mind, the result is the following.

THEOREM 5.26. *If X is second countable and sequentially compact, then it is compact.*

PROOF. We need to show that every open cover of X has a finite subcover. Since X is second countable, we can first use Lemma 5.25 to reduce the given open cover to a *countable* subcover $\mathcal{U}_1, \mathcal{U}_2, \mathcal{U}_3, \ldots \subset X$. Now arguing by contradiction, suppose that X is sequentially compact but the sets $\mathcal{U}_1, \ldots, \mathcal{U}_n$ do not cover X for any $n \in \mathbb{N}$, hence there exists a sequence $x_n \in X$ such that

$$x_n \notin \mathcal{U}_1 \cup \ldots \cup \mathcal{U}_n \tag{5.2}$$

for every $n \in \mathbb{N}$. Some subsequence x_{k_n} then converges to a point $x \in X$, which necessarily lies in $\mathcal{U}_N$ for some $N \in \mathbb{N}$. It follows that x_{k_n} also lies in $\mathcal{U}_N$ for all n sufficiently large, but this contradicts (5.2) as soon as $k_n \geqslant N$. □

EXERCISE 5.27. Consider the space

$$X = \left\{ f \in [0,1]^{\mathbb{R}} \mid f(x) \neq 0 \text{ for at most countably many points } x \in \mathbb{R} \right\},$$

with the subspace topology that it inherits from $[0,1]^{\mathbb{R}}$.

(a) Show that X is sequentially compact.
Hint: For any sequence $f_n \in X$, the set $\bigcup_{n \in \mathbb{N}} \{x \in \mathbb{R} \mid f_n(x) \neq 0\}$ is also countable.
(b) For each $x \in \mathbb{R}$, define $\mathcal{U}_x = \{f \in X \mid -1 < f(x) < 1\}$. Show that the collection $\{U_x \subset X \mid x \in \mathbb{R}\}$ forms an open cover of X that has no finite subcover, hence X is not compact.

Corollary 5.23 and Theorem 5.26 combine to give the following result that is easy to remember:

COROLLARY 5.28. *A second countable space is compact if and only if it is sequentially compact.* □

A loose end: We know from Exercise 5.9 that every separable metric space is second countable, thus Corollary 5.28 implies the equivalence of compactness and sequential compactness for separable metric spaces, which includes most of the metric spaces that one uses in practice. However, more than this was claimed in Lecture 2: the equivalence should hold in *all* metric spaces, and this does not quite follow from what we've proved here. The missing ingredient needed is the notion of *total boundedness*: one can show that every sequentially compact set A in a metric space X is **totally bounded** (*total beschränkt*), meaning that for every $\epsilon > 0$, A is contained in the union of finitely many balls of radius ϵ. Taking $\epsilon = 1/n$ for $n \in \mathbb{N}$ then provides a countable collection of open balls covering A, which can serve as a substitute for the countable subcover we used in the proof of Theorem 5.26. We will not go further into the details here, since this is a topology and not an analysis course, and we will not need the result going forward.

6. Tychonoff's theorem and the separation axioms (May 4, 2023)

Topic 1: Products of compact spaces. Here is a result that may sound less surprising at first than it actually is.

THEOREM 6.1 (Tychonoff's theorem). *For any collection of compact spaces $\{X_\alpha\}_{\alpha \in I}$, the product $\prod_{\alpha \in I} X_\alpha$ is compact.*

NONMATHEMATICAL REMARK. Thinking like an Anglophone may lead you to false assumptions about the pronunciation of the name Tychonoff, e.g. I was mispronouncing it for years until I finally looked up the name on Wikipedia in the context of teaching this course. The original Russian spelling is Тихонов, which would normally get transliterated into English as Tikhonov. The reason he instead became known outside of Russia as Tychonoff is that his papers were published in German, hence different phonetic conventions.

When I is a finite set, Theorem 6.1 says something not at all surprising, and the proof is straightforward, so let's start with that.

Proof of Theorem 6.1 for finite products. By induction, it will suffice to prove that if X and Y are both compact spaces then so is $X \times Y$. We will do so by showing that every net in $X \times Y$ has a convergent subnet. Recall that a net $\{(x_\alpha, y_\alpha)\}_{\alpha \in I}$ in $X \times Y$ converges to $(x, y) \in X \times Y$ if and only if the nets $\{x_\alpha\}_{\alpha \in I}$ in X and $\{y_\alpha\}_{\alpha \in I}$ in Y converge to x and y respectively. (The corresponding fact about sequences was proved in Exercise 4.4—the proof for nets is the same.) Now, since X is compact, $\{x_\alpha\}_{\alpha \in I}$ has a subnet $\{x_{\phi(\beta)}\}_{\beta \in J}$ convergent to some point $x \in X$, where J is some other directed set with a suitable function $\phi : J \to I$. Compactness of Y implies in turn that $\{y_{\phi(\beta)}\}_{\beta \in J}$ has a subnet $\{y_{\phi(\psi(\gamma))}\}_{\gamma \in K}$ convergent to some point $y \in Y$. We therefore obtain a subnet

$$\{(x_{\phi \circ \psi(\gamma)}, y_{\phi \circ \psi(\gamma)})\}_{\gamma \in K}$$

of the original net $\{(x_\alpha, y_\alpha)\}_{\alpha \in I}$ that converges in $X \times Y$ to (x, y). □

The much less obvious aspect of Theorem 6.1 is that it is also true for infinite products, even those for which the index set I is *uncountably* infinite. So it follows for instance that the space

$$[0,1]^{\mathbb{R}} = \{\text{not necessarily continuous functions } f : \mathbb{R} \to [0,1]\} = \prod_{\alpha \in \mathbb{R}} [0,1]$$

with the topology of pointwise convergence is compact, as an immediate consequence of the fact that $[0,1]$ is compact. Of course, this does not mean that every sequence of functions $f_n : \mathbb{R} \to [0,1]$ has a pointwise convergent subsequence! That would be truly surprising, but it is false (see Exercise 6.5); it turns out that $[0,1]^{\mathbb{R}}$ is not a first countable space, so it is allowed to be compact without being sequentially compact.

For a slightly different example, $[-1,1]^{\mathbb{N}}$ is compact. We can identify this space with the set of all sequences in $[-1,1]$, again with the topology of pointwise convergence, i.e. a sequence of sequences $\{x_k^n\}_{k \in \mathbb{N}} \in [-1,1]^{\mathbb{N}}$ converges as $n \to \infty$ to a sequence $\{x_k\}_{k \in \mathbb{N}}$ if $\lim_{n \to \infty} x_k^n = x_k$ for every $k \in \mathbb{N}$. Now observe that $[-1,1]^{\mathbb{N}}$ also contains the unit ball in the infinite-dimensional Hilbert space

$$\ell^2[-1,1] := \left\{ \{x_k \in \mathbb{R}\}_{k \in \mathbb{N}} \;\middle|\; \sum_{k=1}^{\infty} |x_k|^2 < \infty \right\}$$

with metric defined by

$$d(\{x_k\}, \{y_k\})^2 = \sum_{k=1}^{\infty} |x_k - y_k|^2.$$

The unit ball in $\ell^2[-1,1]$ is clearly noncompact since it contains the sequence of sequenes

$$(1,0,0,\ldots), (0,1,0,\ldots), (0,0,1,0,\ldots), \ldots,$$

which converges pointwise to 0 but stays at a constant distance away from 0 with respect to the metric, so it can have no convergent subsequence in the topology of $\ell^2[-1,1]$. It may seem surprising in this case that the *larger* set $[-1,1]^{\mathbb{N}}$ is compact, but the reason is that $[-1,1]^{\mathbb{N}}$ has a much weaker topology than $\ell^2[-1,1]$: since it is easier to converge pointwise than it is to converge in the ℓ^2-norm, $[-1,1]^{\mathbb{N}}$ has more sequences with convergent subsequences (or subnets, as the case may be).

Remark 6.2. One conclusion you should draw from the above discussion is that Tychonoff's theorem depends crucially on the way we defined the product topology on $\prod_{\alpha \in I} X_\alpha$, i.e. it is a result about the topology of pointwise convergence. The result becomes false, for instance, if we replace the usual product topology by the "box" topology from Exercise 4.6. For a concrete example, consider the set $[-1,1]^{\mathbb{N}}$ with the box topology, meaning sets of the form

$$\left\{ f \in [-1,1]^{\mathbb{N}} \;\middle|\; f(k) \in \mathcal{U}_k \text{ for all } k \in \mathbb{N} \right\}$$

for arbitrary collections of open subsets $\{\mathcal{U}_k \subset [-1,1]\}_{k \in \mathbb{N}}$ are open. Then the sequence of constant functions $f_n(k) := 1/n$ converges pointwise to 0, but we claim that it has no cluster point in the box topology. Indeed, the box topology contains the product topology, so if any subnet of f_n converges in the box topology, then it must also converge in the product topology and hence pointwise, meaning the only limit it could possibly converge to is 0, and 0 is therefore the only possible cluster point. But in the box topology,

$$\mathcal{U} := \left\{ f \in [-1,1]^{\mathbb{N}} \;\middle|\; f(k) \in (-1/k, 1/k) \text{ for all } k \in \mathbb{N} \right\}$$

is an open neighborhood of 0 satisfying $f_n \notin \mathcal{U}$ for all $n \in \mathbb{N}$, so 0 is not a cluster point of this sequence.

Let's go ahead and prove another special case of Tychonoff's theorem. The next proof is still relatively straightforward, and it applies for instance to $[-1,1]^{\mathbb{N}}$. Part of the idea is to make our lives easier by dealing with sequences instead of nets, which is made possible by the following simple observation:

Lemma 6.3. *If $X_1, X_2, X_3, \ldots$ is a countably infinite sequence of spaces that are all second countable, then $\prod_{i=1}^\infty X_i$ is also second countable.*

Proof. Fix for each $i = 1, 2, 3, \ldots$ a countable base $\mathcal{B}_i$ for the topology of X_i. Then for each $n \in \mathbb{N}$, the collection of sets

$$\mathcal{O}_n := \left\{ \mathcal{U}_1 \times \ldots \times \mathcal{U}_n \times X_{n+1} \times X_{n+2} \times \ldots \subset \prod_{i=1}^\infty X_i \;\middle|\; \mathcal{U}_i \in \mathcal{B}_i \text{ for each } i = 1, \ldots, n \right\}$$

is countable since $\mathcal{B}_1 \times \ldots \times \mathcal{B}_n$ is countable. Then the countable union of countable sets $\mathcal{O}_1 \cup \mathcal{O}_2 \cup \mathcal{O}_3 \cup \ldots$ is a base for $\prod_{i=1}^\infty X_i$, and it is countable. □

Proof of Theorem 6.1, second countable case. Assume the set I is countable and the spaces X_α are all second countable for $\alpha \in I$. In light of Lemma 6.3 and Theorem 5.26, it will now suffice to prove that for any sequence $X_1, X_2, X_3, \ldots$ of second countable spaces, $\prod_{i=1}^\infty X_i$ is sequentially compact. The idea is to combine the argument above for the case of finite products with Cantor's diagonal method. In order to avoid too many indices, let us denote elements $f \in \prod_{i=1}^\infty X_i$ as functions $f : \mathbb{N} \to \bigcup_{i=1}^\infty X_i$ that satisfy $f(i) \in X_i$ for each $i \in \mathbb{N}$. Now given a sequence $f_n \in \prod_{i=1}^\infty X_i$, the compactness of X_1 guarantees that there is a subsequence f_n^1 of f_n for which the sequence $f_n^1(1)$ in X_1 converges. Continuing inductively, we can construct a sequence of sequences $f_n^k \in \prod_{i=1}^\infty X_i$ for $k, n \in \mathbb{N}$ such that for every $k \geqslant 2$, $\{f_n^k\}_{n=1}^\infty$ is a subsequence of $\{f_n^{k-1}\}_{n=1}^\infty$ and the sequence $f_n^k(k)$ in X_k converges as $n \to \infty$. It follows that for every fixed $k \in \mathbb{N}$, the sequence $\{f_n^n(k)\}_{n=1}^\infty$ in X_k converges, thus $\{f_n^n\}_{n=1}^\infty$ is a convergent subsequence of the original sequence f_n in $\prod_{i=1}^\infty X_i$. □

The ideas in the special cases we've treated so far can be applied toward a general proof of Tychonoff's theorem, but the general case requires one major ingredient that wasn't needed so far: the axiom of choice. This makes e.g. the compactness of $[-1,1]^{[0,1]}$ somewhat harder to grasp intuitively, as invoking the axiom of choice means that the existence of a cluster point for every sequence in $[-1,1]^{[0,1]}$ is guaranteed, but there is nothing even slightly resembling an algorithm for finding one. It is known in fact that this is not just a feature of any particular method of proving the theorem—by a result due to Kelley [Kel50], if one assumes that the usual axioms of set theory (not including choice) hold and that Tychonoff's theorem also holds, then the axiom of choice follows, thus the two are actually equivalent.

Speaking only for myself, I had a Ph.D. in mathematics already for several years before I ever started to find the axiom of choice remotely worrying, so if you've never worried about it before, I don't encourage you to start worrying now. As far as this course is concerned, we actually could have skipped the general case of Tychonoff's theorem with no significant loss of continuity—I am including it here mainly for the sake of cultural education, and because the proof itself is interesting.

The proof given below is based on the characterization of compactness in terms of convergent subnets (Theorem 5.21) and is due to Paul Chernoff [Che92]. Similarly to certain standard results in functional analysis that also depend on the axiom of choice (e.g. the Hahn-Banach theorem), it uses the axiom in a somewhat indirect way, namely via *Zorn's lemma*, which is known to be equivalent to the axiom of choice. I do not want to go far enough into abstract set theory here to explain why it is equivalent: the proof is elementary but somewhat tedious, and you can find it explained e.g. in [Jän05] or [Kel75]. I would recommend reading through that proof exactly once in your life. For our purposes, we will just take the following statement of Zorn's lemma as a black box.

LEMMA 6.4 (Zorn's lemma). *Suppose $(\mathcal{P},\prec)$ is a nonempty partially ordered set in which every totally ordered subset $\mathcal{A}\subset\mathcal{P}$ has an upper bound, i.e. for every subset in which all pairs $x,y\in\mathcal{A}$ satisfy $x\prec y$ or $y\prec x$, there exists an element $p\in\mathcal{P}$ such that $p\succ a$ for all $a\in\mathcal{A}$. Then every totally ordered subset $\mathcal{A}\subset\mathcal{P}$ also has an upper bound $p\in\mathcal{P}$ that is a maximal element, i.e. such that no $q\in\mathcal{P}$ with $q\neq p$ satisfies $q\succ p$.* □

PROOF OF THEOREM 6.1, GENERAL CASE. We shall continue to denote elements of $\prod_{\alpha\in I}X_\alpha$ by functions $f: I\to\bigcup_{\alpha\in I}X_\alpha$ satisfying $f(\alpha)\in X_\alpha$ for each $\alpha\in I$. Assuming all the X_α are compact, it suffices by Theorem 5.21 to prove that every net $\{f_\beta\}_{\beta\in K}$ in $\prod_{\alpha\in I}X_\alpha$ has a cluster point. The idea of Chernoff's proof is as follows: we introduce below the notion of a "partial" cluster point, which may be a function defined only on a subset of I. We will show that the set of all partial cluster points has a partial order for which Zorn's lemma applies and delivers a maximal element. The last step is to show that a maximal element in the set of partial cluster points must in fact be a cluster point of $\{f_\beta\}_{\beta\in K}$.

To define partial cluster points, notice that for any subset $J\subset I$, restricting any function $f\in\prod_{\alpha\in I}X_\alpha$ to the smaller domain J defines an element $f|_J\in\prod_{\alpha\in J}X_\alpha$. We will refer to a pair (J,g) as a **partial cluster point** of the net $\{f_\beta\}_{\beta\in K}$ if J is a subset of I and $g\in\prod_{\alpha\in J}X_\alpha$ is a cluster point of the net $\{f_\beta|_J\}_{\beta\in K}$ in $\prod_{\alpha\in J}X_\alpha$ obtained by restricting the functions $f_\beta: I\to\bigcup_{\alpha\in I}X_\alpha$ to $J\subset I$. Let $\mathcal{P}$ denote the set of all partial cluster points of $\{f_\beta\}_{\beta\in K}$. It is easy to see that $\mathcal{P}$ is nonempty: indeed, for each individual $\alpha\in I$, the compactness of X_α implies that the net $\{f_\beta(\alpha)\}_{\beta\in K}$ in X_α has a cluster point $x_\alpha\in X_\alpha$, hence $(\{\alpha\},x_\alpha)\in\mathcal{P}$.

There is also an obvious partial order on $\mathcal{P}$: we shall write $(J,g)\leqslant(J',g')$ whenever $J\subset J'$ and $g=g'|_J$. In order to satisfy the main hypothesis of Zorn's lemma, we claim that every totally ordered subset $\mathcal{A}\subset\mathcal{P}$ has an upper bound. Being totally ordered means that for any two elements of $\mathcal{A}$, one is obtained from the other by restricting the function to a subset. We can therefore define a set $J_\infty\subset I$ with a function $g_\infty\in\prod_{\alpha\in J_\infty}X_\alpha$ by

$$J_\infty=\bigcup_{\{J\ |\ (J,g)\in\mathcal{A}\}}J,$$

with $g_\infty(\alpha)$ defined as $g(\alpha)$ for any $(J, g) \in \mathcal{A}$ such that $\alpha \in J$. The total ordering condition guarantees that (J_∞, g_∞) is independent of choices, but it is not immediately clear whether it is an element of $\mathcal{P}$, i.e. whether g_∞ is a cluster point of $\{f_\beta|_{J_\infty}\}_{\beta \in K}$. To see this, suppose $\mathcal{U} \subset \prod_{\alpha \in J_\infty} X_\alpha$ is a neighborhood of g_∞, and recall that by the definition of the product topology, this means

$$g_\infty \in \prod_{\alpha \in J_\infty} \mathcal{U}_\alpha \subset \mathcal{U}$$

for some collection of open sets $\mathcal{U}_\alpha \subset X_\alpha$ such that $\mathcal{U}_\alpha = X_\alpha$ for all α outside some finite subset $J_0 \subset J_\infty$. Since J_0 is finite, and $\mathcal{A}$ is totally ordered, there exists some $(J, g) \in \mathcal{A}$ such that $J_0 \subset J$. Then the fact that (J, g) is a partial cluster point means that for every $\beta_0 \in K$, there exists a $\beta \succ \beta_0$ for which

$$f_\beta|_J \in \prod_{\alpha \in J} \mathcal{U}_\alpha.$$

It follows that $f_\beta|_{J_\infty} \in \prod_{\alpha \in J_\infty} \mathcal{U}_\alpha$ as well, hence (J_∞, g_∞) is indeed a partial cluster point.

We can now apply Zorn's lemma and conclude that $\mathcal{P}$ has a maximal element $(J_M, g_M) \in \mathcal{P}$. We claim $J_M = I$, which means g_M is a cluster point of the original net $\{f_\beta\}_{\beta \in K}$ in $\prod_{\alpha \in I} X_\alpha$. Note that since $g_M \in \prod_{\alpha \in J_M} X_\alpha$ is a cluster point of $\{f_\beta|_{J_M}\}_{\beta \in K}$, Lemma 5.20 provides a subnet $\{f_{\phi(\gamma)}\}_{\gamma \in L}$ of $\{f_\beta\}_{\beta \in K}$ in $\prod_{\alpha \in I} X_\alpha$ whose restriction to J_M converges to g_M. But if $J_M \neq I$, then choosing an element $\alpha_0 \in I \backslash J_M$, we can exploit the fact that X_{α_0} is compact and use the same trick as in the proof of Tychonoff for finite products to find a further subnet that also converges at α_0 to some element $x_0 \in X_{\alpha_0}$. We have therefore found a subnet of $\{f_\beta\}_{\beta \in K}$ whose restriction to $J_M \cup \{\alpha_0\}$ converges to the function $g'_M \in \prod_{\alpha \in J_M \cup \{\alpha_0\}} X_\alpha$ defined by $g'_M|_{J_M} = g_M$ and $g'_M(\alpha_0) = x_0$. This means $(J_M \cup \{\alpha_0\}, g'_M) \in \mathcal{P}$ and $(J_M \cup \{\alpha_0\}, g'_M) \succ (J_M, g_M)$, which is a contradiction since (J_M, g_M) is maximal. $\square$

Exercise 6.5. Consider the space $[0,1]^{\mathbb{R}}$ of all functions $f : \mathbb{R} \to [0,1]$, with the topology of pointwise convergence. Tychonoff's theorem implies that $[0,1]^{\mathbb{R}}$ is compact, but one can show that it is not first countable, so it need not be sequentially compact.

(a) For $x \in \mathbb{R}$ and $n \in \mathbb{N}$, let $x_{(n)} \in \{0, \ldots, 9\}$ denote the nth digit to the right of the decimal point in the decimal expansion of x. Now define a sequence $f_n \in [0,1]^{\mathbb{R}}$ by setting $f_n(x) = \frac{x_{(n)}}{10}$. Show that for any subsequence f_{k_n} of f_n, there exists $x \in \mathbb{R}$ such that $f_{k_n}(x)$ does not converge, hence f_n has no pointwise convergent subsequence.
Food for thought: Could you do this if you also had to assume that x is rational? Presumably not, because $[0,1]^{\mathbb{Q}}$ is a product of countably many second countable spaces, and we've proved that such products are second countable (unlike $[0,1]^{\mathbb{R}}$). This implies that since $[0,1]^{\mathbb{Q}}$ is compact, it must also be sequentially compact.

(b) The compactness of $[0,1]^{\mathbb{R}}$ does imply that every sequence has a convergent *subnet*, or equivalently, a cluster point. Use this to deduce that for any given sequence $f_n \in [0,1]^{\mathbb{R}}$, there exists a function $f \in [0,1]^{\mathbb{R}}$ such that for every finite subset $X \subset \mathbb{R}$, some subsequence of f_n converges to f at all points in X.
Achtung: Pay careful attention to the order of quantifiers here. We're claiming that the element f exists independently of the finite set $X \subset \mathbb{R}$ on which we want some subsequence to converge to f. (If you could let f depend on the choice of subset X, this would be easy—but that is not allowed.) On the other hand, the actual choice of subsequence is allowed to depend on the subset X.

Challenge: Find a direct proof of the statement in part (b), without passing through Tychonoff's theorem. I do not know of any way to do this that isn't approximately as difficult as actually proving Tychonoff's theorem and dependent on the axiom of choice.

So much for Tychonoff's theorem. In truth, aside from the easy case of finite products, the general version of this theorem will probably not be mentioned again in this course. You may hear of it again if you take functional analysis since it lies in the background of the Banach-Alaoglu theorem on compactness in the *weak**-topology, and I will have occasion to mention it in *Topologie II* next semester in the context of the Eilenberg-Steenrod axioms for Čech homology. But right now we need to discuss a few more mundane things.

Topic 2: Separation axioms. Recall from Proposition 5.11 that closed subsets of compact spaces are always compact. Your intuition probably tells you that all compact sets are closed, but this in general is false. Here is a counterexample.

EXAMPLE 6.6. Recall from Example 2.2 the so-called "line with two zeroes". We defined it as a quotient $X := (\mathbb{R} \times \{0,1\})/\sim$ by the equivalence relation such that $(x,0) \sim (x,1)$ for all $x \neq 0$, with a topology defined via the pseudometric $d([(x,i)],[(y,j)]) = |x-y|$, i.e. the open balls $B_r(x) := \{y \in X \mid d(y,x) < r\}$ for $x \in X$ and $r > 0$ form a base of the topology. Each $x \in \mathbb{R}\backslash\{0\}$ corresponds to a unique point $[(x,0)] = [(x,1)] \in X$, but for $x = 0$ there are two distinct points, which we shall abbreviate by

$$0_0 := [(0,0)] \in X \quad \text{and} \quad 0_1 := [(0,1)] \in X.$$

As we saw in Exercise 2.3, the one-point subset $\{0_1\} \subset X$ is not closed, but it certainly is compact since finite subsets are always compact (see Example 5.5). The failure of $\{0_1\}$ to be closed results from the fact that since $d(0_0, 0_1) = 0$, every neighborhood of 0_0 also contains 0_1, implying that $X\backslash\{0_1\}$ cannot be open.

The example of the line with two zeroes is pathological in various ways, e.g. it has the property that every sequence convergent to 0_1 also converges to the distinct point 0_0. We would now like to formulate some precise conditions to exclude such behavior. The most important of these will be the *Hausdorff* axiom, but there is a whole gradation of stronger or weaker variations on the same theme, known collectively as the **separation axioms** (*Trennungsaxiome*). Intuitively, they measure the degree to which topological notions such as convergence of sequences and continuity of maps can recognize the difference between two disjoint points or subsets.

DEFINITION 6.7. A space X is said to satisfy axiom T_0 if for every pair of distinct points in X, there exists an open subset of X that contains one of these points but not the other.

Since almost all spaces we want to consider will satisfy the T_0 axiom, we should point out some examples of spaces that do not. One obvious example is any space of more than one element with the trivial topology: if the only open subset other than $\varnothing$ is X, then you clearly cannot find an open set that contains x and not $y \neq x$ or vice versa. A slightly more interesting example is the line with two zeroes as in Example 6.6 above, with the pseudometric topology: it fails to be a T_0 space because every open set that contains 0_0 or 0_1 must contain both of them.

DEFINITION 6.8. A space X is said to satisfy **axiom** T_1 if for every pair of distinct points $x, y \in X$, there exist neighborhoods $\mathcal{U}_x \subset X$ of x and $\mathcal{U}_y \subset X$ of y such that $x \notin \mathcal{U}_y$ and $y \notin \mathcal{U}_x$.

Obviously every T_1 space is also T_0. The following alternative characterization of the T_1 axiom is immediate from the definitions:

PROPOSITION 6.9. *A space X satisfies axiom T_1 if and only if for every point $x \in X$, the subset $\{x\} \subset X$ is closed.* □

DEFINITION 6.10. A space X is said to satisfy **axiom** T_2 (the **Hausdorff** axiom) if for every pair of distinct points $x, y \in X$, there exist neighborhoods $\mathcal{U}_x \subset X$ of x and $\mathcal{U}_y \subset X$ of y such that $\mathcal{U}_x \cap \mathcal{U}_y = \varnothing$.

Every Hausdorff space is clearly also T_1 and T_0. Here is an easy criterion with which to recognize a non-Hausdorff space:

EXERCISE 6.11. Show that if X is Hausdorff, then for any sequence $x_n \in X$ satisfying $x_n \to x$ and $x_n \to y$, we have $x = y$.

Finding an example that is T_1 but not Hausdorff requires only a slight modification of our previous "line with two zeroes".

EXAMPLE 6.12. Consider $X = (\mathbb{R} \times \{0,1\})/\sim$ again with $(x,0) \sim (x,1)$ for every $x \neq 0$, but instead of the pseudometric topology as in Example 6.6, assign it the quotient topology, meaning $\mathcal{U} \subset X$ is open if and only if its preimage under the projection map $\pi : \mathbb{R} \times \{0,1\} \to X : (x,i) \mapsto [(x,i)]$ is open. Recall that the quotient topology is the strongest topology for which π is a continuous map, and in this case, it turns out to be slightly stronger than the pseudometric topology. For example, the open set

$$\mathcal{V} := ((-1,1) \times \{0\}) \cup ((-1,0) \times \{1\}) \cup ((0,1) \times \{1\}) \subset \mathbb{R} \times \{0,1\}$$

is $\pi^{-1}(\mathcal{U})$ for $\mathcal{U} := \pi(\mathcal{V}) \subset X$, thus $\mathcal{U}$ is open in the quotient topology. But $\mathcal{U}$ contains 0_0 and not 0_1, so it is not an open set in the pseudometric topology. The existence of this set implies that X with the quotient topology satisfies T_0. By exchanging the roles of 0 and 1, one can similarly construct an open neighborhood of 0_1 that does not contain 0_0, so the space also satisfies T_1. But it does not satisfy T_2: even in the quotient topology, every neighborhood of 0_0 has nonempty intersection with every neighborhood of 0_1.

Exercise 6.11 has a converse of sorts, which I will state here only for first countable spaces. The countability axiom can be removed at the cost of talking about nets instead of sequences; I will leave the details of this as an exercise for the reader.

PROPOSITION 6.13. *A first countable space X is Hausdorff if and only if the limit of every convergent sequence in X is unique.*

PROOF. In light of Exercise 6.11, we just need to show that if X is a first countable space that is not Hausdorff, we can find a sequence $x_n \in X$ that converges to two distinct points $x, y \in X$. Since X is not Hausdorff, we can pick two distinct points x and y such that every neighborhood of x intersects every neighborhood of y. Fix countable neighborhood bases $X \supset \mathcal{U}_1 \supset \mathcal{U}_2 \supset \ldots \ni x$ and $X \supset \mathcal{V}_1 \supset \mathcal{V}_2 \ldots \ni y$. Then by assumption, for each $n \in \mathbb{N}$ there exists a point $x_n \in \mathcal{U}_n \cap \mathcal{V}_n$. It is now straightforward to verify that $x_n \to x$ and $x_n \to y$. □

The Hausdorff axiom can still be strengthened a bit by talking about neighborhoods of closed sets rather than points. This can be useful, for instance, when considering the quotient space X/A defined by collapsing some closed subset $A \subset X$ to a point; cf. Exercise 6.20 below.

DEFINITION 6.14. A space X is called **regular** (*regulär*) if for every point $x \in X$ and every closed subset $A \subset X$ not containing x, there exist neighborhoods $\mathcal{U}_x \subset X$ of x and $\mathcal{U}_A \subset X$ of A such that $\mathcal{U}_x \cap \mathcal{U}_A = \emptyset$. We say X satisfies **axiom** T_3 if it is regular and also satisfies T_1.

DEFINITION 6.15. A space X is called **normal** if for every pair of disjoint closed subsets $A, B \subset X$, there exist neighborhoods $\mathcal{U}_A \subset X$ of A and $\mathcal{U}_B \subset X$ of B such that $\mathcal{U}_A \cap \mathcal{U}_B = \emptyset$. We say X satisfies **axiom** T_4 if it is normal and also satisfies T_1.

REMARK 6.16. The point of including T_1 in the definitions of T_3 and T_4 is that it makes each one-point subset $\{x\} \subset X$ closed, thus producing obvious implications

$$T_4 \quad \Rightarrow \quad T_3 \quad \Rightarrow \quad T_2 \quad \Rightarrow \quad T_1 \quad \Rightarrow \quad T_0. \tag{6.1}$$

Without assuming T_1, it is possible for spaces to be regular or normal without being Hausdorff, though we will not consider any examples of this. In fact, almost all spaces we actually want to think about in this course will be Hausdorff, and most will also be normal, thus satisfying all of these axioms.

REMARK 6.17. Some of the above definitions, especially for axioms T_3 and T_4, can be found in a few not-quite-equivalent variations in various sources in the literature. One common variation is to interchange the meanings of "regular" with "T_3" and "normal" with "T_4", which destroys the first two implications in (6.1). These discrepancies are matters of convention which are to some extent arbitrary: you are free to choose your favorite convention, but must then be careful about stating your definitions precisely and remaining consistent.

We can now give a better answer to the question of when a compact set must also be closed.

THEOREM 6.18. *If X is Hausdorff, then every compact subset of X is closed.*

PROOF. Given a compact set $K \subset X$, we need to show that $X \backslash K$ is open, or equivalently, that every $x \in X \backslash K$ is contained in an open set disjoint from K. By assumption X is Hausdorff, so for each $y \in K$, we can find open neighborhoods $\mathcal{U}_y \subset X$ of x and $\mathcal{V}_y \subset X$ of y such that $\mathcal{U}_y \cap \mathcal{V}_y = \emptyset$. Then the sets $\{\mathcal{V}_y\}_{y \in K}$ form an open cover of K, and since the latter is compact by assumption, we obtain a finite subset $y_1, \ldots, y_N \in K$ such that

$$K \subset \mathcal{V}_{y_1} \cup \ldots \cup \mathcal{V}_{y_N}.$$

The set $\mathcal{U} := \mathcal{U}_{y_1} \cap \ldots \cap \mathcal{U}_{y_N}$ is then an open neighborhood of x and is disjoint from $\mathcal{V}_{y_1} \cup \ldots \cup \mathcal{V}_{y_N}$, implying in particular that it is disjoint from K. □

EXERCISE 6.19. Prove:

(a) A finite topological space satisfies the axiom T_1 if and only if it carries the discrete topology.
(b) X is a T_2 space (i.e. Hausdorff) if and only if the *diagonal* $\Delta := \{(x, x) \in X \times X\}$ is a closed subset of $X \times X$.
(c) Every compact Hausdorff space is regular, i.e. compact + $T_2 \Rightarrow T_3$.
Hint: The argument needed for this was already used in the proof of Theorem 6.18.
(d) Every metrizable space satisfies the axiom T_4 (in particular it is *normal*).
Hint: Given disjoint closed sets $A, A' \subset X$, each $x \in A$ admits a radius $\epsilon_x > 0$ such that the ball $B_{\epsilon_x}(x)$ is disjoint from A', and similarly for points in A' (why?). The unions of all these balls won't quite produce the disjoint neighborhoods you want, but try cutting their radii in half.

EXERCISE 6.20. Suppose X is a Hausdorff space and $\sim$ is an equivalence relation on X. Let $X/\sim$ denote the quotient space equipped with the quotient topology and denote by $\pi : X \to X/\sim$ the canonical projection. Given a subset $A \subset X$, we will sometimes also use the notation X/A explained in Exercise 5.16.

(a) A map $s : X/\sim \to X$ is called a **section** of π if $\pi \circ s$ is the identity map on $X/\sim$. Show that if a continuous section exists, then $X/\sim$ is Hausdorff.
(b) Show that if X is also regular and $A \subset X$ is a closed subset, then X/A is Hausdorff.
(c) Consider $X = \mathbb{R}$ with the non-closed subset $A = (0, 1]$. Which of the separation axioms $T_0, \ldots, T_4$ does X/A satisfy?

Just for fun: think about some other examples of Hausdorff spaces X with non-Hausdorff quotients $X/\sim$. What stops you from constructing continuous sections $X/\sim \to X$?

REMARK 6.21. In earlier decades, it was common to define compactness slightly differently: what many papers and textbooks from the first half of the 20th centuary call a "compact space" is what we would call a "compact Hausdorff space". You should be aware of this discrepancy if you consult the older literature.

7. Connectedness and local compactness (May 9, 2023)

We would like to formalize the idea that in some spaces, you can find a continuous path connecting any point to any other point, and in other spaces you cannot.

DEFINITION 7.1. A space X is called **path-connected** (*wegzusammenhängend*) if for every pair of points $x, y \in X$, there exists a continuous map $\gamma : [0,1] \to X$ such that $\gamma(0) = x$ and $\gamma(1) = y$.

A subset of X is similarly called path-connected if it is a path-connected space in the subspace topology, which is equivalent to saying that any two points in the subset can be connected by a continuous path in that subset. We will refer to any maximal path-connected subset of a space X as a **path-component** (*Wegzusammenhangskomponente*) of X.

EXERCISE 7.2. Show that any two path-components of a space X must be either identical or disjoint, i.e. the path-components partition X into disjoint subsets. One can also express this by saying that there is a well-defined equivalence relation $\sim$ on X such that $x \sim y$ if and only if x and y belong to the same path-component. (Why is that an equivalence relation?)

The notion of path-connectedness is framed in terms of maps into X, but there is also a "dual" perspective based on functions defined on X. To motivate this, notice that if $f : X \to \{0,1\}$ is any continuous function and $x, y \in X$ belong to the same path-component, then continuity demands $f(x) = f(y)$. (We will formalize this observation in the proof of Theorem 7.13 below.)

DEFINITION 7.3. A space X is **connected** (*zusammenhängend*) if every continuous map $X \to \{0,1\}$ is constant.

In many textbooks one finds a cosmetically different definition of connectedness in terms of subsets that are both open and closed, but the two definitions are equivalent due to the following result.

PROPOSITION 7.4. *A space X is connected if and only if $\varnothing$ and X are the only subsets of X that are both open and closed.*

PROOF. We prove first that the condition in this statement implies connectedness. The key observation is that the sets $\{0\}$ and $\{1\}$ in $\{0,1\}$ are each both open and closed, so if $f : X \to \{0,1\}$ is continuous, the same must hold for both $f^{-1}(0)$ and $f^{-1}(1)$ in X. Then one of these is the empty set and the other is X, so f is constant.

Conversely, suppose X contains a nonempty subset $X_0 \subset X$ that is both open and closed but $X_0 \neq X$. Then $X_1 := X \backslash X_0$ is also a nonempty open and closed subset, implying that X is the union of two disjoint open subsets X_0 and X_1. We can now define a nonconstant continuous function $f : X \to \{0,1\}$ by $f|_{X_0} = 0$ and $f|_{X_1} = 1$. Checking that it is continuous is easy since $\{0,1\}$ only contains four open sets: the main point is that $f^{-1}(0) = X_0$ and $f^{-1}(1) = X_1$ are both open. $\square$

REMARK 7.5. The important fact about $\{0,1\}$ used in the above proof was that it is a space of more than one element with the discrete topology: officially $\{0,1\}$ carries the subspace topology as a subset of $\mathbb{R}$, but this happens to match the discrete topology since 0 and 1 are each centers of open balls in $\mathbb{R}$ that do not touch any other points of $\{0,1\}$. If we preferred, we could have replaced Definition 7.3 with the condition that every continuous map $f : X \to Y$ to any space Y with the discrete topology is constant.

We can of course also talk about **connected subsets** $A \subset X$, meaning subsets that become connected spaces with the subspace topology. Spaces or subsets that are not connected are sometimes called **disconnected**. By analogy with path-components, any maximal connected subset of X will be called a **connected component** (*Zusammenhangskomponente*) of X.

PROPOSITION 7.6. *Any two connected components $A, B \subset X$ are either identical or disjoint.*

PROOF. If A and B are both maximal connected subsets of X and $A \cap B \neq \emptyset$, then we claim that $A \cup B$ is also connected. Indeed, any continuous function $f : A \cup B \to \{0, 1\}$ must restrict to constant functions on both A and B, so if $y \in A \cap B$, then $f(x) = f(y)$ for every $x \in A \cup B$, implying that every continous function $A \cup B \to \{0, 1\}$ is constant. Now if A and B are not identical, then the set $A \cup B$ is strictly larger than either A or B, giving a contradiction to the maximality assumption. □

EXAMPLE 7.7. For any collection $\{X_\alpha\}_{\alpha \in I}$ of connected spaces, the disjoint union $X := \coprod_{\alpha \in I} X_\alpha$ has the individual spaces $X_\alpha \subset X$ for $\alpha \in I$ as its connected components. Indeed, endowing X with the disjoint union topology makes each of the subsets $X_\alpha \subset X$ open, and since $X \backslash X_\alpha = \bigcup_{\beta \neq \alpha} X_\beta$ is then also open, it follows that X_α is also closed. Any strictly larger set $A \subset X$ with $X_\alpha \subset A$ could not then be connected, as it would contain X_α as a nonempty proper open and closed subset; this makes X_α a *maximal* connected subset of X.

EXERCISE 7.8. Show that if the spaces X_α in Example 7.7 are also path-connected, then they also form the path-components of the disjoint union $X = \coprod_{\alpha \in I} X_\alpha$.

For an arbitrary space X, let us choose an index set I with which to label each connected component of X, so the connected components from a collection of spaces $\{X_\alpha\}_{\alpha \in I}$, each of which is a subset $X_\alpha \subset X$ endowed with the subspace topology. Proposition 7.6 shows that $X_\alpha \cap X_\beta = \emptyset$ whenever $\alpha \neq \beta$, and obviously $\bigcup_{\alpha \in I} X_\alpha = X$, so as sets, there is a canonical bijective correspondence between X and the disjoint union $\coprod_{\alpha \in I} X_\alpha$. It is natural to wonder: is this correspondence a homeomorphism? It is easy to see that it is continuous in at least one direction: the individual subsets $X_\alpha \subset X$ come with inclusion maps $i_\alpha : X_\alpha \hookrightarrow X$, and endowing X_α with the subspace topology makes i_α continuous. The canonical bijection from $\coprod_{\alpha \in I} X_\alpha$ to X can then be written as

$$\coprod_{\alpha \in I} i_\alpha : \coprod_{\alpha \in I} X_\alpha \to X, \tag{7.1}$$

meaning it is the unique map whose restriction to each of the subsets $X_\alpha \subset \coprod_{\beta \in I} X_\beta$ is precisely i_α. The definition of the disjoint union topology makes this map automatically continuous. The following example shows however that, in general, its inverse need not be continuous.

EXAMPLE 7.9. The set $\mathbb{Q}$ of rational numbers is a perfectly nice algebraic object, but when endowed with the subspace topology as a subset of $\mathbb{R}$, it becomes a very badly behaved topological space. We claim that if $A \subset \mathbb{Q}$ is any subset with more than one element, then A is disconnected. Indeed, given $x, y \in A$ with $x < y$, we can find an irrational number $r \in \mathbb{R} \backslash \mathbb{Q}$ with $x < r < y$, and the sets $A_- := A \cap (-\infty, r)$ and $A_+ := A \cap (r, \infty)$ are then nonempty open subsets of A which are complements of each other, hence both are open and closed. This proves that the connected components of $\mathbb{Q}$ are simply the one-point subspaces $\{x\} \subset \mathbb{Q}$ for all $x \in \mathbb{Q}$, so the map (7.1) in this case takes the form

$$\coprod_{x \in \mathbb{Q}} \{x\} \to \mathbb{Q}.$$

The domain and target of this map are the same set, and the map itself is the identity, but the two sets are endowed with very different topologies: in particular, the domain carries the discrete topology, while $\mathbb{Q}$ on the right hand side carries the subspace topology that it inherits from the standard topology of $\mathbb{R}$. The identity map is thus continuous—indeed, every map defined on a space with the discrete topology is continuous—but it is not a homeomorphism, because the discrete topology contains many open sets that are not open in the standard topology of $\mathbb{Q}$.

Example 7.9 shows that while every space X has a natural bijective correspondence with the disjont union $\coprod_{\alpha \in I} X_\alpha$ of its connected components, the natural topology on $\coprod_{\alpha \in I} X_\alpha$ may in general be different from the original topology of X. We've seen for instance that each individual X_α is automatically both an open and closed subset of $\coprod_{\beta \in I} X_\beta$, thus there is no hope of (7.1) being a homeomorphism unless X_α is also an open and closed subset of X. The example of $\mathbb{Q}$ shows that the latter is not always true: the 1-point connected components $\{x\} \subset \mathbb{Q}$ are closed subsets, but they are not open. The fact that they are closed turns out to be a completely general phenomenon:

Proposition 7.10. *Every connected component $A \subset X$ of a space X is a closed subset.*

Proof. Assume $A \subset X$ is a maximal connected subset. Recall from Definition 3.1 that the closure $\bar{A} \subset X$ of A is the set of all points $x \in X$ for which every neighborhood of x intersects A. If we equip $\bar{A}$ with the subspace topology and view it as a topological space in itself, with $A \subset \bar{A}$ as a subset, then the closure of A in $\bar{A}$ is still $\bar{A}$: indeed, every neighborhood in $\bar{A}$ of a point $x \in \bar{A}$ takes the form $\mathcal{U} \cap \bar{A}$ for some neighborhood $\mathcal{U}$ of x in X, implying that $\mathcal{U}$ intersects A, and therefore so does $\mathcal{U} \cap \bar{A}$.

Now suppose $f : \bar{A} \to \{0,1\}$ is a continuous function. Its restriction to A is then also continuous, and therefore constant, since A is connected; let us write $f(A) = \{i\} \subset \{0,1\}$. Then since $\{i\}$ is a closed subset of $\{0,1\}$ and f is continuous, $f^{-1}(i)$ is a closed subset of $\bar{A}$ that contains A, and it therefore also contains the closure $\bar{A}$. This implies that f is in fact constant on $\bar{A}$, and thus proves that $\bar{A}$ is connected. Since A is a maximal connected subset, we conclude $A = \bar{A}$, meaning A is closed. □

We note one obvious case in which connected components will necessarily be both closed and open: here openness follows from the fact that the complement of a connected component is a union of disjoint connected components, and finite unions of closed sets are closed.

Corollary 7.11. *If X is a space with only finitely many connected components, then each of them is both closed and open.* □

Exercise 7.12. If $\{X_\alpha \subset X\}_{\alpha \in I}$ are the connected components of a space X, show that the canonical continuous bijection (7.1) from $\coprod_{\alpha \in I} X_\alpha$ to X is a homeomorphism if and only if every X_α is an open subset of X. (In particular, Corollary 7.11 implies that this is always true if I is finite, and we will see in Prop. 7.18 below that it is also true if X is locally connected.)

It is time to clarify the relationship between connectedness and path-connectedness.

Theorem 7.13. *Every path-connected space X is connected.*

Proof. If X is not connected, then there exist points $x, y \in X$ and a continuous function $f : X \to \{0,1\}$ such that $f(x) = 0$ and $f(y) = 1$. But if X is path-connected, then there also exists a continuous map $\gamma : [0,1] \to X$ with $\gamma(0) = x$ and $\gamma(1) = y$. The composition $g := f \circ \gamma$ is then a continuous function $g : [0,1] \to \{0,1\}$ satisfying $g(0) = 0$ and $g(1) = 1$, and this violates the intermediate value theorem. □

Surprisingly, the converse of this theorem is false.

Example 7.14. Define $X \subset \mathbb{R}^2$ to be the subset of $\mathbb{R}^2$ consisting of the vertical line $\{x = 0\}$ and the graph of the equation $\{y = \sin(1/x)\}$ for $x \neq 0$. The latter is a sine curve that oscillates more and more rapidly as $x \to 0$. We claim that

$$X_0 := \{x = 0\}$$

is a path-component of X. It clearly is path-connected, so we need to show that there does not exist any continuous path $\gamma : [0,1] \to X$ that begins on the sine curve $\{y = \sin(1/x)\}$ and ends on the line $\{x = 0\}$. Since $\{x = 0\}$ is a closed subset, the preimage of this set under γ is closed (and therefore compact) in $[0,1]$, implying that it has a minimum $\tau \in (0,1]$. We can therefore restrict our path to $\gamma : [0,\tau] \to X$ and assume that it lies on the sine curve for all $0 \leqslant t < \tau$ but ends on the vertical line at $t = \tau$. Now observe that due to the rapid oscillation as $x \to 0$, we can find for any $y \in [-1,1]$ a sequence $t_n \in [0,\tau)$ with $t_n \to \tau$ such that $\gamma(t_n) \to (0,y)$. The point y here is arbitrary, yet continuity of γ requires $\gamma(t_n) \to \gamma(\tau)$, so this is a contradiction and proves the claim. In particular, this proves that X is not path-connected. The other path-components of X are now easy to identify: they are

$$X_- := X \cap \{x < 0\} \quad \text{and} \quad X_+ := X \cap \{x > 0\},$$

the portions of the sine curve lying to the left and right of X_0, so there are three path-components in total. The path-components are path-connected and therefore (by Theorem 7.13) also connected. But neither X_- nor X_+ is closed, so by Prop. 7.10, neither of these can be a connected component. The maximal connected subset containing X_-, for instance, must be a closed set containing X_- and therefore contains the closure $\overline{X_-}$, which includes points in X_0. Since X_0 is path-connected, it follows that the connected component containing X_- also contains all of X_0. But the same argument applies equally well to X_+, and these two observations together imply that all three path-components are in the same connected component, i.e. X is connected.

The space in Example 7.14 is sometimes called the *topologist's sine curve*. There is a certain "local" character to the pathologies of this space, i.e. part of the reason for its bizarre properties is that one can zoom in on certain points in X arbitrarily far without making it look more reasonable—in particular this is true for the points in X_0 that are in the closure of X_- and X_+. One can use neighborhoods of points to formalize this notion of "zooming in" arbitrarily far.

Definition 7.15. A space X is **locally connected** (*lokal zusammenhängend*) if for all points $x \in X$, every neighborhood of x contains a connected neighborhood of x.

The version of this for path-connectedness is completely analogous.

Definition 7.16. A space X is **locally path-connected** (*lokal wegzusammenhängend*) if for all points $x \in X$, every neighborhood of x contains a path-connected neighborhood of x.

Local path-connectedness obviously implies local connectedness by Theorem 7.13. Since most spaces we can easily imagine will have both properties, it is important at this juncture to look at some examples that do not. The topologist's sine curve in Example 7.14 is one such space: it is not locally connected (even though it is connected), since sufficiently small neighborhoods of points $(0,y) \in X$ for $-1 < y < 1$ always have infinitely many pieces of the sine curve passing through and are thus disconnected. Here is an example that is path-connected, but not locally:

Example 7.17. Let $X \subset \mathbb{R}^2$ denote the compact set

$$X = \left(\bigcup_{n=1}^{\infty} L_n \right) \cup L_\infty,$$

where for each $n \in \mathbb{N}$, L_n denotes the straight line segment from $(0,1)$ to $(1/n, 0)$, and the case $n = \infty$ is included for the vertical segment from $(0,1)$ to $(0,0)$. Then sufficiently small neighborhoods of $(0,0)$ in this space are never connected, so X is not locally connected. Notice however that there are continuous paths along the line segments L_n from any point in X to $(0,1)$, so X is path-connected.

PROPOSITION 7.18. *If X is locally connected, then its connected components are open subsets. Similarly, if X is locally path-connected, then its path-components are open subsets.*

PROOF. If X is locally connected and $A \subset X$ is a maximal connected subset, then for each $x \in A$, fix a connected neighborhood $\mathcal{U}_x \subset X$ of x. Now for $\mathcal{U} := \bigcup_{x \in A} \mathcal{U}_x$, any continuous function $f : \mathcal{U} \to \{0,1\}$ must restrict to a constant on each $\mathcal{U}_x$ and also on A, implying that f is constant, hence $\mathcal{U}$ is connected. The maximality of A thus implies $A = \mathcal{U}$, but $\mathcal{U}$ is also a neighborhood of A and thus contains an open set containing A, therefore A is open.

A completely analogous argument works in the locally path-connected case, taking path-connected neighborhoods $\mathcal{U}_x$ and using the fact that their union must also be path-connected. $\square$

A consequence of this result is that the phenomenon allowing certain spaces to be connected but not path-connected is essentially local:

THEOREM 7.19. *Every space that is connected and locally path-connected is also path-connected.*

PROOF. If X is locally path-connected, then by Prop. 7.18 its path-components are open. Then if $A \subset X$ is a path-component, $X \backslash A$ is a union of path-components and is therefore also open, implying that A is both open and closed. If X is connected, it follows that $A = X$, so X is a path-component. $\square$

EXERCISE 7.20. In this exercise we show that products of (path-)connected spaces are also (path-)connected, so long as one uses the correct topology on the product.

(a) Prove that if X and Y are both connected, then so is $X \times Y$.
Hint: Start by showing that for any $x \in X$ and $y \in Y$, the subsets $\{x\} \times Y$ and $X \times \{y\}$ in $X \times Y$ are connected. Then think about continuous maps $X \times Y \to \{0,1\}$.

(b) Show that for any collection of path-connected spaces $\{X_\alpha\}_{\alpha \in I}$, the space $\prod_{\alpha \in I} X_\alpha$ is path-connected in the usual product topology.
Hint: You might find Exercise 4.5 helpful.

(c) Consider $\mathbb{R}^{\mathbb{N}}$ with the "box topology" which we discussed in Exercise 4.6. Show that the set of all elements $f \in \mathbb{R}^{\mathbb{N}}$ represented as functions $f : \mathbb{N} \to \mathbb{R}$ that satisfy $\lim_{n \to \infty} f(n) = 0$ is both open and closed, hence $\mathbb{R}^{\mathbb{N}}$ in the box topology is not connected (and therefore also not path-connected).

The rest of this exercise is aimed at generalizing part (a) to the statement that for an arbitrary collection $\{X_\alpha\}_{\alpha \in I}$ of connected (but not necessarily path-connected) spaces, $\prod_{\alpha \in I} X_\alpha$ with the product topology is also connected. Choose a point $\{c_\alpha\}_{\alpha \in I} \in \prod_{\alpha \in I} X_\alpha$ and, for each finite subset $J \subset I$ of the index set, consider the set

$$X_J := \left\{ \{x_\alpha\}_{\alpha \in I} \in \prod_{\alpha \in I} X_\alpha \;\middle|\; x_\beta = c_\beta \text{ for all } \beta \in I \backslash J \right\},$$

endowed with the subspace topology that it inherits from the product topology of $\prod_{\alpha \in I} X_\alpha$.

(d) Show that for every choice of finite subset $J \subset I$, X_J is connected.
Hint: This is not really that different from part (a).

(e) Deduce that the union $\bigcup_J X_J \subset \prod_{\alpha \in I} X_\alpha$ is also connected, where J ranges over the set of all finite subsets of I.

(f) Show that the closure of the subset $\bigcup_J X_J \subset \prod_{\alpha\in I} X_\alpha$ is $\prod_{\alpha\in I} X_\alpha$, and deduce that $\prod_{\alpha\in I} X_\alpha$ is also connected.

With the definition of local connectedness in mind, we now briefly revisit the subject of compactness.

Definition 7.21. A space X is **locally compact** (*lokal kompakt*) if every point $x \in X$ has a compact neighorhood.

Local compactness is one of the notions for which one can find multiple inequivalent definitions in the literature, but as we'll see in a moment, all the plausible definitions of this concept are equivalent if we only consider Hausdorff spaces. Let's first note a few examples.

Example 7.22. The Euclidean space $\mathbb{R}^n$ is locally compact, and more generally, so is any closed subset $X \subset \mathbb{R}^n$ endowed with the subspace topology. Indeed, since closed and bounded subsets of $\mathbb{R}^n$ are compact, every $x \in X \subset \mathbb{R}^n$ has a compact neighborhood of the form $\overline{B_r(x)} \cap X$ for any $r > 0$.

Example 7.23. This is a non-example: a Hilbert space is not locally compact if it is infinite dimensional. This is due to the fact that every neighborhood of a point x must contain some closed ball $\overline{B_r(x)}$, but the latter is not compact (cf. Remark 5.8).

Example 7.24. Since a space is a neighborhood of all of its points, every compact space is (trivially) locally compact.

The last example is the one that becomes slightly controversial if you look at alternative definitions of local compactness in the literature, and indeed, if we had phrased Definition 7.21 more analogously to the definition of local (path-)connectedness, it would be easy to imagine spaces that are compact without being locally compact. As it happens, this never happens for Hausdorff spaces, and since we will mainly be interested in Hausdorff spaces, we shall take the following result as an excuse to avoid worrying any further about discrepancies in definitions. It will also be a useful result in its own right.

Theorem 7.25. *If X is Hausdorff, then the following conditions are equivalent:*

(i) X is locally compact (in the sense of Definition 7.21);
(ii) For all $x \in X$, every neighborhood of x contains a compact neighborhood of x;
(iii) If $K \subset \mathcal{U} \subset X$ where K is compact and $\mathcal{U}$ is open, then $K \subset \mathcal{V} \subset \bar{\mathcal{V}} \subset \mathcal{U}$ for some open set $\mathcal{V}$ with compact closure $\bar{\mathcal{V}}$.

Proof. Since single point subsets $\{x\} \subset X$ are always compact, it is clear that (iii) $\Rightarrow$ (ii) $\Rightarrow$ (i). The implication (ii) $\Rightarrow$ (iii) is a relatively straightforward exercise using the finite covering property for the compact set K. We will therefore focus on the implication (i) $\Rightarrow$ (ii).

Assume we are given a neighborhood $\mathcal{U} \subset X$ of x and would like to find a compact neighborhood inside $\mathcal{U}$. By assumption, x also has a compact neighborhood $K \subset X$. It will do no harm to replace $\mathcal{U}$ with a smaller neighorhood such as the interior of $\mathcal{U} \cap K$, so without loss of generality, let us assume $\mathcal{U}$ is open and contained in K, in which case (since X is Hausdorff and K is therefore closed) its closure $\bar{\mathcal{U}}$ is also contained in K and is thus compact. We define the *boundary* of $\bar{\mathcal{U}}$ by

$$\partial\bar{\mathcal{U}} = \bar{\mathcal{U}} \cap \overline{X\backslash\mathcal{U}}.$$

This is a closed subset of $\bar{\mathcal{U}}$ and is therefore also compact, and we observe that since x is contained in a neighborhood disjoint from $X\backslash\mathcal{U}$, x is not in the closure $\overline{X\backslash\mathcal{U}}$ and thus

$$x \notin \partial\bar{\mathcal{U}}.$$

Since X is Hausdorff, for every $y \in \partial\bar{\mathcal{U}}$ there exists a pair of open neighborhoods

$$x \in A_y \subset X, \quad y \in B_y \subset X \quad \text{such that} \quad A_y \cap B_y = \emptyset.$$

Then the sets B_y for $y \in \partial\bar{\mathcal{U}}$ form an open cover of the compact set $\partial\bar{\mathcal{U}}$, hence there exists a finite subset $\{y_1, \ldots, y_N\} \subset \partial\bar{\mathcal{U}}$ such that

$$\partial\bar{\mathcal{U}} \subset \bigcup_{i=1}^{N} B_{y_i}.$$

Now the set

$$\mathcal{V} := \mathcal{U} \cap \left(\bigcap_{i=1}^{N} A_{y_i} \right)$$

is an open neighborhood of x contained in $\mathcal{U}$ and disjoint from the neighborhood $\bigcup_{i=1}^{N} B_{y_i}$ of $\partial\bar{\mathcal{U}}$. The latter implies that for any $y \in \partial\bar{\mathcal{U}}$, y has a neighborhood disjoint from $\mathcal{V}$, hence $y \notin \bar{\mathcal{V}}$. Similarly, $\mathcal{V} \subset \mathcal{U}$ implies y cannot be in the closure of $\mathcal{V}$ if it is in the interior of $\overline{X \backslash \mathcal{U}}$, so we conclude $\bar{\mathcal{V}} \subset \mathcal{U}$. The compactness of $\bar{\mathcal{V}}$ follows because it is a closed subset of $\bar{\mathcal{U}}$ and the latter is compact. $\square$

EXERCISE 7.26. Prove the implication that was skipped in the proof of Theorem 7.25 above, namely: if X is locally compact and Hausdorff, then for any nested pair of subsets $K \subset \mathcal{U} \subset X$ with K compact and $\mathcal{U}$ open, there exists an open set $\mathcal{V} \subset X$ with compact closure $\bar{\mathcal{V}}$ such that $K \subset \mathcal{V} \subset \bar{\mathcal{V}} \subset \mathcal{U}$.

EXERCISE 7.27. There is a cheap trick to view any topological space as a compact space with a single point removed. For a space X with topology $\mathcal{T}$, let $\{\infty\}$ denote a set consisting of one element that is not in X, and define the **one point compactification** of X as the set $X^* = X \cup \{\infty\}$ with topology $\mathcal{T}^*$ consisting of all subsets in $\mathcal{T}$ plus all subsets of the form $(X \backslash K) \cup \{\infty\} \subset X^*$ where $K \subset X$ is closed and compact.

(a) Verify that $\mathcal{T}^*$ is a topology and that X^* is always compact.
(b) Show that if X is first countable and Hausdorff, a sequence in $X \subset X^*$ converges to $\infty \in X^*$ if and only if it has no convergent subsequence with a limit in X. Conclude that if X is first countable and Hausdorff, X^* is sequentially compact.
(c) Show that for $X = \mathbb{R}$, X^* is homeomorphic to S^1. *(More generally, one can use stereographic projection to show that the one point compactification of $\mathbb{R}^n$ is homeomorphic to S^n.)*
(d) Show that if X is already compact, then X^* is homeomorphic to the disjoint union $X \amalg \{\infty\}$.
(e) Show that X^* is Hausdorff if and only if X is both Hausdorff and locally compact.

Notice that $\mathbb{Q}$ is not locally compact, since every neighborhood of a point $x \in \mathbb{Q}$ contains sequences without convergent subsequences, e.g. any sequence of rational numbers that converges to an irrational number sufficiently close to x. The one point compactification $\mathbb{Q}^*$ is a compact space, and by part (b) it is also sequentially compact, but those are practically the only nice things we can say about it.

(f) Show that for any $x \in \mathbb{Q}$, every neighborhood of x in $\mathbb{Q}^*$ intersects every neighborhood of ∞, so in particular, $\mathbb{Q}^*$ is not Hausdorff.
Advice: Do not try to argue in terms of sequences with non-unique limits (cf. part (g) below), and do not try to describe precisely what arbitrary compact subsets of $\mathbb{Q}$ can look like (the answer is not nice). One useful thing you can say about arbitrary compact subsets of $\mathbb{Q}$ is that they can never contain the intersection of $\mathbb{Q}$ with any open interval. (Why not?)

(g) Show that every convergent sequence in $\mathbb{Q}^*$ has a unique limit. *(Since $\mathbb{Q}^*$ is not Hausdorff, this implies via Proposition 6.13 that $\mathbb{Q}^*$ is not first countable—in particular, ∞ does not have a countable neighborhood base.)*

(h) Find a point in $\mathbb{Q}^*$ with a neighborhood that does not contain any compact neighborhood.

EXERCISE 7.28. Given spaces X and Y, let $C(X,Y)$ denote the set of all continuous maps from X to Y, and consider the natural *evaluation map*

$$\mathrm{ev} : C(X,Y) \times X \to Y : (f,x) \mapsto f(x).$$

It is easy to show that ev is a continuous map if we assign the discrete topology to $C(X,Y)$, but usually one can also find more interesting topologies on $C(X,Y)$ for which ev is continuous. The **compact-open topology** is defined via a subbase consisting of all subsets of the form

$$\mathcal{U}_{K,V} := \left\{ f \in C(X,Y) \;\middle|\; f(K) \subset V \right\},$$

where K ranges over all compact subsets of X, and V ranges over all open subsets of Y. Prove:

(a) If Y is a metric space, then convergence of a sequence $f_n \in C(X,Y)$ in the compact-open topology means that f_n converges uniformly on all compact subsets of X.

(b) If $C(X,Y)$ carries the topology of pointwise convergence (i.e. the subspace topology defined via the obvious inclusion $C(X,Y) \subset Y^X$), then ev is not sequentially continuous in general.

(c) If $C(X,Y)$ carries the compact-open topology, then ev is always sequentially continuous.

(d) If $C(X,Y)$ carries the compact-open topology and X is locally compact and Hausdorff, then ev is continuous.

(e) Every topology on $C(X,Y)$ for which ev is continuous contains the compact-open topology. (This proves that if X is locally compact and Hausdorff, the compact-open topology is the weakest topology for which the evaluation map is continuous.)
Hint: If $(f_0,x_0) \in \mathrm{ev}^{-1}(V)$ where $V \subset Y$ is open, then $(f_0,x_0) \in \mathcal{O} \times \mathcal{U} \subset \mathrm{ev}^{-1}(V)$ for some open $\mathcal{O} \subset C(X,Y)$ and $\mathcal{U} \subset X$. Is $\mathcal{U}_{K,V}$ a union of sets $\mathcal{O}$ that arise in this way?

(f) For the compact-open topology on $C(\mathbb{Q},\mathbb{R})$, $\mathrm{ev} : C(\mathbb{Q},\mathbb{R}) \times \mathbb{Q} \to \mathbb{R}$ is not continuous.

EXERCISE 7.29. One of the good reasons to use the notation X^Y for the set of all functions $f : Y \to X$ between two sets is that there is an obvious bijection

$$Z^{X \times Y} \to (Z^Y)^X$$

sending a function $F : X \times Y \to Z$ to the function $\Phi : X \to Z^Y$ defined by

$$\Phi(x)(y) = F(x,y). \tag{7.2}$$

The existence of this bijection is sometimes called the *exponential law* for sets. In this exercise we will explore to what extent the exponential law carries over to topological spaces and continuous maps. We will see that this is also related to the question of how to define a natural topology on the group of homeomorphisms of a space.

If X and Y are topological spaces, let us denote by $C(X,Y)$ the space of all continuous maps $X \to Y$, with the compact-open topology, which has a subbase consisting of all sets of the form

$$\mathcal{U}_{K,V} := \left\{ f \in C(X,Y) \;\middle|\; f(K) \subset V \right\}$$

for $K \subset X$ compact and $V \subset Y$ open (see Exercise 7.28 above). Assume Z is also a topological space.

(a) Prove that if $F : X \times Y \to Z$ is continuous, then the correspondence (7.2) defines a continuous map $\Phi : X \to C(Y,Z)$.

(b) Prove that if Y is locally compact and Hausdorff, then the converse also holds: any continuous map $\Phi : X \to C(Y,Z)$ defines a continuous map $F : X \times Y \to Z$ via (7.2).

Let's pause for a moment to observe what these two results imply for the case $X := I = [0,1]$. First, here is a quick definition of a notion that will appear very often in the remainder of this course: given two continuous maps $f_0, f_1 : Y \to Z$, a continuous map

$$h : I \times Y \to Z \qquad \text{such that} \qquad h(0,\cdot) = f_0 \text{ and } h(1,\cdot) = f_1$$

is called a **homotopy** (*Homotopie*) between f_0 and f_1, and we call f_0 and f_1 **homotopic** (*homotop*) if a homotopy between them exists. According to part (a), a homotopy between two maps $Y \to Z$ can always be regarded as a continuous path in $C(Y,Z)$, and part (b) says that the converse is also true if Y is locally compact and Hausdorff, hence two maps $Y \to Z$ are homotopic if and only if they lie in the same path-component of $C(Y,Z)$.[5]

(c) Deduce from part (b) a new proof of the following result from Exercise 7.28(d): if X is locally compact and Hausdorff, then the *evaluation map* $\mathrm{ev} : C(X,Y) \times X \to Y : (f,x) \mapsto f(x)$ is continuous.
Hint: This is very easy if you look at it from the right perspective.
Remark: If you were curious to see a counterexample to part (b) in a case where Y is not locally compact, you could now extract one from Exercise 7.28(f).

(d) The following cannot be deduced directly from part (b), but it is a similar result and requires a similar proof: show that if Y is locally compact and Hausdorff, then

$$C(X,Y) \times C(Y,Z) \to C(X,Z) : (f,g) \mapsto g \circ f$$

is a continuous map.
Hint: Exercise 7.26 is useful here.

Now let's focus on maps from a space X to itself. A group G with a topology is called a **topological group** if the maps

$$G \times G \to G : (g,h) \mapsto gh \quad \text{ and } \quad G \to G : g \mapsto g^{-1}$$

are both continuous. Common examples include the standard matrix groups $\mathrm{GL}(n,\mathbb{R})$, $\mathrm{GL}(n,\mathbb{C})$ and their subgroups, which have natural topologies as subsets of the vector space of (real or complex) n-by-n matrices. Another natural example to consider is the group

$$\mathrm{Homeo}(X) = \left\{ f \in C(X,X) \;\middle|\; f \text{ is bijective and } f^{-1} \in C(X,X) \right\}$$

for any topological space X, where the group operation is defined via composition of maps. We would like to know what topologies can be assigned to $C(X,X)$ so that $\mathrm{Homeo}(X) \subset C(X,X)$, with the subspace topology, becomes a topological group. Notice that the discrete topology clearly works; this is immediate because all maps between spaces with the discrete topology are automatically continuous, so there is nothing to check. But the discrete topology is not very interesting. Let $\mathcal{T}_H$ denote the topology on $C(X,X)$ with subbase consisting of all sets of the form $\mathcal{U}_{K,V}$ and $\mathcal{U}_{X \setminus V, X \setminus K}$, where again $K \subset X$ can be any compact subset and $V \subset X$ any open subset. Notice that if X is compact and Hausdorff, then for any V open and K compact, $X \setminus V$ is compact and $X \setminus K$ is open, thus $\mathcal{T}_H$ is again simply the compact-open topology. But if X is not compact or Hausdorff, $\mathcal{T}_H$ may be stronger than the compact-open topology.

[5] Since $C(X \times Y, Z)$ and $C(X, C(Y,Z))$ both have natural topologies in terms of the compact-open topology, you may be wondering whether the correspondence (7.2) defines a homeomorphism between them. The answer to this is more complicated than one would like, but Steenrod showed in a famous paper in 1967 [Ste67] that the answer is "yes" if one restricts attention to spaces that are *compactly generated*, a property that most respectable spaces have. The caveat is that $C(X,Y)$ in the compact-open topology will not always be compactly generated if X and Y are, so one must replace the compact-open topology by a slightly stronger one that is compactly generated but otherwise has the same properties for most practical purposes. If you want to know what "compactly generated" means and why it is a useful notion, see [Ste67]. These issues are somewhat important in homotopy theory at more advanced levels, though it is conventional to worry about them as little as possible.

(e) Show that if X is locally compact and Hausdorff, then $\mathrm{Homeo}(X)$ with the topology $\mathcal{T}_H$ is a topological group.
Hint: Notice that $f(K) \subset V$ if and only if $f^{-1}(X \backslash V) \subset X \backslash K$. Use this to show directly that $f \mapsto f^{-1}$ is continuous, and reduce the rest to what was proved already in part (d).

Conclusion: We've shown that if X is compact and Hausdorff, then $\mathrm{Homeo}(X)$ with the compact-open topology is a topological group. This is actually true under somewhat weaker hypotheses, e.g. it suffices to know that X is Hausdorff, locally compact and locally connected. (If you're interested, a quite clever proof of this fact may be found in [**Are46**]*.)*

Just for fun, here's an example to show that just being locally compact and Hausdorff is not enough: let $X = \{0\} \cup \{e^n \mid n \in \mathbb{Z}\} \subset \mathbb{R}$ with the subspace topology, and notice that X is neither compact (since it is unbounded) nor locally connected (since every neighborhood of 0 is disconnected). Consider the sequence $f_k \in \mathrm{Homeo}(X)$ defined for $k \in \mathbb{N}$ by $f_k(0) = 0$, $f_k(e^n) = e^{n-1}$ for $n \leqslant -k$ or $n > k$, $f_k(e^n) = e^n$ for $-k < n < k$, and $f_k(e^k) = e^{-k}$. It is not hard to show that in the compact-open topology on $C(X, X)$, $f_k \to \mathrm{Id}$ but $f_k^{-1} \not\to \mathrm{Id}$ as $k \to \infty$, hence the map $\mathrm{Homeo}(X) \to \mathrm{Homeo}(X) : f \mapsto f^{-1}$ is not continuous.

8. Paths, homotopy and the fundamental group (May 11, 2023)

The rest of this course will concentrate on *algebraic* topology. The class of spaces we consider will often be more restrictive than up to this point, e.g. we will usually (though not always) require them to be Hausdorff, second countable, locally path-connected and one or two other conditions that are satisfied in all interesting examples.[6] It will happen often from now on that the best way to prove any given result is with a picture, but I might not always have time to produce the relevant picture in these notes. I'll do what I can.

As motivation, let us highlight two examples of questions that the tools of algebraic topology are designed to answer.

Sample question 8.1. The following figures show two examples of **knots** K and K_0 in $\mathbb{R}^3$:

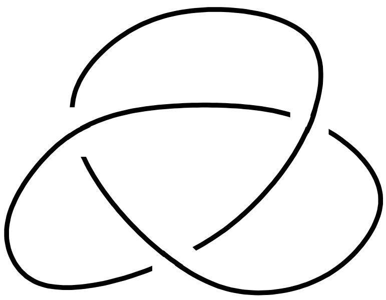

$K \subset \mathbb{R}^3$

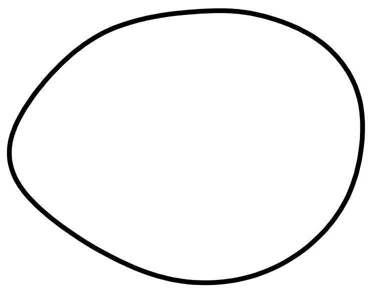

$K_0 \subset \mathbb{R}^3$

The first knot K is known as the **trefoil** knot (*Kleeblattknoten*), and the second K_0 is the *trivial* knot or **unknot** (*Unknoten*). Roughly speaking, a knot is a subset in $\mathbb{R}^3$ that is homeomorphic to S^1 and satisfies some additional condition to avoid overly "wild" behavior, e.g. one could sensibly require each of K and K_0 to be the image of some infinitely differentiable 1-periodic map $\mathbb{R} \to \mathbb{R}^3$. The question then is: can K be deformed continuously to K_0? Let us express this more precisely. If you imagine K and K_0 as physical knots in space, then when you move them around, you don't

[6] The question of which examples are considered "interesting" depends highly on context, of course. In functional analysis, one encounters many interesting spaces of functions that do not have all of the properties we just listed. But this is not a course in functional analysis.

move only the knots—you also displace the air around them, and the motion of this collection of air particles over time can be viewed as defining a continuous family of homeomorphisms on $\mathbb{R}^3$. Mathematically, the question is then, does there exists a continuous map

$$\varphi : [0,1] \times \mathbb{R}^3 \to \mathbb{R}^3$$

such that $\varphi(t,\cdot) : \mathbb{R}^3 \to \mathbb{R}^3$ is a homeomorphism for every $t \in [0,1]$, $\varphi(0,\cdot)$ is the identity map on $\mathbb{R}^3$ and $\varphi(1,\cdot) : \mathbb{R}^3 \to \mathbb{R}^3$ sends K_0 to K?

It turns out that the answer is no: in particular, if a homeomorphism $\varphi(1,\cdot)$ on $\mathbb{R}^3$ sending K_0 to K exists, then there must also be a homeomorphism between $\mathbb{R}^3\backslash K$ and $\mathbb{R}^3\backslash K_0$, and we will see that the latter is impossible. The reason is because we can associate to these spaces groups $\pi_1(\mathbb{R}^3\backslash K)$ and $\pi_1(\mathbb{R}^3\backslash K_0)$, which would need to be isomorphic if $\mathbb{R}^3\backslash K$ and $\mathbb{R}^3\backslash K_0$ were homeomorphic, and we will be able to compute enough information about both groups to show that they are not isomorphic.

Sample question 8.2. Here is another pair of spaces defined as subsets of $\mathbb{R}^3$:

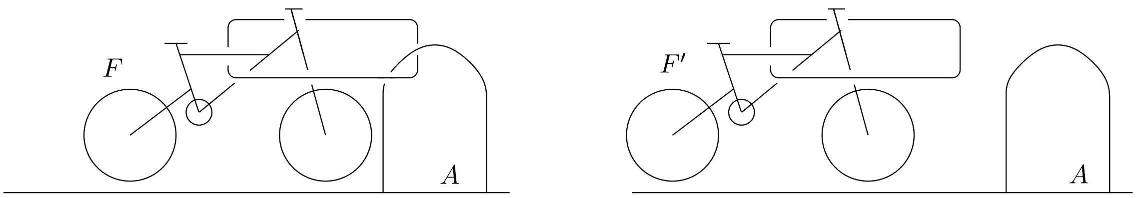

A question of tremendous practical import: can the set F in the picture at the left be shifted continuously to match the set F' in the picture at the right, but without "passing through" A, i.e. is there a continuous family of embeddings $F \hookrightarrow \mathbb{R}^3\backslash A$ that begins as the natural inclusion and ends by sending F to F'? If there is, then you may want to adjust your bike lock.

Of course there is no such continuous family of embeddings, and to see why, you could just delete the bicycle from the picture and pay attention only to the loop representing the bike lock, which is shown "linked" with A in the left picture and not in the right picture. The precise way to express the impossibility of deforming one picture to the other is that this loop is parametrized by a "noncontractible loop" $\gamma : S^1 \to \mathbb{R}^3\backslash A$, meaning γ represents a nontrivial element in the fundamental group $\pi_1(\mathbb{R}^3\backslash A)$.

Our task in this lecture is to define what the fundamental group is for an arbitrary space. We will then develop a few more of its general properties in the next lecture and spend the next four or five weeks developing methods to compute it.

We must first discuss paths in a space X. Since the unit interval $[0,1]$ will appear very often in the rest of this course, let us abbreviate it from now on by

$$I := [0,1].$$

For two points $x, y \in X$, a **path** (*Pfad*) from x to y is a map $\gamma : I \to X$ satisfying $\gamma(0) = x$ and $\gamma(1) = y$.[7] We will sometimes use the notation

$$x \overset{\gamma}{\rightsquigarrow} y$$

to indicate that γ is a path from x to y.

The **inverse** of a path $x \overset{\gamma}{\rightsquigarrow} y$ is the path

$$y \overset{\gamma^{-1}}{\rightsquigarrow} x$$

[7]This seems a good moment to emphasize that all maps in this course are assumed continuous unless otherwise noted.

defined by $\gamma^{-1}(t) := \gamma(1-t)$. The reason for this terminology and notation will become clearer when we give the definition of the fundamental group below. The same goes for the notion of the **product** of two paths: there is no natural multiplication defined for a pair of paths between arbitrary points, but given $x \overset{\alpha}{\rightsquigarrow} y$ and $y \overset{\beta}{\rightsquigarrow} z$, we can define the product path $x \overset{\alpha\cdot\beta}{\rightsquigarrow} z$ by

$$(\alpha \cdot \beta)(t) = \begin{cases} \alpha(2t) & \text{if } 0 \leqslant t \leqslant 1/2, \\ \beta(2t-1) & \text{if } 1/2 \leqslant t \leqslant 1. \end{cases} \tag{8.1}$$

This operation is also called a **concatenation** of paths. The **trivial path** at a point $x \in X$ is defined as the constant path $x \overset{e_x}{\rightsquigarrow} x$, i.e.

$$e_x(t) = x.$$

The idea is for this to play the role of the identity element in some kind of group structure.

If we want to turn concatenation into a product structure on a group, then we have one immediate problem: it is not associative. In fact, given paths $x \overset{\alpha}{\rightsquigarrow} y$, $y \overset{\beta}{\rightsquigarrow} z$ and $z \overset{\gamma}{\rightsquigarrow} a$, we have

$$\alpha \cdot (\beta \cdot \gamma) \neq (\alpha \cdot \beta) \cdot \gamma,$$

though clearly the images of these two concatenations are the same, and their difference is only in the way they are parametrized. We would like to introduce an equivalence relation on the set of paths that forgets this distinction in parametrizations.

DEFINITION 8.3. Two maps $f, g : X \to Y$ are **homotopic** (*homotop*) if there exists a map

$$H : I \times X \to Y \quad \text{such that } H(0, \cdot) = f \text{ and } H(1, \cdot) = g.$$

The map H is in this case called a **homotopy** (*Homotopie*) from f to g, and when a homotopy exists, we shall write

$$f \underset{h}{\sim} g.$$

It is straightforward to show that $\underset{h}{\sim}$ is an equivalence relation. In particular, if there are homotopies from f to g and from g to h, then by reparametrizing the parameter in $I = [0,1]$ we can "glue" the two homotopies together to form a homotopy from f to h. The definition of the new homotopy is analogous to the definition of the concatenation of paths in (8.1).

For paths in particular we will need a slightly more restrictive notion of homotopy that fixes the end points.

DEFINITION 8.4. For two paths α and β from x to y, we write

$$\alpha \underset{h+}{\sim} \beta$$

and say α is **homotopic with fixed end points** to β if there exists a map $H : I \times I \to X$ satisfying $H(0, \cdot) = \alpha$, $H(1, \cdot) = \beta$, $H(s, 0) = x$ and $H(s, 1) = y$ for all $s \in I$.

EXERCISE 8.5. Show that for any two points $x, y \in X$, $\underset{h+}{\sim}$ defines an equivalence relation on the set of all paths from x to y.

We will now prove several easy results about paths and homotopies. In most cases we will give precise formulas for the necessary homotopies, but one can also represent the main idea quite easily in pictures (see e.g. [**Hat02**, pp. 26–27]). We adopt the following convenient terminology: if $H : I \times X \to Y$ is a homotopy from $f_0 := H(0, \cdot) : X \to Y$ to $f_1 := H(1, \cdot) : X \to Y$, then we obtain a **continuous family of maps** $f_s := H(s, \cdot) : X \to Y$ for $s \in I$. The words "continuous family" will be understood as synonymous with "homotopy" in this sense.

PROPOSITION 8.6. *If $\alpha \underset{h+}{\sim} \alpha'$ are homotopic paths from x to y and $\beta \underset{h+}{\sim} \beta'$ are homotopic paths from y to z, then*

$$\alpha \cdot \beta \underset{h+}{\sim} \alpha' \cdot \beta'.$$

PROOF. By assumption, there exist continuous families of paths $x \overset{\alpha_s}{\rightsquigarrow} y$ and $y \overset{\beta_s}{\rightsquigarrow} z$ for $s \in I$ with $\alpha_0 = \alpha$, $\alpha_1 = \alpha'$, $\beta_0 = \beta$ and $\beta_1 = \beta'$. Then a homotopy with fixed end points from $\alpha \cdot \beta$ to $\alpha' \cdot \beta'$ can be defined via the continuous family

$$x \overset{\alpha_s \cdot \beta_s}{\rightsquigarrow} z \quad \text{for} \quad s \in I.$$

□

We next show that while concatenation of paths is not an associative operation, it is associative "up to homotopy".

PROPOSITION 8.7. *Given paths $x \overset{\alpha}{\rightsquigarrow} y$, $y \overset{\beta}{\rightsquigarrow} z$ and $z \overset{\gamma}{\rightsquigarrow} a$,*

$$(\alpha \cdot \beta) \cdot \gamma \underset{h+}{\sim} \alpha \cdot (\beta \cdot \gamma).$$

PROOF. A suitable homotopy $H : I \times I \to X$ can be defined as a family of linear reparametrizations of the sequence of paths α, β, γ:

$$H(s,t) = \begin{cases} \alpha\left(\frac{4t}{s+1}\right) & \text{if } 0 \leqslant t \leqslant \frac{s+1}{4}, \\ \beta(4t - (s+1)) & \text{if } \frac{s+1}{4} \leqslant t \leqslant \frac{s+2}{4}, \\ \gamma\left(\frac{4}{2-s}(t-1)+1\right) & \text{if } \frac{s+2}{4} \leqslant t \leqslant 1. \end{cases}$$

□

And finally, a result that allows us to interpret the constant paths e_x as "identity elements" and γ and γ^{-1} as "inverses":

PROPOSITION 8.8. *For any path $x \overset{\gamma}{\rightsquigarrow} y$, the following relations hold:*

(i) $e_x \cdot \gamma \underset{h+}{\sim} \gamma$
(ii) $\gamma \underset{h+}{\sim} \gamma \cdot e_y$
(iii) $\gamma \cdot \gamma^{-1} \underset{h+}{\sim} e_x$
(iv) $\gamma^{-1} \cdot \gamma \underset{h+}{\sim} e_y$

PROOF. For (i), we define a family of reparametrizations of the concatenated path $e_x \cdot \gamma$ that shrinks the amount of time spent on e_x from $1/2$ to 0:

$$H(s,t) = \begin{cases} x & \text{if } 0 \leqslant t \leqslant \frac{1-s}{2}, \\ \gamma\left(\frac{2}{s+1}(t-1)+1\right) & \text{if } \frac{1-s}{2} \leqslant t \leqslant 1. \end{cases}$$

The homotopy for (ii) is analogous.

For (iii), the idea is to define a family of paths that traverse only part of γ up to some time depending on s, then stay still for a suitable length of time and, in a third step, follow γ^{-1} back to x:

$$H(s,t) = \begin{cases} \gamma(2t) & \text{if } 0 \leqslant t \leqslant \frac{1-s}{2}, \\ \gamma(1-s) & \text{if } \frac{1-s}{2} \leqslant t \leqslant \frac{1+s}{2}, \\ \gamma(2-2t) & \text{if } \frac{1+s}{2} \leqslant t \leqslant 1. \end{cases}$$

The last relation follows from this by interchanging the roles of γ and γ^{-1}. □

The last three propositions combine to imply that the group structure in the following definition is a well-defined associative product which admits an identity element and inverses.

Definition 8.9. Given a space X and a point $p \in X$, the **fundamental group** (*Fundamentalgruppe*) of X with **base point** (*Basispunkt*) p is defined as the set of equivalence classes of paths $p \rightsquigarrow p$ up to homotopy with fixed end points:

$$\pi_1(X, p) := \left\{\text{paths } p \overset{\gamma}{\rightsquigarrow} p\right\} \Big/ \underset{h+}{\sim} .$$

The product of two equivalence classes $[\alpha], [\beta] \in \pi_1(X, p)$ is defined via concatenation:

$$[\alpha][\beta] := [\alpha \cdot \beta],$$

and the identity element is represented by the constant path $[e_p]$. The inverse element for $[\gamma] \in \pi_1(X, p)$ is represented by the reversed path γ^{-1}.

Before exploring the further properties of the group $\pi_1(X, p)$, let us clarify in what sense it is a "topological invariant" of the space X. Intuitively, we would like this to mean that whenever X and Y are two homeomorphic spaces, their fundamental groups should be isomorphic groups. What makes this statement a tiny bit more complicated is that the fundamental group of X doesn't just depend on X alone, but also on a choice of base point, so in order to make precise and correct statements about topological invariance, we will need to carry around a base point as extra data. The following definition is intended to formalize this notion.

Definition 8.10. A **pointed space** (*punktierter Raum*) is a pair (X, p) consisting of a topological space X and a point $p \in X$. The point $p \in X$ is in this case called the **base point** (*Basispunkt*) of X. Given pointed spaces (X, p) and (Y, q), any continuous map $f : X \to Y$ satisfying $f(p) = q$ is called a **pointed map** or **map of pointed spaces**, and can be denoted by

$$f : (X, p) \to (Y, q).$$

We also sometimes refer to such objects as **base-point preserving** maps. Finally, given two pointed maps $f, g : (X, p) \to (Y, q)$, a homotopy $H : I \times X \to Y$ from f to g that satisfies $H(s, p) = q$ for all $s \in I$ is called a **pointed homotopy**, or **homotopy of pointed maps**, or **base-point preserving homotopy**. One can equivalently describe such a homotopy as a continuous 1-parameter family of pointed maps $f_s := H(s, \cdot) : (X, p) \to (Y, q)$ defined for $s \in I$.

Here is the first main result about the topological invariance of π_1:

Theorem 8.11. *One can associate to every pointed map $f : (X, p) \to (Y, q)$ a group homomorphism*

$$f_* : \pi_1(X, p) \to \pi_1(Y, q) : [\gamma] \mapsto [f \circ \gamma],$$

which has the following properties:

(i) For any pointed maps $(X, p) \xrightarrow{f} (Y, q)$ and $(Y, q) \xrightarrow{g} (Z, r)$, $(g \circ f)_* = g_* \circ f_*$.
(ii) The map associated to the identity map $(X, p) \xrightarrow{\mathrm{Id}} (X, p)$ is the identity homomorphism $\pi_1(X, p) \xrightarrow{\mathbb{1}} \pi_1(X, p)$.
(iii) Each homomorphism f_* depends only on the pointed homotopy class of f.

Proof. It is clear that up to homotopy (with fixed end points), the path $q \overset{f\circ\gamma}{\rightsquigarrow} q$ in Y depends only on the path $p \overset{\gamma}{\rightsquigarrow} p$ only up to homotopy with fixed end points; indeed, if $H : I \times I \to X$ defines a homotopy with fixed end points between two paths α and β based at p, then $f \circ H : I \times I \to Y$ defines a corresponding homotopy between $f \circ \alpha$ and $f \circ \beta$. Similarly, if $[\gamma] \in \pi_1(X, p)$ and $f, g : (X, p) \to (Y, q)$ are homotopic via a base-point preserving homotopy $H : I \times X \to Y$, then

$h : I \times I \to Y : (s,t) \mapsto H(s, \gamma(t))$ defines a homotopy with fixed end points between $f \circ \gamma$ and $g \circ \gamma$. This shows that $f_* : \pi_1(X,p) \to \pi_1(Y,q)$ is a well-defined map that depends on f only up to base-point preserving homotopy. It is similarly easy to check that f_* is a homomorphism and satisfies the first two stated properties: e.g. for any two paths $p \overset{\alpha,\beta}{\rightsquigarrow} p$, we have

$$f_*([\alpha][\beta]) = [f \circ (\alpha \cdot \beta)] = [(f \circ \alpha) \cdot (f \circ \beta)] = f_*[\alpha] f_*[\beta]$$

and

$$f_*[e_p] = [e_q].$$

□

Corollary 8.12. *If X and Y are spaces admitting a homeomorphism $f : X \to Y$, then for any choice of base point $p \in X$, the groups $\pi_1(X,p)$ and $\pi_1(Y, f(p))$ are isomorphic.*

Proof. Abbreviate $q := f(p)$, so $f : (X,p) \to (Y,q)$ is a pointed map, and since its inverse is continuous, $f^{-1} : (Y,q) \to (X,p)$ is also a pointed map. Using Theorem 8.11, the commutative diagram (see Remark 8.14 below) of continuous maps

$$\begin{tikzcd} & (Y,q) \arrow[dr, "f^{-1}"] & \\ (X,p) \arrow[ur, "f"] \arrow[rr, "\mathrm{Id}"] & & (X,p) \end{tikzcd} \tag{8.2}$$

then gives rise to a similar commutative diagram of group homomorphisms

$$\begin{tikzcd} & \pi_1(Y,q) \arrow[dr, "f_*^{-1}"] & \\ \pi_1(X,p) \arrow[ur, "f_*"] \arrow[rr, "\mathbb{1}"] & & \pi_1(X,p) \end{tikzcd} \tag{8.3}$$

Reversing the roles of (X,p) and (Y,q) produces similar diagrams to show that f_* and f_*^{-1} are inverse homomorphisms, hence both are isomorphisms. □

Remark 8.13. The fancy way to summarize Theorem 8.11 is that π_1 defines a "covariant functor" from the category of pointed spaces and pointed homotopy classes to the category of groups and homomorphisms. We will discuss categories and functors more next semester in *Topologie II.*

Remark 8.14. Commutative diagrams such as (8.2) and (8.3) will appear more and more often as we get deeper into algebraic topology. When we say that such a diagram **commutes**, it means that any two maps obtained by composing a sequence of arrows along different paths from one place in the diagram to another must match, so e.g. the message carried by (8.2) is the relation $f^{-1} \circ f = \mathrm{Id}$, and (8.3) means $f_*^{-1} \circ f_* = \mathbb{1}$. These were especially simple examples, but later we will also encounter larger diagrams like

$$\begin{tikzcd} A \arrow[r, "f"] \arrow[d, "\alpha"] & B \arrow[r, "g"] \arrow[d, "\beta"] & C_* \arrow[d, "\gamma"] \\ A \arrow[r, "f'"] & B' \arrow[r, "g'"] & C' \end{tikzcd}$$

The purpose of this one is to communicate the two relations $\beta \circ f = f' \circ \alpha$ and $\gamma \circ g = g' \circ \beta$, along with all the more complicated relations that follow from these, such as $g' \circ f' \circ \alpha = \gamma \circ g \circ f$.

Since the paths representing elements of $\pi_1(X,p)$ have the same fixed starting and ending point, we often think of them as *loops* in X. We will establish some general properties of $\pi_1(X,p)$ in the next lecture, starting with the observation that whenever X is path-connected, $\pi_1(X,p)$ up to isomorphism does not actually depend on the choice of the base point $p \in X$, thus we can sensibly write it as $\pi_1(X)$. Computing $\pi_1(X)$ for a given space X is not always easy or possible, but we will develop some methods that are very effective on a wide class of spaces. I can already mention two simple examples: first, $\pi_1(\mathbb{R}^n)$ is the trivial group, resulting from the relatively obvious fact that (by linear interpolation) every path in $\mathbb{R}^n$ from a point to itself is homotopic with fixed end points to the constant path. In contrast, we will see that $\pi_1(S^1)$ and $\pi_1(\mathbb{R}^2\backslash\{0\})$ are both isomorphic to the integers, and this simple result already has many useful applications, e.g. we will derive from it a very easy proof of the fundamental theorem of algebra.

9. Some properties of the fundamental group (May 16, 2023)

We would now like to clarify to what extent $\pi_1(X,p)$ depends on p in addition to X.

Theorem 9.1. *Given $p,q \in X$, any homotopy class (with fixed end points) of paths $p \overset{\gamma}{\rightsquigarrow} q$ determines a group isomorphism*

$$\Phi_\gamma : \pi_1(X,q) \to \pi_1(X,p) : [\alpha] \mapsto [\gamma \cdot \alpha \cdot \gamma^{-1}].$$

Proof. Note that in writing the formula above for $\Phi_\gamma([\alpha])$, we are implicitly using the fact (Proposition 8.7) that concatenation of paths is an associative operation up to homotopy, so one can represent $\Phi_\gamma([\alpha])$ by either of the paths $\gamma\cdot(\alpha\cdot\gamma^{-1})$ or $(\gamma\cdot\alpha)\cdot\gamma^{-1}$ without the result depending on this choice. Similarly, Proposition 8.6 implies that the homotopy class of $\gamma\cdot\alpha\cdot\gamma^{-1}$ with fixed end points only depends on the homotopy classes of γ and α (also with fixed end points).[8] This proves that Φ_γ is a well-defined map as written. The propositions in the previous lecture imply in a similarly straightforward manner that Φ_γ is a homomorphism, i.e.

$$\Phi_\gamma([\alpha][\beta]) = [\gamma\cdot\alpha\cdot\beta\cdot\gamma^{-1}] = [\gamma\cdot\alpha\cdot\gamma^{-1}\cdot\gamma\cdot\beta\cdot\gamma^{-1}] = \Phi_\gamma([\alpha])\Phi_\gamma([\beta]),$$

and

$$\Phi_\gamma([e_q]) = [\gamma\cdot e_q\cdot\gamma^{-1}] = [\gamma\cdot\gamma^{-1}] = [e_p].$$

It remains only to observe that Φ_γ and $\Phi_{\gamma^{-1}}$ are inverses of each other, hence both are isomorphisms. ☐

Corollary 9.2. *If X is path-connected, then $\pi_1(X,p)$ up to isomorphism is independent of the choice of base point $p \in X$.* ☐

Due to this corollary, it is conventional to abbreviate the fundamental group by

$$\pi_1(X) := \pi_1(X,p)$$

whenever X is path-connected, and we will see many theorems about $\pi_1(X)$ in situations where the base point plays no important role. If X is not path-connected but $X_0 \subset X$ denotes the path-component containing p, then $\pi_1(X,p) = \pi_1(X_0,p) \cong \pi_1(X_0)$, so in practice it is sufficient to restrict our attention to path-connected spaces. Some caution is nonetheless warranted in using the notation $\pi_1(X)$: strictly speaking, $\pi_1(X)$ is not a concrete group but only an isomorphism class of groups, and the subtle distinction between these two notions occasionally leads to trouble. You should always keep in the back of your mind that even if the base point is not mentioned, it is an essential piece of the definition of $\pi_1(X)$.

[8]Note that the homotopy class of γ determines that of γ^{-1}. (Why?)

We next discuss some alternative ways to interpret $\pi_1(X,p)$. Recall the following useful notational device: given a space X with subset $A \subset X$, we define

$$X/A := X/\sim$$

with the quotient topology, where the equivalence relation defines $a \sim b$ for all $a, b \in A$. In other words, this is the quotient space obtained from X by "collapsing" the subset A to a single point. For example, it is straightforward (see Exercise 5.16) to show that $\mathbb{D}^n/S^{n-1}$ is homeomorphic to S^n for every $n \in \mathbb{N}$, and if we replace $\mathbb{D}^1 = [-1,1]$ by the unit interval $I = [0,1]$, we obtain the special case

$$[0,1]/\{0,1\} = I/\partial I \cong S^1.$$

Here we have used the notation

$$\partial X := \text{“boundary of } X\text{”},$$

which comes from differential geometry, so for instance $\partial\mathbb{D}^n = S^{n-1}$ and we can therefore also identify S^n with $\mathbb{D}^n/\partial\mathbb{D}^n$. A specific homeomorphism $I/\partial I \to S^1$ can be written most easily by thinking of S^1 as the unit circle in $\mathbb{C}$:

$$I/\partial I \to S^1 : [t] \mapsto e^{2\pi i t}.$$

LEMMA 9.3. *For any space X and subset $A \subset X$, there is a canonical bijection between the set of all continuous maps $f : X \to Y$ that are constant on A and the set of all continuous maps $g : X/A \to Y$. For any two maps f and g that correspond under this bijection, the diagram*

$$\begin{tikzcd} X \arrow[rr, "\pi"] \arrow[dr, "f"'] & & X/A \arrow[dl, "g"] \\ & Y & \end{tikzcd}$$

commutes, where $\pi : X \to X/A$ denotes the quotient projection; in other words, $g \circ \pi = f$.

PROOF. The diagram determines the correspondence: given $g : X/A \to Y$, we can define $f := g \circ \pi$ to obtain a map $X \to Y$ that is automatically constant on A, and conversely, if $f : X \to Y$ is given and is constant on A, then there is a well-defined map $g : X/A \to Y : [x] \mapsto f(x)$. Our main task is to show that f is continuous if and only if g is continuous. In one direction this is immediate: if g is continuous, then $f = g \circ \pi$ is the composition of two continuous maps and is therefore also continuous. Conversely, if f is continuous, then for every open set $\mathcal{U} \subset Y$, we know $f^{-1}(\mathcal{U}) \subset X$ is open. A point $[x] \in X/A$ is then in $g^{-1}(\mathcal{U})$ if and only if $x \in f^{-1}(\mathcal{U})$, so $g^{-1}(\mathcal{U}) = \pi(f^{-1}(\mathcal{U}))$ and thus $\pi^{-1}(g^{-1}(\mathcal{U})) = f^{-1}(\mathcal{U})$ is open. By the definition of the quotient topology, this means that $g^{-1}(\mathcal{U}) \subset X/A$ is open, so g is continous. $\square$

Lemma 9.3 gives a canonical bijection between the set of all paths $p \overset{\gamma}{\rightsquigarrow} p$ in X beginning and ending at the base point and the set of all continuous pointed maps

$$(I/\partial I, [0]) \to (X, p).$$

It is easy to check moreover that two paths $p \overset{\gamma}{\rightsquigarrow} p$ are homotopic with fixed end points if and only if they correspond to maps $(I/\partial I, [0]) \to (X, p)$ in the same *pointed* homotopy class. Under the aforementioned homeomorphism $I/\partial I \cong S^1 \subset \mathbb{C}$ that identifies $[0] = [1]$ with 1, this gives us an alternative description of $\pi_1(X,p)$ as

$$\pi_1(X,p) = \left\{\text{pointed maps } \gamma : (S^1, 1) \to (X,p)\right\} \big/ \underset{h+}{\sim},$$

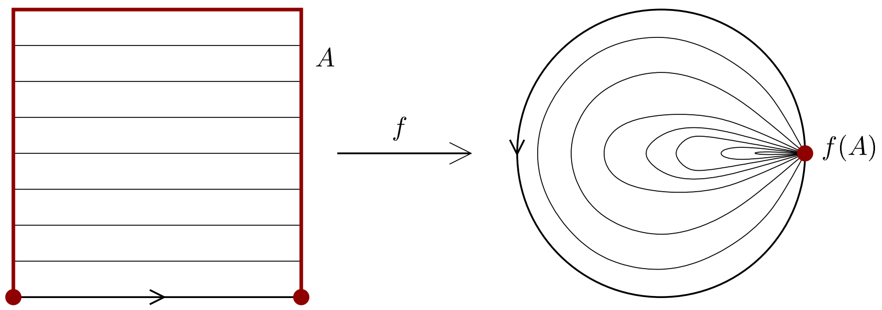

FIGURE 1. A map $f : I^2 \to \mathbb{D}^2$ which descends to a homeomorphism $g : I^2/A \to \mathbb{D}^2$ in the proof of Theorem 9.4.

where $\underset{h+}{\sim}$ now denotes the equivalence relation defined by pointed homotopy. The group structure of $\pi_1(X, p)$ is less easy to see from this perspective, but it will nonetheless be extremely useful to think of elements of $\pi_1(X)$ as represented by *loops* $\gamma : S^1 \to X$.

THEOREM 9.4. *A loop $\gamma : (S^1, 1) \to (X, p)$ represents the identity element in $\pi_1(X, p)$ if and only if there exists a continuous map $u : \mathbb{D}^2 \to X$ with $u|_{\partial\mathbb{D}^2} = \gamma$.*

PROOF. I can't explain this proof without a picture, so to start with, have a look at Figure 1. It depicts a map $f : I^2 \to \mathbb{D}^2 \subset \mathbb{C}$ that collapses the red region consisting of three sides of the square

$$A := (\partial I \times I) \cup (I \times \{1\}) \subset I^2$$

to the single point $f(A) = \{1\} \subset \mathbb{D}^2$, but is bijective everywhere else, and maps the path $I \times \{0\} \subset I^2$ to the loop $\partial\mathbb{D}^2$. By Lemma 9.3, f determines a map

$$g : I^2/A \to \mathbb{D}^2$$

which is continuous and bijective, and it is also an open map (i.e. it maps open sets to open sets), hence its inverse is also continuous and g is therefore a homeomorphism. Now, a path $\gamma : I \to X$ with $\gamma(0) = \gamma(1) = p$ represents the identity in $\pi_1(X, p)$ if and only if there exists a homotopy $H : I^2 \to X$ with $H(0, \cdot) = \gamma$ and $H|_A \equiv p$. Applying Lemma 9.3 again, such a map is equivalent to a map $h : I^2/A \to X$ which sends the equivalence class represented by every point in A to the base point p. In this case, $h \circ g^{-1}$ is a map $\mathbb{D}^2 \to X$ whose restriction to $\partial\mathbb{D}^2$ is the loop $S^1 \cong I/\partial I \to X$ determined by $\gamma : I \to X$. □

REMARK 9.5. Maps $\gamma : S^1 \to X$ that admit extensions over $\mathbb{D}^2$ as in the above theorem are called **contractible loops** (*zusammenziehbare Schleifen*).

DEFINITION 9.6. A space X is called **simply connected** (*einfach zusammenhängend*) if it is path-connected and its fundamental group is trivial.

It is common to denote the trivial group by "0", so for path-connected spaces, we can write

$$X \text{ is simply connected} \quad \Leftrightarrow \quad \pi_1(X) = 0.$$

By Theorem 9.4, this is equivalent to the condition that every map $\gamma : S^1 \to X$ admits a continuous extension $u : \mathbb{D}^2 \to X$ satisfying $u|_{\partial\mathbb{D}^2} = \gamma$. Note that there was no need to mention the base point in this formulation: if X is path-connected, then $\pi_1(X) = 0$ means $\pi_1(X, p) = 0$ for every p, so for a given loop $\gamma : S^1 \to X$ we are free to choose $p := \gamma(1) \in X$ as the base point and then apply Theorem 9.4.

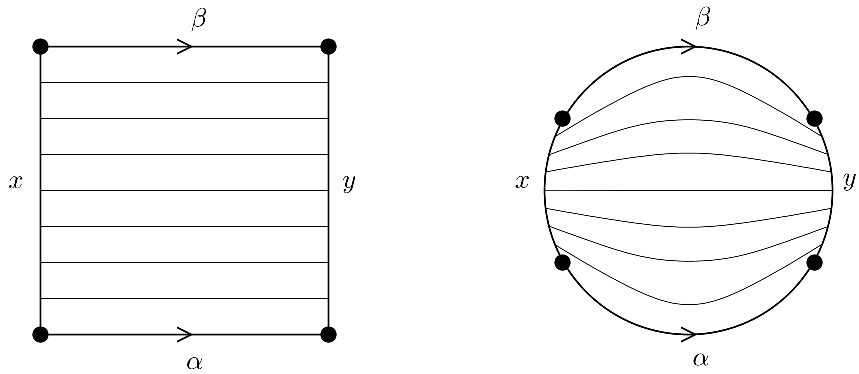

FIGURE 2. Two equivalent pictures of the same homotopy with fixed end points x and y between two paths α and β, using a homeomorphism $I^2 \cong \mathbb{D}^2$.

EXAMPLES 9.7. Though we will need to develop a few more tools before we can prove it, the sphere S^2 is simply connected. (Try to imagine a loop in S^2 that cannot be filled in by a disk—but do not try too hard!)

In contrast, $\mathbb{R}^2\backslash\{0\}$ is not simply connected: we will see that the natural inclusion map $\gamma : S^1 \hookrightarrow \mathbb{R}^2\backslash\{0\}$ is an example of a loop that cannot be extended to a map $u : \mathbb{D}^2 \to \mathbb{R}^2\backslash\{0\}$. Of course, it *can* be extended to a map $\mathbb{D}^2 \to \mathbb{R}^2$, but it will turn out that such an extension must always hit the origin somewhere—in other words, the loop is contractible in $\mathbb{R}^2$, but not contractible in $\mathbb{R}^2\backslash\{0\}$. This observation has many powerful implications, e.g. we will see in the next lecture that it is the key idea behind one of the simplest proofs of the *fundamental theorem of algebra*, that every nonconstant complex polynomial has a root.

Another example with nontrivial fundamental group is the **torus** $\mathbb{T}^2 := S^1 \times S^1$. Pictures of this space embedded in $\mathbb{R}^3$ typically depict it as the surface of a tube (or a doughnut or a bagel—depending on your cultural preferences). Can you visualize a loop on this surface that is contractible in $\mathbb{R}^3$ but not in $\mathbb{T}^2$?

One can also use the fundamental group to gain insight into homotopy classes of non-closed paths:

THEOREM 9.8. *Two paths $x \overset{\alpha,\beta}{\rightsquigarrow} y$ in X are homotopic with fixed end points if and only if the concatenated path $x \overset{\alpha\cdot\beta^{-1}}{\rightsquigarrow} x$ represents the identity element in $\pi_1(X, x)$.*

PROOF. The condition $\alpha \underset{h+}{\sim} \beta$ means the existence of a homotopy $H : I^2 \to X$ with certain properties as depicted at the left in Figure 2, but by a suitable choice of homeomorphism $I^2 \cong \mathbb{D}^2$ as shown to the right of that picture, we can equally well regard H as a map $\mathbb{D}^2 \to X$. The loop $\gamma := H|_{\partial\mathbb{D}^2} : S^1 \to X$ can then be viewed as the concatenation $\alpha \cdot e_y \cdot \beta^{-1} \cdot e_x$, which by Proposition 8.8 is homotopic with fixed end points to $\alpha \cdot \beta^{-1}$. The result then follows directly from Theorem 9.4. $\square$

COROLLARY 9.9. *A space X is simply connected if and only if for every pair of points $p, q \in X$, there exists a path from p to q and it is unique up to homotopy with fixed end points.* $\square$

Let us finally work out a few concrete examples.

EXAMPLE 9.10. For each $n \geqslant 0$, the Euclidean space $\mathbb{R}^n$ is simply connected. Indeed, since it is path-connected, we are free to choose the base point $0 \in \mathbb{R}^n$, and can then observe that every loop $0 \overset{\gamma}{\rightsquigarrow} 0$ is homotopic to the constant loop via the continuous family of loops

$$\gamma_s : I \to \mathbb{R}^n : t \mapsto s\gamma(t) \quad \text{for} \quad s \in I.$$

EXAMPLE 9.11. Since every open ball $B_r(x)$ in $\mathbb{R}^n$ is homeomorphic to $\mathbb{R}^n$ itself, Corollary 8.12 implies that $\pi_1(B_r(x))$ also vanishes, i.e. $B_r(x)$ is simply connected. One could also give a direct proof of this, analogously to Example 9.10: just choose $x \in B_r(x)$ as the base point and define γ_s via linear interpolation between γ and the constant loop at x. A similar trick works in fact for any *convex* subset $K \subset \mathbb{R}^n$, i.e. any set K with the property that the straight line segment connecting any two points $x, y \in K$ is also contained in K. It follows that all convex subsets of finite-dimensional vector spaces are simply connected.

EXAMPLE 9.12. Our first example of a nontrivial fundamental group (and probably also the most important one to take note of in this course) is the circle: we claim that

$$\pi_1(S^1) \cong \mathbb{Z}.$$

The proof is based on a pair of lemmas that we will prove (in more general forms) in a few weeks, though I suspect you will already find them easy to believe. Regarding S^1 as the unit circle in $\mathbb{C}$, consider the map

$$f : \mathbb{R} \to S^1 : t \mapsto e^{2\pi i t}.$$

This is our first interesting example of a so-called **covering map** (*Überlagerung*): it is surjective, and it looks like a homeomorphism *on the small scale* (i.e. if you zoom in close enough on any particular point in $\mathbb{R}$), but it is not injective, in fact it "wraps" the line $\mathbb{R}$ around S^1 infinitely many times. The next two statements are special cases of results that we will later prove about a much more general class of covering spaces:

(1) Given a path $x \overset{\gamma}{\rightsquigarrow} y$ in S^1 and a point $\tilde{x} \in f^{-1}(x)$, there exists a unique path $\tilde{x} \overset{\tilde{\gamma}}{\rightsquigarrow} \tilde{y}$ in $\mathbb{R}$ that is a "lift" of γ in the sense that $f \circ \tilde{\gamma} = \gamma$.
(2) Given a homotopy $H : I \times I \to S^1$ of paths $x \overset{\gamma}{\rightsquigarrow} y$ (with fixed end points) and a point $\tilde{x} \in f^{-1}(x)$, there exists a unique homotopy $\widetilde{H} : I \times I \to \mathbb{R}$ of lifted paths $\tilde{x} \overset{\tilde{\gamma}}{\rightsquigarrow} \tilde{y}$ which lifts H in the sense that $f \circ \widetilde{H} = H$.

Now for any $[\gamma] \in \pi_1(S^1, 1)$ represented by a path $1 \overset{\gamma}{\rightsquigarrow} 1$, there is a unique lift to a path $0 \overset{\tilde{\gamma}}{\rightsquigarrow} \tilde{\gamma}(1)$ in $\mathbb{R}$. Unlike γ, the end point of the lift need not match its starting point, but the fact that it is a lift implies $\tilde{\gamma}(1) \in f^{-1}(1) = \mathbb{Z}$, and the fact that homotopies can be lifted implies that this integer does not change if we replace γ with any other representative of $[\gamma] \in \pi_1(S^1, 1)$. We therefore obtain a well-defined map

$$\Phi : \pi_1(S^1, 1) \to \mathbb{Z} : [\gamma] \mapsto \tilde{\gamma}(1).$$

It is easy to show that Φ is a group homomorphism by lifting concatenated paths. Moreover, Φ is surjective since $\Phi([\gamma_k]) = k$ for each of the loops $\gamma_k(t) = e^{2\pi i k t}$ with $k \in \mathbb{Z}$, as these have lifts $\tilde{\gamma}(t) = kt$. Injectivity amounts to the statement that γ must be homotopic to a constant whenever its lift satisfies $\tilde{\gamma}(1) = 0$, and this follows from the fact that $\pi_1(\mathbb{R}) = 0$: indeed, in this case $\tilde{\gamma}$ is not just a path in $\mathbb{R}$ but is also a loop, thus it represents an element of $\pi_1(\mathbb{R}, 0) = 0$ and is therefore homotopic to the constant loop. Composing that homotopy with $f : \mathbb{R} \to S^1$ gives a homotopy of the original loop γ to a constant.

EXERCISE 9.13. In this exercise we show that the fundamental group of a product is a product of fundamental groups.

(a) Given two pointed spaces (X, x) and (Y, y), prove that $\pi_1(X \times Y, (x, y))$ is isomorphic to the product group $\pi_1(X, x) \times \pi_1(Y, y)$.

Hint: Use the projections $p^X : X \times Y \to X$ and $p^Y : X \times Y \to Y$ to define a natural map from π_1 of the product to the product of π_1's, then prove that it is an isomorphism.

(b) Generalize part (a) to the case of an infinite product of pointed spaces (with the product topology).

EXERCISE 9.14. Let us regard $\pi_1(X, p)$ as the set of base-point preserving homotopy classes of maps $(S^1, \mathrm{pt}) \to (X, p)$, and let $[S^1, X]$ denote the set of homotopy classes of maps $S^1 \to X$, with no conditions on base points. (The elements of $[S^1, X]$ are called **free homotopy classes** of loops in X). There is a natural map

$$F : \pi_1(X, p) \to [S^1, X]$$

defined by ignoring base points. Prove:

(a) F is surjective if X is path-connected.

(b) $F([\alpha]) = F([\beta])$ if and only if $[\alpha]$ and $[\beta]$ are conjugate in $\pi_1(X, p)$.
Hint: If $H : [0,1] \times S^1 \to X$ is a homotopy with $H(0, \cdot) = \alpha$ and $H(1, \cdot) = \beta$, and $t_0 \in S^1$ is the base point in S^1, then $\gamma := H(\cdot, t_0) : [0,1] \to X$ begins and ends at p, and therefore also defines a loop. Compare α and the concatenation $\gamma \cdot \beta \cdot \gamma^{-1}$.

The conclusion is that if X is path-connected, F induces a bijection between $[S^1, X]$ and the set of conjugacy classes in $\pi_1(X)$. In particular, $\pi_1(X) \cong [S^1, X]$ whenever $\pi_1(X)$ is abelian.

10. Retractions and homotopy equivalence (May 23, 2023)

Having proved that two homeomorphic spaces always have isomorphic fundamental groups, it is natural to wonder whether the converse is true. The answer is an emphatic *no*, but this will turn out to be more of an advantage than a disadvantage: it becomes much easier to compute $\pi_1(X)$ if we are free to replace X with another space X' that is not homeomorphic to X but still has certain features in common. This idea leads us naturally to the notion of *homotopy equivalence*, another equivalence relation on topological spaces that is strictly weaker than homeomorphism.

Let us first discuss conditions that make the homomorphisms $f_* : \pi_1(X, p) \to \pi_1(Y, q)$ injective or surjective.

DEFINITION 10.1. For a space X with subset $A \subset X$, a map $f : X \to A$ is called a **retraction** (*Retraktion*) if $f|_A$ is the identity map $A \to A$. Equivalently, if $i : A \hookrightarrow X$ denotes the natural inclusion map, then f being a retraction means that the following diagram commutes:

$$\begin{tikzcd} A \arrow[rr, "\mathrm{Id}"] \arrow[dr, hook, "i"'] & & A \\ & X \arrow[ur, "f"'] & \end{tikzcd} \tag{10.1}$$

We say in this case that A is a **retract** of X.

EXAMPLE 10.2. For $A := \mathbb{R} \times \{0\} \subset \mathbb{R}^2$, the map $f : \mathbb{R}^2 \to A : (x, y) \mapsto (x, 0)$ is a retraction.

A wide class of examples of retractions arises from the following general construction.

DEFINITION 10.3. The **wedge sum** of two pointed spaces (X, p) and (Y, q) is the space

$$X \vee Y := (X \amalg Y)/\sim$$

where the equivalence relation sets $p \in X$ equivalent to $q \in Y$ and is otherwise trivial. More generally, any (potentially infinite) collection of pointed spaces $\{(X_\alpha, p_\alpha)\}_{\alpha \in J}$ has a wedge sum

$$\bigvee_{\alpha \in J} X_\alpha := \coprod_{\alpha \in J} X_\alpha \Big/ \sim,$$

where the equivalence relation identifies all the base points $p_\alpha \sim p_\beta$ for $\alpha, \beta \in J$. The wedge sum is naturally also a pointed space, with base point $[p_\alpha] \in \bigvee_\beta X_\beta$.

REMARK 10.4. I did not specify the topology on $X \vee Y$ or $\bigvee_\alpha X_\alpha$, but by now you know enough to deduce from context what it must be: e.g. for the wedge of two spaces, we assign the disjoint union topology to $X \amalg Y$ and then endow $(X \amalg Y)/\sim$ with the resulting quotient topology. We will see many more constructions of this sort that involve a combination of quotients with disjoint unions and/or products, so you should always assume unless otherwise specified that the topology is whatever arises naturally from disjoint union, product and/or quotient topologies.

The notation for wedge sums is slightly nonideal since the definition of $\bigvee_\alpha X_\alpha$ depends not just on the spaces X_α but also on their base points $p_\alpha \in X_\alpha$, and it is not true in general that changing base points always produces homeomorphic wedge sums. It is true however for most examples that arise in practice, so the ambiguity in notation will usually not cause a problem. Note that since each of the individual spaces X_α are naturally subspaces of $\coprod_\beta X_\beta$, they can equally well be regarded as subspaces of $\bigvee_\beta X_\beta$, and it is straightforward to show that the obvious inclusion $X_\alpha \hookrightarrow \bigvee_\beta X_\beta$ for each α is a homeomorphism onto its image. As subspaces of a disjoint union $\coprod_\alpha X_\alpha$, the individual spaces X_β and X_γ for $\beta \neq \gamma$ are by definition disjoint, whereas in $\bigvee_\alpha X_\alpha$, they intersect each other at the base point, and only there.

EXERCISE 10.5. Show that for any collection of pointed maps $\{f_\alpha : (X_\alpha, p_\alpha) \to (Y, q)\}_{\alpha \in J}$, the unique map $f : \bigvee_{\alpha \in J} X_\alpha \to Y$ determined by the condition $f|_{X_\alpha} = f_\alpha$ for each $\alpha \in J$ is continuous.

EXAMPLE 10.6. For the wedge sum $X \vee Y$ of two pointed spaces (X, p) and (Y, q), there is a natural base-point preserving retraction

$$f : X \vee Y \to X : [x] \mapsto \begin{cases} x & \text{if } x \in X, \\ p & \text{if } x \in Y. \end{cases}$$

In words, f maps $X \subset X \vee Y$ to itself as the identity map while collapsing all of $Y \subset X \vee Y$ to the base point. One can analogously define a natural retraction $X \vee Y \to Y$, and for a wedge sum of arbitrarily many spaces, a natural retraction $\bigvee_{\beta \in J} X_\beta \to X_\alpha$ for each $\alpha \in J$.

EXERCISE 10.7. Convince yourself that the map $f : X \vee Y \to X$ in Example 10.6 is continuous.

EXAMPLE 10.8. For $X = Y = S^1$, the wedge sum $S^1 \vee S^1$ is a space homeomorphic to the symbols "8" and "∞", i.e. a so-called *figure eight*. Note that in this case, we did not need to specify the base points on the two copies of S^1 because choosing different base points leads to wedge sums that are homeomorphic. As a special case of Example 10.6, there are two retractions $S^1 \vee S^1 \to S^1$ that collapse either the top half or the bottom half of the "8" to a point.

The next example originates in the proof of the Brouwer fixed point theorem that we sketched at the end of Lecture 1 (cf. Theorem 1.13).

EXAMPLE 10.9. As explained in Lecture 1, if there exists a continuous map $f : \mathbb{D}^n \to \mathbb{D}^n$ with no fixed point, then one can use it to define a map $g : \mathbb{D}^n \to \partial\mathbb{D}^n = S^{n-1}$ that satisfies $g(x) = x$ for all $x \in \partial\mathbb{D}^n$. The idea is to follow the unique line from x through $f(x)$ until arriving at some point of the boundary, which is defined to be $g(x)$. This makes g a retraction of $\mathbb{D}^n$ to $\partial\mathbb{D}^n$. The main step in the proof of Brouwer's fixed point theorem is to show that no such retraction exists. We will carry this out for $n = 2$ in a moment.

THEOREM 10.10. *If $f : X \to A$ is a retraction and $i : A \hookrightarrow X$ denotes the inclusion, then for any choice of base point $a \in A$, the induced homomorphism $i_* : \pi_1(A, a) \to \pi_1(X, a)$ is injective, while $f_* : \pi_1(X, a) \to \pi_1(A, a)$ is surjective.*

PROOF. Since the maps in the commutative diagram (10.1) all send the base point $a \in A$ to itself, Theorem 8.11 produces a corresponding commutative diagram of homomorphisms:

$$\begin{tikzcd} \pi_1(A,a) \arrow[rr, "\mathbb{1}"] \arrow[dr, "i_*"'] & & \pi_1(A,a) \\ & \pi_1(X,a) \arrow[ur, "f_*"'] & \end{tikzcd}$$

In particular, $f_* \circ i_*$ is both injective and surjective, which is only possible if i_* is injective and f_* is surjective. ☐

PROOF OF THE BROUWER FIXED POINT THEOREM FOR $n = 2$. If there is a map $f : \mathbb{D}^2 \to \mathbb{D}^2$ with no fixed point, then there is also a retraction $g : \mathbb{D}^2 \to \partial\mathbb{D}^2 = S^1$ as explained in Example 10.9, so Theorem 10.10 implies that the induced homomorphism $g_* : \pi_1(\mathbb{D}^2) \to \pi_1(S^1)$ is surjective. As we saw at the end of the previous lecture, $\pi_1(S^1) \cong \mathbb{Z}$, and an easy modification of Example 9.10 shows that $\pi_1(\mathbb{D}^2) = 0$. (In fact, the same argument proves that every convex subset of $\mathbb{R}^n$ is simply connected—this will also follow from the more general Corollary 10.24 below.) But there is no surjective homomorphism from the trivial group to $\mathbb{Z}$, so this is a contradiction. ☐

DEFINITION 10.11. Assume X is a space with subset $A \subset X$ and $i : A \hookrightarrow X$ denotes the inclusion. A **deformation retraction** (*Deformationsretraktion*) of X to A is a homotopy $H : I \times X \to X$ such that $H(s, \cdot)|_A = \mathrm{Id}_A$ for every $s \in I$, $H(1, \cdot) = \mathrm{Id}_X$ and $H(0, \cdot) = i \circ f$ for some retraction $f : X \to A$. If a deformation retraction exists, we say that A is a a **deformation retract** (*Deformationsretrakt*) of X.

You should imagine a deformation retraction as a gradual "pulling" of all points in X toward the subset A until eventually all of them end up in A.

EXAMPLE 10.12. We call $X \subset \mathbb{R}^n$ a **star-shaped domain** (*sternförmige Menge*) if for every $x \in X$, the rescaled vector tx is also in X for every $t \in [0, 1]$. In this case $H(t, x) := tx$ defines a deformation retraction of X to the one-point subset $\{0\}$.

EXAMPLE 10.13. This is actually a non-example: while the maps $f : S^1 \vee S^1 \to S^1$ in Example 10.8 are retractions, $i \circ f$ in this case is not homotopic to the identity on $S^1 \vee S^1$, so S^1 is not a deformation retract of $S^1 \vee S^1$. We are not yet in a position to prove this, as it will require more knowledge of $\pi_1(S^1 \vee S^1)$ than we presently have, but the necessary results will be proved within the next four lectures. For now, feel free to try to imagine how you might define a homotopy of maps $S^1 \vee S^1 \to S^1 \vee S^1$ that starts with the identity and ends with a retraction collapsing one of the circles. (Keep in mind however that it is not possible, so don't try too hard.)

EXAMPLE 10.14. The sphere $S^{n-1} \subset \mathbb{R}^n \backslash \{0\}$ is a deformation retract of the punctured Euclidean space. A suitable homotopy $H : I \times (\mathbb{R}^n \backslash \{0\}) \to \mathbb{R}^n \backslash \{0\}$ can be defined by

$$H(t, x) = \frac{x}{t + (1 - t)|x|},$$

which makes $H(1, \cdot)$ the identity map, while $H(0, x) := x/|x|$ retracts $\mathbb{R}^n \backslash \{0\}$ to S^{n-1} and $H(t, x) = x$ for $x \in S^{n-1}$. It is important to observe that no continuous map can be defined in this way with all of $\mathbb{R}^n$ as its domain: the removal of one point changes the topology of $\mathbb{R}^n$ in an essential way that makes the deformation retraction to S^{n-1} possible. (We will later be able to prove that $\mathbb{R}^n$ does not admit any retraction to S^{n-1}. When $n = 2$, this already follows from Theorem 10.10 since $\pi_1(S^1) \cong \mathbb{Z}$ and $\pi_1(\mathbb{R}^2) = 0$.)

EXAMPLE 10.15. Writing $S^n = \{(\mathbf{x}, z) \in \mathbb{R}^n \times \mathbb{R} \mid |\mathbf{x}|^2 + z^2 = 1\}$, define the two "poles" $p_\pm = (0, \pm 1)$. Removing these poles produces a space that can be decomposed into a 1-parameter family of $(n-1)$-spheres, i.e. there is a homeomorphism

$$S^n \backslash \{p_+, p_-\} \xrightarrow{\cong} S^{n-1} \times (-1, 1) : (\mathbf{x}, z) \mapsto \left(\frac{\mathbf{x}}{|\mathbf{x}|}, z \right).$$

If we identify $S^n \backslash \{p_+, p_-\}$ with $S^{n-1} \times (-1,1)$ in this way, then we see that the "equator" $S^{n-1} \times \{0\} \subset S^n$ is a deformation retract of $S^n \backslash \{p_+, p_-\}$. This follows from the fact that $\{0\}$ is a deformation retract of $(-1,1)$.

DEFINITION 10.16. A map $f : X \to Y$ is a **homotopy equivalence** (*Homotopieäquivalenz*) if there exists a map $g : Y \to X$ such that $g \circ f$ and $f \circ g$ are each homotopic to the identity map on X and Y respectively. When this exists, we say that g is a **homotopy inverse** (*Homotopieinverse*) of f, and that the spaces X and Y are **homotopy equivalent** (*homotopieäquivalent*). This defines an equivalence relation on topological spaces which we shall denote in these notes by

$$X \underset{h.e.}{\simeq} Y.$$

EXERCISE 10.17. Verify that homotopy equivalence defines an equivalence relation.

REMARK 10.18. The notation "$\underset{h.e.}{\simeq}$" for homotopy equivalence is not universal, and there are several similar but slightly different standards that frequently appear in the literature. This one happens to be my current favorite, but I may change to something else next year.

EXAMPLE 10.19. A homeomorphism $f : X \to Y$ is obviously also a homotopy equivalence, with homotopy inverse f^{-1}.

EXAMPLE 10.20. If $H : I \times X \to X$ is a deformation retraction with $H(0, \cdot) = f \circ i$ for a retraction $f : X \to A$, then the inclusion $i : A \hookrightarrow X$ is a homotopy inverse of f, so that both f and i are homotopy equivalences and thus $X \underset{h.e.}{\simeq} A$. Indeed, the retraction condition implies that $f \circ i$ is not just homotopic but also equal to Id_A, and adding the word "deformation" provides the condition $i \circ f \underset{h}{\sim} \mathrm{Id}_X$.

DEFINITION 10.21. We say that a space X is **contractible** (*zusammenziehbar* or *kontrahierbar*) if it is homotopy equivalent to a one-point space.

REMARK 10.22. The above definitions imply immediately that any space admitting a deformation retraction to a one-point subset (as in Example 10.12) is contractible. The converse is not quite true. Indeed, suppose $\{x\}$ is a one-point space and $f : X \to \{x\}$ is a homotopy equivalence with homotopy inverse $g : \{x\} \to X$ and a homotopy $H : I \times X \to X$ from Id_X to $g \circ f$. (We do not need to discuss any homotopy of $f \circ g$ since there is only one map $\{x\} \to \{x\}$.) Then if $p := g(x) \in X$, $F : X \to \{p\}$ denotes the constant map at p and $i : \{p\} \hookrightarrow X$ is the inclusion, we have $F \circ i = \mathrm{Id}_{\{p\}}$, and H is a homotopy from Id_X to $i \circ F$. Unfortunately, the definition of homotopy equivalence does not guarantee that this homotopy will satisfy $H(t,p) = p$ for all $t \in I$, so H might not be a deformation retraction in the strict sense of Definition 10.11. It turns out that this distinction matters, but only for fairly strange spaces: see [**Hat02**, p. 18, Exercise 6] for an example of a space that is contractible but does not admit a deformation retraction to any point.

We can now state the main theorem of this lecture.

THEOREM 10.23. *If $f : X \to Y$ is a homotopy equivalence with $f(p) = q$, then the induced homomorphism $f_* : \pi_1(X,p) \to \pi_1(Y,q)$ is an isomorphism.*

Since a one-point space contains only one path and therefore has trivial fundamental group, this implies:

COROLLARY 10.24. *For every contractible space X, $\pi_1(X) = 0$.* □

PROOF OF THEOREM 10.23. Here is a preliminary remark: if you're only half paying attention, then you might reasonably think this theorem follows immediately from Theorem 8.11. Indeed, we stated in that theorem that the homomorphism $f_* : \pi_1(X, p) \to \pi_1(Y, q)$ depends only on the pointed homotopy class of f, and the same is of course true of the compositions $g \circ f$ and $f \circ g$, which ought to make $g_* \circ f_*$ and $f_* \circ g_*$ both the identity if $g \circ f$ and $f \circ g$ are homotopic to the identity. The problem however is that we are not paying attention to the base point: the definition of homotopy equivalence never mentions any base point and says "homotopy" rather than "pointed homotopy," while in Theorem 8.11, maps and homotopies are always required to preserve base points. In particular, if $f(p) = q$ and $g : Y \to X$ is a homotopy inverse of f, then there is no reason to expect $g(q) = p$, in which case $g_* : \pi_1(Y, q) \to \pi_1(X, g(q))$ cannot be an inverse of $f_* : \pi_1(X, p) \to \pi_1(Y, q)$, as its target is not even the same group as the domain of f_*. The main content of the following proof is an argument to cope with this annoying detail.

With that out of the way, assume $f : X \to Y$ is a map with homotopy inverse $g : Y \to X$, satisfying $f(p) = q$ and $g(q) = r$, so we have a sequence of pointed maps

$$(X, p) \xrightarrow{f} (Y, q) \xrightarrow{g} (X, r)$$

and induced homomorphisms

$$\pi_1(X, p) \xrightarrow{f_*} \pi_1(Y, q) \xrightarrow{g_*} \pi_1(X, r). \tag{10.2}$$

By assumption there exists a homotopy $H : I \times X \to X$, which we shall write as a 1-parameter family of maps

$$h_s := H(s, \cdot) : X \to X \quad \text{for} \quad s \in I,$$

satisfying $h_0 = \mathrm{Id}_X$ and $h_1 = g \circ f$. We can therefore define a path $p \overset{\gamma}{\rightsquigarrow} r$ by

$$\gamma(t) := h_t(p),$$

and by Theorem 9.1, this gives rise to an isomorphism

$$\Phi_\gamma : \pi_1(X, r) \to \pi_1(X, p) : [\alpha] \mapsto [\gamma \cdot \alpha \cdot \gamma^{-1}].$$

We claim that the diagram

$$\begin{tikzcd} \pi_1(X,p) \arrow[r, "f_*"] \arrow[dr, "\Phi_\gamma^{-1}"'] & \pi_1(Y,q) \arrow[d, "g_*"] \\ & \pi_1(X,r) \end{tikzcd}$$

commutes, or equivalently, $\Phi_\gamma \circ g_* \circ f_*$ is the identity map on $\pi_1(X, p)$. Given a loop $p \overset{\alpha}{\rightsquigarrow} p$, the element $\Phi_\gamma \circ g_* \circ f_*[\alpha] = \Phi_\gamma \circ (g \circ f)_*[\alpha]$ is represented by $\gamma \cdot (g \circ f \circ \alpha) \cdot \gamma^{-1}$, so we need to show that the latter is homotopic with fixed end points to α. A precise formula for such a homotopy is provided by the following 1-parameter family of loops: for $s \in I$, let

$$\alpha_s := \gamma_s \cdot (h_s \circ \alpha) \cdot \gamma_s^{-1},$$

where $p \overset{\gamma_s}{\rightsquigarrow} \gamma(s)$ denotes the path $\gamma_s(t) := \gamma(st)$. (For a visualization of what this homotopy is actually doing, I recommend the picture on page 37 of [Hat02].) This proves the claim, and since Φ_γ is an isomorphism, it implies that $g_* \circ f_* = \Phi_\gamma^{-1}$ is also an isomorphism, from which we deduce that f_* is injective and g_* is surjective.

The preceding argument was based on the assumption that $g \circ f : X \to X$ is homotopic to the identity. We have not yet used the assumption that $f \circ g : Y \to Y$ is also homotopic to the identity, but we can use it now to carry out the same argument again with the roles of f and g reversed. The conclusion is that $f_* \circ g_*$ is also an isomorphism, implying g_* is injective and f_* is surjective. We conclude that f_* and g_* are in fact both isomorphisms. $\square$

EXAMPLE 10.25. Here are some examples of contractible spaces, which therefore have isomorphic (trivial) fundamental groups even though they are not all homeomorphic: $\mathbb{R}^n$, $\mathbb{D}^n$ (not homeomorphic to $\mathbb{R}^n$ since it is compact), any convex subset or star-shaped domain in $\mathbb{R}^n$ as in Example 10.12. A quite different type of example comes from *graph theory*: a **graph** is a combinatorial object consisting of a set V (called the **vertices**) and a set E whose elements (the **edges**) are unordered pairs of vertices. A graph is typically represented by depicting the vertices as points and the edges $\{x, y\} \in E$ as curves connecting the corresponding vertices x and y to each other. One can thus naturally view a graph as a topological space in which each vertex is a point and each edge is a subset homeomorphic to $[0, 1]$ (possibly with its end points identified if its two vertices are the same one). A graph is called a **tree** if there is exactly one path (up to parametrization) connecting any two of its vertices. It is not hard to show that any finite graph with this property is a contractible space: pick your favorite vertex $v \in V$, draw the unique path from v to every other vertex, then define a deformation retraction to v by pulling everything back along these paths.

EXAMPLE 10.26. Viewing S^1 as the unit circle in $\mathbb{C}$, associate to each $z \in \mathbb{C}$ the loop $\gamma_z : S^1 \hookrightarrow \mathbb{C}\backslash\{z\} : e^{i\theta} \mapsto z + e^{i\theta}$. Since these are pointed maps $(S^1, 1) \to (\mathbb{C}\backslash\{z\}, z + 1)$, they represent elements $[\gamma_z] \in \pi_1(\mathbb{C}\backslash\{z\}, z+1)$. We claim in fact that this group is isomorphic to $\mathbb{Z}$, and that $[\gamma_z]$ generates it. The proof is mainly the observation that $\gamma_z(S^1)$ is a deformation retract of $\mathbb{C}\backslash\{z\}$, by a construction analogous to Example 10.14, hence γ_z is a homotopy equivalence and therefore induces an isomorphism $\pi_1(S^1, 1) \to \pi_1(\mathbb{C}\backslash\{z\}, z + 1)$. Since the identity map $(S^1, 1) \to (S^1, 1)$ represents a generator of $\pi_1(S^1, 1)$, composing this with γ_z now represents a generator of $\pi_1(\mathbb{C}\backslash\{z\}, z + 1)$ as claimed.

EXERCISE 10.27. For a point $z \in \mathbb{C}$ and a continuous map $\gamma : [0, 1] \to \mathbb{C}\backslash\{z\}$ with $\gamma(0) = \gamma(1)$, one defines the **winding number** of γ about z as

$$\text{wind}(\gamma; z) = \theta(1) - \theta(0) \in \mathbb{Z}$$

where $\theta : [0, 1] \to \mathbb{R}$ is any choice of continuous function such that

$$\gamma(t) = z + r(t)e^{2\pi i\theta(t)}$$

for some function $r : [0, 1] \to (0, \infty)$. Notice that since $\gamma(t) \neq z$ for all t, the function $r(t)$ is uniquely determined, and requiring $\theta(t)$ to be continuous makes it unique up to the addition of a constant integer, hence $\theta(1) - \theta(0)$ depends only on the path γ and not on any additional choices. One of the fundamental facts about winding numbers is their important role in the computation of $\pi_1(S^1)$: as we saw in Example 9.12, viewing S^1 as $\{z \in \mathbb{C} \mid |z| = 1\}$, the map

$$\pi_1(S^1, 1) \to \mathbb{Z} : [\gamma] \mapsto \text{wind}(\gamma; 0)$$

is an isomorphism to the abelian group $(\mathbb{Z}, +)$. Assume in the following that $\Omega \subset \mathbb{C}$ is an open set and $f : \Omega \to \mathbb{C}$ is a continuous function.

(a) Suppose $f(z) = w$ and $w \notin f(\mathcal{U}\backslash\{z\})$ for some neighborhood $\mathcal{U} \subset \Omega$ of z. This implies that the loop $f \circ \gamma_\epsilon$ for $\gamma_\epsilon : [0, 1] \to \Omega : t \mapsto z + \epsilon e^{2\pi it}$ has image in $\mathbb{C}\backslash\{w\}$ for all $\epsilon > 0$ sufficiently small, hence $\text{wind}(f \circ \gamma_\epsilon; w)$ is well defined. Show that for some $\epsilon_0 > 0$, $\text{wind}(f \circ \gamma_\epsilon; w)$ does not depend on ϵ as long as $0 < \epsilon \leqslant \epsilon_0$.

(b) Show that if the ball $B_r(z_0)$ of radius $r > 0$ about $z_0 \in \Omega$ has its closure contained in Ω, and the loop $\gamma(t) = z_0 + re^{2\pi it}$ satisfies $\mathrm{wind}(f \circ \gamma; w) \neq 0$ for some $w \in \mathbb{C}$, then there exists $z \in B_r(z_0)$ with $f(z) = w$.
Hint: Recall that if we regard elements of $\pi_1(X, p)$ as pointed homotopy classes of maps $S^1 \to X$, then such a map represents the identity in $\pi_1(X, p)$ if and only if it admits a continuous extension to a map $\mathbb{D}^2 \to X$. Define X in the present case to be $\mathbb{C}\backslash\{w\}$.

(c) Prove the Fundamental Theorem of Algebra: every nonconstant complex polynomial has a root.
Hint: Consider loops $\gamma(t) = Re^{2\pi it}$ with $R > 0$ large.

(d) We call $z_0 \in \Omega$ an **isolated zero** of $f : \Omega \to \mathbb{C}$ if $f(z_0) = 0$ but $0 \notin f(\mathcal{U}\backslash\{z_0\})$ for some neighborhood $\mathcal{U} \subset \Omega$ of z_0. Let us say that such a zero has **order** $k \in \mathbb{Z}$ if $\mathrm{wind}(f \circ \gamma_\epsilon; 0) = k$ for $\gamma_\epsilon(t) = z_0 + \epsilon e^{2\pi it}$ and $\epsilon > 0$ small (recall from part (a) that this does not depend on the choice of ϵ if it is small enough). Show that if $k \neq 0$, then for any neighborhood $\mathcal{U} \subset \Omega$ of z_0, there exists $\delta > 0$ such that every continuous function $g : \Omega \to \mathbb{C}$ satisfying $|f - g| < \delta$ everywhere has a zero somewhere in $\mathcal{U}$.

(e) Find an example of the situation in part (d) with $k = 0$ such that f admits arbitrarily close perturbations g that have no zeroes in some fixed neighborhood of $\mathcal{U}$.
Hint: Write f as a continuous function of x and y where $x + iy \in \Omega$. You will not be able to find an example for which f is holomorphic—they do not exist!

General advice: Throughout this problem, it is important to remember that $\mathbb{C}\backslash\{w\}$ is homotopy equivalent to S^1 for every $w \in \mathbb{C}$. Thus all questions about $\pi_1(\mathbb{C}\backslash\{w\})$ can be reduced to questions about $\pi_1(S^1)$.

11. The easy part of van Kampen's theorem (May 25, 2023)

The main question of this lecture is the following: If X is the union of two subsets $A \cup B$ and we know both $\pi_1(A)$ and $\pi_1(B)$, what can we say about $\pi_1(X)$?

EXAMPLE 11.1. The sphere S^n can be viewed as the union of two subsets A and B that are both homeomorphic to $\mathbb{D}^n$, e.g. when $n = 2$, we would take the northern and southern "hemispheres" of the globe. Since $\mathbb{D}^n$ is contractible, $\pi_1(A) = \pi_1(B) = 0$. We will see below that this is almost enough information to compute $\pi_1(S^n)$.

The next lemma is the "easy" first half of an important result about fundamental groups known as the *Seifert-van Kampen theorem*, or often simply *van Kampen's theorem.* The much more powerful "hard" part of the theorem will be dealt with in the two subsequent lectures, though the easy part already has several impressive applications. We will state it here in somewhat greater generality than is needed for most applications: on first reading, you are free to replace the arbitrary open covering $X = \bigcup_{\alpha \in J} A_\alpha$ with a covering by *two* open subsets $X = A \cup B$, which will be the situation in all of the examples below.

LEMMA 11.2. *Suppose $X = \bigcup_{\alpha \in J} A_\alpha$ for a collection of open subsets $\{A_\alpha \subset X\}_{\alpha \in J}$ satisfying the following conditions:*

(1) A_α is path-connected for every $\alpha \in J$;
(2) $A_\alpha \cap A_\beta$ is path-connected for every pair $\alpha, \beta \in J$;
(3) $\bigcap_{\alpha \in J} A_\alpha \neq \varnothing$.

Let $A_\alpha \overset{i_\alpha}{\hookrightarrow} X$ denote the natural inclusion maps. Then for any base point $p \in \bigcap_{\alpha \in J} A_\alpha$, $\pi_1(X, p)$ is generated by the subgroups

$$(i_\alpha)_* (\pi_1(A_\alpha, p)) \subset \pi_1(X, p),$$

i.e. every element of $\pi_1(X, p)$ *is a product of elements of the form* $(i_\alpha)_*[\gamma]$ *for some* $\alpha \in J$ *and* $[\gamma] \in \pi_1(A_\alpha, p)$.

Before proving the lemma, let's look at several more examples, starting with a rehash of Example 11.1 above.

EXAMPLE 11.3. Denote points in the unit sphere S^n by $(\mathbf{x}, z) \in \mathbb{R}^n \times \mathbb{R}$ such that $|\mathbf{x}|^2 + z^2 = 1$, and define the open subsets

$$A := \{z > -\epsilon\} \subset S^n, \qquad B := \{z < \epsilon\} \subset S^n$$

for some $\epsilon > 0$ small. Then $A \cong B \cong \mathbb{R}^n$, so both have trivial fundamental group. Moreover, $A \cap B \cong S^{n-1} \times (-\epsilon, \epsilon)$ is path-connected if $n \geqslant 2$. (Note that this is not true if $n = 1$: the 0-sphere S^0 is just the set of two points $\{1, -1\} \subset \mathbb{R}$, so it is not path-connected.) The lemma therefore implies that for any $p \in A \cap B$, $\pi_1(S^n, p)$ is generated by images of homomorphisms into $\pi_1(S^n, p)$ from the groups $\pi_1(A, p)$ and $\pi_1(B, p)$, both of which are trivial, therefore $\pi_1(S^n, p)$ is trivial.

We just proved:

COROLLARY 11.4. *For all* $n \geqslant 2$, S^n *is simply connected.* □

Here is an easy application:

THEOREM 11.5. *For every* $n \geqslant 3$, $\mathbb{R}^2$ *is not homeomorphic to* $\mathbb{R}^n$.

PROOF. The complement of one point in $\mathbb{R}^n$ is homotopy eqivalent to S^{n-1}, thus $\pi_1(\mathbb{R}^n \backslash \{\mathrm{pt}\}) \cong \pi_1(S^{n-1}) = 0$ if $n \geqslant 3$, while $\pi_1(\mathbb{R}^2 \backslash \{\mathrm{pt}\}) \cong \pi_1(S^1) \cong \mathbb{Z}$. It follows that $\mathbb{R}^2 \backslash \{\mathrm{pt}\}$ and $\mathbb{R}^n \backslash \{\mathrm{pt}\}$ for $n \geqslant 3$ are not homeomorphic, hence neither are $\mathbb{R}^2$ and $\mathbb{R}^n$. □

A wider class of examples comes from the following general construction known as *gluing* of spaces. Assume X, Y and A are spaces and we have inclusions[9]

$$i_X : A \hookrightarrow X, \qquad i_Y : A \hookrightarrow Y.$$

We then define the space

$$X \cup_A Y := (X \amalg Y)/\sim$$

where the equivalence relation identifies $i_X(a) \in X$ with $i_Y(a) \in Y$ for every $a \in A$. As usual in such constructions, we assign to $X \amalg Y$ the disjoint union topology and then give $X \cup_A Y$ the quotient topology. We say that $X \cup_A Y$ is the space obtained by **gluing** X **to** Y **along** A. Note that we can regard X and Y both as subspaces of $X \cup_A Y$, and their intersection is a subspace homeomorphic to A. The wedge sum of two spaces (see Example 10.3) is the special case of this construction where A is a single point. (The notation is slightly non-ideal since $X \cup_A Y$ depends on the inclusions of A into X and Y, not just on the three spaces themselves, but in most interesting examples the inclusions are obvious, so the notation is easy to interpret.)

EXAMPLE 11.6. If $X = Y = \mathbb{D}^n$ and $A = S^{n-1}$ is included in both as the boundary $\partial\mathbb{D}^n$, then the descriptions of S^n in Examples 11.1 and 11.3 translates into

$$\mathbb{D}^n \cup_{S^{n-1}} \mathbb{D}^n \cong S^n.$$

[9]The technical meaning of the word **inclusion** in this context is a map $A \hookrightarrow X$ which is injective and is a homeomorphism onto its image (with the subspace topology). Such a map is also sometimes called a **topological embedding**.

EXAMPLE 11.7. In Example 1.2 we gave a description of $\mathbb{RP}^2$ as the space obtained by gluing a disk $\mathbb{D}^2$ to a Möbius strip

$$\mathbb{M} := \left\{ (e^{i\theta}, t\cos(\theta/2), t\sin(\theta/2)) \in S^1 \times \mathbb{R}^2 \;\middle|\; e^{i\theta} \in S^1,\ t \in [-1,1] \right\}$$

along their boundaries, which are both homeomorphic to S^1. Choose a particular inclusion of S^1 as the boundary of $\mathbb{M}$, e.g.

$$S^1 \hookrightarrow \mathbb{M} : e^{i\theta} \mapsto (e^{2i\theta}, \cos(\theta), \sin(\theta)).$$

Then our picture of $\mathbb{RP}^2$ can be expressed succinctly as

$$\mathbb{RP}^2 \cong \mathbb{D}^2 \cup_{S^1} \mathbb{M}.$$

Lemma 11.2 can now be applied to this as follows. There is an obvious deformation retraction of $\mathbb{M}$ to the "central" circle $S^1 \times \{0\} \subset \mathbb{M}$, defined via the homotopy

$$H : I \times \mathbb{M} \to \mathbb{M} : (s, (e^{i\theta}, t\cos(\theta/2), t\sin(\theta/2))) \mapsto (e^{i\theta}, st\cos(\theta/2), st\sin(\theta/2)),$$

thus $\mathbb{M} \underset{h.e.}{\simeq} S^1$. The gluing construction allows us to view both $\mathbb{D}^2$ and $\mathbb{M}$ as subsets of $\mathbb{RP}^2$, but they are not *open* subsets as required by the lemma. This can easily be fixed by slightly expanding both of them. Concretely, by adding a neighborhood of $\partial\mathbb{M}$ in $\mathbb{M}$ to $\mathbb{D}^2$, we obtain an open neighborhood $A \subset \mathbb{RP}^2$ of $\mathbb{D}^2$ that is homeomorphic to an open disk, and similarly, adding a neighborhood of $\partial\mathbb{D}^2$ in $\mathbb{D}^2$ to $\mathbb{M}$ gives an open neighborhood $B \subset \mathbb{RP}^2$ of $\mathbb{M}$ that admits a deformation retraction to $\mathbb{M}$ and thus also to the central circle $S^1 \times \{0\} \subset \mathbb{M}$. We now have

$$\pi_1(A) \cong \pi_1(\mathring{\mathbb{D}}^2) = 0 \quad \text{and} \quad \pi_1(B) \cong \pi_1(\mathbb{M}) \cong \pi_1(S^1) \cong \mathbb{Z},$$

and notice also that A and B are both path connected, and so is $A \cap B$ since we can arrange for the latter to be homeomorphic to $S^1 \times (-1,1)$, i.e. it is the union of an annular neighborhood of $\partial\mathbb{D}^2$ in $\mathbb{D}^2$ with another annular neighborhood of $\partial\mathbb{M}$ in $\mathbb{M}$. The lemma thus implies that for any $p \in A \cap B$, $\pi_1(\mathbb{RP}^2, p)$ is generated by the element $i^B_*[\gamma] \in \pi_1(\mathbb{RP}^2, p)$, where $i^B : B \hookrightarrow \mathbb{RP}^2$ is the inclusion and $\gamma : (S^1, 1) \to (B, p)$ is any loop such that $[\gamma]$ generates $\pi_1(B, p) \cong \mathbb{Z}$. In light of the deformation retraction to the central circle, the inclusion of that circle into B induces an isomorphism of fundamental groups, thus we can take γ to be the obvious inclusion of S^1 into B as the central circle:

$$\begin{aligned} \gamma : S^1 &\xrightarrow{\cong} S^1 \times \{0\} \subset \mathbb{M} \subset \mathbb{RP}^2, \\ e^{i\theta} &\mapsto (e^{i\theta}, 0). \end{aligned} \tag{11.1}$$

The conclusion is that if we regard γ in this way as a loop in $\mathbb{RP}^2$, then $[\gamma]$ generates $\pi_1(\mathbb{RP}^2, p)$. The loop γ is not hard to visualize if you translate from our picture of $\mathbb{RP}^2$ as $\mathbb{D}^2 \cup_{S^1} \mathbb{M}$ back to the usual definition of $\mathbb{RP}^2$ as a quotient of S^2 (see Example 1.2): in the latter picture you can realize γ as a path along the equator of S^2 that goes exactly halfway around. Note that this is not a loop in S^2, but it becomes a loop when you project it to $\mathbb{RP}^2$ since its starting and end point are antipodal.

A word of caution is in order: we have not yet actually computed $\pi_1(\mathbb{RP}^2)$, we have only shown that every element in $\pi_1(\mathbb{RP}^2)$ is a power of a single element $[\gamma]$. It is still possible that $\pi_1(\mathbb{RP}^2)$ is trivial because γ is contractible—this will turn out not to be the case, but we are not in a position to prove it just yet. We can say one more thing, however: $[\gamma]^2$ is the identity element in $\pi_1(\mathbb{RP}^2, p)$. Indeed, $[\gamma]^2$ is represented by the concatenation of γ with itself, which can also be realized as the projection through $S^2 \xrightarrow{\pi} \mathbb{RP}^2$ of a path that goes *all the way* around the equator in S^2, i.e. it is the concatenation of two paths that go halfway around. But if $\alpha : S^1 \to S^2$ parametrizes this loop around the equator, then there is obviously an extension of α to a map $u : \mathbb{D}^2 \to S^2$ satisfying $u|_{\partial\mathbb{D}^2} = \alpha$, namely the inclusion of either the northern or southern hemisphere of S^2.

The map $\pi \circ u : \mathbb{D}^2 \to \mathbb{RP}^2$ is then an extension over the disk of our loop representing $[\gamma]^2$, which proves via Theorem 9.4 that $[\gamma]^2$ is trivial. This proves that $\pi_1(\mathbb{RP}^2)$ is either the trivial group or is isomorphic to $\mathbb{Z}_2$; we will see that it is the latter when we prove that the generator $[\gamma]$ is nontrivial.

Here is another pair of general constructions that produce many more examples.

DEFINITION 11.8. Given a space X, the **cone** (*Kegel*) of X is the space

$$CX := (X \times I)/(X \times \{1\}).$$

The single point in CX represented by $(x, 1)$ for every $x \in X$ is sometimes called the "summit" or "node" of the cone.

EXERCISE 11.9. Show that CS^{n-1} is homeomorphic to $\mathbb{D}^n$.

LEMMA 11.10. *For every space X, the cone CX is contractible.*

PROOF. There is an obvious deformation retraction of $X \times I$ to $X \times \{1\}$ defined by pushing every $(x,t) \in X \times I$ upward in the t-coordinate. Writing down this same deformation retraction on the quotient $(X \times I)/(X \times \{1\})$, the result is that everything gets pushed to a single point, the summit of the cone. □

DEFINITION 11.11. Given a space X, the **suspension** (*Einhängung*) of X is the space

$$SX := C_+X \cup_{X \times \{0\}} C_-X,$$

where $C_+X := CX$ as above, and C_-X is the "reversed" cone $(X \times [-1, 0])/(X \times \{-1\})$. Equivalently, the suspension can be written as

$$SX = (X \times [-1, 1])/\sim$$

where $(x, 1) \sim (y, 1)$ and $(x, -1) \sim (y, -1)$ for every $x, y \in X$.

EXERCISE 11.12. Show that $SS^{n-1} \cong S^n$.

We can now generalize the result that $\pi_1(S^n) = 0$ for $n \geqslant 2$ as follows.

THEOREM 11.13. *If X is path-connected, then its suspension SX is simply connected.*

PROOF. We define $A, B \subset SX$ to be open neighborhoods of C_+X and C_-X respectively, e.g.

$$A := (X \times (-\epsilon, 1])/(X \times \{1\}), \qquad B := (X \times [-1, \epsilon))/(X \times \{-1\})$$

for any $\epsilon \in (0, 1)$. The subspaces are both contractible for the same reason that C_+X and C_-X are: one can define deformation retractions to a point by pushing upward in A and downward in B. Moreover, $A \cap B = X \times (-\epsilon, \epsilon)$ is path-connected if and only if X is path-connected, and in that case, Lemma 11.2 implies that $\pi_1(SX)$ is generated by the images of homomorphisms from $\pi_1(A)$ and $\pi_1(B)$, both of which are trivial, therefore $\pi_1(SX)$ is trivial. □

Let us finally prove the lemma.

PROOF OF LEMMA 11.2. We assume $X = \bigcup_{\alpha \in J} A_\alpha$ and $p \in \bigcap_{\alpha \in J} A_\alpha$, where the sets $A_\alpha \subset X$ are open and path-connected, and $A_\alpha \cap A_\beta$ is also path-connected for every pair $\alpha, \beta \in J$. What we need to show is that every loop $p \overset{\gamma}{\rightsquigarrow} p$ in X is homotopic with fixed end points to a concatenation of finitely many loops based at p that are each contained in one of the subsets A_α. To start with, observe that since $\gamma : I \to X$ is continuous, $I_\alpha := \gamma^{-1}(A_\alpha)$ is an open subset of I for every α, and is therefore a union of open subintervals of I.[10] The union of all these open subintervals for all

[10]Remember that since sets like $[0, \epsilon) \subset I$ that include an end point are open subsets of I, they are included in the term "open subinterval of I".

$\alpha \in J$ thus forms an open covering of I, which has a finite subcovering since I is compact, giving rise to a finite collection of open subintervals

$$I = I_1 \cup \ldots \cup I_N$$

such that for each $j = 1, \ldots, N$, $\gamma(I_j) \subset A_{\alpha_j}$ for some $\alpha_j \in J$. After relabeling the α_j's if necessary, we can then find a finite increasing sequence

$$0 =: t_0 < t_1 < \ldots < t_{N-1} < t_N := 1$$

such that $\gamma([t_{j-1}, t_j]) \subset A_{\alpha_j}$ for each $j = 1, \ldots, N$. In particular, for $j = 1, \ldots, N-1$, each $\gamma(t_j)$ lies in both A_{α_j} and $A_{\alpha_{j+1}}$. The intersection of these two sets is path-connected by assumption, so choose a path β_j in $A_{\alpha_j} \cap A_{\alpha_{j+1}}$ from $\gamma(t_j)$ to the base point p. Then if we write $\gamma_j := \gamma|_{[t_{j-1}, t_j]}$ and reparametrize each of these paths to define them on the usual interval I, we have

$$\gamma = \gamma_1 \cdot \ldots \cdot \gamma_N \underset{h+}{\sim} \gamma_1 \cdot \beta_1 \cdot \beta_1^{-1} \cdot \gamma_2 \cdot \beta_2 \cdot \beta_2^{-1} \cdot \ldots \cdot \beta_{N-2} \cdot \beta_{N-2}^{-1} \cdot \gamma_{N-1} \cdot \beta_{N-1} \cdot \beta_{N-1}^{-1} \cdot \gamma_N.$$

The latter is the concatenation we were looking for since $\gamma_1 \cdot \beta_1$ is a loop from p to itself in A_{α_1}, $\beta_1^{-1} \cdot \gamma_2 \cdot \beta_2$ is a loop from p to itself in A_{α_2}, and so forth up to $\beta_{N-2}^{-1} \cdot \gamma_{N-1} \cdot \beta_{N-1}$ in $A_{\alpha_{N-1}}$ and $\beta_{N-1}^{-1} \cdot \gamma_N$ in A_{α_N}. □

To conclude this lecture, we would like to restate Lemma 11.2 in more precise terms. This requires a few notions from combinatorial group theory.

Definition 11.14. Suppose $\{G_\alpha\}_{\alpha \in J}$ is a collection of groups, with the identity element in each denoted by $e_\alpha \in G_\alpha$. For any integer $N \geqslant 0$, an ordered set $b_1 b_2 \ldots b_N$ together with a corresponding ordered set $\alpha_1, \alpha_2, \ldots, \alpha_N \in J$ is called a **word in** $\{G_\alpha\}_{\alpha \in J}$ if $b_i \in G_{\alpha_i}$ for each $i = 1, \ldots, N$. Informally, we call the elements of the sequence *letters*, and denote the word by $b_1 \ldots b_N$ even though, strictly speaking, the set of indices $\alpha_1, \ldots, \alpha_N \in J$ is also part of the data defining the word.[11] Note that this definition includes the so-called *empty word*, with $N = 0$, i.e. the word with no letters. A word $a_1 \ldots a_N$ is called a **reduced word** if:

- none of the letters b_i are the identity element $e_{\alpha_i} \in G_{\alpha_i}$ in the corresponding group, and
- no two adjacent letters b_i and b_{i+1} satisfy $\alpha_i = \alpha_{i+1}$, i.e. the groups that appear in adjacent positions are distinct.

Note that the empty word trivially satisfies both conditions, thus it is a reduced word.

There is an obvious map called **reduction** from the set of all words to the set of all reduced words: it acts on a given word $b_1 \ldots b_N$ by replacing all adjacent pairs $b_i b_{i+1}$ with their product in G_α whenever $\alpha_i = \alpha_{i+1} = \alpha$, and removing all e_α's.

Definition 11.15. The **free product** (*freies Produkt*) $\circledast_{\alpha \in J} G_\alpha$ of a collection of groups $\{G_\alpha\}_{\alpha \in J}$ is defined as the set of all reduced words in $\{G_\alpha\}_{\alpha \in J}$. The product of two reduced words $w = b_1 \ldots b_N$ and $w' = b_1' \ldots b_{N'}'$ in this group is defined to be the reduction of the concatenated word $ww' = b_1 \ldots b_N b_1' \ldots b_{N'}'$. The identity element is the empty word, and will be denoted by

$$e \in \underset{\alpha \in J}{\circledast} G_\alpha.$$

We will typically deal with collections of only finitely many groups $G_1, \ldots, G_N$, in which case the free product is usually denoted by

$$G_1 * \ldots * G_N.$$

[11]This is important to remember in case some G_α and G_β contain common elements for $\alpha \neq \beta$, e.g. if they are both subgroups of a single larger group. If not, then this detail is safe to ignore and the notation $b_1 \ldots b_N$ for a word is completely unambiguous.

In general, this is an enormous group, e.g. it is always infinite if there are at least two nontrivial groups in the collection, no matter how small those groups are. It is also always nonabelian in those cases. Let us see some examples.

EXAMPLE 11.16. Consider two copies of the same group $G = H = \mathbb{Z}_2$, with the unique nontrivial elements of G and H denoted by $a \in G$ and $b \in H$. Then $G * H$ consists of all possible reduced words built out of these two letters, plus the empty word e, so

$$\mathbb{Z}_2 * \mathbb{Z}_2 \cong G * H = \{e, a, b, ab, ba, aba, bab, abab, baba, \dots\}.$$

For an example of how multiplication in $\mathbb{Z}_2 * \mathbb{Z}_2$ works, the product of aba and ab is a, i.e. this is the result of reducing the unreduced word $abaab$ since aa and bb are both identity elements.

EXAMPLE 11.17. Let $G = \mathbb{Z}$ with a generator denoted by $a \in G$, and $H = \mathbb{Z}_2$ with nontrivial element b. If we write G as a multiplicative group so that its elements are all of the form a^p for $p \in \mathbb{Z}$, then

$$\mathbb{Z} * \mathbb{Z}_2 \cong G * H = \left\{e, a^p, b, a^p b, ba^p, a^p ba^q, ba^p ba^q, a^p ba^q ba^r, \dots \;\middle|\; p, q, r, \dots \in \mathbb{Z}\right\}.$$

For an example of a product, $a^p ba^r$ times $a^{-1}b$ gives $a^p ba^{r-1}b$.

With this terminology understood, here is what we actually proved when we proved Lemma 11.2.

LEMMA 11.18. *Given $X = \bigcup_{\alpha\in J} A_\alpha$ and $p \in \bigcap_{\alpha\in J} A_\alpha$ as in Lemma 11.2, there exists a natural group homomorphism*

$$\mathop{\ast}_{\alpha\in J} \pi_1(A_\alpha, p) \xrightarrow{\Phi} \pi_1(X, p)$$

sending each reduced word $[\gamma_1]\dots[\gamma_N] \in \ast_{\alpha\in J}\pi_1(A_\alpha, p)$ with $[\gamma_i] \in \pi_1(A_{\alpha_i}, p)$ to the concatenation $[\gamma_1 \cdot \dots \cdot \gamma_N] \in \pi_1(X, p)$, and Φ is surjective. □

The existence of the homomorphism Φ is an easy and purely algebraic fact, which we'll expand on a bit in the next lecture. The truly nontrivial statement here is that Φ is surjective. If we can now identify the kernel of Φ, then Φ descends to an isomorphism from the quotient of the free product by $\ker\Phi$ to $\pi_1(X, p)$, and we will thus have a formula for $\pi_1(X, p)$. Identifying the kernel and then using the resulting formula in applications will be our main topic for the next two lectures.

12. Normal subgroups, generators and relations (May 30, 2023)

Before stating the general version of the Seifert-van Kampen theorem, we need to collect a few more useful algebraic facts about groups and the free product. Recall from the previous lecture that the free product $\ast_{\alpha\in J} G_\alpha$ of an arbitrary collection of groups $\{G_\alpha\}_{\alpha\in J}$ is defined to consist of all so-called *reduced words* $g_1 \dots g_N$ in which each "letter" g_i is an element of one of the groups G_{α_i}, and the choice of $\alpha_i \in J$ such that $g_i \in G_{\alpha_i}$ for each $i = 1, \dots, N$ is considered part of the data defining the word.[12] The word "reduced" means that the sequence of letters in the word cannot be simplified by computing products in any of the individual groups, hence no consecutive letters $g_i g_{i+1}$ with $\alpha_i = \alpha_{i+1} =: \alpha$ appear—if such a pair appeared then it could be replaced by a single letter formed from the product $g_i g_{i+1} \in G_\alpha$—and similarly, none of the letters is the identity element in any of the groups. Products in $\ast_{\alpha\in J} G_\alpha$ are formed by concatenating words and then

[12]This latter detail is unimportant if the groups G_α are all disjoint sets in the first place, but if any of them have elements in common, e.g. if some G_α and G_β for $\alpha \neq \beta$ are copies of the same group, then we regard them as *separate* copies and always keep track of which letter belongs to which copy. The idea is somewhat analogous to constructing the disjoint union $\coprod_{\alpha\in J} X_\alpha$ of sets, in which X_β and X_γ for $\beta \neq \gamma$ always become disjoint subsets of $\coprod_{\alpha\in J} X_\alpha$, even if they are originally defined as the same set, e.g. $\mathbb{R} \amalg \mathbb{R}$ is by definition two disjoint copies of $\mathbb{R}$, which is different from the ordinary union $\mathbb{R} \cup \mathbb{R} = \mathbb{R}$.

reducing them if necessary, so for example, if G and H are two groups containing elements $g \in G$ and $h, k \in H$, then the product of the reduced words $gh \in G * H$ and $h^{-1}k \in G * H$ is

$$(gh)(h^{-1}k) = gk \in G * H,$$

since the concatenated word $ghh^{-1}k$ can be reduced by replacing hh^{-1} with the identity element $e \in H$ and then removing e from the word. The identity element in $\ast_{\alpha \in J} G_\alpha$ itself is the so-called "empty" word, with zero letters, which we will usually denote by e; there should be no danger of confusing this with the identity elements of the individual groups G_α, since they never appear in reduced words.

The following result is easy to prove directly from the definitions.

Proposition 12.1. *Assume $\{G_\alpha\}_{\alpha \in J}$ is a collection of groups. Then:*

(1) *For each $\alpha \in J$, the free product $\ast_{\beta \in J} G_\beta$ contains a distinguished subgroup isomorphic to G_α: it consists of the empty word plus all reduced words of exactly one letter which is in G_α.*
(2) *If we regard each G_α as a subgroup of $\ast_{\gamma \in J} G_\gamma$ as described above, then for every $\alpha, \beta \in J$ with $\alpha \neq \beta$, the intersection $G_\alpha \cap G_\beta$ in $\ast_{\gamma \in J} G_\gamma$ consists only of the identity element e (i.e. the empty word), and any two nontrivial elements $g \in G_\alpha$ and $h \in G_\beta$ satisfy $gh \neq hg$ in $\ast_{\gamma \in J} G_\gamma$.*
(3) *For any group H with a collection of homomorphisms $\{\Phi_\alpha : G_\alpha \to H\}_{\alpha \in J}$, there exists a unique homomorphism*

$$\Phi : \mathop{\ast}_{\alpha \in J} G_\alpha \to H$$

whose restriction to each of the subgroups $G_\alpha \subset \ast_{\beta \in J} G_\beta$ is Φ_α.

The third item in this list deserves brief comment: the homomorphism $\Phi : \ast_{\alpha \in J} G_\alpha \to H$ exists and is unique because every element of $\ast_{\alpha \in J} G_\alpha$ is uniquely expressible as a reduced word $g_1 \dots g_N$ with $g_i \in G_{\alpha_i}$ for some specified $\alpha_1, \dots, \alpha_N \in J$, hence the definition of Φ can only be

$$\Phi(g_1 \dots g_N) = \Phi_{\alpha_1}(g_1) \dots \Phi_{\alpha_N}(g_N) \in H.$$

It is similarly straightfoward to verify that Φ by this definition is a homomorphism.

Remark 12.2. In Lemma 11.18 at the end of the previous lecture the homomorphism

$$\mathop{\ast}_{\alpha \in J} \pi_1(A_\alpha, p) \xrightarrow{\Phi} \pi_1(X, p) \tag{12.1}$$

is determined as in the proposition above by the homomorphisms $(i_\alpha)_* : \pi_1(A_\alpha, p) \to \pi_1(X, p)$ induced by the inclusions $i_\alpha : A_\alpha \hookrightarrow X$.

We now address the previously unanswered question about the homomorphism (12.1) from Lemma 11.18: what is its kernel?

We can make two immediate observations about this: first, for any group homomorphism $\Psi : G \to H$, $\ker \Psi$ is a normal subgroup of G. Recall that a subgroup $K \subset G$ is called **normal** if it is invariant under conjugation with arbitrary elements of G, i.e.

$$gkg^{-1} \in K \quad \text{for all } k \in K \text{ and } g \in G.$$

This condition is abbreviated by "$gKg^{-1} = K$". It is obviously satisfied if $K = \ker \Psi$ since $\Psi(k) = e$ implies $\Psi(gkg^{-1}) = \Psi(g)\Psi(k)\Psi(g^{-1}) = \Psi(g)e\Psi(g)^{-1} = e$. Recall further that for any subgroup $K \subset G$, the **quotient** G/K is defined as the set of all **left cosets** of K, meaning subsets of the form $gK := \{gh \mid h \in K\}$ for fixed elements $g \in G$. For arbitrary subgroups $K \subset G$, the quotient

G/K does not have a natural group structure, but it does when K is a *normal* subgroup: indeed, the condition $gKg^{-1} = K$ gives rise to a well-defined product

$$(aK)(bK) := (ab)K \in G/K$$

since, as subsets of G, $aKbK = a(bKb^{-1})bK = abKK = abK$. In particular, any homomorphism $\Psi : G \to H$ between groups G and H gives rise to a normal subgroup $K := \ker \Psi \subset G$ and thus a quotient group G/K, such that Ψ determines a a well-defined map

$$G/\ker \Psi \to H : gK \mapsto \Psi(g),$$

meaning that the value $\Psi(g)$ of this map does not depend on the choice of element $g \in G$ representing the coset $gK \in G/K$. It is easy to check that this map is also a group homomorphism, in which case we say that Ψ **descends** to a homomorphism $G/K \to H$, and moreover, it is injective since $\Psi(g) = e$ means $g \in \ker \Psi = K$ and thus $gK = K = eK$, which is the identity element of G/K. It follows that the induced map $G/\ker \Psi \to H$ is an isomorphism whenever the original homomorphism Ψ is surjective. (A standard reference for these basic notions from group theory is [**Art91**].)

The second observation concerns certain specific elements that obviously belong to the kernel of the map (12.1). Consider the inclusions

$$j_{\alpha\beta} : A_\alpha \cap A_\beta \hookrightarrow A_\alpha$$

for each pair $\alpha, \beta \in J$, and recall that $i_\alpha : A_\alpha \hookrightarrow X$ denotes the inclusion of $A_\alpha \subset X$. Then the following diagram commutes,

$$\begin{tikzcd}
 & A_\alpha \arrow[dr, "i_\alpha"] & \\
A_\alpha \cap A_\beta \arrow[ur, "j_{\alpha\beta}"] \arrow[dr, "j_{\beta\alpha}"'] & & X \\
 & A_\beta \arrow[ur, "i_\beta"'] &
\end{tikzcd}$$

meaning $i_\alpha \circ j_{\alpha\beta} = i_\beta \circ j_{\beta\alpha}$, since both are just the inclusion of $A_\alpha \cap A_\beta$ into X. This trivial observation has a nontrivial consequence for the homomorphism Φ. Indeed, for any loop $p \overset{\gamma}{\rightsquigarrow} p$ in $A_\alpha \cap A_\beta$ representing a nontrivial element of $\pi_1(A_\alpha \cap A_\beta, p)$, the two elements $(j_{\alpha\beta})_*[\gamma] \in \pi_1(A_\alpha, p)$ and $(j_{\beta\alpha})_*[\gamma] \in \pi_1(A_\beta, p)$ belong to distinct subgroups in the free product $*_{\gamma \in J}\, \pi_1(A_\gamma, p)$, yet clearly

$$(i_\alpha)_*(j_{\alpha\beta})_*[\gamma] = (i_\beta)_*(j_{\beta\alpha})_*[\gamma] \in \pi_1(X, p)$$

since $i_\alpha \circ j_{\alpha\beta} = i_\beta \circ j_{\beta\alpha}$. It follows that $\Phi((j_{\alpha\beta})_*[\gamma]) = \Phi((j_{\beta\alpha})_*[\gamma])$, hence $\ker \Phi$ must contain the reduced word formed by the two letters $(j_{\alpha\beta})_*[\gamma] \in \pi_1(A_\alpha, p)$ and $(j_{\beta\alpha})_*[\gamma]^{-1} \in \pi_1(A_\beta, p)$:

$$(j_{\alpha\beta})_*[\gamma](j_{\beta\alpha})_*[\gamma]^{-1} \in \ker \Phi.$$

Combining this with the first observation, $\ker \Phi$ must contain the smallest normal subgroup of $*_{\gamma \in J}\, \pi_1(A_\gamma, p)$ that contains all elements of this form.

Definition 12.3. For any group G and subset $S \subset G$, we denote by

$$\langle S \rangle \subset G$$

the smallest subgroup of G that contains S, i.e. $\langle S \rangle$ is the set of all products of elements $g \in S$ and their inverses g^{-1}. Similarly,

$$\langle S \rangle_N \subset G$$

denotes the smallest *normal* subgroup of G that contains S. Concretely, this means $\langle S \rangle_N$ is the set of all conjugates of products of elements of S and their inverses.

We are now in a position to state the complete version of the Seifert-van Kampen theorem. The first half of the statement is just a repeat of Lemma 11.18, which we have proved already. The second half tells us what $\ker \Phi$ is, and thus gives a formula for $\pi_1(X, p)$.

THEOREM 12.4 (Seifert-van Kampen). *Suppose* $X = \bigcup_{\alpha \in J} A_\alpha$ *for a collection of open and path-connected subsets* $\{A_\alpha \subset X\}_{\alpha \in J}$ *with nonempty intersection, denote by* $i_\alpha : A_\alpha \hookrightarrow X$ *and* $j_{\alpha\beta} : A_\alpha \cap A_\beta \hookrightarrow A_\alpha$ *the inclusion maps for* $\alpha, \beta \in J$*, and fix* $p \in \bigcap_{\alpha \in J} A_\alpha$*.*

(1) If $A_\alpha \cap A_\beta$ *is path-connected for every pair* $\alpha, \beta \in J$*, then the natural homomorphism*
$$\Phi : \mathop{\ast}_{\alpha \in J} \pi_1(A_\alpha, p) \to \pi_1(X, p)$$
induced by the homomorphisms $(i_\alpha)_* : \pi_1(A_\alpha, p) \to \pi_1(X, p)$ *is surjective.*

(2) If additionally $A_\alpha \cap A_\beta \cap A_\gamma$ *is path-connected for every triple* $\alpha, \beta, \gamma \in J$*, then*
$$\ker \Phi = \left\langle \left\{ (j_{\alpha\beta})_*[\gamma](j_{\beta\alpha})_*[\gamma]^{-1} \;\middle|\; \alpha, \beta \in J,\ [\gamma] \in \pi_1(A_\alpha \cap A_\beta, p) \right\} \right\rangle_N .$$
In particular, Φ *then descends to an isomorphism*
$$\mathop{\ast}_{\alpha \in J} \pi_1(A_\alpha, p) \Big/ \ker \Phi \xrightarrow{\cong} \pi_1(X, p).$$

REMARK 12.5. In most applications, we will consider coverings of X by only two subsets $X = A \cup B$, and the condition on triple intersections in the second half of the statement then merely demands that $A \cap B$ be path-connected, which we already needed for the first half. (One can take the third subset in that condition to be either A or B; we never said that α, β and γ need to be distinct!)

I will give you the remaining part of the proof of this theorem in the next lecture. Let's now discuss some simple applications.

EXAMPLE 12.6. Consider the figure-eight $S^1 \vee S^1$ with its natural base point $p \in S^1 \vee S^1$, i.e. $S^1 \vee S^1$ is the union of two circles $A, B \subset S^1 \vee S^1$ with $A \cap B = \{p\}$. These are not open subsets, but since a neighborhood of p in $S^1 \vee S^1$ has a fairly simple structure, we can get away with the usual trick (cf. Examples 11.3 and 11.7) of replacing both with homotopy equivalent open neighborhoods: define $A' \subset S^1 \vee S^1$ as a small open neighborhood of A and $B' \subset S^1 \vee S^1$ as a small open neighborhood of B such that there exist deformation retractions of A' to A and B' to B. The inclusions $A \hookrightarrow A'$ and $B \hookrightarrow B'$ then induce isomorphisms $\mathbb{Z} \cong \pi_1(A, p) \xrightarrow{\cong} \pi_1(A', p)$ and $\mathbb{Z} \cong \pi_1(B, p) \xrightarrow{\cong} \pi_1(B', p)$. The intersection $A' \cap B'$ is now a pair of line segments with one intersection point at p, so it admits a deformation retraction to p and is thus contractible, implying $\pi_1(A' \cap B', p) = 0$. This makes $\ker \Phi$ in Theorem 12.4 trivial, hence the map
$$\pi_1(A, p) * \pi_1(B, p) \to \pi_1(S^1 \vee S^1, p)$$
determined by the homomorphisms of $\pi_1(A, p)$ and $\pi_1(B, p)$ to $\pi_1(S^1 \vee S^1, p)$ induced by the inclusions $A, B \hookrightarrow S^1 \vee S^1$ is an isomorphism. To see more concretely what this group looks like, fix generators $\alpha \in \pi_1(A, p) \cong \mathbb{Z}$ and $\beta \in \pi_1(B, p) \cong \mathbb{Z}$, each of which can also be identified with elements of $\pi_1(S^1 \vee S^1, p)$ via the inclusions of A and B into $S^1 \vee S^1$. Then
$$\pi_1(S^1 \vee S^1, p) \cong \mathbb{Z} * \mathbb{Z} = \{e, \alpha^p, \beta^q, \alpha^p\beta^q, \beta^p\alpha^q, \alpha^p\beta^q\alpha^r, \ldots \mid p, q, r, \ldots \in \mathbb{Z}\}.$$
These elements are easy to visualize: α and β are represented by loops that start and end at p and run once around the circles A or B respectively, so each element in the above list is a concatenation of finitely many repetitions of these two loops and their inverses. Notice that $\alpha\beta \neq \beta\alpha$, so $\pi_1(S^1 \vee S^1)$ is our first example of a nonabelian fundamental group.

EXAMPLE 12.7. Recall from Exercise 7.27 that for each $n \in \mathbb{N}$, one can identify S^n with the *one point compactification* of $\mathbb{R}^n$, a space defined by adjoining a single point called "∞" to $\mathbb{R}^n$:

$$S^n \cong \mathbb{R}^n \cup \{\infty\}.$$

This gives rise to an inclusion map $\mathbb{R}^n \overset{i}{\hookrightarrow} S^n$ with image $S^n \backslash \{\infty\}$. We claim that for any compact subset $K \subset \mathbb{R}^3$ such that $\mathbb{R}^3 \backslash K$ is path-connected, and any choice of base point $p \in \mathbb{R}^3 \backslash K$,

$$i_* : \pi_1(\mathbb{R}^3 \backslash K, p) \to \pi_1(S^3 \backslash K, p)$$

is an isomorphism. To see this, define the open subset $A := \mathbb{R}^3 \backslash K \subset S^3 \backslash K$, and choose $B_0 \subset S^3 \backslash K$ to be an open ball about ∞, i.e. a set of the form $(\mathbb{R}^3 \backslash \overline{B_R(0)}) \cup \{\infty\}$ where $\overline{B_R(0)} \subset \mathbb{R}^3$ is any closed ball large enough to contain K. Since p might not be contained in B_0 but $\mathbb{R}^3 \backslash K$ is path-connected, we can then define a larger set B by adjoining to B_0 the neighborhood in $\mathbb{R}^3 \backslash K$ of some path from a point in B_0 to p: this can be done so that both B_0 and B are homeomorphic to an open ball, so in particular they are contractible. The intersection $A \cap B$ is then $B \backslash \{\infty\}$ and is thus homoemorphic to $\mathbb{R}^3 \backslash \{0\}$ and homotopy equivalent to S^2, implying $\pi_1(A \cap B) = 0$. The Seifert-van Kampen theorem therefore gives an isomorphism $\pi_1(\mathbb{R}^3 \backslash K, p) * \pi_1(B, p) \to \pi_1(S^3 \backslash K, p)$, but $\pi_1(B, p)$ is the trivial group, so this proves the claim.

A frequently occuring special case of this example is when $K \subset \mathbb{R}^3$ is a knot, i.e. the image of an embedding $S^1 \hookrightarrow \mathbb{R}^3$. The fundamental group $\pi_1(\mathbb{R}^3 \backslash K)$ is then called the **knot group** of K, and the argument above shows that we are free to adjoin a point at infinity and thus replace the knot group with $\pi_1(S^3 \backslash K)$. This will be convenient for certain computations.

As in the previous lecture, we shall conclude this one by introducing some more terminology from combinatorial group theory in order to state a more usable variation on the Seifert-van Kampen theorem.

DEFINITION 12.8. Given a set S, the **free group on** S is defined as

$$F_S := \mathop{\ast}_{\alpha \in S} \mathbb{Z},$$

or in other words, the set of all reduced words $a_1^{p_1} a_2^{p_2} \ldots a_N^{p_N}$ for $N \geqslant 0$, $p_i \in \mathbb{Z}$ with $p_i \neq 0$, $a_i \in S$ and $a_i \neq a_{i+1}$ for every i, with the product defined by concatenation of words followed by reduction. The elements of S are called the **generators** of F_S.

EXAMPLE 12.9. The computation in Example 12.6 gives $\pi_1(S^1 \vee S^1) \cong F_{\{\alpha,\beta\}} \cong \mathbb{Z} * \mathbb{Z}$, where the set generating $F_{\{\alpha,\beta\}}$ consists of the two loops α and β parametrizing the two circles that form $S^1 \vee S^1$.

PROPOSITION 12.10. *For any set S, group G and map $\phi : S \to G$, there is a unique group homomorphism $\Phi : F_S \to G$ satisfying $\Phi(a) = \phi(a)$ for single-letter words $a \in F_S$ defined by elements $a \in S$.*

PROOF. Writing elements of F_S in the form $a_1^{p_1} a_2^{p_2} \ldots a_N^{p_N}$, there is clearly only one formula for $\Phi : F_S \to G$ that will match ϕ on single-letter words and also be a homomorphism, namely

$$\Phi(a_1^{p_1} \ldots a_N^{p_N}) = \phi(a_1)^{p_1} \ldots \phi(a_N)^{p_N}.$$

It is straightforward to check that this defines a homomorphism. $\square$

PROPOSITION 12.11. *Every group is isomorphic to a quotient of a free group by some normal subgroup.*

PROOF. Pick any subset $S \subset G$ that generates G, e.g. one can choose $S := G$, though smaller subsets are usually also possible. Then the unique homomorphism $\Phi : F_S \to G$ sending each $g \in S \subset F_S$ to $g \in G$ is surjective, thus Φ descends to an isomorphism $F_S/\ker \Phi \to G$. $\square$

DEFINITION 12.12. Given a set S, a **relation** in S is defined to mean any equation of the form "$a = b$" where $a, b \in F_S$.

DEFINITION 12.13. For any set S and a set R consisting of relations in S, we define the group

$$\{S \mid R\} := F_S \Big/ \langle R' \rangle_N$$

where R' is the set of all elements of the form $ab^{-1} \in F_S$ for relations "$a = b$" in R. The elements of S are called the **generators** of this group, and elements of R are its **relations**.

Let us pause a moment to interpret this definition. By a slight abuse of notation, we can write each element of $\{S \mid R\}$ as a reduced word w formed out of letters in S, with the understanding that w represents an equivalence class in the quotient $F_S/\langle R' \rangle_N$, thus it is possible to have $w = w'$ in $\{S \mid R\}$ even if w and w' are distinct elements of F_S. This will happen if and only if $w^{-1}w'$ belongs to the normal subgroup $\langle R' \rangle_N$, and in particular, it happens whenever "$w = w'$" is one of the relations in R. The relations are usually necesary because most groups are not free groups: while free groups are easy to describe (they depend only on their generators), most groups have more interesting structure than free groups, and this structure is encoded by relations. Proposition 12.11 implies that *every* group can be presented in this way, i.e. every group is isomorphic to $\{S \mid R\}$ for some set of generators S and relations R. Indeed, if $G = F_S/\ker \Phi$ for a set S and a surjective homomorphism $\Phi : F_S \to G$, then we can take S as the set of generators and define R to consist of all relations of the form "$a = b$" such that $ab^{-1} \in \ker \Phi$; the latter is equivalent to the condition $\Phi(a) = \Phi(b)$, so the relations tell us precisely when two products of generators give us the same element in G.

DEFINITION 12.14. Given a group G, a **presentation** of G consists of a subset $S \subset G$ together with a set R of relations in S such that the unique homomorphism $F_S \to G$ matching the inclusion $S \hookrightarrow G$ on single-letter words descends to a group isomorphism

$$\{S \mid R\} \xrightarrow{\cong} G.$$

We say that G is **finitely presented** if it admits a presentation such that S and R are both finite sets.

EXAMPLE 12.15. The group $\{a\} := \{a \mid \varnothing\}$ consisting of a single generator a with no relations is isomorphic to the free group $F_{\{a\}}$ on one element. The isomorphism $a^p \mapsto p$ identifies this with the integers $\mathbb{Z}$.

EXAMPLE 12.16. The group $\{a, b \mid ab = ba\}$ has two generators and is abelian, so it is isomorphic to $\mathbb{Z}^2$. An explicit isomorphism is defined by $a^p b^q \mapsto (p, q)$. To see that this is an isomorphism, observe first that since $F_{\{a,b\}}$ is free, there exists a unique homomorphism $\Phi : F_{\{a,b\}} \to \mathbb{Z}^2$ with $\Phi(a) = (1, 0)$ and $\Phi(b) = (0, 1)$, and Φ is clearly surjective since it necesarily sends $a^p b^q$ to (p, q). Since $\mathbb{Z}^2$ is abelian, we also have

$$\Phi(ab(ba)^{-1}) = \Phi(aba^{-1}b^{-1}) = \Phi(a) + \Phi(b) - \Phi(a) - \Phi(b) = 0,$$

so $\ker \Phi$ contains $ab(ba)^{-1}$ and therefore also contains the smallest normal subgroup containing $ab(ba)^{-1}$, which is the group $\langle R' \rangle_N$ appearing in the quotient $\{a, b \mid ab = ba\} = F_{\{a,b\}}/\langle R' \rangle_N$. This proves that Φ descends to a surjective homomorphism $\{a, b \mid ab = ba\} \to \mathbb{Z}^2$. Finally, observe that since $ab = ba$ in the quotient $\{a, b \mid ab = ba\}$, every reduced word in $F_{\{a,b\}}$ is equivalent in this quotient to a word of the form $a^p b^q$ for some $(p, q) \in \mathbb{Z}^2$, and $\Phi(a^p b^q)$ then vanishes if and only if $a^p b^q = e$, proving that Φ is also injective.

EXAMPLE 12.17. The group $\{a \mid a^p = e\}$ is isomorphic to $\mathbb{Z}_p := \mathbb{Z}/p\mathbb{Z}$, with an explicit isomorphism defined in terms of the unique homomorphism $F_{\{a\}} \to \mathbb{Z}_p$ that sends a to $[1]$.

EXAMPLE 12.18. We will prove in Lecture 14 that for the trefoil knot $K \subset \mathbb{R}^3 \subset S^3$, (see Lecture 8), $\pi_1(S^3 \backslash K) \cong \{a, b \mid a^2 = b^3\}$, and Exercise 12.20 below proves that this group is not abelian. By contrast, we will also see that the unknot $K_0 \subset \mathbb{R}^3 \subset S^3$ has $\pi_1(S^3 \backslash K_0) \cong \mathbb{Z}$, which is abelian. This implies via Example 12.7 that $\pi_1(\mathbb{R}^3 \backslash K) \not\cong \pi_1(\mathbb{R}^3 \backslash K_0)$, so $\mathbb{R}^3 \backslash K$ and $\mathbb{R}^3 \backslash K_0$ are not homeomorphic, hence the trefoil cannot be deformed continuously to the unknot.

Note that for any given set of generators S and relations R, it is often possible to reduce these to smaller sets without changing the isomorphism class of the group that they define. For the relations in particular, it is easy to imagine multiple distinct choices of the subset $R' \subset F_S$ that will produce the same normal subgroup $\langle R' \rangle_N$. In general, it is a very hard problem to determine whether or not two groups described via generators and relations are isomorphic; in fact, it is known that there does not exist any algorithm to decide whether a given presentation defines the trivial group. Nonetheless, generators and relations provide a very convenient way to describe many simple groups that arise in practice, especially in the context of van Kampen's theorem. This is due to the following reformulation of Theorem 12.4 for the case of two open subsets when all fundamental groups are finitely presented.

COROLLARY 12.19 (Seifert-van Kampen for finitely-presented groups). *Suppose $X = A \cup B$ where $A, B \subset X$ are open and path-connected subsets such that $A \cap B$ is also path-connected, and $j_A : A \cap B \hookrightarrow A$ and $j_B : A \cap B \hookrightarrow B$ denote the inclusions. Suppose moreover that there exist finite presentations*

$$\pi_1(A) \cong \big\{\{a_i\} \bigm| \{R_j\}\big\}, \qquad \pi_1(B) \cong \big\{\{b_k\} \bigm| \{S_\ell\}\big\}, \qquad \pi_1(A \cap B) \cong \big\{\{c_p\} \bigm| \{T_q\}\big\},$$

with the indices i, j, k, ℓ, p, q each ranging over finite sets. Then

$$\pi_1(X) \cong \big\{\{a_i\} \cup \{b_k\} \bigm| \{R_j\} \cup \{S_\ell\} \cup \{(j_A)_* c_p = (j_B)_* c_p\}\big\}.$$

□

In other words, as generators for $\pi_1(X)$, one can take all generators of $\pi_1(A)$ together with all generators of $\pi_1(B)$. The relations must then include all of the relations among the generators of $\pi_1(A)$ and $\pi_1(B)$ separately, but there may be additional relations that mix the generators from $\pi_1(A)$ and $\pi_1(B)$: these extra relations set $(j_A)_* c_p \in \pi_1(A)$ equal to $(j_B)_* c_p \in \pi_1(B)$ for each of the generators c_p of $\pi_1(A \cap B)$. These extra relations are exactly what is needed to describe the normal subgroup $\ker \Phi$ in the statement of Theorem 12.4. The relations in $\pi_1(A \cap B)$ do not play any role.

EXERCISE 12.20. Let us prove that the finitely-presented group $G = \{x, y \mid x^2 = y^3\}$ mentioned in Example 12.18 is nonabelian.

(a) Denoting the identity element by e, consider the related group

$$H = \{x, y \mid x^2 = y^3,\ y^3 = e,\ xyxy = e\}.$$

Show that every element of H is equivalent to one of the six elements $e, x, y, y^2, xy, xy^2 \in H$. This proves that H has order at most six, though in theory it could be less, since some of those six elements might still be equivalent to each other. To prove that this is not the case, construct (by writing down a multiplication table) a nonabelian group H' of order six that is generated by two elements a, b satisfying the relations $a^2 = b^3 = e$ and $abab = e$. Show that there exists a surjective homomorphism $H \to H'$, which is therefore an isomorphism since $|H| \leqslant 6$.
Remark: You don't need this fact, but if you've seen some of the standard examples of finite groups before, you might in any case notice that H is isomorphic to the dihedral group (Diedergruppe) of order 6.

(b) Show that H is a quotient of G by some normal subgroup, and deduce that G is also nonabelian.

EXERCISE 12.21. Given a group G, the **commutator subgroup** $[G,G] \subset G$ is the subgroup generated by all elements of the form

$$[x,y] := xyx^{-1}y^{-1}$$

for $x, y \in G$.

(a) Show that $[G,G] \subset G$ is always a normal subgroup, and it is trivial if and only if G is abelian.
(b) The **abelianization** (Abelisierung) of G is defined as the quotient group $G/[G,G]$. Show that this group is always abelian, and it is equal to G if G is already abelian.[13]
(c) Given any two abelian groups G, H, find a natural isomorphism from the abelianization of the free product $G * H$ to the Cartesian product $G \times H$.
(d) Prove that the abelianization of $\{x, y \mid x^2 = y^3\}$ is isomorphic to $\mathbb{Z}$.
Hint: An isomorphism φ from the abelianization to $\mathbb{Z}$ will be determined by two integers, $\varphi(x)$ and $\varphi(y)$. If φ exists, how must these two integers be related to each other?

13. Proof of the Seifert-van Kampen theorem (June 1, 2023)

We have put off the proof of the Seifert-van Kampen theorem long enough. Here again is the statement.

THEOREM 13.1 (Seifert-van Kampen). *Suppose $X = \bigcup_{\alpha \in J} A_\alpha$ for a collection of open and path-connected subsets $\{A_\alpha \subset X\}_{\alpha \in J}$, $i_\alpha : A_\alpha \hookrightarrow X$ and $j_{\alpha\beta} : A_\alpha \cap A_\beta \hookrightarrow A_\alpha$ denote the natural inclusion maps for $\alpha, \beta \in J$, and $p \in \bigcap_{\alpha \in J} A_\alpha$.*

(1) If $A_\alpha \cap A_\beta$ is path-connected for every pair $\alpha, \beta \in J$, then the unique homomorphism

$$\Phi : \mathop{\ast}_{\alpha \in J} \pi_1(A_\alpha, p) \to \pi_1(X, p)$$

that restricts to each subgroup $\pi_1(A_\alpha, p) \subset \ast_{\beta \in J} \pi_1(A_\beta, p)$ as $(i_\alpha)_*$ is surjective.

(2) If additionally $A_\alpha \cap A_\beta \cap A_\gamma$ is path-connected for every triple $\alpha, \beta, \gamma \in J$, then

$$\ker \Phi = \langle S \rangle_N,$$

meaning $\ker \Phi$ is the smallest normal subgroup containing the set

$$S := \left\{ (j_{\alpha\beta})_*[\gamma](j_{\beta\alpha})_*[\gamma]^{-1} \;\middle|\; \alpha, \beta \in J,\ [\gamma] \in \pi_1(A_\alpha \cap A_\beta, p) \right\}.$$

In particular, if we abbreviate $F := \ast_{\alpha \in J} \pi_1(A_\alpha, p)$, then Φ descends to an isomorphism

$$F \big/ \langle S \rangle_N \to \pi_1(X, p).$$

PROOF. We proved the first statement already in Lecture 11, so assume the hypothesis of the second statement holds. As observed in the previous lecture, $\Phi((j_{\alpha\beta})_*\gamma) = \Phi((j_{\beta\alpha})_*\gamma)$ for every $\alpha, \beta \in J$ and $\gamma \in \pi_1(A_\alpha \cap A_\beta, p)$, thus $\ker \Phi$ clearly contains $\langle S \rangle_N$, and in particular, Φ descends to a surjective homomorphism $F/\langle S \rangle_N \to \pi_1(X, p)$. We need to show that this homomorphism is injective, or equivalently, that whenever $\Phi(w) = \Phi(w')$ for a pair of reduced words $w, w' \in F$, their equivalence classes in $F/\langle S \rangle_N$ must match.

[13]Note that if $G = \{S \mid R\}$ is a finitely-presented group with generators S and relations R, then its abelianization is $\{S \mid R'\}$ where R' is the union of R with all relations of the form "$ab = ba$" for $a, b \in S$.

Given a loop $p \overset{\gamma}{\rightsquigarrow} p$ in X, let us say that a *factorization of* γ is any finite sequence $\{(\gamma_i, \alpha_i)\}_{i=1}^N$ such that $\alpha_i \in J$ and $p \overset{\gamma_i}{\rightsquigarrow} p$ is a loop in A_{α_i} for each $i = 1, \dots, N$, and

$$\gamma \underset{h+}{\sim} \gamma_1 \cdot \ldots \cdot \gamma_N.$$

The first half of the theorem follows from the fact (proved in Lemma 11.2) that every γ has a factorization. Now observe that any factorization as described above determines a reduced word $w \in F$, defined as the reduction of the word $[\gamma_1] \dots [\gamma_N]$ with $[\gamma_i] \in \pi_1(A_{\alpha_i}, p)$ for $i = 1, \dots, N$, and this word satisfies $\Phi(w) = [\gamma]$. Conversely, every reduced word $w \in \Phi^{-1}([\gamma])$ can be realized as a factorization of γ by choosing specific loops to represent the letters in w. The theorem will then follow if we can show that any two factorizations of γ can be related to each other by a finite sequence of the following operations and their inverses:

(A) Given two adjacent loops γ_i and γ_{i+1} such that $\alpha_i = \alpha_{i+1}$, replace them with their concatenation $p \overset{\gamma_i \cdot \gamma_{i+1}}{\rightsquigarrow} p$. (This does not change the corresponding reduced word in F, as it just implements a step in the reduction of an unreduced word.)

(B) Replace some γ_i with any loop γ_i' that is homotopic (with fixed end points) in A_{α_i}. (This also does not change the corresponding reduced word in F; in fact it doesn't even change the unreduced word from which it is derived.)

(C) Given a loop γ_i that lies in $A_{\alpha_i} \cap A_\beta$ for some $\beta \in J$, replace α_i with β. (In the corresponding reduced word in F, this replaces a letter of the form $(j_{\alpha_i \beta})_* [\gamma_i] \in \pi_1(A_{\alpha_i}, p)$ with one of the form $(j_{\beta \alpha_i})_* [\gamma_i] \in \pi_1(A_\beta, p)$, thus it changes the word but does not change its equivalence class in $F / \langle S \rangle_N$.)

We now prove that any two factorizations $\{(\gamma_i, \alpha_i)\}_{i=1}^N$ and $\{(\gamma_i', \alpha_i')\}_{i=1}^{N'}$ of γ are related by these operations. By assumption $\gamma_1 \cdot \ldots \cdot \gamma_N \underset{h+}{\sim} \gamma_1' \cdot \ldots \cdot \gamma_{N'}'$, so after choosing suitable parametrizations of both of these concatenations on the unit interval I,[14] there exists a homotopy

$$H : I^2 \to X$$

with $H(0, \cdot) = \gamma_1 \cdot \ldots \cdot \gamma_N$, $H(1, \cdot) = \gamma_1' \cdot \ldots \cdot \gamma_N'$ and $H(s, 0) = H(s, 1) = p$ for all $s \in I$. Since I^2 is compact, one can find a number $\epsilon > 0$ such that for every $(s, t) \in I^2$,[15] the intersection of I^2 with the box

$$[s - 2\epsilon, s + 2\epsilon] \times [t - 2\epsilon, t + 2\epsilon] \subset \mathbb{R}^2$$

is contained in $H^{-1}(A_\alpha)$ for some $\alpha \in J$. For suitably small $\epsilon = 1/n$ with $n \in \mathbb{N}$, we can therefore break up I^2 into n^2 boxes of side length ϵ which are each contained in $H^{-1}(A_\alpha)$ for some $\alpha \in J$ (possibly a different α for each box), forming a grid in I^2. For each box in the diagram there may be multiple $\alpha \in J$ that satisfy this condition, but let us choose a specific one to associate to each box. (These choices are indicated by the three colors in Figure 3.) Notice that each vertex in the grid is contained in the intersection of $H^{-1}(A_\alpha)$ for each of the $\alpha \in J$ associated to boxes that it touches. We can now perturb this diagram slightly to fill I^2 with a collection of boxes of slightly varying sizes such that every vertex in the interior touches only three of them (see the right side of Figure 3). We can similarly assume after such a perturbation that the vertices in $\{s = 0\}$ and $\{s = 1\}$ never coincide with the starting or ending times of the loops γ_i, γ_i' in the concatenations

[14] Recall that concatenation of paths is associative up to homotopy, so the N-fold concatenation $\gamma_1 \cdot \ldots \cdot \gamma_N$ is not a uniquely determined path $I \to X$ if $N > 2$, but it is unique up to homotopy with fixed end points.

[15] I do not consider this statement completely obvious, but it is a not very difficult exercise in point-set topology, and since that portion of the course is now over, I would rather leave it as an exercise than give the details here. Here is a hint: if the claim is not true, one can find a sequence $(s_k, t_k) \in I^2$ such that for each k, the intersection of I^2 with the box of side length $1/k$ about (s_k, t_k) is not fully contained in any of the subsets $H^{-1}(A_\alpha)$. This sequence has a convergent subsequence. What can you say about its limit?

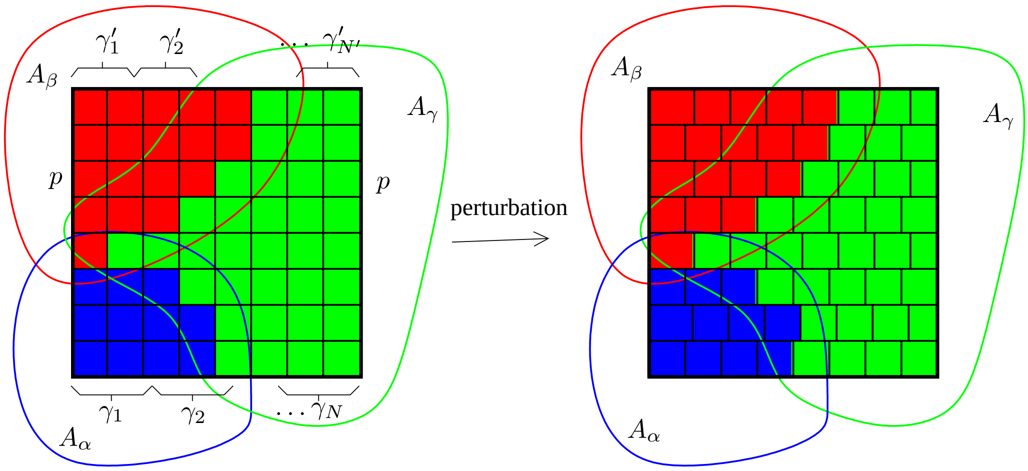

FIGURE 3. A grid on the domain of the homotopy $H : I^2 \to X$ between two factorizations $\gamma_1 \cdot \ldots \cdot \gamma_N$ and $\gamma'_1 \cdot \ldots \cdot \gamma'_{N'}$ of a loop $p \overset{\gamma}{\rightsquigarrow} p$ in X. In this example, there are three open sets $A_\alpha, A_\beta, A_\gamma \subset X$, and colors are used to indicate that each of the small boxes filling I^2 has image lying in (at least) one of these subsets. In the perturbed picture at the right, every vertex in the interior touches exactly three boxes.

$\gamma_1 \cdot \ldots \cdot \gamma_N$ and $\gamma'_1 \cdot \ldots \cdot \gamma'_{N'}$. Moreover, each vertex still lies in the same intersection of sets $H^{-1}(A_\alpha)$ as before, assuming the perturbation is sufficiently small.

Now suppose $(s,t) \in I^2$ is a vertex in the interior of the perturbed grid. Then (s,t) is on the boundary of exactly three boxes in the diagram, each of which belongs to one of the sets $H^{-1}(A_\alpha)$, $H^{-1}(A_\beta)$ and $H^{-1}(A_\gamma)$ for three associated elements $\alpha, \beta, \gamma \in J$ (they need not necessarily be distinct). If $(0,t)$ is a vertex with $t \notin \{0,1\}$, then it is on the boundary of exactly two boxes and thus lies in $H^{-1}(A_\alpha \cap A_\beta)$ for two associated elements $\alpha, \beta \in J$, but it also lies in $H^{-1}(A_\gamma)$ where $\gamma := \alpha_i$ is associated to the particular path γ_i whose domain as part of the concatenation $H(0,\cdot) = \gamma_1 \cdot \ldots \cdot \gamma_N$ contains $(0,t)$. For vertices $(1,t)$ with $t \notin \{0,1\}$, choose $A_\gamma := A_{\alpha'_i}$ similarly in terms of the concatenation $\gamma'_1 \cdot \ldots \cdot \gamma'_{N'}$. In any of these cases, we have associated to each vertex (s,t) a path-connected set $A_\alpha \cap A_\beta \cap A_\gamma$ that contains $H(s,t)$, thus we can choose a path[16]

$$H(s,t) \overset{\delta_{(s,t)}}{\rightsquigarrow} p \quad \text{in} \quad A_\alpha \cap A_\beta \cap A_\gamma.$$

Since $H(s,t) = p$ for $t \in \{0,1\}$, this definition can be extended to vertices with $t \in \{0,1\}$ by defining $\delta_{(s,t)}$ as the trivial path. Now if E is any edge in the diagram, i.e. a side of one of the boxes, connecting two neighboring vertices (s_0,t_0) and (s_1,t_1), then we can identify E with the unit interval in order to regard $H|_E : E \to X$ as a path, and thus associate to E a loop

$$p \overset{\gamma_E}{\rightsquigarrow} p \quad \text{in} \quad A_\alpha \cap A_\beta, \qquad \gamma_E := \delta^{-1}_{(s_0,t_0)} \cdot H|_E \cdot \delta_{(s_1,t_1)},$$

where $\alpha, \beta \in J$ are the two (not necessarily distinct) elements associated to the boxes bordered by E.

[16] This is the specific step where we need the assumption that triple intersections are path-connected. If you're curious to see an example of the second half of the theorem failing without this assumption, I refer you to [**Hat02**, p. 44].

With these choices in place, any path through I^2 that follows a sequence of edges $E_1, \ldots, E_k$ starting at some vertex in $(s_0, 0)$ and ending at a vertex $(s_1, 1)$ produces various factorizations of γ in the form $\{(\gamma_{E_i}, \beta_i)\}_{i=1}^k$. Here there is some freedom in the choices of $\beta_i \in J$: whenever a given edge E_i lies in $H^{-1}(A_\beta) \cap H^{-1}(A_\gamma)$, we can choose β_i to be either β or γ and thus produce two valid factorizations, which are related to each other by operation (C) in the list above.

We can now describe a procedure to modify the factorization $\{(\gamma_i, \alpha_i)\}_{i=1}^N$ to $\{(\gamma_i', \alpha_i')\}_{i=1}^{N'}$. We show first that $\{(\gamma_i, \alpha_i)\}_{i=1}^N$ is equivalent via our three operations to the factorization corresponding to the sequence of edges in $\{s = 0\}$ moving from $t = 0$ to $t = 1$. This is not so obvious because, although $H(0, \cdot)$ is a parametrization of the concatenated path $\gamma_1 \cdot \ldots \cdot \gamma_N$, the times that mark the boundaries between one path and the next in this concatenation need not have anything to do with the vertices of our chosen grid. Instead, our perturbation of the grid ensured that each γ_i in the concatenation hits vertices only in the interior of its domain, not at starting or end points. Denote by $(0, t_1), \ldots, (0, t_{m-1})$ the particular grid vertices in the domain of γ_i, thus splitting up γ_i into a concatenation of paths $\gamma_i = \gamma_i^1 \cdot \ldots \cdot \gamma_i^m$ which have these vertices as starting and/or end points. Then

$$\gamma_i \underset{h+}{\sim} (\gamma_i^1 \cdot \delta_{(0,t_1)}) \cdot (\delta_{(0,t_1)}^{-1} \cdot \gamma_i^2 \cdot \delta_{(0,t_2)}) \cdot \ldots \cdot (\delta_{(0,t_{m-1})}^{-1} \cdot \gamma_i^m) \quad \text{in } A_{\alpha_i}.$$

We can now apply operations (B) and (A) in that order to replace γ_i with the sequence of loops of the form $\delta_{(0,t_{j-1})}^{-1} \cdot \gamma_i^j \cdot \delta_{(0,t_j)}$ in A_{α_i} as indicated above. The result is a new factorization that has more loops in the sequence, but the resulting concatenation is broken up along points that include all vertices in $\{s = 0\}$. It is also broken along more points, corresponding to the pieces of the original concatenation $\gamma_1 \cdot \ldots \cdot \gamma_N$, but after applying operation (C) if necessary, we can now apply operation (A) to combine all adjacent loops whose domains belong to the same edge. The result is precisely the factorization corresponding to the sequence of edges in $\{s = 0\}$. The same procedure can be used to modify $\{(\gamma_i', \alpha_i')\}_{i=1}^{N'}$ to the factorization corresponding to the sequence of edges in $\{s = 1\}$.

To finish, we need to show that the factorization given by the edges in $\{s = 0\}$ can be transformed into the corresponding factorization at $\{s = 1\}$ by applying our three operations. The core of the idea for this is shown in Figure 4, where the purple curves show two sequences of edges which represent two factorizations. In this case the difference between one path and the other consists only of replacing two edges on adjacent sides of a particular box $Q \subset I^2$ with their two opposite sides, and we can change from one to the other as follows. First, if the box Q is in $H^{-1}(A_\alpha)$, apply the operation (C) to both factorizations until all the loops corresponding to sides of Q are regarded as loops in A_α. Having done this, both factorizations now contain two consecutive loops in A_α that correspond to two sides of Q, so we can apply the operation (A) to concatenate each of these pairs, reducing two loops to one distinguished loop through A_α in each factorization. Those two distinguished loops are also homotopic in A_α, as one can see by choosing a homotopy of paths through the square Q that connects two adjacent sides to their two opposite sides (Figure 4, right). This therefore applies the operation (B) to change one factorization to the other.

We note finally that for any sequence of edges that includes edges in $\{t = 0\}$ or $\{t = 1\}$, those edges represent the constant path at the base point p, and since concatenation with constant paths produces homotopic paths, adding these edges or removing them from the diagram changes the factorization by a combination of operations (A) and (B). It now only remains to observe that the path of edges along $\{s = 0\}$ can always be modified to the path of edges along $\{s = 1\}$ by a finite sequence of the modifications just described.

$\square$

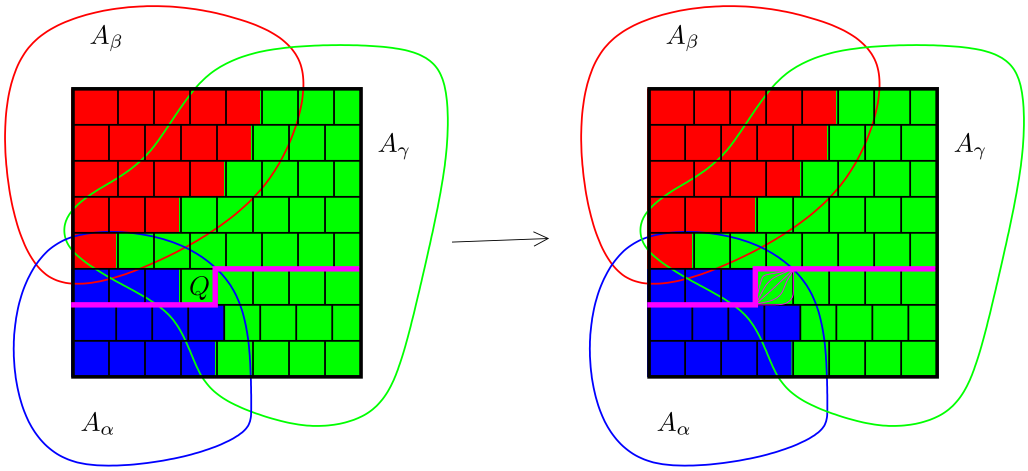

FIGURE 4. The magenta paths in both pictures are sequences of edges that define factorizations of γ, differing only at pairs of edges that surround a particular box Q. We can change one to the other by applying the three operations in our list.

EXERCISE 13.2. Recall that the wedge sum of two pointed spaces (X, x) and (Y, y) is defined as $X \vee Y = (X \amalg Y)/\sim$ where the equivalence relation identifies the two base points x and y. It is commonly said that whenever X and Y are both path-connected and are otherwise "reasonable" spaces, the formula

$$\pi_1(X \vee Y) \cong \pi_1(X) * \pi_1(Y) \tag{13.1}$$

holds. We saw for instance in Example 12.6 that this is true when X and Y are both circles. The goal of this problem is to understand slightly better what "reasonable" means in this context, and why such a condition is needed.

(a) Show by a direct argument (i.e. without trying to use Seifert-van Kampen) that if X and Y are both Hausdorff and simply connected, then $X \vee Y$ is simply connected.
Hint: Hausdorff implies that $X\backslash\{x\}$ and $Y\backslash\{y\}$ are both open subsets. Consider loops $\gamma : [0, 1] \to X \vee Y$ based at $[x] = [y]$ and decompose $[0, 1]$ into subintervals in which $\gamma(t)$ stays in either X or Y.

(b) Call a pointed space (X, x) *nice*[17] if x has an open neighborhood that admits a deformation retraction to x. Show that the formula (13.1) holds whenever (X, x) and (Y, y) are both nice, and more generally, the formula

$$\pi_1\left(\bigvee_{\alpha \in J} X_\alpha\right) \cong \mathop{\ast}_{\alpha \in J} \pi_1(X_\alpha)$$

holds for any (possibly infinite) collection of nice pointed spaces $\{(X_\alpha, x_\alpha)\}_{\alpha \in J}$.

[17]Not a standardized term, I made it up.

(c) Here is an example of a space that is not "nice" in the sense of part (b): for each $n \in \mathbb{N}$, let $S^1_n \subset \mathbb{R}^2$ denote the circle of radius $1/n$ centered at $(1/n, 0)$. The union of all these circles is a space known informally as the *Hawaiian earring*

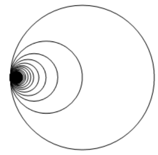

$$H := \bigcup_{n \in \mathbb{N}} S^1_n \subset \mathbb{R}^2.$$

As usual, we assign to H the subspace topology induced by the standard topology of $\mathbb{R}^2$. Show that in this space, the point $(0, 0)$ does not have any simply connected open neighborhood.

(d) It is tempting to liken the Hawaiian earring H to the infinite wedge sum of circles $X := \bigvee_{n=1}^{\infty} S^1$, defined as above by choosing a base point in each copy of the circle and then identifying all the base points in the infinite disjoint union $\coprod_{n=1}^{\infty} S^1$. Both are unions of infinite collections of circles that all intersect each other at one point. Show in fact that there exists a continuous map

$$f : X \to H$$

that is a bijection sending the natural base point of $\bigvee_n S^1$ to $(0, 0) \in \bigcap_n S^1_n$, but that X (unlike H) is a "nice" space, hence $f : X \to H$ cannot be a homeomorphism.
Hint: Continuity of maps defined on wedge sums is easy to check—see Exercise 10.5.

(e) Show that there exists a surjective continuous map $S^1 \to H$, but continuous maps $S^1 \to X$ are never surjective.
Hint: In H, start at $(0, 0)$ and traverse the largest circle first, then continue to smaller circles.

(f) Show that for any finite subset $J \subset \mathbb{N}$, there exists a retraction

$$r_J : H \to \bigcup_{n \in J} S^1_n \subset H,$$

and deduce from this that the map $f_* : \pi_1(X) \to \pi_1(H)$ is injective.
Hint: Unlike H, $\bigcup_{n \in J} S^1_n$ really is homeomorphic to a wedge sum of circles, the crucial detail in this case being that there are only finitely many.

(g) Writing $r_n := r_{\{n\}} : H \to S^1_n$ for each individual value of $n \in \mathbb{N}$, show that the homomorphism

$$\pi_1(H) \to \prod_{n \in \mathbb{N}} \pi_1(S^1_n) \cong \prod_{n \in \mathbb{N}} \mathbb{Z}$$

determined by the maps $(r_n)_* : \pi_1(H) \to \pi_1(S^1_n)$ is surjective, and deduce from this that $f_* : \pi_1(X) \to \pi_1(H)$ is not injective.
Remark: The direct product $\prod_{n \in \mathbb{N}} \mathbb{Z}$ of infinitely many groups (or in this case copies of the same group) is much larger than the direct sum $\bigoplus_{n \in \mathbb{N}} \mathbb{Z}$, and in fact, the standard "Cantor diagonal trick" that is typically used for proving the uncountability of $\mathbb{R}$ implies that $\prod_{n \in \mathbb{N}} \mathbb{Z}$ is likewise an uncountable set. It follows that $\pi_1(H)$ itself is uncountable, whereas $\pi_1(X) \cong *_{n \in \mathbb{N}} \mathbb{Z}$, being generated by countably many countable groups, is countable.

14. Surfaces and torus knots (June 6, 2023)

We will discuss two applications of the Seifert-van Kampen theorem in this lecture: one to the study of surfaces, and the other to knots. Let's begin with surfaces.

Someday, when we talk about topological manifolds in this course (namely in Lecture 18), I will give you a precise mathematical definition of what the word "surface" means, but that day is not today. For now, we're just going to consider a class of specific examples that can be presented in a way that is convenient for computing their fundamental groups. A theorem we will discuss later in the semester implies that *all* compact surfaces can be presented in this way, but that is rather far from obvious.

We are going to consider pictures of polygons such as the following:

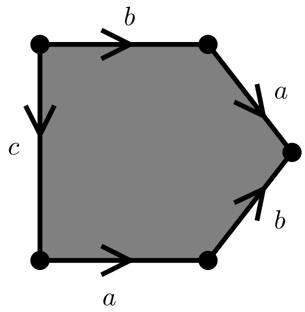

Suppose in general that $P \subset \mathbb{R}^2$ is a compact region bounded by some collection of N smooth curves that are arranged in a cyclic sequence with matching end points and do not intersect each other except at the matching end points. We will refer to these curves as *edges*, and label each of them with a letter a_i and an arrow. The letters $a_1, \dots, a_N$ need not all be distinct. We then define a topological space

$$X := P/\sim,$$

where the equivalence relation is trivial on the interior of P but identifies all vertices with each other, thus collapsing the set of vertices to a single point, and it also identifies any pair of edges labeled by the same letter with each other via a homeomorphism that matches the directions of the arrows. (The exact choice of this homeomorphism will not matter.) In the picture above, this means the two edges labeled with "a" get identified, and so do the two edges labeled with "b". (By the time you've read to the end of this lecture, you should be able to form a fairly clear picture of this surface in your mind, but I suggest reading somewhat further before you try this.)

Example 14.1. Take P to be a square whose sides have two labels a and b such that opposite sides of the square have matching letters and arrows pointing in the same direction. You could then build a physical model of $X = P/\sim$ in two steps: take a square piece of paper and bend it until you can tape together the two opposite sides labeled a, producing a cylinder. The two boundary components of this cylinder are circles labeled b, so if you were doing this with a sufficiently stretchable material (paper is not stretchable enough), you could then bend the cylinder around and tape together its two circular boundary components. The result is what's depicted in the picture at the right, a space conventionally known as the **2-torus** (or just "the torus" for short) and denoted by $\mathbb{T}^2$. It is homeomorphic to the product $S^1 \times S^1$.

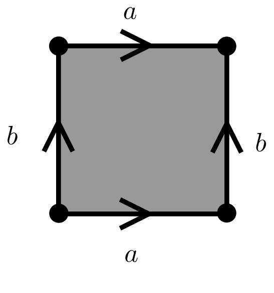

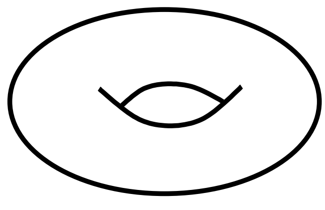

Example 14.2. If you relax your usual understanding of what a "polygon" is, you can also allow edges of the polygon to be curved as in the following example with only two edges:

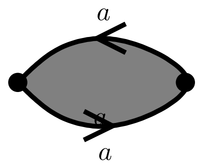

The polygon itself is homeomorphic to the disk $\mathbb{D}^2$, but identifying the two edges via a homeomorphism matching the arrows means we identify each point on $\partial\mathbb{D}^2$ with its antipodal point. The result matches the second description of $\mathbb{RP}^2$ that we saw in the first lecture, see Example 1.2.

THEOREM 14.3. *Suppose $X = P/\sim$ is a space defined as described above by a polygon P with N edges labeled by (possibly repeated) letters $a_1, \ldots, a_N$, where we are listing them in the order in which they appear as the boundary is traversed once counterclockwise. Let G denote the set of all letters that appear in this list, and for each $i = 1, \ldots, N$, write $p_i = 1$ if the arrow at edge i points counterclockwise around the boundary and $p_i = -1$ otherwise. Then $\pi_1(X)$ is isomorphic to the group with generators G and exactly one relation $a_1^{p_1} \ldots a_N^{p_N} = e$:*

$$\pi_1(X) \cong \left\{ G \mid a_1^{p_1} \ldots a_N^{p_N} = e \right\}.$$

PROOF. Let $P^1 := \partial P/\sim \subset X$. Since all vertices are identified to a point, P^1 is homeomorphic to a wedge sum of circles, one for each of the letters that appear as labels of edges, hence by an easy application of the Seifert-van Kampen theorem (cf. Exercise 13.2(b)),

$$\pi_1(P^1) \cong \pi_1(S^1) * \ldots * \pi_1(S^1) \cong \mathbb{Z} * \ldots * \mathbb{Z} = F_G,$$

the free group generated by the set G. Now decompose X into two open subsets A and B, where A is the interior of the polygon (not including its boundary) and B is an open neighborhood of P^1. We can arrange this so that $A \cap B$ is homeomorphic to an annulus $S^1 \times (-1, 1)$ occupying a neighborhood of ∂P in the interior of P, so for any choice of base point $p \in A \cap B$, $\pi_1(A \cap B, p) \cong \mathbb{Z}$ is generated by a loop that circles around parallel to ∂P. Since the neighborhood of ∂P admits a deformation retraction to ∂P, there is similarly a deformation retraction of B to P^1, giving $\pi_1(B, p) \cong \pi_1(P^1) = F_G$. Likewise, A is homeomorphic to an open disk, hence $\pi_1(A) = 0$. The Seifert-van Kampen theorem then idenifies $\pi_1(X, p)$ with a quotient of the free product $\pi_1(A, p) * \pi_1(B, p) \cong \pi_1(P^1) = F_G$, modulo the normal subgroup generated by the relation that if $j_A : A \cap B \hookrightarrow A$ and $j_B : A \cap B \hookrightarrow B$ denote the inclusion maps and $[\gamma] \in \pi_1(A \cap B, p) \cong \mathbb{Z}$ is a generator, then $(j_A)_*[\gamma] = (j_B)_*[\gamma]$. The left hand side of this equation is the trivial element since $\pi_1(A) = 0$. On the right hand side, we have the element of $\pi_1(B, p)$ represented by a loop $p \overset{\gamma}{\rightsquigarrow} p$ in the annulus $A \cap B$ that is parallel to the boundary of the polygon. Under the deformation retraction of $A \cap B$ to P^1, γ becomes the concatenated loop $a_1^{p_1} \ldots a_N^{p_N}$ defined by composing a traversal of ∂P with the quotient projection $\partial P \to P^1$, thus producing the relation $a_1^{p_1} \ldots a_N^{p_N} = e$. □

EXAMPLE 14.4. Applying the theorem to the torus in Example 14.1 gives

$$\pi_1(\mathbb{T}^2) \cong \{a, b \mid aba^{-1}b^{-1} = e\} = \{a, b \mid ab = ba\} \cong \mathbb{Z}^2.$$

Notice that this matches the result of applying Exercise 9.13(a), which gives $\pi_1(S^1 \times S^1) \cong \pi_1(S^1) \times \pi_1(S^1) \cong \mathbb{Z} \times \mathbb{Z}$.

EXAMPLE 14.5. For the picture of $\mathbb{RP}^2$ in Example 14.2, we obtain

$$\pi_1(\mathbb{RP}^2) \cong \{a \mid a^2 = e\} \cong \mathbb{Z}_2.$$

We already saw in Example 11.7 that $\pi_1(\mathbb{RP}^2)$ is generated by a single loop $\gamma : S^1 \to \mathbb{RP}^2$, the projection to $\mathbb{RP}^2 = S^2/\sim$ of a path that goes halfway around the equator of the sphere from one point to its antipodal point. We have now shown that $[\gamma]$ really is a nontrivial element of $\pi_1(\mathbb{RP}^2)$, but its square is trivial. The latter was also observed in Example 11.7, where it followed essentially from the fact that S^2 is simply connected: the concatenation of γ with itself is the projection to $\mathbb{RP}^2$ of a path that goes *all the way* around the equator in S^2, i.e. it is a loop, and can then be filled in with a map $\mathbb{D}^2 \to S^2$ since $\pi_1(S^2) = 0$. Composing the map $\mathbb{D}^2 \to S^2$ with the projection $S^2 \to \mathbb{RP}^2$ then contracts the loop γ^2 in $\mathbb{RP}^2$. However, we could not have deduced so easily from our knowledge of S^2 the fact that γ itself is *not* a contractible loop in $\mathbb{RP}^2$; that required the full strength of the Seifert-van Kampen theorem.

In Lecture 1, I drew you some pictures of topological spaces that I called "surfaces of genus g" for various values of a nonnegative integer g. I will now give you a precise definition of this space which, unfortunately, looks completely different from the original pictures, but we will soon see that it is equivalent.

DEFINITION 14.6. For any integer $g \geqslant 0$, the **closed orientable surface** Σ_g of **genus** (*Geschlecht*) g is defined to be S^2 if $g = 0$, and otherwise $\Sigma_g := P/\sim$ where P is a polygon with $4g$ edges labeled by $2g$ distinct letters $\{a_i, b_i\}_{i=1}^{g}$ in the order

$$a_1, b_1, a_1, b_1, a_2, b_2, a_2, b_2, \ldots, a_g, b_g, a_g, b_g,$$

such that the arrows point counterclockwise on the first instance of each letter in this sequence and clockwise on the second instance.

Once you've fully digested this definition, you may recognize that Σ_1 is defined by the square in Example 14.1, i.e. it is the torus $\mathbb{T}^2$. The diagram for Σ_2 is shown at the bottom of Figure 5. The projective plane $\mathbb{RP}^2$ is not an "orientable" surface, so it is not Σ_g for any g, though it is sometimes called a "non-orientable surface of genus 1". This terminology will make more sense when we later discuss the classification of surfaces.

In order to understand what Σ_g has to do with pictures we've seen before, we consider an operation on surfaces called the *connected sum*. It can be defined on any pair of surfaces Σ and Σ', or more generally, on any pair of n-dimensional topological manifolds, though for now we will consider only the case $n = 2$. Since I haven't yet actually given you precise definitions of the terms "surface" and "topological manifold," for now you should just assume Σ and Σ' come from the list of specific examples $\Sigma_0 = S^2$, $\Sigma_1 = \mathbb{T}^2$, $\Sigma_2, \Sigma_3, \ldots$ defined above.

Given a pair of inclusions $\mathbb{D}^2 \hookrightarrow \Sigma$ and $\mathbb{D}^2 \hookrightarrow \Sigma'$, the **connected sum** (*zusammenhängende Summe*) of Σ and Σ' is defined as the space

$$\Sigma \# \Sigma' := \left(\Sigma \backslash \mathring{\mathbb{D}}^2\right) \cup_{S^1} \left(\Sigma' \backslash \mathring{\mathbb{D}}^2\right).$$

The result of this operation is not hard to visualize in many concrete examples, see e.g. Figure 6.

More generally, for topological n-manifolds M and M', one defines the connected sum $M \# M'$ by choosing inclusions of $\mathbb{D}^n$ into M and M', then removing the interiors of these disks and gluing together $M \backslash \mathring{\mathbb{D}}^n$ and $M' \backslash \mathring{\mathbb{D}}^n$ along $S^{n-1} = \partial \mathbb{D}^n$. The notation $M \# M'$ obscures the fact that the definition of the connected sum depends explicitly on choices of inclusions of $\mathbb{D}^n$ into both spaces, and it is not entirely true in general that $M \# M'$ up to homeomorphism is independent of this choice. It is true however for surfaces:

LEMMA 14.7 (slightly nontrivial). *Up to homeomorphism, the connected sum $\Sigma \# \Sigma'$ of two closed connected surfaces Σ and Σ' does not depend on the choices of inclusions $\mathbb{D}^2 \hookrightarrow \Sigma$ and $\mathbb{D}^2 \hookrightarrow \Sigma'$.*

SKETCH OF A PROOF. A complete proof of this would be too much of a digression and require more knowledge about the classification of surfaces than is presently safe to assume, but I can

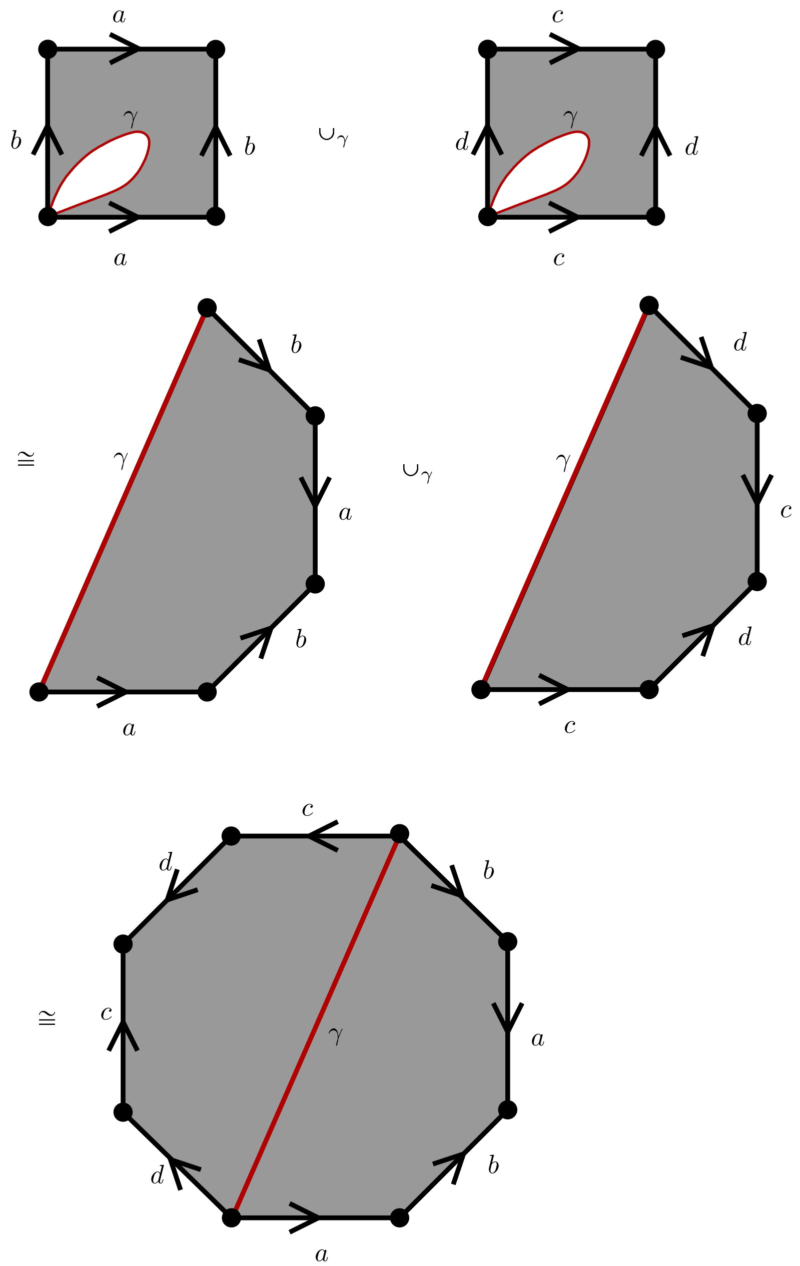

FIGURE 5. The connected sum $\mathbb{T}^2 \# \mathbb{T}^2$ is formed by cutting holes $\mathbb{D}^2$ out of two copies of $\mathbb{T}^2$ along some loop γ, and then gluing together the two copies of $\mathbb{T}^2 \backslash \mathbb{D}^2$. The result is Σ_2, the closed orientable surface of genus 2.

give the rough idea. The main thing you need to believe is that "up to orientation" (I'll come back to that detail in a moment), any inclusion $i_0 : \mathbb{D}^2 \hookrightarrow \Sigma$ can be deformed into any other inclusion $i_1 : \mathbb{D}^2 \hookrightarrow \Sigma$ through a continuous family of inclusions $i_t : \mathbb{D}^2 \hookrightarrow \Sigma$ for $t \in I$. You should imagine this roughly as follows: since $\mathbb{D}^2$ is homeomorphic via the obvious rescalings to the disk

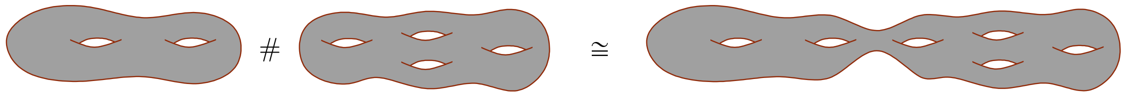

FIGURE 6. The connected sum of two surfaces is defined by cutting a hole out of each of them and gluing the rest together along the resulting boundary circle.

$\mathbb{D}^2_r$ of radius r for every $r > 0$, one can first deform i_0 and i_1 to inclusions whose images lie in arbitrarily small neighborhoods of two points $z_0, z_1 \in \Sigma$. Now since Σ is connected (and therefore also path-connected, as all topological manifolds are locally path-connected), we can choose a path γ from z_0 to z_1, and the idea is then to define i_t as a continuous family of inclusions $\mathbb{D}^2 \hookrightarrow \Sigma$ such that the image of i_t lies in an arbitrarily small neighborhood of $\gamma(t)$ for each t. You should be able to imagine concretely how to do this in the special case $\Sigma = \mathbb{R}^2$. That it can be done on arbitrary connected surfaces Σ depends on the fact that every point in Σ has a neighborhood homeomorphic to $\mathbb{R}^2$ (in other words, Σ is a topological 2-manifold).

Now for the detail that was brushed under the rug in the previous paragraph: even if $i_0, i_1 : \mathbb{D}^2 \hookrightarrow \Sigma$ are two inclusions that send 0 to the same point $z \in \Sigma$ and have images in an arbitrarily small neighborhood of z, it is not always true that i_0 can be deformed to i_1 through a continuous family of inclusions. For example, if we take $\Sigma = \mathbb{R}^2$, it is not true for the two inclusions $i_0, i_1 : \mathbb{D}^2 \hookrightarrow \mathbb{R}^2$ defined by $i_0(x, y) = (\epsilon x, \epsilon y)$ and $i_1(x, y) = (\epsilon x, -\epsilon y)$. In this example, both inclusions are defined as restrictions of injective linear maps $\mathbb{R}^2 \to \mathbb{R}^2$, but one has positive determinant and the other has negative determinant, so one cannot deform from one to the other through injective linear maps. One can use the technology of *local homology groups* (which we'll cover next semester) to remove the linearity from this argument and show that there also is no deformation from i_0 to i_1 through continuous inclusions. The issue here is one of *orientations*: i_0 is an orientation-preserving map, while i_1 is orientation-reversing. It turns out that two inclusions of $\mathbb{D}^2$ into $\mathbb{R}^2$ can be deformed to each other through inclusions if and only if they are either both orientation preserving or both orientation reversing. This obstruction sounds like bad news for our proof, but the situation is saved by the following corollary of the classification of surfaces: every closed orientable surface admits an orientation-reversing homeomorphism to itself. For example, if you picture the torus as the usual tube embedded in $\mathbb{R}^3$ and you embed it so that it is *symmetric* about some 2-dimensional coordinate plane, then the linear reflection through that plane restricts to a homeomorphism of $\mathbb{T}^2$ that is orientation reversing. Once we see what all the other closed orientable surfaces look like, it will be easy to see that one can do that with all of them. Actually, it is also not so hard to see this for the surfaces Σ_g defined as polygons: you just need to choose a sufficiently clever axis in the plane containing the polygon and reflect across it. Once this is understood, you realize that the orientation of your inclusion $\mathbb{D}^2 \hookrightarrow \Sigma$ does not really matter, as you can always replace it with an inclusion having the opposite orientation, and the picture you get in the end will be homeomorphic to the original.

With this detail out of the way, you just have to convince yourself that if you have a pair of continuous families of inclusions $i_t : \mathbb{D}^2 \hookrightarrow \Sigma$ and $j_t : \mathbb{D}^2 \hookrightarrow \Sigma'$ defined for $t \in [0, 1]$, then the resulting glued surfaces

$$\Sigma \#_t \Sigma' := \left(\Sigma \backslash i_t(\mathring{\mathbb{D}}^2)\right) \cup_{S^1} \left(\Sigma' \backslash j_t(\mathring{\mathbb{D}}^2)\right)$$

are homeomorphic for all t. It suffices in fact to prove that this is true just for t varying in an arbitrarily small interval $(t_0 - \epsilon, t_0 + \epsilon)$, since $[0, 1]$ is compact and can therefore be covered by finitely many such intervals. A homeomorphism $\Sigma \#_t \Sigma' \to \Sigma \#_s \Sigma'$ for $t \neq s$ is easy to define if we can first find a homeomorphism $\Sigma \to \Sigma$ that sends $i_t(z) \mapsto i_s(z)$ for every $z \in \mathbb{D}^2$ and similarly on Σ'. This is not hard to construct if t and s are sufficiently close. $\square$

Now we are in a position to relate Σ_g with the more familiar pictures of surfaces.

THEOREM 14.8. *For any nonnegative integers g, h, $\Sigma_g \# \Sigma_h \cong \Sigma_{g+h}$. In particular, Σ_g is the connected sum of g copies of the torus:*

$$\Sigma_g \cong \underbrace{\mathbb{T}^2 \# \ldots \# \mathbb{T}^2}_{g}$$

PROOF. The result becomes obvious if one makes a sufficiently clever choice of hole to cut out of Σ_g and Σ_h, and Lemma 14.7 tells us that the resulting space up to homeomorphism is independent of this choice. The example of $g = h = 1$ is shown in Figure 5, and the same idea works (but is more effort to draw) for any values of g and h. $\square$

Now that we know how to draw pretty pictures of the surfaces Σ_g, we can also observe that we have already proved something quite nontrivial about them: we have computed their fundamental groups!

COROLLARY 14.9 (of Theorem 14.3). *The closed orientable surface Σ_g of genus $g \geqslant 0$ has a fundamental group with $2g$ generators and one relation, namely*

$$\pi_1(\Sigma_g) \cong \left\{ a_1, b_1, \ldots, a_g, b_g \;\middle|\; a_1 b_1 a_1^{-1} b_1^{-1} a_2 b_2 a_2^{-1} b_2^{-1} \ldots a_g b_g a_g^{-1} b_g^{-1} = e \right\}.$$

$\square$

Using the commutator notation from Exercise 12.21, the relation in Corollary 14.9 can be conveniently abbreviated as

$$\prod_{i=1}^{g} [a_i, b_i] = e.$$

EXERCISE 14.10. Show that the abelianization (cf. Exercise 12.21) of $\pi_1(\Sigma_g)$ is isomorphic to the additive group $\mathbb{Z}^{2g}$.
Hint: $\pi_1(\Sigma_g)$ is a particular quotient of the free group on $2g$ generators. Observe that the abelianization of that free group is identical to the abelianization of $\pi_1(\Sigma_g)$. (Why?)

By the classification of finitely generated abelian groups, $\mathbb{Z}^m$ and $\mathbb{Z}^n$ are never isomorphic unless $m = n$, so Exercise 14.10 implies that $\pi_1(\Sigma_g)$ and $\pi_1(\Sigma_h)$ are not isomorphic unless $g = h$. This completes the first step in the classification of closed surfaces:

COROLLARY 14.11. *For two nonnegative integers $g \neq h$, Σ_g and Σ_h are not homeomorphic.* $\square$

EXERCISE 14.12. Assume X and Y are path-connected topological manifolds of dimension n.

(a) Use the Seifert-Van Kampen theorem to show that if $n \geqslant 3$, then $\pi_1(X \# Y) \cong \pi_1(X) * \pi_1(Y)$. Where does your proof fail in the cases $n = 1$ and $n = 2$?
(b) Show that the formula of part (a) is false in general for $n = 1, 2$.

EXERCISE 14.13. For integers $g, m \geqslant 0$, let $\Sigma_{g,m}$ denote the compact surface obtained by cutting m disjoint disk-shaped holes out of the closed orientable surface with genus g. (By this convention, $\Sigma_g = \Sigma_{g,0}$.) The boundary $\partial \Sigma_{g,m}$ is then a disjoint union of m circles, e.g. the case with $g = 1$ and $m = 3$ is shown in Figure 7.

(a) Show that $\pi_1(\Sigma_{g,1})$ is a free group with $2g$ generators, and if $g \geqslant 1$, then any simple closed curve parametrizing $\partial \Sigma_{g,1}$ represents a nontrivial element of $\pi_1(\Sigma_{g,1})$.[18]
Hint: Think of Σ_g as a polygon with some of its edges identified. If you cut a hole in the middle of the polygon, what remains admits a deformation retraction to the edges. Prove it with a picture.

[18] Terminology: one says in this case that $\partial \Sigma_{g,1}$ is **homotopically nontrivial** or **essential**, or equivalently, $\partial \Sigma_{g,1}$ is not **nullhomotopic**.

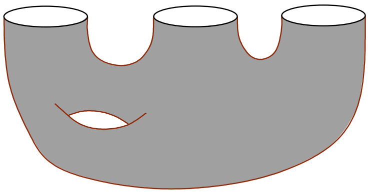

FIGURE 7. The surface $\Sigma_{1,3}$ as in Exercise 14.13.

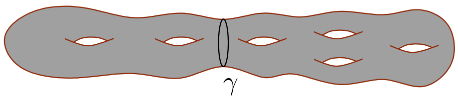

(b) Assume γ is a simple closed curve separating Σ_g into two pieces homeomorphic to $\Sigma_{h,1}$ and $\Sigma_{k,1}$ for some $h, k \geqslant 0$. (The picture at the right shows an example with $h = 2$ and $k = 4$.) Show that the image of $[\gamma] \in \pi_1(\Sigma_g)$ under the natural projection to the abelianization of $\pi_1(\Sigma_g)$ is trivial.
Hint: What does γ look like in the polygonal picture from part (a)? What is it homotopic to?

(c) Prove that if $g \geqslant 2$ and G denotes the group $\{a_1, b_1, \ldots, a_g, b_g \mid \prod_{i=1}^g [a_i, b_i] = e\}$, then for any proper subset $J \subset \{1, \ldots, g\}$, $\prod_{i \in J}[a_i, b_i]$ is a nontrivial element of G.
Hint: Given $j \in J$ and $\ell \in \{1, \ldots, g\} \backslash J$, there is a homomorphism $\Phi : F_{\{a_1, b_1, \ldots, a_g, b_g\}} \to F_{\{x,y\}}$ that sends $a_j \mapsto x$, $b_j \mapsto y$, $a_\ell \mapsto y$, $b_\ell \mapsto x$ and maps all other generators to the identity. Show that Φ descends to the quotient G and maps $\prod_{i \in J}[a_i, b_i] \in G$ to something nontrivial.

(d) Deduce from part (c) that if $h > 0$ and $k > 0$, then the curve γ in part (b) represents a nontrivial element of $\pi_1(\Sigma_g)$.

(e) Generalize part (a): show that if $m \geqslant 1$, $\pi_1(\Sigma_{g,m})$ is a free group with $2g + m - 1$ generators.

Now let's talk about knots. Back in Lecture 8, I showed you two simple examples of knots $K \subset \mathbb{R}^3$: the *trefoil* and the *unknot*. I claimed that it is impossible to deform one of these knots into the other, and in fact that the complements of both knots in $\mathbb{R}^3$ are not homeomorphic. It is time to prove this.

We will consider both as special cases of a more general class of knots called *torus knots*. Fix the standard embedding of the torus

$$f : \mathbb{T}^2 = S^1 \times S^1 \hookrightarrow \mathbb{R}^3,$$

where by "standard," I mean the one that you usually picture when you imagine a torus embedded in $\mathbb{R}^3$ (see the surface bounding the grey region in Figure 9). Given any two relatively prime integers $p, q \in \mathbb{Z}$, the (p, q)**-torus knot** is defined by

$$K_{p,q} := \left\{ f(e^{pi\theta}, e^{qi\theta}) \mid \theta \in \mathbb{R} \right\} \subset \mathbb{R}^3.$$

In other words, $K_{p,q}$ is a knot lying on the image of the embedded torus $f(\mathbb{T}^2) \subset \mathbb{R}^3$, obtained from a loop that rotates p times around one of the dimensions of $\mathbb{T}^2 = S^1 \times S^1$ while rotating q times around the other. It is conventional to assume p and q are relatively prime, since the definition of $K_{p,q}$ above would not change if both p and q were multiplied by the same nonzero constant.

EXAMPLE 14.14. $K_{2,3}$ is the trefoil knot (Figure 8, left).

EXAMPLE 14.15. $K_{1,0}$ is the unknot (Figure 8, right).

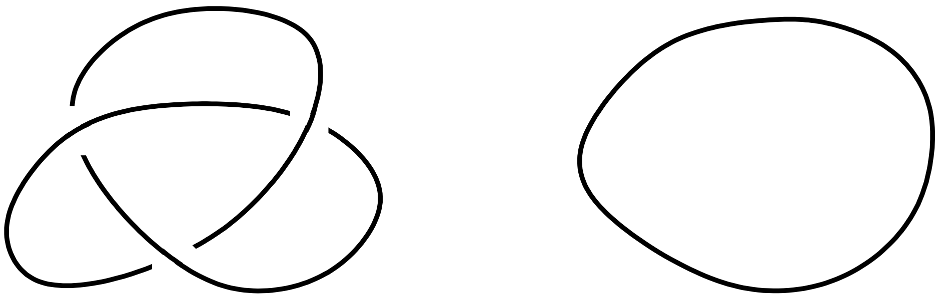

$K_{2,3} \subset \mathbb{R}^3$ $K_{1,0} \subset \mathbb{R}^3$

FIGURE 8. The trefoil knot $K_{2,3}$ and unknot $K_{1,0}$.

The **knot group** of a knot $K \subset \mathbb{R}^3$ is defined as the fundamental group of the so-called *knot complement*, $\pi_1(\mathbb{R}^3 \backslash K)$. We saw in Example 12.7 that the natural inclusion $\mathbb{R}^3 \hookrightarrow S^3$ defined by identifying S^3 with the one-point compactification $\mathbb{R}^3 \cup \{\infty\}$ induces an isomorphism of $\pi_1(\mathbb{R}^3 \backslash K)$ to $\pi_1(S^3 \backslash K)$, thus in order to compute knot groups, we may as well regard the knot $K \subset \mathbb{R}^3$ as a subset of the slightly larger but *compact* space S^3 and compute $\pi_1(S^3 \backslash K)$. We shall now answer the question: given relatively prime integers p and q, what is $\pi_1(S^3 \backslash K_{p,q})$?

Here is a useful trick for picturing S^3. By definition, $S^3 = \partial \mathbb{D}^4$, but notice that $\mathbb{D}^4$ is also homeomorphic to the "box" $\mathbb{D}^2 \times \mathbb{D}^2$, whose boundary consists of the two pieces $\partial \mathbb{D}^2 \times \mathbb{D}^2$ and $\mathbb{D}^2 \times \partial \mathbb{D}^2$, intersecting each other along $\partial \mathbb{D}^2 \times \partial \mathbb{D}^2$. The latter is a copy of $\mathbb{T}^2$, and the pieces $S^1 \times \mathbb{D}^2$ and $\mathbb{D}^2 \times S^1$ are called **solid tori** since we usually picture them as the region in $\mathbb{R}^3$ bounded by the standard embedding of the torus. The homeomorphism $\mathbb{D}^4 \cong \mathbb{D}^2 \times \mathbb{D}^2$ thus allows us to identify S^3 with the space constructed by gluing together these two solid tori along the obvious identification of their boundaries:

$$S^3 \cong (S^1 \times \mathbb{D}^2) \cup_{\mathbb{T}^2} (\mathbb{D}^2 \times S^1).$$

A picture of this decomposition is shown in Figure 9. Here the 2-torus along which the two solid tori are glued together is depicted as the standard embedding of $\mathbb{T}^2$ in $\mathbb{R}^3$, so this is where we will assume $K_{p,q}$ lies. The region bounded by this torus is $S^1 \times \mathbb{D}^2$, shown in the picture as an S^1-parametrized family of disks $\mathbb{D}^2$. It requires a bit more imagination to recognize $\mathbb{D}^2 \times S^1$ in the picture: instead of a family of disks, we have drawn it as a $\mathbb{D}^2$-parametrized family of circles, where it is important to understand that one of those circles passes through $\infty \in S^3$ and thus looks like a line instead of a circle in the picture. This picture will now serve as the basis for a Seifert-van Kampen decomposition of $S^3 \backslash K_{p,q}$ into two open subsets. They will be defined as open neighborhoods of the two subsets

$$A_0 := (S^1 \times \mathbb{D}^2) \backslash K_{p,q}, \qquad B_0 := (\mathbb{D}^2 \times S^1) \backslash K_{p,q}.$$

In order to define suitable neighborhoods, let us identify a neighborhood of $f(\mathbb{T}^2)$ in $\mathbb{R}^3$ with $(-1,1) \times \mathbb{T}^2$ such that $f(\mathbb{T}^2)$ becomes $\{0\} \times \mathbb{T}^2 \subset \mathbb{R}^3$. We then define

$$A := \left(S^1 \times \mathring{\mathbb{D}}^2\right) \cup \left((-1,1) \times (\mathbb{T}^2 \backslash f^{-1}(K_{p,q}))\right),$$

and

$$B := \left(\mathring{\mathbb{D}}^2 \times S^1\right) \cup \left((-1,1) \times (\mathbb{T}^2 \backslash f^{-1}(K_{p,q}))\right).$$

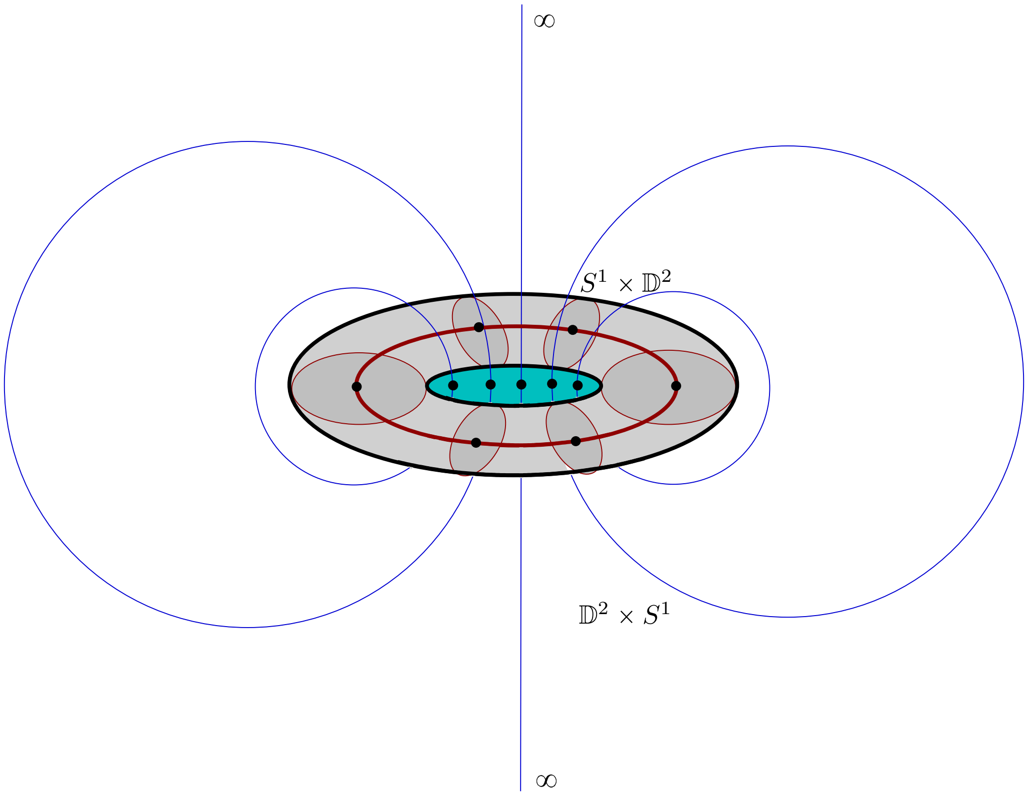

FIGURE 9. The sphere $S^3 = \mathbb{R}^3 \cup \{\infty\}$ decomposed as a union of two solid tori whose common boundary is the "standard" embedding of $\mathbb{T}^2$ in $\mathbb{R}^3$: $S^3 \cong \partial(\mathbb{D}^2 \times \mathbb{D}^2) = (S^1 \times \mathbb{D}^2) \cup_{\mathbb{T}^2} (\mathbb{D}^2 \times S^1)$. The vertical blue line passing through the middle is actually a circle in S^3 passing through the point at ∞.

By contracting the interval $(-1,1)$, we can define a deformation retraction of A to A_0 and then retract further by contractng the disk $\mathbb{D}^2$ to its center, eventually producing a deformation retraction of A to the circle $S^1 \times \{0\}$ at the center of the inner solid torus—this is the red circle in Figure 9 that passes through the center of each disk. In an analogous way, there is a deformation retraction of B to the center $\{0\} \times S^1$ of the outer solid torus, which is the blue line through ∞ in the picture, though you might prefer to perturb this to one of the parallel circles $\{z\} \times S^1 \subset \mathbb{D}^2 \times S^1$ for $z \neq 0$, since these actually look like circles in the picture. We can now regard $\pi_1(A)$ and $\pi_1(B)$ as separate copies of the integers whose generators we shall call a and b respectively,

$$\pi_1(A) \cong \{a \mid \varnothing\}, \qquad \pi_1(B) \cong \{b \mid \varnothing\}.$$

The intersection is

$$A \cap B = (-1,1) \times \left(\mathbb{T}^2 \backslash f^{-1}(K_{p,q})\right) \underset{h.e.}{\simeq} \mathbb{T}^2 \backslash f^{-1}(K_{p,q}) \underset{h.e.}{\simeq} S^1.$$

That last homotopy equivalence deserves an explanation: if you draw $\mathbb{T}^2$ as a square with its sides identified, then $f^{-1}(K_{p,q})$ looks like a straight line that periodically exits one side of the square and reappears at the opposite side. Now draw another straight path parallel to this one (I recommend using a different color), and you will easily see that after removing $f^{-1}(K_{p,q})$ from $\mathbb{T}^2$, what remains admits a deformation retraction to the parallel path, which is an embedded copy of S^1. We will call the generator of its fundamental group c,

$$\pi_1(A \cap B) \cong \{c \mid \varnothing\}.$$

According to the Seifert-van Kampen theorem (in particular Corollary 12.19, the version for finitely-presented groups), we can now write

$$\pi_1(S^3 \backslash K_{p,q}) \cong \left\{a, b \;\middle|\; (j_A)_* c = (j_B)_* c\right\},$$

where j_A and j_B denote the inclusions of $A \cap B$ into A and B respectively. To interpret this properly, we should choose a base point in $A \cap B$ and picture a, b and c as represented by specific loops through this base point, so without loss of generality, a is a loop near the boundary $\mathbb{T}^2$ of $S^1 \times \mathbb{D}^2$ that wraps once around the S^1 direction, and b is another loop near $\mathbb{T}^2$ that wraps once around the S^1-direction of $\mathbb{D}^2 \times S^1$, which is the other dimension of $\mathbb{T}^2 = S^1 \times S^1$. The interesting part is c, as it is represented by a loop in $\mathbb{T}^2$ that is parallel to $K_{p,q}$, thus it wraps p times around the direction of a and q times around the direction of b. This means $(j_A)_* c = a^p$ and $(j_B)_* c = b^q$, so putting all of this together yields:

THEOREM 14.16. $\pi_1(S^3 \backslash K_{p,q}) \cong \{a, b \mid a^p = b^q\}$. □

EXAMPLE 14.17. For $(p, q) = (1, 0)$, we obtain the knot group of the unknot: $\pi_1(S^3 \backslash K_{1,0}) \cong \{a, b \mid a = e\} = \{b \mid \varnothing\} = \mathbb{Z}$. In particular, this is an abelian group.

EXAMPLE 14.18. The knot group of the trefoil is $\pi_1(S^3 \backslash K_{2,3}) \cong \{a, b \mid a^2 = b^3\}$. We proved in Exercise 12.20 that this group is not abelian, in contrast to Example 14.17, hence $\pi_1(S^3 \backslash K_{2,3})$ and $\pi_1(S^3 \backslash K_{1,0})$ are not isomorphic.

COROLLARY 14.19. *The knot complements $\mathbb{R}^3 \backslash K_{1,0}$ and $\mathbb{R}^3 \backslash K_{2,3}$ are not homeomorphic.* □

Before moving on[19] from the Seifert-van Kampen theorem, I would like to sketch one more application, which answers the question, "which groups can be fundamental groups of nice spaces?" If we are only interested in finitely-presented groups and decide that "nice" should mean "compact and Hausdorff", then the answer turns out to be that there is no restriction at all.

THEOREM 14.20. *Every finitely-presented group is the fundamental group of some compact Hausdorff space.*

PROOF. The following lemma will be used as an inductive step. Suppose X_0 is a compact Hausdorff space with a finitely-presented fundamental group

$$\pi_1(X_0, p) \cong \left\{\{a_i\} \;\middle|\; \{R_j\}\right\}.$$

Then for any loop $\gamma : (S^1, 1) \to (X_0, p)$, we claim that the space

$$X := \mathbb{D}^2 \cup_\gamma X_0 := \left(\mathbb{D}^2 \amalg X_0\right) \Big/ z \sim \gamma(z) \in X_0 \text{ for all } z \in \partial\mathbb{D}^2$$

is compact and Hausdorff with

$$\pi_1(X, p) \cong \left\{\{a_i\} \;\middle|\; \{R_j\},\ [\gamma] = e\right\},$$

i.e. its fundamental group has the same generators and one new relation, defined by setting $[\gamma] \in \pi_1(X_0, p)$ equal to the trivial element. This claim follows easily[20] from the Seifert-van Kampen

[19] We ran out of time in the actual lecture before we could talk about Theorem 14.20, but I am including it in the notes just because it is interesting.

[20] I am glossing over the detail where we need to prove that X is also compact and Hausdorff. This is not completely obvious, but it is yet another exercise in point-set topology that I feel justified in not explaining now that that portion of the course is finished.

theorem using the decomposition $X = A \cup B$ where $A = \mathring{\mathbb{D}}^2$ and B is an open neighborhood of X_0 obtained by adding a small annulus near the boundary of $\partial\mathbb{D}^2$. Since the annulus admits a deformation retraction to $\partial\mathbb{D}^2$, we have $B \underset{h.e.}{\simeq} X_0$, while $A \cap B \underset{h.e.}{\simeq} S^1$ and A is contractible. According to Corollary 12.19, $\pi_1(X,p)$ then inherits all the generators and relations of $\pi_1(B) \cong \pi_1(X_0)$, no new generators from $\pi_1(A) = 0$, and one new relation from the generator of $\pi_1(A \cap B) \cong \mathbb{Z}$, whose inclusion into A is trivial, so the relation says that its inclusion into B must become the trivial element. That inclusion is precisely $[\gamma] \in \pi_1(X_0,p)$, hence the claim is proved.

Now suppose G is a finitely-presented group with generators $x_1,\dots,x_N$ and relations $w_1 = e,\dots,w_m = e$ for $w_i \in F_{\{x_1,\dots,x_N\}}$. We start with a space X_0 whose fundamental group is the free group on $\{x_1,\dots,x_N\}$: the wedge sum of N circles will do. As the previous paragraph demonstrates, we can then attach a 2-disk for each individual relation we would like to add to the fundamental group, and doing this finitely many times produces a compact Hausdorff space with the desired fundamental group. □

15. Covering spaces and the lifting theorem (June 8, 2023)

We now leave the Seifert-van Kampen theorem behind and introduce the second major tool for computing fundamental groups: the theory of covering spaces.

Definition 15.1. A map $f : Y \to X$ is called a **covering map** (*Überlagerung*), or simply a **cover** of X, if for every $x \in X$, there exists an open neighborhood $\mathcal{U} \subset X$ such that

$$f^{-1}(\mathcal{U}) = \bigcup_{\alpha\in J} \mathcal{V}_\alpha$$

for a collection of disjoint open subsets $\{\mathcal{V}_\alpha \subset Y\}_{\alpha\in J}$ such that $f|_{\mathcal{V}_\alpha} : \mathcal{V}_\alpha \to \mathcal{U}$ is a homeomorphism for each $\alpha \in J$. The domain Y of this map is called a **covering space** (*Überlagerungsraum*) of X. Any subset $\mathcal{U} \subset X$ satisfying the conditions stated above is said to be **evenly covered**.

Example 15.2. The map $f : \mathbb{R} \to S^1 : \theta \mapsto e^{i\theta}$ is a covering map of S^1.

Example 15.3. The map $S^1 \to S^1$ sending $e^{i\theta}$ to $e^{ki\theta}$ for any nonzero $k \in \mathbb{Z}$ is also a covering map of S^1.

Example 15.4. The n-**dimensional torus** $\mathbb{T}^n := \underbrace{S^1 \times \ldots \times S^1}_{n}$ admits a covering map

$$\mathbb{R}^n \to \mathbb{T}^n : (\theta_1,\dots,\theta_n) \mapsto (e^{i\theta_1},\dots,e^{i\theta_n}).$$

More generally, it is straightforward to show that given any two covering maps $f_i : Y_i \to X_i$ for $i = 1,2$, there is a "product" cover

$$Y_1 \times Y_2 \xrightarrow{f_1\times f_2} X_1 \times X_2 : (x_1,x_2) \mapsto (f_1(x_1), f_2(x_2)).$$

Example 15.5. For any space X, the identity map $X \to X$ is trivially a covering map.

Example 15.6. Another trivial example of a covering map can be defined for any space X and any set J by setting $X_\alpha := X$ for every $\alpha \in J$ and defining $f : \coprod_{\alpha\in J} X_\alpha \to X$ as the unique map that restricts to each $X_\alpha = X$ as the identity map on X. This is a *disconnected* covering map. We will usually restrict our attention to covering spaces that are connected.

Example 15.7. For each $n \in \mathbb{N}$, the quotient projection $S^n \to \mathbb{RP}^n = S^n/\!\sim$ is a covering map.

Theorem 15.8. *If X is connected and $f : Y \to X$ is a cover, then the number (finite or infinite) of points in $f^{-1}(x) \subset Y$ does not depend on the choice of a point $x \in X$.*

Proof. Given $x \in X$, choose an evenly covered neighborhood $\mathcal{U} \subset X$ of x and write $f^{-1}(\mathcal{U}) = \bigcup_{\alpha \in J} \mathcal{V}_\alpha$. Then for every $y \in \mathcal{U}$, $|f^{-1}(y)| = |J|$, and it follows that for every $n \in \{0, 1, 2, 3, \ldots, \infty\}$, the subset $X_n := \{x \in X \mid |f^{-1}(x)| = n\} \subset X$ is open. If $x \in X_n$, notice that $\bigcup_{m \neq n} X_m$ is also open, thus X_n is also closed, so connectedness implies $X_n = X$. $\square$

In the setting of the above theorem, the number of points in $f^{-1}(x)$ is called the **degree** (*Grad*) of the cover. If $\deg(f) = n$, we sometimes call f an n**-fold** cover.

Examples 15.9. The cover $S^1 \to S^1 : z \mapsto z^k$ from Example 15.3 has degree $|k|$, while the quotient projection $S^n \to \mathbb{RP}^n$ has degree 2 and the cover $\mathbb{R} \to S^1$ from Example 15.2 has infinite degree.

Remark 15.10. Some authors strengthen the definition of a covering map $f : Y \to X$ by requiring f to be surjective. We did not require this in Definition 15.1, but notice that if X is connected, then it follows immediately from Theorem 15.8. In practice, it is only sensible to consider covers of connected spaces, and we shall always assume connectedness.

Note that in Definition 15.1, one should explicitly require the sets $\mathcal{V}_\alpha \subset f^{-1}(\mathcal{U})$ to be open. This is important, as part of the point of that definition is that X can be covered by open neighborhoods $\mathcal{U}$ whose preimages are homeomorphic to *disjoint unions* of copies of $\mathcal{U}$, i.e.

$$f^{-1}(\mathcal{U}) \cong \coprod_{\alpha \in J} \mathcal{U}.$$

This is true specifically because each of the sets $\mathcal{V}_\alpha$ is open, and therefore (as the complement of $\bigcup_{\beta \neq \alpha} \mathcal{V}_\beta$) also closed in $f^{-1}(\mathcal{U})$. To put it another way, in a covering map, every point $x \in X$ has a neighborhood $\mathcal{U}$ such that $f^{-1}(\mathcal{U})$ is the disjoint union of homeomorphic neighborhoods of the individual points in $f^{-1}(x)$. An important consequence of this definition is that every covering map $f : Y \to X$ is also a *local homeomorphism*, meaning that for each $y \in Y$ and $x := f(y)$, f maps some neighborhood of y homeomorphically to some neighborhood of x.

Almost every result in covering space theory is based on the answer to the following question: given a map $f : A \to X$ and a covering map $p : Y \to X$, can f be "lifted" to a map $\tilde{f} : A \to Y$ satisfying $p \circ \tilde{f} = f$? This problem can be summarized with the diagram

$$\begin{tikzcd} & Y \arrow[d, "p"] \\ A \arrow[ur, dashed, "\tilde{f}"] \arrow[r, "f"] & X \end{tikzcd} \tag{15.1}$$

in which the maps f and p are given, but the dashed arrow for $\tilde{f}$ indicates that we do not know whether such a map exists. If it does, then we call $\tilde{f}$ a **lift** of f to the cover. It is easy to see that lifts do not always exist: take for instance the cover $p : \mathbb{R} \to S^1 : \theta \mapsto e^{i\theta}$ and let $f : S^1 \to S^1$ be the identity map. A lift $\tilde{f} : S^1 \to \mathbb{R}$ would need to associate to every $e^{i\theta} \in S^1$ some point $\phi := \tilde{f}(e^{i\theta})$ such that $e^{i\phi} = e^{i\theta}$. It is easy to define a function that does this, but can we make it *continuous*? If it were continuous, then $\tilde{f}(e^{i\theta})$ would have to increase by 2π as $e^{i\theta}$ turns around the circle from $\theta = 0$ to $\theta = 2\pi$, producing *two* values $\tilde{f}(e^{2\pi i}) = \tilde{f}(1) + 2\pi$ even though $e^{2\pi i} = 1$. The goal for the remainder of this lecture is to determine precisely which maps can be lifted to which covering spaces and which cannot.

We start with the following observation: choose base points $a \in A$ and $x \in X$ to make $f : (A, a) \to (X, x)$ into a pointed map. Then if a lift $\tilde{f} : A \to Y$ exists and we set $y := \tilde{f}(a)$ to make $\tilde{f}$ a pointed map, p now becomes one as well since $p(y) = p(\tilde{f}(a)) = f(a) = x$, hence (15.1) becomes a diagram of pointed maps and induces a corresponding diagram of group homomorphisms

$$\begin{tikzcd} & \pi_1(Y,y) \arrow[d, "p_*"] \\ \pi_1(A,a) \arrow[ur, "\tilde{f}_*"] \arrow[r, "f_*"] & \pi_1(X,x). \end{tikzcd} \tag{15.2}$$

The existence of this diagram implies a nontrivial condition that relates the homomorphisms f_* and p_* but has nothing intrinsically to do with the lift: it implies $\operatorname{im} f_* \subset \operatorname{im} p_*$, i.e. these are two subgroups of $\pi_1(X,x)$, and one of them must be contained in the other. The lifting theorem states that under some assumptions that are satisfied by most reasonable spaces, this necessary condition is also sufficient.

THEOREM 15.11 (lifting theorem). *Assume X, Y, A are all path-connected spaces, A is also locally path-connected, $p : Y \to X$ is a covering map and $f : (A, a_0) \to (X, x_0)$ is a base-point preserving map. Then for any choice of base point $y_0 \in f^{-1}(x_0) \subset Y$, f admits a base-point preserving lift $\tilde{f} : (A, a_0) \to (Y, y_0)$ if and only if*

$$f_*\left(\pi_1(A, a_0)\right) \subset p_*\left(\pi_1(Y, y_0)\right),$$

and the point $y_0 = \tilde{f}(a_0)$ uniquely determines the lift $\tilde{f}$.

Let us discuss some applications before we get to the proof.

COROLLARY 15.12. *For any covering map $p : Y \to X$ between path-connected spaces and any space A that is simply connected and locally path-connected, every map $f : A \to X$ can be lifted to Y.* □

COROLLARY 15.13. *For every base-point preserving covering map $p : (Y, y_0) \to (X, x_0)$ between path-connected spaces, the homomorphism $p_* : \pi_1(Y, y_0) \to \pi_1(X, x_0)$ is injective.*

PROOF. Suppose $\tilde{\gamma} : (S^1, 1) \to (Y, y_0)$ is a loop such that $p_*[\tilde{\gamma}] = e \in \pi_1(X, x_0)$. Then $\gamma := p \circ \tilde{\gamma} : (S^1, 1) \to (X, x_0)$ admits an extension $u : (\mathbb{D}^2, 1) \to (X, x_0)$ with $u|_{\partial \mathbb{D}^2} = \gamma$. But $\mathbb{D}^2$ is simply connected, so u admits a lift $\tilde{u} : (\mathbb{D}^2, 1) \to (Y, y_0)$ satisfying $p \circ \tilde{u} = u$, thus $p \circ \tilde{u}|_{\partial \mathbb{D}^2} = \gamma$ implies that $\tilde{u}|_{\partial \mathbb{D}^2} : (S^1, 1) \to (Y, y_0)$ is a lift of γ. Uniqueness of lifts then implies $\tilde{u}|_{\partial \mathbb{D}^2} = \tilde{\gamma}$ and thus $[\tilde{\gamma}] = e \in \pi_1(Y, y_0)$. □

COROLLARY 15.14. *If X is simply connected, then every path-connected covering space of X is also simply connected.* □

EXAMPLE 15.15. Corollary 15.14 implies that there does not exist any covering map $S^1 \to \mathbb{R}$.

Here is an application important in complex analysis. Observe that

$$p : \mathbb{C} \to \mathbb{C}^* := \mathbb{C} \backslash \{0\} : z \mapsto e^z$$

is a covering map. Writing $p(x + iy) = e^x e^{iy}$, we can picture p as a transformation from Cartesian to polar coordinates: it maps every horizontal line $\{\operatorname{Im} z = \text{const}\}$ to a ray in $\mathbb{C}^*$ emanating from the origin, and every vertical line $\{\operatorname{Re} z = \text{const}\}$ to a circle in $\mathbb{C}^*$, which it covers infinitely many times. This shows that p is not bijective, so it has no global inverse, but it will admit inverses if we restrict it to suitably small domains, and it is useful to know what domains will generally suffice for this. In other words, we would like to know which open subsets $\mathcal{U} \subset \mathbb{C}^*$ can be the domain of a continuous function

$$\log : \mathcal{U} \to \mathbb{C} \quad \text{such that} \quad e^{\log z} = z \text{ for all } z \in \mathcal{U}.$$

For simplicity, we will restrict our attention to path-connected[21] domains and also assume $1 \in \mathcal{U}$, so that we can adopt the convention $\log(1) := 0$. Defining $f : (\mathcal{U}, 1) \hookrightarrow (\mathbb{C}^*, 1)$ as the inclusion, the desired function $\log : (\mathcal{U}, 1) \to (\mathbb{C}, 0)$ will then be the unique solution to the lifting problem

$$\begin{tikzcd} & (\mathbb{C}, 0) \arrow[d, "p"] \\ (\mathcal{U}, 1) \arrow[ur, dashed, "\log"] \arrow[r, hook, "f"] & (\mathbb{C}^*, 1) \end{tikzcd}$$

Theorem 15.11 now gives the answer: $\log : \mathcal{U} \to \mathbb{C}$ exists if and only if $f_*(\pi_1(\mathcal{U}, 1)) \subset p_*(\pi_1(\mathbb{C}, 0)) = 0$, or in other words, if every loop in $\mathcal{U}$ can be extended to a map $\mathbb{D}^2 \to \mathbb{C}^*$. Using the notion of the *winding number* from Exercise 10.27, this is the same as saying every loop $\gamma : S^1 \to \mathcal{U}$ satisfies $\operatorname{wind}(\gamma; 0) = 0$. For example, $\log : \mathcal{U} \to \mathbb{C}$ can be defined whenever $\mathcal{U}$ is simply connected, or if $\mathcal{U}$ has the shape of an annulus whose outer circle does not enclose the origin. Examples that do not work include any annulus whose inner circle encloses the origin: this will always contain a loop that winds nontrivially around the origin, so that trying to define log along this loop produces a function that shifts by $2\pi i$ as one rotates fully around the loop. Notice that when $\log : \mathcal{U} \to \mathbb{C}$ exists, it is uniquely determined by the condition $\log(1) = 0$; without this one could equally well modify any given definition of log by adding integer multiples of $2\pi i$.

The proof of the lifting theorem requires two lemmas that are also special cases of the theorem. We assume for the remainder of this lecture that $(Y, y_0) \xrightarrow{p} (X, x_0)$ is a covering map and X, Y and A are all path-connected.

Lemma 15.16 (the path lifting property). *Every path $\gamma : (I, 0) \to (X, x_0)$ has a unique lift $\tilde{\gamma} : (I, 0) \to (Y, y_0)$.*

Proof. Since I is compact, we can find a finite partition $0 =: t_0 < t_1 < \ldots < t_{N-1} < t_N := 1$ such that for each $j = 1, \ldots, N$, the image of $\gamma_j := \gamma|_{[t_{j-1}, t_j]}$ lies in an evenly covered open subset $\mathcal{U}_j \subset X$ with $p^{-1}(\mathcal{U}_j) = \bigcup_{\alpha \in J} \mathcal{V}_\alpha$. Now given any $y \in p^{-1}(\gamma(t_{j-1}))$, we have $y \in \mathcal{V}_\alpha$ for a unique $\alpha \in J$, and γ_j has a unique lift $\tilde{\gamma}_j : [t_{j-1}, t_j] \to Y$ with $\tilde{\gamma}_j(t_{j-1}) = y$, defined by

$$\tilde{\gamma}_j = (p|_{\mathcal{V}_\alpha})^{-1} \circ \gamma_j.$$

With this understood, the unique lift $\tilde{\gamma}$ of γ with $\tilde{\gamma}(0) = y_0$ can be constructed by lifting $\tilde{\gamma}_1$ as explained above, then lifting $\tilde{\gamma}_2$ with starting point $\tilde{\gamma}_2(t_1) := \tilde{\gamma}_1(t_1)$, and continuing in this way to cover the entire interval. □

Lemma 15.17 (the homotopy lifting property). *Suppose $H : I \times A \to X$ is a homotopy with $H(0, \cdot) = f : A \to X$, and $\tilde{f} : A \to Y$ is a lift of f. Then there exists a unique lift $\widetilde{H} : I \times A \to Y$ of H satisfying $\widetilde{H}(0, \cdot) = \tilde{f}$.*

Proof. The previous lemma implies that each of the paths $s \mapsto H(s, a) \in X$ for $a \in A$ have unique lifts $s \mapsto \widetilde{H}(s, a) \in Y$ with $\widetilde{H}(0, a) = \tilde{f}(a)$. One should then check that the map $\widetilde{H} : I \times A \to Y$ defined in this way is continuous; I leave this as an exercise. □

Proof of Theorem 15.11. We shall first define an appropriate map $\tilde{f} : A \to Y$ and then show that the definition is independent of choices. Its uniqueness will be immediately clear, but its continuity will not be: in the final step we will use the hypothesis that A is locally path-connected in showing that $\tilde{f}$ is continuous.

[21]Since $\mathcal{U} \subset \mathbb{C}^*$ is open, it is locally path-connected, thus it will automatically be path-connected if it is connected.

Given $a \in A$, choose a path $a_0 \overset{\alpha}{\rightsquigarrow} a$, giving a path $x_0 \overset{f\circ\alpha}{\rightsquigarrow} f(a)$, which lifts via Lemma 15.16 to a unique path $\widetilde{f \circ \alpha}$ in Y that starts at y_0. If a lift $\tilde{f}$ exists, it clearly must satisfy

$$\tilde{f}(a) = \widetilde{f \circ \alpha}(1).$$

We claim that this point in Y does not depend on the choice of the path α, and thus gives a well-defined (though not necessarily continuous) map $\tilde{f} : A \to Y$. Indeed, suppose $a_0 \overset{\beta}{\rightsquigarrow} a$ is another path. Then $\alpha \cdot \beta^{-1}$ is a loop based at a_0 and thus represents an element of $\pi_1(A, a_0)$, and $f_*[\alpha \cdot \beta^{-1}] \in \pi_1(X, x_0)$ is represented by the loop $(f \circ \alpha) \cdot (f \circ \beta^{-1})$. The hypothesis $\operatorname{im} f_* \subset \operatorname{im} p_*$ then implies the existence of a loop $y_0 \overset{\tilde{\gamma}}{\rightsquigarrow} y_0$ in Y such that

$$[(f \circ \alpha) \cdot (f \circ \beta^{-1})] = p_*[\tilde{\gamma}] = [p \circ \tilde{\gamma}],$$

so there is a homotopy $H : I^2 \to X$ with $H(0, \cdot) = \gamma := p \circ \tilde{\gamma}$, $H(1, \cdot) = (f \circ \alpha) \cdot (f \circ \beta^{-1})$, and $H(s, 0) = H(s, 1) = x_0$ for all $s \in I$. Notice that $\tilde{\gamma}$ is a lift of $\gamma : (I, 0) \to (X, x_0)$. Now Lemma 15.17 provides a lift $\widetilde{H} : I^2 \to Y$ of H with $\widetilde{H}(0, \cdot) = \tilde{\gamma}$. In this homotopy, the paths $s \mapsto \widetilde{H}(s, 0)$ and $s \mapsto \widetilde{H}(s, 1)$ are lifts of the constant path $H(\cdot, 0) = H(\cdot, 1) \equiv x_0$ starting at $\tilde{\gamma}(0) = \tilde{\gamma}(1) = y_0$, so the uniqueness in Lemma 15.16 implies that both are also constant paths, hence $\widetilde{H}(s, 0) = \widetilde{H}(s, 1) = y_0$ for all $s \in I$. This shows that the unique lift of $(f \circ \alpha) \cdot (f \circ \beta^{-1})$ to a path in Y starting at y_0 is actually a loop, i.e. its end point is also y_0: indeed, this lift is $\widetilde{H}(1, \cdot)$. This lift is necessarily the concatenation of the lift $\widetilde{f \circ \alpha}$ of $f \circ \alpha$ starting at y_0 with the lift of $f \circ \beta^{-1}$ starting at $\widetilde{f \circ \alpha}(1)$. Since it ends at y_0, we conclude that this second lift is simply the inverse of $\widetilde{f \circ \beta}$, implying that

$$\widetilde{f \circ \alpha}(1) = \widetilde{f \circ \beta}(1),$$

which proves the claim.

It remains to show that $\tilde{f} : A \to Y$ as defined by the above procedure is continuous. Given $a \in A$ with $x = f(a) \in X$ and $y = \tilde{f}(a) \in Y$, choose any neighborhood $\mathcal{V} \subset Y$ of y that is small enough for $\mathcal{U} := p(\mathcal{V}) \subset X$ to be an evenly covered neighborhood of x, with $p|_{\mathcal{V}} : \mathcal{V} \to \mathcal{U}$ a homeomorphism. It will suffice to show that a has a neighborhood $\mathcal{O} \subset A$ with $\tilde{f}(\mathcal{O}) \subset \mathcal{V}$. Since A is locally path-connected, we can choose $\mathcal{O} \subset f^{-1}(\mathcal{U})$ to be a path-connected neighborhood of a, fix a path $a_0 \overset{\gamma}{\rightsquigarrow} a$ in A and, for any $a' \in \mathcal{O}$, choose a path $a \overset{\beta}{\rightsquigarrow} a'$ in $\mathcal{O}$. Now $\gamma \cdot \beta$ is a path from a_0 to a', so

$$\tilde{f}(a) = \widetilde{f \circ \gamma}(1) = y \in \mathcal{V} \quad \text{and} \quad \tilde{f}(a') = \widetilde{f \circ \gamma} \cdot \widetilde{f \circ \beta}(1),$$

where $\widetilde{f \circ \beta}$ is the unique lift of $f \circ \beta$ starting at y. Since $f \circ \beta$ lies entirely in the evenly covered neighborhood $\mathcal{U}$, this second lift is simply $(p|_{\mathcal{V}})^{-1} \circ (f \circ \beta)$, which lies entirely in $\mathcal{V}$, proving $\tilde{f}(a') \in \mathcal{V}$. $\square$

EXAMPLE 15.18. If the local path-connectedness assumption on A is dropped, then the proof above gives a procedure for defining a unique lift $\tilde{f} : A \to Y$, but it may fail to be continuous. A concrete example is depicted in [**Hat02**, p. 79], Exercise 7. The idea is to define A as a space that mostly consists of the usual circle $S^1 \subset \mathbb{R}^2$, but replace a portion just to the right of the top point $(0, 1)$ with a curve resembling the graph of the function $y = \sin(1/x) + 1$. The point $(0, 1)$ is included in A, along with every point of the usual circle just to the left of it, but on the right, A consists of an infinitely long curve that is compressed into a compact space and has accumulation points along an interval but no well-defined limit. This space is path-connected, because one can start from $(0, 1)$ and go around the circle to reach any other point, including any point on the infinitely long compressed sine curve; it is also simply connected, due to the fact that continuous paths along the compressed sine curve can never actually reach the end of it, but must instead go back the other way around the circle before they can reach $(0, 1)$. But A is not locally path-connected, because sufficiently small neighborhoods of $(0,1)$ in A always contain many disjoint segments of the compressed sine curve and thus cannot be path-connected. Now consider the covering map $\mathbb{R} \to S^1 : \theta \mapsto e^{i\theta}$ and a continuous map $f : A \to S^1$ defined as the identity on most of A, but projecting the graph of $y = \sin(1/x) + 1$ to the circle in the obvious way near $(0,1)$. One can define a lift $\tilde{f} : A \to \mathbb{R}$ by choosing $\tilde{f}(0,1)$ to be any point in $p^{-1}(f(0,1))$ and then lifting paths to define $\tilde{f}$ everywhere else. But since every neighborhood of $(0,1)$ contains some points that cannot be reached except by paths rotating almost all the way around the circle, this neighborhood will contain points $a \in A$ for which $\tilde{f}(a)$ differs from $\tilde{f}(0,1)$ by nearly 2π. In particular, $\tilde{f}$ cannot be continuous at $(0,1)$.

16. Classification of covers (June 13, 2023)

Throughout this lecture, all spaces should be assumed path-connected and locally path-connected unless otherwise noted. We will occasionally need a slightly stronger condition, which we will abbreviate with the word "reasonable":[22]

Definition 16.1. We will say that a space X is **reasonable** if it is path-connected and locally path-connected, and every point $x \in X$ has a simply connected neighborhood.

For the purposes of the theorems in this lecture, the definition of the term "reasonable" can be weakened somewhat at the expense of making it more complicated, but we will stick with the above definition since it is satisfied by almost all spaces we would ever like to consider. A popular example of an "unreasonable" space is the so-called *Hawaiian earring*, see Exercise 13.2(c).

We will state several theorems in this lecture related to the problem of classifying covers of a given space. All of them are in some way applications of the lifting theorem (Theorem 15.11). Before stating them, we need to establish what it means for two covers of the same space to be "equivalent".

Definition 16.2. Given two covers $p_i : Y_i \to X$ for $i = 1, 2$, a **map of covers** from p_1 to p_2 is a map $f : Y_1 \to Y_2$ such that $p_2 \circ f = p_1$, i.e. the following diagram commutes:

$$\begin{tikzcd} Y_1 \arrow[rr, "f"] \arrow[dr, "p_1"'] & & Y_2 \arrow[dl, "p_2"] \\ & X & \end{tikzcd} \tag{16.1}$$

Additionally, we call f an **isomorphism of covers** if there also exists a map of covers from p_2 to p_1 that inverts f; this is true if and only if the map $f : Y_1 \to Y_2$ is a homeomorphism, since its inverse $f^{-1} : Y_2 \to Y_1$ is then automatically a map of covers from p_2 to p_1. If such an isomorphism exists, we say that the two covers p_1 and p_2 are **isomorphic** (or **equivalent**). If base points $x \in X$ and $y_i \in Y_i$ are specified such that $p_i : (Y_i, y_i) \to (X, x)$ and $f : (Y_1, y_1) \to (Y_2, y_2)$ are also pointed maps, then we call f an **isomorphism of pointed covers**. In the case where p_1 and p_2 are both the same cover $p : Y \to X$, an isomorphism of covers from p to itself is called a **deck transformation**[23] (*Decktransformation*) of $p : Y \to X$.

The terms **covering translation** and **automorphism** are also sometimes used as synonyms for "deck transformation". The set of all deck transformations of a given cover $p : Y \to X$ forms a

[22]This is not a universally standard term.

[23]This terminology gives you a hint that some portion of this subject was developed by German mathematicians in the time before English was fully established as an international language. I don't happen to know who invented the term.

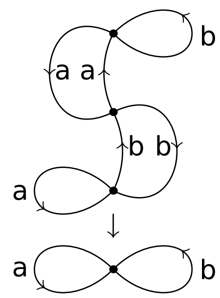

FIGURE 10. A 3-fold cover of $S^1 \vee S^1$ with trivial automorphism group.

group, called the **automorphism group**

$$\mathrm{Aut}(p) := \left\{ f : Y \to Y \;\middle|\; f \text{ is a homeomorphism such that } p \circ f = p \right\},$$

where the group operation is defined by composition of maps.

EXAMPLE 16.3. For the cover $p : \mathbb{R} \to S^1 : \theta \mapsto e^{i\theta}$, $\mathrm{Aut}(p)$ consists of all maps $f_k : \mathbb{R} \to \mathbb{R}$ of the form $f_k(\theta) = \theta + 2\pi k$ for $k \in \mathbb{Z}$, so in particular, $\mathrm{Aut}(p)$ is isomorphic to $\mathbb{Z}$.

EXAMPLE 16.4. Figure 10 illustrates a covering map $p : Y \to S^1 \vee S^1$ of degree 3. If we label the base point of $S^1 \vee S^1$ as x, then the three elements of $p^{-1}(x) \subset Y$ are the three dots in the top portion of the diagram: label them y_1, y_2 and y_3 from bottom to top. The covering map is defined such that each loop or path beginning and ending at any of the points y_1, y_2, y_3 is sent to the loop in $S^1 \vee S^1$ labeled by the same letter with the orientations of the arrows matching. Suppose $f : Y \to Y$ is a deck transformation satisfying $f(y_1) = y_2$. Then since f is a homeomorphism, it must map the loop labeled a based at y_1 to a loop based at y_2 that also must be labeled a. But no such loop exists, so we conclude that there is no deck transformation sending y_1 to y_2. By similar arguments, it is not hard to show that the only deck transformation of this cover is the identity map, in other words, $\mathrm{Aut}(p)$ is the trivial group.

Almost everything we will be able to prove about maps of covers is based on the following observation: if the diagram (16.1) commutes, it means that $f : Y_1 \to Y_2$ is a *lift* of the map $p_1 : Y_1 \to X$ to the cover Y_2, i.e. in our previous notation for lifts, $f = \tilde{p}_1$. The fact that p_1 itself is a covering map is irrelevant for this observation. Now if all the spaces involved are path-connected and locally path-connected, the lifting theorem gives us a condition characterizing the existence and uniqueness of a map of covers: for any choices of base points $x \in X$, $y_1 \in p_1^{-1}(x) \subset Y_1$ and $y_2 \in p_2^{-1}(x) \subset Y_2$, a map of covers $f : Y_1 \to Y_2$ satisfying $f(y_1) = y_2$ exists (and is unique) if and only if

$$(p_1)_*\pi_1(Y_1, y_1) \subset (p_2)_*\pi_1(Y_2, y_2).$$

This map will then be an isomorphism if and only if there exists a map of covers going the other direction, and the latter exists if and only if the reverse inclusion holds. This proves:

THEOREM 16.5. *Two covers $p_i : Y_i \to X$ for $i = 1, 2$ are isomorphic if and only if for some choice of base points $x \in X$ and $y_i \in p_i^{-1}(x) \subset Y_i$ for $i = 1, 2$, the subgroups $(p_1)_*\pi_1(Y_1, y_1)$ and $(p_2)_*\pi_1(Y_2, y_2)$ in $\pi_1(X, x)$ are identical.* □

Next we use the same perspective to study deck transformations of a single cover $p : Y \to X$. Given $x \in X$ and $y_1, y_2 \in p^{-1}(x) \subset Y$, the uniqueness of lifts implies that there exists at most one deck transformation $f : Y \to Y$ sending y_1 to y_2. We've seen in Example 16.4 that this transformation might not always exist.

DEFINITION 16.6. A cover $p : Y \to X$ is called **regular** (or equivalently **normal**) if for every $x \in X$ and all $y_1, y_2 \in p^{-1}(x) \subset Y$, there exists a deck transformation sending y_1 to y_2.

The following exercise says that in order to check whether a cover of a path-connected space is regular, it suffices to choose a base point $x \in X$ and investigate whether deck transformations can be used to relate arbitrary points in the preimage of *that particular point.* The proof is an easy application of the path lifting property (Lemma 15.16).

EXERCISE 16.7. Show that if $p : Y \to X$ is a covering map and X is path-connected, then p is also regular if the following slightly weaker condition holds: for some fixed $x \in X$, any two elements $y_1, y_2 \in p^{-1}(x) \subset X$ satisfy $y_2 = f(y_1)$ for some deck transformation $f \in \operatorname{Aut}(p)$.

If $\deg(p) < \infty$, the previous remarks about uniqueness of deck transformations imply $|\operatorname{Aut}(p)| \leqslant \deg(p)$, and equality is satisfied if and only if p is regular. By the lifting theorem, the desired deck transformation sending y_1 to y_2 will exist if and only if

$$p_*\pi_1(Y, y_1) = p_*\pi_1(Y, y_2). \tag{16.2}$$

Let us try to translate this into a condition for recognizing when p is regular. Recall that any path $y_1 \overset{\tilde{\gamma}}{\rightsquigarrow} y_2$ in Y determines an isomorphism

$$\Phi_{\tilde{\gamma}} : \pi_1(Y, y_2) \to \pi_1(Y, y_1) : [\alpha] \mapsto [\tilde{\gamma} \cdot \alpha \cdot \tilde{\gamma}^{-1}].$$

Since y_1 and y_2 are both in $p^{-1}(x)$, the projection of this concatenation down to X gives a concatenation of *loops*, i.e. $\gamma := p \circ \tilde{\gamma}$ is a loop $x \rightsquigarrow x$ and thus represents an element $[\gamma] \in \pi_1(X, x)$. Now in order to check whether (16.2) holds, we can represent an arbitrary element of $\pi_1(Y, y_1)$ as $\Phi_{\tilde{\gamma}}[\alpha]$ for some loop $y_2 \overset{\alpha}{\rightsquigarrow} y_2$, and then observe

$$p_*\Phi_{\tilde{\gamma}}[\alpha] = [p \circ (\tilde{\gamma} \cdot \alpha \cdot \tilde{\gamma}^{-1})] = [\gamma \cdot (p \circ \alpha) \cdot \gamma^{-1}] = [\gamma]p_*[\alpha][\gamma]^{-1}.$$

This proves that the subgroup $p_*\pi_1(Y, y_1) \subset \pi_1(X, x)$ is the conjugate of $p_*\pi_1(Y, y_2) \subset \pi_1(X, x)$ by the specific element $[\gamma] \in \pi_1(X, x)$, so the desired deck transformation exists if and only if $p_*\pi_1(Y, y_2)$ is invariant under conjugation with $[\gamma]$. We could now ask the same question about deck transformations sending y_i to y_2 for arbitrary $y_i \in p^{-1}(x)$, and the answer in each case can be expressed in terms of conjugation of $p_*\pi_1(Y, y_2)$ by some element $[\gamma] \in \pi_1(X, x)$ for which the loop γ lifts to a path $y_i \overset{\tilde{\gamma}}{\rightsquigarrow} y_2$. Now observe: *any* loop $x \overset{\gamma}{\rightsquigarrow} x$ can arise in this way for some choice of $y_i \in p^{-1}(x)$. Indeed, if γ is given, then γ^{-1} has a unique lift to a path from y_2 to some other point in $p^{-1}(x)$, and the inverse of this path is then a lift of γ. Using Exercise 16.7 above, the question of regularity therefore reduces to the question of whether $p_*\pi_1(Y, y_2)$ is invariant under arbitrary conjugations, and we have thus proved:

THEOREM 16.8. *If Y and X are path-connected and locally path-connected, then a cover $p : (Y, y_0) \to (X, x_0)$ is regular if and only if the subgroup $p_*\pi_1(Y, y_0) \subset \pi_1(X, x_0)$ is normal.* □

Notice that while the algebraic condition in this theorem appears to depend on a choice of base points, the condition of p being regular clearly does not. It follows that if $p_*\pi_1(Y, y_0) \subset \pi_1(X, x_0)$ is a normal subgroup, then this condition will remain true for any other choice of base points $x \in X$ and $y \in p^{-1}(x) \subset Y$.

The next two results require the restriction to "reasonable" spaces in the sense of Definition 16.1.

THEOREM 16.9 (the Galois correspondence). *If X is a reasonable space with base point $x_0 \in X$, there is a natural bijection from the set of all isomorphism classes of pointed covers $p : (Y, y_0) \to (X, x_0)$ to the set of all subgroups of $\pi_1(X, x_0)$: it is defined by*

$$[p : (Y, y_0) \to (X, x_0)] \mapsto p_*\pi_1(Y, y_0).$$

It is easy to verify from the definition of isomorphism for covers that the map in this theorem is well defined, and we proved in Theorem 16.5 that it is injective. Surjectivity will be a consequence of the following result, which will be proved in the next lecture.

THEOREM 16.10. *Every reasonable space admits a simply connected covering space.*

Notice that if $p_i : (Y_i, y_i) \to (X, x_0)$ for $i = 1, 2$ are two reasonable covers satisfying $\pi_1(Y_1) = \pi_1(Y_2) = 0$, then Theorem 16.5 implies that they are isomorphic covers. For this reason it is conventional to abuse terminology slightly by referring to any simply connected cover of a given space X as "the" **universal cover** (*universelle Überlagerung*) of X. It is often denoted by $\widetilde{X}$.

EXAMPLES 16.11. The universal cover $\widetilde{S^1}$ of S^1 is $\mathbb{R}$, due to the covering map $\mathbb{R} \to S^1 : \theta \mapsto e^{i\theta}$. Similarly, $\widetilde{\mathbb{RP}^n} \cong S^n$ for $n \geqslant 2$, and $\widetilde{\mathbb{T}^n} \cong \mathbb{R}^n$.

A substantially less obvious class of examples is given by the surfaces Σ_g of genus $g \geqslant 2$: these have universal cover $\widetilde{\Sigma}_g \cong \mathbb{R}^2$. It would take us too far afield to explain why, but one standard way of constructing this cover comes from hyperbolic geometry, where instead of $\mathbb{R}^2$ we consider the open disk $\mathring{\mathbb{D}}^2$ with a Riemannian metric that has constant negative curvature. One can identify each of the surfaces Σ_g with the quotient of $\mathring{\mathbb{D}}^2$ by a suitable group of isometries and then define a covering map $\mathring{\mathbb{D}}^2 \to \Sigma_g$ as the quotient projection.

For the remainder of this lecture, fix a base-point preserving covering map $p : (Y, y_0) \to (X, x_0)$ where X and Y are assumed reasonable, and denote

$$G := \pi_1(X, x_0), \qquad H := p_*\pi_1(Y, y_0) \subset G.$$

If H is not a normal subgroup, then there is no natural notion of a quotient group G/H, but we can still define G/H as the *set* of left cosets

$$G\big/H = \{gH \subset G \mid g \in G\},$$

where gH denotes the subset $\{gh \mid h \in H\} \subset G$. One can similarly consider the set of right cosets

$$H\big\backslash G = \{Hg \subset G \mid g \in G\}.$$

These two sets are identical if and only if H is normal, in which case both are denoted by G/H and they form a group. With or without this condition, $G\big/H$ and $H\big\backslash G$ have the same number (finite or infinite) of elements, which is called the **index** of H in G and denoted by

$$[G : H] := \left|G\big/H\right| = \left|H\big\backslash G\right|.$$

In the following we will make repeated use of the fact that for any $y \in p^{-1}(x_0)$, any path $y_0 \overset{\tilde{\gamma}}{\rightsquigarrow} y$ gives rise to a loop $\gamma := p \circ \tilde{\gamma}$ based at x_0, and conversely, any such loop gives rise to a path that starts at y_0 and ends at some point in $p^{-1}(x_0)$.

LEMMA 16.12. *There is a natural bijection*

$$\Phi : p^{-1}(x_0) \to H\big\backslash G : y \mapsto H[\gamma],$$

where $x_0 \overset{\gamma}{\rightsquigarrow} x_0$ is any loop that lifts to a path $y_0 \overset{\tilde{\gamma}}{\rightsquigarrow} y$.

COROLLARY 16.13. $\deg(p) = [G : H]$. $\square$

Proof of Lemma 16.12. We first show that Φ is well defined. Given two choices of paths $\tilde{\alpha}, \tilde{\beta}$ from y_0 to y, we have loops $\alpha := p \circ \tilde{\alpha}$ and $\beta := p \circ \tilde{\beta}$ based at x_0, and $\tilde{\alpha} \cdot \tilde{\beta}^{-1}$ is a loop based at y_0. We therefore have

$$[\alpha][\beta]^{-1} = [p \circ (\tilde{\alpha} \cdot \tilde{\beta}^{-1})] = p_*[\tilde{\alpha} \cdot \tilde{\beta}^{-1}] \in H,$$

implying $H[\alpha] = H[\beta]$.

The surjectivity of Φ is obvious: given $[\gamma] \in G$, there exists a lift $\tilde{\gamma}$ of γ to a path from y_0 to some point $y \in p^{-1}(x_0)$, so $\Phi(y) = H[\gamma]$.

To see that Φ is injective, suppose $\Phi(y) = \Phi(y')$, choose paths $y_0 \overset{\tilde{\alpha}}{\rightsquigarrow} y$ and $y_0 \overset{\tilde{\beta}}{\rightsquigarrow} y'$, giving rise to loops $\alpha := p \circ \tilde{\alpha}$ and $\beta := p \circ \tilde{\beta}$ based at x_0 such that

$$H[\alpha] = \Phi(y) = \Phi(y') = H[\beta],$$

thus $[\alpha][\beta]^{-1} \in H$. It follows that there exists a loop $y_0 \overset{\tilde{\gamma}}{\rightsquigarrow} y_0$ projecting to $\gamma := p \circ \tilde{\gamma}$ such that $[\alpha \cdot \beta^{-1}] = [\gamma]$, hence $[\alpha] = [\gamma] \cdot [\beta]$, so α is homotopic to $\gamma \cdot \beta$ with fixed end points. Since γ lifts to a loop $\tilde{\gamma}$ and homotopies can also be lifted, we conclude that $\tilde{\alpha}$ is homotopic to $\tilde{\gamma} \cdot \tilde{\beta}$ with fixed end points, implying $y = \tilde{\alpha}(1) = \tilde{\beta}(1) = y'$. □

If the cover is regular so $H \subset G$ is normal, then $\deg(p) = |\operatorname{Aut}(p)|$, and Corollary 16.13 therefore implies that $\operatorname{Aut}(p)$ has the same order as the quotient group G/H. The next result should then seem relatively unsurprising.

Theorem 16.14. *For a regular cover $p : (Y, y_0) \to (X, x_0)$ of reasonable spaces with $\pi_1(X, x_0) = G$ and $p_*\pi_1(Y, y_0) = H \subset G$, there exists a group isomorphism*

$$\Psi : \operatorname{Aut}(p) \to G/H : f \mapsto [\gamma]H,$$

where $x_0 \overset{\gamma}{\rightsquigarrow} x_0$ is any loop that has a lift to a path from y_0 to $f(y_0)$.

Notice that the universal cover $p : (\widetilde{X}, \tilde{x}_0) \to (X, x_0)$ is automatically regular since the trivial subgroup of $\pi_1(X, x_0)$ is always normal, so applying this theorem to the universal cover gives:

Corollary 16.15. *For the universal cover $p : (\widetilde{X}, \tilde{x}_0) \to (X, x_0)$, there is an isomorphism $\operatorname{Aut}(p) \to \pi_1(X, x_0)$ sending each deck transformation f to the homotopy class of any loop $x_0 \rightsquigarrow x_0$ that lifts to a path $\tilde{x}_0 \rightsquigarrow f(\tilde{x}_0)$.* □

Proof of Theorem 16.14. Regularity implies that the map $\operatorname{Aut}(p) \to p^{-1}(x_0) : f \mapsto f(y_0)$ is bijective, so Ψ is then well defined and bijective due to Lemma 16.12. For the identity element $\operatorname{Id} \in \operatorname{Aut}(p)$, we have $\Psi(\operatorname{Id}) = [\gamma]H$ for any loop γ that lifts to a loop from y_0 to $\operatorname{Id}(y_0) = y_0$, which means $[\gamma] \in H$, so $[\gamma]H$ is the identity element in G/H.

It remains to show that $\Psi(f \circ g) = \Psi(f)\Psi(g)$ for any two deck transformations $f, g \in \operatorname{Aut}(p)$. Choose loops α, β based at x_0 which lift to paths $y_0 \overset{\tilde{\alpha}}{\rightsquigarrow} f(y_0)$ and $y_0 \overset{\tilde{\beta}}{\rightsquigarrow} g(y_0)$. Then $f \circ \tilde{\beta}$ is a path from $f(y_0)$ to $f \circ g(y_0)$ and can thus be concatenated with $\tilde{\alpha}$, forming a path

$$y_0 \overset{\tilde{\alpha} \cdot (f \circ \tilde{\beta})}{\rightsquigarrow} f \circ g(y_0).$$

Now since $f \in \operatorname{Aut}(p)$, $p \circ f = p$ implies $p \circ (f \circ \tilde{\beta}) = p \circ \tilde{\beta} = \beta$, thus

$$\Psi(f \circ g) = [p \circ (\tilde{\alpha} \cdot (f \circ \tilde{\beta}))] = [\alpha][\beta] = \Psi(f)\Psi(g).$$

□

Corollary 16.15 says that we can compute the fundamental group of any reasonable space X if we can understand the deck transformations of its universal cover. Combining this with the natural bijection $\mathrm{Aut}(p) \to p^{-1}(x_0)$ that sends each deck transformation to its image on the base point, we also obtain from this an intuitively appealing interpretation of the meaning of $\pi_1(X, x_0)$: every loop γ based at x_0 lifts uniquely to a path starting at $\tilde{x}_0$ and ending at some point in $p^{-1}(x_0)$. As far as $\pi_1(X, x_0)$ is concerned, all that matters is the end point of the lift: two loops are equivalent in $\pi_1(X, x_0)$ if and only if their lifts to $\widetilde{X}$ have the same end point, and a loop is trivial in $\pi_1(X, x_0)$ if and only if its lift to $\widetilde{X}$ is also a loop.

EXAMPLE 16.16. Applying Corollary 16.15 to the cover $p : \mathbb{R} \to S^1 : \theta \mapsto e^{i\theta}$ reproduces the isomorphism $\pi_1(S^1, 1) \cong \mathbb{Z}$ we discussed at the end of Lecture 9. The loop $\gamma_k(t) := e^{2\pi i k t}$ in S^1 for each $k \in \mathbb{Z}$ lifts to $\mathbb{R}$ with base point 0 as the path $\tilde{\gamma}_k(t) = 2\pi k t$.

EXAMPLE 16.17. For each $n \geqslant 2$, Corollary 16.15 implies $\pi_1(\mathbb{RP}^n) \cong \mathbb{Z}_2$, as this is the automorphism group of the universal cover $p : S^n \to \mathbb{RP}^n$, defined as the natural quotient projection. Concretely, after fixing base points $x_0 \in \mathbb{RP}^n$ and $y_0 \in p^{-1}(x_0) \subset S^n$, each loop in $\mathbb{RP}^n$ based at x_0 lifts to S^n as a path that starts at y_0 and ends at either y_0 or its antipodal point $-y_0$. The nontrivial element of $\pi_1(\mathbb{RP}^n, x_0)$ is thus represented by any loop whose lift to S^n starts and ends at antipodal points.

17. The universal cover and group actions (June 15, 2023)

In Theorem 16.14, we saw a formula that can be used to compute the automorphism group of any regular cover as a quotient of two fundamental groups. I want to mention how this generalizes for non-regular covers, though I will leave most of the details as an exercise. One way to approach the problem is as follows: any pointed covering map $p : (Y, y_0) \to (X, x_0)$ of reasonable spaces can be fit into a diagram

$$\begin{tikzcd} (Z, z_0) \arrow[r, "q"] \arrow[rr, bend right, "P"'] & (Y, y_0) \arrow[r, "p"] & (X, x_0), \end{tikzcd} \tag{17.1}$$

in which q and P are also pointed covering maps and are both *regular*. For example, if you already believe that every reasonable space has a universal cover (and we shall prove this below), then we can always take $q : Z \to Y$ to be the universal cover of Y, which makes $P : Z \to X$ the universal cover of X since $\pi_1(Z) = 0$, and universal covers are always regular because the trivial subgroup is always normal. In this case, Corollary 16.15 gives us natural isomorphisms $\mathrm{Aut}(P) \cong \pi_1(X, x_0)$ and $\mathrm{Aut}(q) \cong \pi_1(Y, y_0)$. This is not true if Z is not simply connected, and we will not assume this in the following exercise, but it turns out that if P and q are nonetheless regular, then we can derive a formula for $\mathrm{Aut}(p)$ in terms of the other two automorphism groups.

EXERCISE 17.1. Assuming the spaces in (17.1) are all reasonable, let us abbreviate the automorphism groups of P and q by

$$G := \mathrm{Aut}(P), \qquad \text{and} \qquad H := \mathrm{Aut}(q).$$

(a) Use the path-lifting property to prove the following lemma: If $\Psi \in G$ and $\psi \in \mathrm{Aut}(p)$ are deck transformations for which the relation $q \circ \Psi = \psi \circ q$ holds at the base point $z_0 \in Z$, then it holds everywhere.
Hint: For any $z \in Z$, choose a path from z_0 to z, then use Ψ, ψ and the covering projections to cook up other paths in Z, Y and X. Some of them are lifts of others, and two important ones will turn out to be the same.

(b) Deduce from part (a) that H is the subgroup of G consisting of all deck transformations $\Psi : Z \to Z$ for P that satisfy $\Psi(z_0) \in q^{-1}(y_0)$.
(c) Show that if $P : Z \to X$ is regular then so is $q : Z \to Y$. Give two proofs: one using the result of part (b), and another using the characterization of regularity in terms of normal subgroups.
(d) The **normalizer** (*Normalisator*) $N(H) \subset G$ of the subgroup H is by definition the largest subgroup of G that contains H as a normal subgroup, i.e.

$$N(H) := \left\{ g \in G \;\middle|\; gHg^{-1} = H \right\}.$$

Show that if the cover $q : Z \to Y$ is regular, then for any $\Psi \in N(H)$, there exists a deck transformation $\psi : Y \to Y$ of p satisfying the relation $q \circ \Psi = \psi \circ q$, and it is unique. Moreover, the correspondence $\Psi \mapsto \psi$ defines a group homomorphism $N(H) \to \mathrm{Aut}(p)$ whose kernel is H.
(e) Show that if the cover $P : Z \to X$ is also regular, then the homomorphism $N(H) \to \mathrm{Aut}(p)$ in part (d) is also surjective, and thus descends to an isomorphism

$$N(H)/H \xrightarrow{\cong} \mathrm{Aut}(p).$$

Applying Exercise 17.1 with Z simply connected now gives:

Corollary 17.2. *For any covering map $p : (Y, y_0) \to (X, x_0)$ of reasonable spaces with $\pi_1(X, x_0) = G$ and $p_*\pi_1(Y, y_0) = H \subset G$, there is a natural isomorphism* $\mathrm{Aut}(p) \cong N(H)/H$. □

Notice that there always exists a subgroup of G in which H is normal, e.g. H itself is such a subgroup, and it may well happen that no larger subgroup satisfies this condition, in which case $N(H) = H$ and $\mathrm{Aut}(p)$ is therefore trivial. If H is normal in G, then $N(H) = G$ and the cover is therefore regular, hence Corollary 17.2 reduces to Theorem 16.14.

Moving on from non-regular covers, we have some unfinished business from the previous lecture: it remains to prove the surjectivity of the Galois correspondence (Theorem 16.9), and the existence of the universal cover (Theorem 16.10). The latter is actually a special case of the former: recall from Corollary 15.13 that the homomorphism $p_* : \pi_1(Y, y_0) \to \pi_1(X, x_0)$ induced by a covering map $p : (Y, y_0) \to (X, x_0)$ is always injective, thus the existence of a universal cover amounts to the statement that the image of the Galois correspondence includes the trivial subgroup of $\pi_1(X, x_0)$. We will prove this first, and then use it to deduce the Galois correspondence in full generality.

As before, we need to restrict our attention to "reasonable spaces," meaning spaces that are path-connected and locally path-connected, and in which every point has a simply connected neighborhood. The first two conditions are needed in order to apply the lifting theorem, which we used several times in the previous lecture. The third condition has not yet been used, but this is the moment where we will need it. In constructing a universal cover $p : (\widetilde{X}, \tilde{x}_0) \to (X, x_0)$, the theorems at the end of the previous lecture give some useful intuition on what to aim for: in particular, there needs to be a one-to-one correspondence between $p^{-1}(x_0) \subset \widetilde{X}$ and $\pi_1(X, x_0)$. What we will actually construct is a cover for which these two sets are not just in bijective correspondence but are literally the same set. In set-theoretic terms, the construction is quite straightforward, but giving it a topology that makes it a covering map is a bit subtle—that is where we will need to assume that simply connected neighborhoods exist.

Proof of Theorem 16.10 (the universal cover). We will not give every detail but sketch the main idea. Given a reasonable space X with base point $x_0 \in X$, define the set

$$\widetilde{X} := \{\text{paths } \gamma : (I, 0) \to (X, x_0)\} \big/ \underset{h+}{\sim},$$

i.e. it is the set of all equivalence classes of paths that start at the base point, with equivalence defined as homotopy with fixed end points. Since this definition does not specify the end point of any path but the equivalence relation leaves these end points unchanged, we obtain a natural map

$$p : \widetilde{X} \to X : [\gamma] \mapsto \gamma(1),$$

which is obviously surjective since X is path-connected. Notice that $p^{-1}(x_0) = \pi_1(X, x_0)$.

We claim that $\widetilde{X}$ can be assigned a topology that makes $p : \widetilde{X} \to X$ into a covering map. To see this, suppose $\mathcal{U} \subset X$ is a path-connected subset and $i^{\mathcal{U}} : \mathcal{U} \hookrightarrow X$ denotes its inclusion. For any point $x \in \mathcal{U}$, the induced homomorphism $i^{\mathcal{U}}_* : \pi_1(\mathcal{U}, x) \to \pi_1(X, x)$ is trivial if and only if every loop $S^1 \to \mathcal{U}$ based at x can be extended to a map $\mathbb{D}^2 \to X$. Notice that this is weaker in general than demanding an extension $\mathbb{D}^2 \to \mathcal{U}$; the latter would mean that $\mathcal{U}$ is simply connected, but we do not want to assume this. Notice also that if this condition holds for some choice of base point $x \in \mathcal{U}$, then the usual change of base-point arguments imply that it will hold for any other base point $y \in \mathcal{U}$, thus we can sensibly speak of the condition that $i^{\mathcal{U}}_* : \pi_1(\mathcal{U}) \to \pi_1(X)$ is trivial. With this understood, consider the collection of sets

$$\mathcal{B} := \left\{ \mathcal{U} \subset X \;\middle|\; \mathcal{U} \text{ is open and path-connected and } i^{\mathcal{U}}_* : \pi_1(\mathcal{U}) \to \pi_1(X) \text{ is trivial} \right\}.$$

It is a straightforward exercise to verify the following properties:

(1) $\mathcal{U} \in \mathcal{B}$ if and only if for every pair of paths α, β in $\mathcal{U}$ with the same end points, α and β are homotopic in X with fixed end points (cf. Corollary 9.9).
(2) If $\mathcal{U} \in \mathcal{B}$ and $\mathcal{V} \subset \mathcal{U}$ is a path-connected open subset, then $\mathcal{V} \in \mathcal{B}$.
(3) $\mathcal{B}$ is a base for the topology of X.

In particular, the third property holds because X is reasonable: every point $x \in X$ has a simply connected neighborhood, which contains an open neighborhood that necessarily belongs to $\mathcal{B}$, and it follows that every open subset of X is a union of such sets.

Now for any $\mathcal{U} \in \mathcal{B}$ with a point $x \in \mathcal{U}$ and a path γ in X from x_0 to x, let

$$\mathcal{U}_{[\gamma]} := \left\{ [\gamma \cdot \alpha] \in \widetilde{X} \;\middle|\; \alpha \text{ is a path in } \mathcal{U} \text{ starting at } x \right\}.$$

Notice that $\mathcal{U}_{[\gamma]}$ depends only on the homotopy class $[\gamma] \in \widetilde{X}$; this relies on the fact that since $\mathcal{U} \in \mathcal{B}$, the path α in the definition above is uniquely determined up to homotopy in X by its end point. It follows in fact that $p : \widetilde{X} \to X$ restricts to a bijection

$$\mathcal{U}_{[\gamma]} \xrightarrow{p} \mathcal{U}.$$

With all this in mind, one can now show that

$$\widetilde{\mathcal{B}} := \left\{ \mathcal{U}_{[\gamma]} \subset \widetilde{X} \;\middle|\; \mathcal{U} \in \mathcal{B} \text{ and } [\gamma] \in \widetilde{X} \text{ with } \gamma(1) \in \mathcal{U} \right\}$$

is a base for a topology on $\widetilde{X}$ such that each $\mathcal{U} \in \mathcal{B}$ is evenly covered by $p : \widetilde{X} \to X$. We leave the details of this as an exercise.

There is an obvious choice of base point in $\widetilde{X}$: define $\tilde{x}_0 \in \widetilde{X}$ as the homotopy class of the constant path at x_0. It remains to prove that $\pi_1(\widetilde{X}, \tilde{x}_0) = 0$. Since we now know that $p : (\widetilde{X}, \tilde{x}_0) \to (X, x_0)$ is a covering map, Corollary 15.13 implies that $p_* : \pi_1(\widetilde{X}, \tilde{x}_0) \to \pi_1(X, x_0)$ is injective, thus it will suffice to show that the subgroup $p_*\pi_1(\widetilde{X}, \tilde{x}_0)$ in $\pi_1(X, x_0)$ is trivial. This subgroup is the set of homotopy classes $[\gamma] \in \pi_1(X, x_0)$ for which the loop γ lifts to a loop $\tilde{\gamma}$ based at $\tilde{x}_0$. The lift of γ to $\widetilde{X}$ can be written as

$$\tilde{\gamma}(t) = [\gamma_t] \in \widetilde{X},$$

where for each $t \in I$ we define

$$\gamma_t(s) := \begin{cases} \gamma(s) & \text{for } 0 \leqslant s \leqslant t, \\ \gamma(t) & \text{for } t \leqslant s \leqslant 1. \end{cases}$$

Then assuming $\tilde{\gamma}$ is a loop, we find $\tilde{\gamma}(1) = [\gamma] = \tilde{\gamma}(0) = [\text{const}]$, which is simply the statement that γ is homotopic with fixed end points to a constant loop, hence $[\gamma] \in \pi_1(X, x_0)$ is the trivial element. □

I do not have the energy to draw the picture myself, but I highly recommend looking at the picture of the universal cover of $S^1 \vee S^1$ on page 59 of [**Hat02**]. The idea here is that for every homotopically nontrivial loop in $S^1 \vee S^1$, one obtains a non-closed path in the universal cover $\widetilde{X}$. One can thus construct $\widetilde{X}$ one path at a time if one denotes by a and b the generators of $\pi_1(S^1 \vee S^1, x) \cong F_{\{a,b\}}$: at each step, the loops a, b, a^{-1} and b^{-1} furnish four homotopically distinct choices of loops to traverse, which lift to four distinct paths in $\widetilde{X}$ from one copy of the base point to another. Starting at the natural base point $\tilde{x}_0$ and following this procedure recursively produces the fractal picture in [**Hat02**, p. 59].

The application to the Galois correspondence requires a brief digression on topological groups and group actions.

Definition 17.3. A **topological group** (*topologische Gruppe*) is a group G with a topology such that the maps

$$G \times G \to G : (g, h) \mapsto gh \quad \text{and} \quad G \to G : g \mapsto g^{-1}$$

are both continuous.

Popular examples of topological groups include the various subgroups of the real or complex general linear groups $\mathrm{GL}(n, \mathbb{R})$ and $\mathrm{GL}(n, \mathbb{C})$, e.g. the orthogonal group $\mathrm{O}(n)$ and unitary group $\mathrm{U}(n)$, the special linear groups $\mathrm{SL}(n, \mathbb{R})$ and $\mathrm{SL}(n, \mathbb{C})$, and so forth. We saw in Exercise 7.29 that for any locally compact and locally connected Hausdorff space X, the group of homeomorphisms $\mathrm{Homeo}(X)$ is a topological group with the group operation defined by composition. Finally, *any* group can be regarded as a topological group if we assign to it the discrete topology; this follows from the fact that every map on a space with the discrete topology is continuous. Topological groups with the discrete topology are often referred to as **discrete groups**.

Definition 17.4. Given a topological group G and a space X, a (continuous) G-action (*Wirkung*) on X is a (continuous) map

$$G \times X \to X : (g, x) \mapsto g \cdot x$$

such that the identity element $e \in G$ satisfies $e \cdot x = x$ for all $x \in X$ and $(gh) \cdot x = g \cdot (h \cdot x)$ holds for all $g, h \in G$ and $x \in X$.

Notice that for any G-action on X, there is a natural group homomorphism $G \to \mathrm{Homeo}(X)$ sending $g \in G$ to the homeomorphism $\varphi_g : X \to X$ defined by $\varphi_g(x) = g \cdot x$. If G is a discrete group then the converse is also true: every group homomorphism $G \to \mathrm{Homeo}(X)$ comes from a G-action on X. This is true because as long as the topology of G is discrete, the map $G \times X \to X : (g, x) \mapsto g \cdot x$ is continuous if and only if the map $X \to X : x \mapsto g \cdot x$ is continuous for every fixed $g \in G$. If G has a more interesting topology, then continuity of the map $(g, x) \mapsto g \cdot x$ with respect to $g \in G$ is also a nontrivial condition that would need to be checked—but we have no need to worry about this right now, as most of the groups we will deal with below are discrete.

Example 17.5. For any covering map $p : Y \to X$, $\mathrm{Aut}(p)$ acts as a discrete group on Y by $f \cdot y := f(y)$.

EXAMPLE 17.6. Regarding $\mathbb{Z}_2$ as a discrete group, a $\mathbb{Z}_2$-action on any space X is determined by the homeomorphism $\varphi_1 : X \to X$ associated to the nontrivial element $[1] \in \mathbb{Z}/2\mathbb{Z} =: \mathbb{Z}_2$, and this is necessarily an **involution**, i.e. it is its own inverse. A frequently occurring example is the action of $\mathbb{Z}_2$ on S^n defined via the antipodal map $\mathbf{x} \mapsto -\mathbf{x}$.

EXAMPLE 17.7. Here is a non-discrete example: any subgroup of the orthogonal group $\mathrm{O}(n)$ acts on $S^{n-1} \subset \mathbb{R}^n$ by matrix-vector multiplication, $A \cdot \mathbf{x} = A\mathbf{x}$.

For any G-action on X and a subset $\mathcal{U} \subset X$, we denote

$$g \cdot \mathcal{U} := \{g \cdot x \mid x \in \mathcal{U}\} \subset X.$$

Similarly, for each point $x \in X$, we define its **orbit** (*Bahn*) as the subset

$$G \cdot x := \{g \cdot x \mid g \in G\} \subset X.$$

One can easily check that for any two points $x, y \in X$, their orbits $G \cdot x$ and $G \cdot y$ are either identical or disjoint, thus there is an equivalence relation $\sim$ on X such that $x \sim y$ if and only if $G \cdot x = G \cdot y$. The quotient topological space defined by this equivalence relation is denoted by

$$X/G := X/\sim = \{\text{orbits } G \cdot x \subset X \mid x \in X\}.$$

EXAMPLE 17.8. The quotient $S^n/\mathbb{Z}_2$ arising from the action in Example 17.6 is $\mathbb{RP}^n$.

PROPOSITION 17.9. *Regarding $\pi_1(X, x_0)$ as a discrete group, any covering map $p : (Y, y_0) \to (X, x_0)$ of reasonable spaces with $\pi_1(Y) = 0$ gives rise to a natural action of $\pi_1(X, x_0)$ on Y.*

PROOF. There are at least two ways to see the action of $\pi_1(X, x_0)$ on a simply connected cover. First, Corollary 16.15 identifies $\pi_1(X, x_0)$ with $\mathrm{Aut}(p)$, and the latter acts on Y as explained in Example 17.5.

Alternatively, one can appeal to the uniqueness of the universal cover, so $p : (Y, y_0) \to (X, x_0)$ is necessarily isomorphic to the specific cover $\widetilde{X} = \{\text{paths } x_0 \rightsquigarrow x\}/\underset{h+}{\sim}$ that we constructed in the proof of Theorem 16.10. Then the obvious way for homotopy classes of loops $[\alpha] \in \pi_1(X, x_0)$ to act on homotopy classes of paths $[\gamma] \in \widetilde{X}$ is by concatenation:

$$[\alpha] \cdot [\gamma] := [\alpha \cdot \gamma].$$

It is easy to verify that this also defines a group action. $\square$

EXERCISE 17.10. Show that the two actions of $\pi_1(X, x_0)$ on the universal cover constructed in the above proof are the same.

DEFINITION 17.11. A G-action on X is **free** (*frei*) if the only element $g \in G$ satisfying $g \cdot x = x$ for some $x \in X$ is the identity $g = e$.

The action is called **properly discontinuous** (*eigentlich diskontinuierlich*) if every $x \in X$ has a neighborhood $\mathcal{U} \subset X$ such that

$$(g \cdot \mathcal{U}) \cap \mathcal{U} = \varnothing$$

for every $g \in G$ with $g \cdot x \neq x$.

EXERCISE 17.12. Show that if a G-action is free and properly discontinuous, then G is discrete.

EXERCISE 17.13. Show that for any covering map $p : Y \to X$, the action of $\mathrm{Aut}(p)$ on Y as in Example 17.5 is free and properly discontinuous.

The observation that actions of deck transformation groups are free already has some nontrivial consequences, for instance:

PROPOSITION 17.14. *There exists no covering map $p : \mathbb{D}^2 \to X$ with $\deg(p) > 1$.*

Proof. If $\deg(p) > 1$, then since $\pi_1(\mathbb{D}^2) = 0$, we observe that the cover $p : \mathbb{D}^2 \to X$ must be regular and therefore has a nontrivial deck transformation group $\mathrm{Aut}(p)$ which acts freely on $\mathbb{D}^2$. But the Brouwer fixed point theorem rules out the existence of any nontrivial free group action on $\mathbb{D}^2$. □

The main purpose of the above definitions is that they lead to the following theorem, whose proof is now an easy exercise.

Theorem 17.15. *If G acts on X freely and properly discontinuously, then the quotient projection*

$$q : X \to X/G : x \mapsto G \cdot x$$

is a regular covering map with $\mathrm{Aut}(q) = G$. □

Now we are ready to finish the proof of the Galois correspondence.

Proof of Theorem 16.9. We have already shown that the correspondence is well defined and injective, so we need to prove surjectivity, in other words: given a reasonable space X with base point $x_0 \in X$ and any subgroup $H \subset G := \pi_1(X, x_0)$, we need to find a reasonable space Y with a covering map $p : (Y, y_0) \to (X, x_0)$ such that $p_*\pi_1(Y, y_0) = H$. Since X is reasonable, there exists a universal cover $f : (\widetilde{X}, \tilde{x}_0) \to (X, x_0)$, whose automorphism group is isomorphic to G, so this isomorphism defines a free and properly discontinuous action of G on $\widetilde{X}$. It also defines a free and properly discontinuous action of every subgroup of G on $\widetilde{X}$, and in particular an H-action. Define

$$Y := \widetilde{X}/H \quad \text{and} \quad p : Y \to X : H \cdot \tilde{x} \mapsto f(\tilde{x}).$$

It is straightforward to check that this is a covering map, and it is base-point preserving if we define $y_0 := H \cdot \tilde{x}_0$ as the base point of Y. Moreover, the quotient projection $q : (\widetilde{X}, \tilde{x}_0) \to (Y, y_0)$ is now the universal cover of Y, and it fits into the following commutative diagram:

$$\begin{tikzcd}
(\widetilde{X}, \tilde{x}_0) \arrow[r, "f"] \arrow[d, "q"] & (X, x_0) \\
(Y, y_0) \arrow[ur, "p"] &
\end{tikzcd}$$

Given a loop γ in X based at x_0, let γ' denote its lift to a path in Y starting at y_0, and let $\tilde{\gamma}$ denote the lift to a path in $\widetilde{X}$ starting at $\tilde{x}_0$, The subgroup $p_*\pi_1(Y, y_0) \subset \pi_1(X, x_0)$ is precisely the set of all homotopy classes $[\gamma] \in \pi_1(X, x_0)$ for which γ' is a loop. Notice that since all maps in the diagram are covering maps, $\tilde{\gamma}$ is also a lift of γ' via the covering map q. Then $[\gamma] \in H$ so that γ' is a loop if and only if the end point of $\tilde{\gamma}$ is in $q^{-1}(y_0) = H \cdot \tilde{x}_0$. Under the natural bijection between $\pi_1(X, x_0)$ and $f^{-1}(x_0) = G \cdot \tilde{x}_0$, this just means $[\gamma] \in H$, hence $p_*\pi_1(Y, y_0) = H$. □

18. Manifolds (June 20, 2023)

I have mentioned manifolds already a few times in this course, but now it is time to discuss them somewhat more precisely. While we do not plan to go to deeply into this subject this semester, the goal is in part to understand what the main definitions are and why, forming the basis of the subject known as "geometric topology". In so doing, we will also establish an inventory of examples and concepts that will serve as useful intuition when we start to talk about homology next week.

Definition 18.1. A **topological manifold** (*Mannigfaltigkeit*) of dimension $n \geqslant 0$ (often abbreviated with the term "n-manifold") is a second countable Hausdorff space M such that every point $p \in M$ has a neighborhood homeomorphic to $\mathbb{R}^n$.

More generally, a **topological n-manifold with boundary** (*Mannigfaltigkeit mit Rand*) is a second countable Hausdorff space M such that every point $p \in M$ has a neighborhood homeomorphic to either $\mathbb{R}^n$ or the so-called "n-dimensional half-space"

$$\mathbb{H}^n := [0, \infty) \times \mathbb{R}^{n-1}.$$

The third condition in each of these definitions is probably the most intuitive and is the most distinguishing feature of manifolds: we abbreviate it by saying that manifolds are "locally Euclidean". It means in effect that sufficiently small open subsets of a manifold can be described via *local coordinate systems*. The technical term for this is "chart": a **chart** (*Karte*) on an n-manifold with boundary is a homeomorphism

$$\varphi : \mathcal{U} \to \Omega$$

where $\mathcal{U} \subset M$ and $\Omega \subset \mathbb{H}^n$ are open subsets. As special cases, Ω may be the whole of $\mathbb{H}^n$, or an open ball in $\mathbb{H}^n$ disjoint from

$$\partial\mathbb{H}^n := \{0\} \times \mathbb{R}^{n-1},$$

in which case Ω is also homeomorphic to $\mathbb{R}^n$. It follows that on any n-manifold (with or without boundary), every point is in the domain of a chart. Conversely, if we are given a collection of charts $\{\varphi_\alpha : \mathcal{U}_\alpha \to \Omega_\alpha\}_{\alpha \in J}$ such that $M = \bigcup_{\alpha \in J} \mathcal{U}_\alpha$, then after shrinking the domains and targets of these charts if necessary, we can assume every point $p \in M$ is in the domain of some chart $\varphi_\alpha : \mathcal{U}_\alpha \to \Omega_\alpha$ such that Ω_α is either an open ball in $\mathbb{H}^n \backslash \partial\mathbb{H}^n$ or a half-ball with boundary on $\partial\mathbb{H}^n$, so that Ω is homeomorphic to either $\mathbb{R}^n$ or $\mathbb{H}^n$. This means M is locally Euclidean, so both versions of the third condition in our definition can be rephrased as the condition that M is covered by charts. The **boundary** of a manifold M with boundary can now be defined as the subset

$$\partial M := \left\{ p \in M \;\middle|\; \varphi(p) \in \partial\mathbb{H}^n \text{ for some chart } \varphi \right\},$$

which is clearly an $(n-1)$-manifold (without boundary).

The word "topological" is included before "manifold" in order to make the distinction between topological manifolds and *smooth manifolds*, which we will discuss a little bit below. By default in this course, you should assume that everything we refer to simply as a "manifold" is actually a *topological* manifold unless otherwise specified. (If this were a differential geometry course, you would instead want to assume that "manifold" always means *smooth* manifold.) One can regard manifolds without boundary as being special cases of manifolds M with boundary such that $\partial M = \varnothing$, so we shall also use "manifold" as an abbreviation for the term "manifold with boundary" and will generally specify "without boundary" when we want to assume $\partial M = \varnothing$. You should be aware that some books adopt different conventions for such details, e.g. some authors assume $\partial M = \varnothing$ always unless the words "with boundary" are explicitly included.

REMARK 18.2. The following detail deserves emphasis: the way we have expressed the definition of the boundary $\partial M \subset M$ above makes sense in part because when we defined the notion of a chart $\varphi : \mathcal{U} \to \Omega$, we required[24] its image Ω to be an open subset of the half-space $\mathbb{H}^n$, and not necessarily an open subset of $\mathbb{R}^n$. If we were allowing arbitrary open subsets $\Omega \subset \mathbb{R}^n$, then every point $p \in M$ would be a boundary point, because e.g. one could take any chart $\varphi : \mathcal{U} \to \Omega$ with $p \in \mathcal{U}$ and compose it with a translation on $\mathbb{R}^n$ so that $\varphi(p) = 0 \in \partial\mathbb{H}^n$. Requiring $\Omega \subset \mathbb{H}^n$ prevents this in general, because if we start with a chart $\varphi : \mathcal{U} \to \Omega$ whose image contains an open ball around $\varphi(p)$, then translating it to achieve $\varphi(p) = 0$ will produce something whose image cannot be contained in $\mathbb{H}^n$. In fact, the translation trick works only for points $p \in \mathcal{U}$ with $\varphi(p) \in \partial\mathbb{H}^n$, as

[24]This convention is not universal: many books allow charts to have images that are arbitrary open subsets of $\mathbb{R}^n$. The latter is a sensible convention especially if one only wants to consider manifolds with empty boundary, and even if nonempty boundaries are allowed, one can work with charts defined in this way, but the definition of $\partial M \subset M$ would need to be expressed a bit differently.

these are precisely the points for which Ω does not contain any ball around $\varphi(p)$. It can happen that $\Omega \subset \mathbb{H}^n$ is *also* an open subset of $\mathbb{R}^n$: this is true if and only if $\Omega \cap \partial\mathbb{H}^n = \varnothing$, and in that case, none of the points in the domain of the chart are boundary points. One can show that whenever $\varphi(p) \in \partial\mathbb{H}^n$ for some chart $\varphi : \mathcal{U} \to \Omega$ with $p \in \mathcal{U}$, the same must hold for all other charts whose domains contain p; in other words, no point of M can be simultaneously a boundary point and an *interior* point, where the latter means that some chart maps it into $\mathbb{H}^n \backslash \partial\mathbb{H}^n$. For $n \leqslant 2$, this can be proved using methods that we have already developed (see Exercise 19.13); the proof for $n > 2$ requires some other methods that we haven't developed yet, but will soon, e.g. singular homology.

Manifolds are usually what we have in mind when we think of spaces that are "nice" or "reasonable". In particular, the following is an immediate consequence of the observation that every point in $\mathbb{R}^n$ or $\mathbb{H}^n$ has a neighborhood homeomorphic to the closed n-disk:

PROPOSITION 18.3. *For an n-manifold M and a point $p \in M$, every neighborhood of p contains one that is homeomorphic to $\mathbb{D}^n$.* ☐

COROLLARY 18.4. *Manifolds are locally compact and locally path-connected. They are also **locally contractible**, meaning every neighborhood of every point in M contains a contractible neighborhood. In particular, they are "reasonable" in the sense of Definition 16.1.* ☐

It follows via Theorem 7.19 that a manifold M is connected if and only if it is path-connected. More generally, the path-components of M are the same as its connected components (cf. Prop. 7.18), each of which are open and closed subsets, hence M is homeomorphic to the disjoint union of its connected components. It is similarly easy to show that these connected components are also manifolds.

DEFINITION 18.5. A manifold M is **closed** (*geschlossen*) if it is compact and $\partial M = \varnothing$. It is **open** (*offen*) if none of its connected components are closed, i.e. all of them either are noncompact or have nonempty boundary.

You need to be aware that these usages of the words "closed" and "open" are different from the notions of closed or open subsets in a topological space. The distinction between a "closed manifold" and a "closed subset" is at least more explicit in German: the former is a *geschlossene Mannigfaltigkeit*, while the latter is an *abgeschlossene Teilmenge*. For openness there is the same ambiguity in German and English, but it is rarely a problem: you just need to pay attention to the context in which these adjectives are used and what kinds of nouns they are modifying. We will not have much occasion to talk about open manifolds in this course, and many authors apparently dislike seeing the word "open" used in this way, but it has some advantages, e.g. in differential topology, there are some elegant theorems that can be stated most naturally for open manifolds but are not true for manifolds that are not open.

EXAMPLE 18.6. Any discrete space with only countably many points is a 0-manifold. (Discrete spaces with uncountably many points are excluded because they are not second countable.) Conversely, this is an accurate description of every 0-manifold, and the closed ones are those that are finite. Note that a 0-manifold can never have boundary.

EXAMPLE 18.7. The line $\mathbb{R}$, the interval $(-1, 1)$ and the circle S^1 are all examples of 1-manifolds without boundary, where S^1 is closed and the others are open. Further examples without boundary are obtained by taking arbitrary countable disjoint unions of these examples, e.g. $S^1 \amalg \mathbb{R}$ is a 1-manifold without boundary, though it is neither closed nor open since it has one closed component and one that is not closed. Some examples of 1-manifolds with nonempty boundary include the interval $I = [0, 1]$, whose boundary is the compact 0-manifold $\partial I = \{0, 1\}$, and $[0, 1)$, whose boundary is $\partial[0, 1) = \{0\}$.

EXAMPLE 18.8. The word **surface** (*Fläche*) refers in general to a 2-dimensional manifold. Examples without boundary include S^2, $\mathbb{T}^2 = S^1 \times S^1$, the surfaces Σ_g of genus $g \geqslant 0$, $\mathbb{RP}^2$, $\mathbb{R}^2$, and arbitrary countable disjoint unions of any of these. One can also take connected sums of these examples to obtain more, though as we've seen, not all of the examples that arise in this way are new, e.g. Σ_g for $g \geqslant 1$ is the g-fold connected sum of copies of $\mathbb{T}^2$. Some compact examples with boundary include $\mathbb{D}^2$ (with $\partial\mathbb{D}^2 = S^1$) and the surface $\Sigma_{g,m}$ of genus g with $m \geqslant 1$ holes cut out, which has $\partial\Sigma_{g,m} \cong \coprod_{i=1}^m S^1$. An obvious noncompact example with nonempty boundary is the half-plane $\mathbb{H}^2$, with $\partial\mathbb{H}^2 \cong \mathbb{R}$.

EXAMPLE 18.9. Some examples of arbitrary dimension n without boundary are S^n, $\mathbb{RP}^n$, $\mathbb{R}^n$, $\mathbb{T}^n := S^1 \times \ldots \times S^1$, any open subset of any of these, and anything obtained from these by (countable) disjoint unions or connected sums.[25] Some obvious examples with nonempty boundary are $\mathbb{D}^n$ (with $\partial\mathbb{D}^n = S^{n-1}$), and $[-1,1] \times \mathbb{T}^{n-1}$, whose boundary is the disjoint union of two copies of $\mathbb{T}^{n-1}$.

While we don't plan to do very much with it in this course, we now make a brief digression on the subject of *smooth* manifolds, which are the main object of study in differential geometry and differential topology. As preparation, observe that if $\varphi_\alpha : \mathcal{U}_\alpha \to \Omega_\alpha$ and $\varphi_\beta : \mathcal{U}_\beta \to \Omega_\beta$ are two charts on the same manifold M, then on any region $\mathcal{U}_\alpha \cap \mathcal{U}_\beta$ where they overlap, we can think of them as describing two alternative coordinate systems, so that there is a well-defined "coordinate transformation" map switching from one to the other. To be more precise, $\varphi_\alpha(\mathcal{U}_\alpha \cap \mathcal{U}_\beta)$ and $\varphi_\beta(\mathcal{U}_\alpha \cap \mathcal{U}_\beta)$ are open subsets of Ω_α and Ω_β respectively, and there is a homeomorphism from one to the other defined via the following diagram:

$$\begin{tikzcd}[column sep=small] & \mathcal{U}_\alpha \cap \mathcal{U}_\beta \arrow[dl, "\varphi_\alpha"'] \arrow[dr, "\varphi_\beta"] & \\ \varphi_\alpha(\mathcal{U}_\alpha \cap \mathcal{U}_\beta) \arrow[rr, "\varphi_\beta \circ \varphi_\alpha^{-1}"] & & \varphi_\beta(\mathcal{U}_\alpha \cap \mathcal{U}_\beta) \end{tikzcd}$$

The map $\varphi_\beta \circ \varphi_\alpha^{-1}$ is called the **transition map** (*Übergang*) relating φ_α and φ_β. The key point about a transition map is that its domain and target are open subsets of a Euclidean space (or half-space), thus we know what it means for such a map to be "differentiable". This observation makes it possible to do differential calculus on manifolds and to speak of functions $f : M \to \mathbb{R}$ as being differentiable or not: the idea is that f should be called differentiable if it appears differentiable whenever it is written in a local coordinate system. But for this to be well defined, we need to be assured that the answer to the differentiability question will not change if we change coordinate systems, i.e. if we compose our local coordinate expression for f with a transition map. If all conceivable charts for M are allowed, then the answer will indeed sometimes change, because the composition of a differentiable function with a non-differentiable map is not usually differentiable. We therefore need to be able to assume that transition maps are always differentiable, and since this is not true if all conceivable charts are allowed, we need to restrict the class of charts that we consider. This restriction introduces a bit of structure on M that is not determined by its topology, but is something extra:

DEFINITION 18.10. A **smooth structure** (*glatte Struktur*) on an n-dimensional topological manifold M is a maximal collection of charts $\{\varphi_\alpha : \mathcal{U}_\alpha \to \Omega_\alpha\}_{\alpha \in J}$ for which $M = \bigcup_{\alpha \in J} \mathcal{U}_\alpha$ and the corresponding transition maps $\varphi_\beta \circ \varphi_\alpha^{-1}$ for all $\alpha, \beta \in J$ are of class C^∞. A topological manifold endowed with a smooth structure is called a **smooth manifold** (*glatte Mannigfaltigkeit*).

[25] Recall from Lecture 13 the connected sum of two n-manifolds M and N: it is defined by deleting the interiors of two embedded n-disks from M and N and then gluing them together along the spheres S^{n-1} at the boundaries of these disks.

It is easy to see that a single topological manifold can have multiple distinct smooth structures, e.g. on $M = \mathbb{R}$, the functions $\varphi_\alpha(t) = t$ and $\varphi_\beta(t) = t^3$ are homeomorphisms $\mathbb{R} \to \mathbb{R}$ and can thus be regarded as charts, but $\varphi_\alpha \circ \varphi_\beta^{-1}$ is not everywhere differentiable, hence φ_α and φ_β can each be regarded as belonging to smooth structures on $\mathbb{R}$, but they are distinct smooth structures. That is a relatively uninteresting example, but there are also known examples of topological manifolds admitting multiple smooth structures that are not even equivalent up to *diffeomorphism* (the smooth version of homeomorphism), as well as topological manifolds that do not admit any smooth structure at all. Such things are very hard to prove, but you should not worry about them right now, because the basic fact is that most manifolds we encounter in nature have natural smooth structures. A very high proportion of them come from the following geometric version of the implicit function theorem.

THEOREM 18.11 (implicit function theorem). *Suppose $\mathcal{U} \subset \mathbb{R}^n$ is an open subset, $F : \mathcal{U} \to \mathbb{R}^k$ is a C^∞-map and $q \in \mathbb{R}^k$ is a point such that for all $p \in F^{-1}(q)$, the derivative $dF(p) : \mathbb{R}^n \to \mathbb{R}^k$ is surjective (we say in this case that q is a* ***regular value*** *of F). Then $F^{-1}(q) \subset \mathbb{R}^n$ is a smooth manifold of dimension $n - k$.* □

The above theorem is provided "for your information," meaning we do not plan to either prove or use it in any serious way in this course, but you should be aware that it exists because it provides many examples of manifolds that arise naturally in various applications. For instance:

EXAMPLE 18.12. The n-sphere $S^n = F^{-1}(1)$, where $F : \mathbb{R}^{n+1} \to \mathbb{R} : \mathbf{x} \mapsto |\mathbf{x}|^2$, which has 1 as a regular value.

EXAMPLE 18.13. The special linear group $\mathrm{SL}(n, \mathbb{R}) = \det^{-1}(1)$ for the determinant map $\det : \mathbb{R}^{n\times n} \to \mathbb{R}$. One can show that 1 is a regular value of det by relating the derivative of the determinants of a family of matrices passing through $\mathbb{1}$ to the trace of the derivative of that family of matrices. Thus $\mathrm{SL}(n, \mathbb{R})$ is a smooth manifold of dimension $n^2 - 1$.

Now let's look at a couple of non-examples.

EXAMPLE 18.14. The wedge sum $S^1 \vee S^1$ is *not* a manifold of any dimension. It does look like a 1-manifold in the complement of the base point $x \in S^1 \vee S^1$, but x does not have any neighborhood homeomorphic to Euclidean space. Indeed, sufficiently small neighborhoods $\mathcal{U} \subset S^1 \vee S^1$ of x all look like two line segments intersecting, so that if we delete the point x, we obtain a space $\mathcal{U}\backslash\{x\}$ with four path-components. This cannot happen in an n-manifold for any n, as deleting a point from $\mathbb{R}$ produces two path-components, while deleting a point from $\mathbb{R}^n$ with $n \geqslant 2$ leaves a space that is still path-connected.

EXAMPLE 18.15. Here is a space that is locally Euclidean and second countable, but not Hausdorff: the line with two zeroes, i.e. $X := (\mathbb{R} \times \{0, 1\})/\sim$ with $(x, 0) \sim (x, 1)$ for all $x \neq 0$. If we endow X with the quotient topology induced by the natural topology of $\mathbb{R} \times \{0, 1\} \cong \mathbb{R} \amalg \mathbb{R}$, then a subset $\mathcal{U} \subset X$ is open if and only if its preimage under the quotient projection $\mathbb{R} \times \{0, 1\} \to X$ is open, and it follows in particular that the images of $\mathbb{R} \times \{0\}$ and $\mathbb{R} \times \{1\}$ under this projection are open subsets of X that are each (in obvious ways) homeomorphic to $\mathbb{R}$. The two zeroes $0_0 := [(0, 0)]$ and $0_1 := [(0, 1)]$ therefore each have neighborhoods homeomorphic to $\mathbb{R}$, and so (for more obvious reasons) does every other point, so the line with two zeroes would count as a 1-manifold if we did not require manifolds to be Hausdorff. We should emphasize that we are considering the quotient topology on X, not the pseudometric topology (cf. Example 6.12); X with the pseudometric topology is not locally homeomorphic to $\mathbb{R}$, because every neighborhood of 0_0 must also contain 0_1 and vice versa, so the two subsets described above would no longer be open.

EXAMPLE 18.16. The following is a compact variation on the previous example: writing X for the line with two zeroes, its one point compactification X^* is obtained by adding a single point called ∞, which is the limit of any sequence in X that has no bounded subsequence. Just as the one point compactification $\mathbb{R} \cup \{\infty\}$ of $\mathbb{R}$ is homeomorphic to S^1, we can think of X^* as the result of replacing one point $0 \in \mathbb{R} \subset S^1$ with a pair of points $0_0, 0_1 \in X^*$ that each have neighborhoods homeomorphic to $\mathbb{R}$, but with every neighborhood of 0_0 intersecting every neighborhood of 0_1. This would also be a 1-manifold if manifolds were not required to be Hausdorff.

You probably don't need much convincing by this point that spaces which are Hausdorff and second countable are "good," while those that lack either of these properties are "bad". Nonetheless, it's worth taking a moment to consider *why* it would be bad if we dropped either of these conditions from the definition of a manifold. The first answer is clearly that if we dropped the Hausdorff axiom, then Example 18.15 would be a manifold, and we don't like Example 18.15. But there are better reasons. One of them is related to the implicit function theorem, Theorem 18.11 above, which produces many examples of manifolds that are subsets of larger-dimensional Euclidean spaces. Notice that in this situation, it is completely unnecessary to verify whether those subsets are Hausdorff or second countable, because every subset of a finite-dimensional Euclidean space is both. (See Exercise 5.9 if you've forgotten how we know that $\mathbb{R}^n$ is second countable.) Now, it is reasonable to ask whether *all* conceivable manifolds arise from something similar to Theorem 18.11, i.e. are all of them embeddable into $\mathbb{R}^N$ for some $N \in \mathbb{N}$? The answer is yes, though clearly it would not be if the Hausdorff and second countability conditions were not included:

THEOREM 18.17. *Every topological manifold is homeomorphic to a closed subset of $\mathbb{R}^N$ for $N \in \mathbb{N}$ sufficiently large.* □

This is another theorem that I am providing "for your information," as I do not intend to use it for anything and therefore will not prove it. A readable proof for the case of a compact manifold appears in [**Hat02**, Corollary A.9]. The noncompact case is significantly harder and proofs typically do not appear in textbooks, but the idea is outlined and some precise references given in [**Lee11**, p. 116]. I would caution you in any case against taking this theorem more seriously than it deserves: while it's nice to know that all manifolds are in some sense *submanifolds* of some $\mathbb{R}^N$, many of them do not come with any canonical choice of embedding into $\mathbb{R}^N$, so this property is not in any way intrinsic to their structure and one should (and usually can) avoid using it to prove things about manifolds. It might also be argued that Theorem 18.17 undermines my point about the Hausdorff and second countability assumptions being indispensable, since it may seem desirable to be able to consider "manifolds" that are *more general* than just submanifolds of Euclidean spaces.

As a general principle, mathematicians consider a definition to be a "good" definition if it appears as the hypothesis for a good theorem. I'm not sure if Theorem 18.17 truly qualifies as a good theorem. But I want to talk about another one that I think is better.

THEOREM 18.18. *Every connected nonempty 1-manifold without boundary is homeomorphic to either S^1 or $\mathbb{R}$.*

If this statement sounds at first too restrictive, it makes up for it by being extremely useful. In combination with the implicit function theorem, one can deduce from it e.g. the possible topologies of regular level sets of arbitrary smooth functions $F : \mathbb{R}^n \to \mathbb{R}^{n-1}$. This ability has a surprising number of beautiful applications in differential topology and related fields; one example is the definition of the "mapping degree," sketched in Exercise 19.14. Those applications are typically based on the following corollary for compact manifolds with boundary.

COROLLARY 18.19. *Every compact 1-manifold M with boundary is homeomorphic to a disjoint union of finitely many copies of S^1 and $[0, 1]$. In particular, ∂M consists of evenly many points.*

PROOF. Since M is compact, it can have at most finitely many connected components (otherwise we can find a noncompact closed subset by choosing one point from every component). Restricting to connected components, it will therefore suffice to show that every connected compact 1-manifold M is either S^1 or $[0,1]$. Theorem 18.18 implies that $M \cong S^1$ if $\partial M = \varnothing$, so assume otherwise. Then ∂M is a closed subset and therefore is compact, and it is also a 0-manifold, which means it is a nonempty finite set. Let us modify M by attaching a half-line $[0,\infty)$ to each boundary point, that is, let

$$\widehat{M} := M \cup_{\partial M} \left(\coprod_{p \in \partial M} [0,\infty) \right).$$

This makes $\widehat{M}$ a noncompact connected 1-manifold with empty boundary, so by Theorem 18.18, $\widehat{M} \cong \mathbb{R}$. It follows that $M \subset \widehat{M}$ is homeomorphic to a path-connected compact subset of $\mathbb{R}$. All such subsets are compact intervals $[a,b]$, hence $M \cong [0,1]$. □

The proof of Theorem 18.18 given below is based on a series of exercises outlined in [Gal87]. I will not go through every step in exhaustive detail, as my main objective is just to point out explicitly where the Hausdorff and second countability conditions are needed. You saw already from Examples 18.15 and 18.16 that the theorem becomes false if the Hausdorff condition is dropped, and after the proof we will look at an even stranger example to see what can happen without second countability.

Here is a lemma that depends explicitly on the Hausdorff property, e.g. you will find if you look again at the line with two zeroes (Example 18.15) that it is not satisfied in that particular example.

LEMMA 18.20. *Suppose M is a Hausdorff space with two overlapping open subsets $\mathcal{U}_\alpha, \mathcal{U}_\beta \subset M$ that are each homeomorphic to $\mathbb{R}$, and neither is contained in the other. Then each connected component $\mathcal{W}$ of $\mathcal{U}_\alpha \cap \mathcal{U}_\beta$ is homeomorphic to $\mathbb{R}$ and has compact closure $\overline{\mathcal{W}} \subset M$ homeomorphic to $[0,1]$, whose boundary consists of a point $p_\alpha \in \mathcal{U}_\alpha$ that is not in $\mathcal{U}_\beta$ and a point $p_\beta \in \mathcal{U}_\beta$ that is not in $\mathcal{U}_\alpha$.*

PROOF. Choose explicit homeomorphisms $\varphi_\alpha : \mathcal{U}_\alpha \to \mathbb{R}$ and $\varphi_\beta : \mathcal{U}_\beta \to \mathbb{R}$. The image $\varphi_\beta(\mathcal{W}) \subset \mathbb{R}$ is necesarily a connected open subset of $\mathbb{R}$, and is therefore an open interval, implying $\mathcal{W} \cong \mathbb{R}$. But $\varphi_\beta(\mathcal{W})$ cannot be the entirety of $\mathbb{R}$, as that would imply $\mathcal{W} = \mathcal{U}_\beta$ since φ_β is a homeomorphism, and thus $\mathcal{U}_\beta \subset \mathcal{U}_\alpha$, which was excluded in the hypotheses. For the same reasons, $\varphi_\alpha(\mathcal{W})$ is an open interval in $\mathbb{R}$, but not the entirety of $\mathbb{R}$.

Let us show that the closure $\overline{\mathcal{W}} \subset M$ contains two boundary points p_α, p_β with the stated properties. To find p_α, choose a point $t \in \mathbb{R}$ that is in the closure of $\varphi_\alpha(\mathcal{W}) \subset \mathbb{R}$ but not in $\varphi_\alpha(\mathcal{W})$. Since φ_α is a homeomorphism, there must then exist a sequence $x_n \in \mathcal{W}$ converging to a point $p_\alpha := \varphi_\alpha^{-1}(t) \in \mathcal{U}_\alpha$, and p_α cannot belong to $\mathcal{U}_\beta$ since this would imply $p_\alpha \in \mathcal{W}$ and thus $t \in \varphi_\alpha(\mathcal{W})$. We claim: $|\varphi_\beta(x_n)| \to \infty$. Indeed, if this does not hold, then after replacing x_n with a suitable subsequence, we can assume $\varphi_\beta(x_n)$ converges to some point $y \in \mathbb{R}$, in which case x_n also converges to $x := \varphi_\beta^{-1}(y) \in \mathcal{U}_\beta$ since φ_β is a homeomorphism. But we already know $x_n \to p_\alpha$, so the assumption that M is Hausdorff implies $x = p_\alpha$ and gives a contradiction, since $p_\alpha \notin \mathcal{U}_\beta$.

It follows from the claim above that $\varphi_\beta(\mathcal{W}) \subset \mathbb{R}$ is an unbounded interval, and since it is not the entirety of $\mathbb{R}$, it is therefore an infinite half-interval of the form $(-\infty, a)$ or (b, ∞) for some $a, b \in \mathbb{R}$. Reversing the roles of α and β, a similar conclusion holds for $\varphi_\alpha(\mathcal{W})$, so for concreteness, let us suppose

$$\varphi_\alpha(\mathcal{W}) = (-\infty, a) \qquad \text{and} \qquad \varphi_\beta(\mathcal{W}) = (b, \infty),$$

in which case the recipe described above for defining $p_\alpha, p_\beta \in \overline{\mathcal{W}}$ gives

$$p_\alpha = \varphi_\alpha^{-1}(a), \qquad p_\beta = \varphi_\beta^{-1}(b).$$

(Only minor modifications to this discussion are necessary if $\varphi_\alpha(\mathcal{W})$ is instead bounded below or $\varphi_\beta(\mathcal{W})$ bounded above.) Moreover, the transition map

$$\mathbb{R} \supset \varphi_\alpha(\mathcal{W}) = (-\infty, a) \xrightarrow{\varphi_\beta \circ \varphi_\alpha^{-1}} (b, \infty) = \varphi_\beta(\mathcal{W}) \subset \mathbb{R},$$

being a homeomorphism between two open intervals in $\mathbb{R}$, is a monotone function whose value approaches $\pm\infty$ at the bounded end of its domain, and the same applies to its inverse, implying that this transition map also has a *finite* limit at the unbounded end of its domain. Now if $x_n \in \mathcal{W}$ is any sequence that has no subsequence converging to any point in $\mathcal{W}$ or to p_β, it follows that $|\varphi_\beta(x_n)| \to \infty$ and thus $\varphi_\alpha(x_n) \to a$, implying $x_n \to p_\alpha$. This proves that the union of $\mathcal{W}$ with the two points p_α, p_β is compact, as claimed. Putting the obvious topology on the extended interval $[b, \infty]$, φ_β now has a unique extension to a homeomorphism $\overline{\mathcal{W}} \to [b, \infty]$ that sends $p_\alpha \mapsto \infty$, so $\overline{\mathcal{W}}$ has the topology of a compact interval. □

Note that in the setting of the lemma, $\mathcal{U}_\alpha \cap \mathcal{U}_\beta$ may in general have multiple connected components, but the proof showed that a homeomorphism $\varphi_\alpha : \mathcal{U}_\alpha \to \mathbb{R}$ sends each of them to an unbounded half-interval. Here's a useful fact we know about $\mathbb{R}$: you can't fit more than two disjoint unbounded half-intervals into it!

Corollary 18.21. *In the setting of Lemma 18.20, $\mathcal{U}_\alpha \cap \mathcal{U}_\beta$ has either one or two connected components.* □

Exercise 18.22. Show that the compact non-Hausdorff space in Example 18.16 admits an open covering by two sets homeomorphic to $\mathbb{R}$ whose intersection with each other has three connected components.

Proof of Theorem 18.18. Given a nonempty connected 1-manifold M without boundary, every point has an open neighborhood homeomorphic to $\mathbb{R}$, and since M is second countable, we can cover M with a *finite or countable* collection $\{\mathcal{U}_n \subset M\}_{n=1}^N$ of such neighborhoods with homeomorphisms $\varphi_n : \mathcal{U}_n \to \mathbb{R}$; here N is either a natural number or ∞. After removing some of these sets from the collection, we can assume without loss of generality that none of them are contained in any one of the others.

If $N = 1$, then M is homeomorphic to $\mathbb{R}$, and we are done.

If $N \geqslant 2$, then since M is also Hausdorff and connected, we can appeal to Lemma 18.20 and Corollary 18.21 in order to relabel the subsets $\{\mathcal{U}_n\}_{n=1}^N$ in the following manner. Choose $\mathcal{U}_1$ to be an arbitrary set in the collection. By definition $\mathcal{U}_1$ is an open subset of M, but it might also be a closed subset—if it is, then since M is connected, we can conclude that $M = \mathcal{U}_1 \cong \mathbb{R}$, so again we are done. If however $\mathcal{U}_1 \subset M$ is not a closed subset, then it is not the complement of any open set, and in particular it is not the complement of the union of the rest of the sets in our collection, which means at least one of them—which we shall now call $\mathcal{U}_2$—must intersect $\mathcal{U}_1$. There are now three possibilities:

1. If $\mathcal{U}_1 \cap \mathcal{U}_2$ has two connected components, one can deduce from Lemma 18.20 that $\mathcal{U}_1 \cup \mathcal{U}_2$ is homeomorphic to S^1, which is compact and is therefore (since M is Hausdorff) a closed subset of M. Since it is clearly also an open subset and M is connected, this implies $M = \mathcal{U}_1 \cup \mathcal{U}_2 \cong S^1$, so we are done.
2. If $\mathcal{U}_1 \cap \mathcal{U}_2$ has only one connected component, then $\mathcal{U}_1 \cup \mathcal{U}_2$ must be homeomorphic to $\mathbb{R}$. If $\mathcal{U}_1 \cup \mathcal{U}_2$ is also a closed subset of M, then connectedness again implies $M = \mathcal{U}_1 \cup \mathcal{U}_2 \cong \mathbb{R}$, and we are done.

(3) If $\mathcal{U}_1 \cap \mathcal{U}_2$ has only one connected component and the subset $\mathcal{U}_1 \cup \mathcal{U}_2 \subset M$ is not closed, then appealing again to the fact that M is connected, $\mathcal{U}_1 \cup \mathcal{U}_2$ must intersect one of the remaining subsets in our collection, which we shall now call $\mathcal{U}_3$.

Now repeat the previous step like so: if $(\mathcal{U}_1 \cup \mathcal{U}_2) \cap \mathcal{U}_3$ has two connected components, we can conclude $M = \mathcal{U}_1 \cup \mathcal{U}_2 \cup \mathcal{U}_3 \cong S^1$, and if not, then $\mathcal{U}_1 \cup \mathcal{U}_2 \cup \mathcal{U}_3 \cong \mathbb{R}$ and either this is all of M or it has nonempty intersection with one of the remaining sets in the collection. If the latter happens, repeat. And so on.

If N is finite, this process eventually exhausts all the sets $\mathcal{U}_1, \dots, \mathcal{U}_N$ and produces a homeomorphism of M to either S^1 or $\mathbb{R}$, the former if an intersection with two connected components ever occurs, and the latter otherwise.

If N is infinite, the process may still terminate if an intersection with two connected components appears, implying that finitely many of the sets $\mathcal{U}_n$ cover M and it is homeomorphic to S^1.

The remaining possibility is that the process never terminates, but instead produces a countable sequence of nested open subsets

$$I_1 \subset I_2 \subset I_3 \subset \dots \bigcup_{n=1}^{\infty} I_n = M,$$

where each $I_n := \mathcal{U}_1 \cup \dots \cup \mathcal{U}_n$ is homeomorphic to $\mathbb{R}$ and is obtained from I_{n-1} by gluing two copies of $\mathbb{R}$ together along a pair of connected half-intervals of infinite length. Up to homeomorphism, we could instead describe this process as follows: identify I_1 with $(0, 1)$, and by induction, if I_{n-1} for some $n \geqslant 2$ has been identified with a finite interval (a, b), then I_n is identified with the union of (a, b) and another finite open interval that contains either a or b in its interior and has an end point in (a, b). Up to homeomorphism, we can thus assume $I_{n-1} = (a, b)$ and I_n is either $(a - 1, b)$ or $(a, b + 1)$. Continuing this process indefinitely, the union $\bigcup_{n=1}^{\infty} I_n$ gets identified with some subinterval in $\mathbb{R}$, and is thus homeomorphic to $\mathbb{R}$. □

The second countability axiom became relevant in the last step of this proof because M was presented as the union of a *countable* collection of intervals; if we had been forced to assume that the collection of Euclidean neighborhoods covering M was uncountable, we would not have been able to conclude in the same manner that M is homeomorphic to $\mathbb{R}$. I would now like to describe an example showing that this danger is serious, and that something other than S^1 or $\mathbb{R}$ can indeed arise if the second countability axiom is dropped. We will need to appeal to a rather non-obvious result from elementary set theory. Recall that a **totally ordered set** $(I, \prec)$ consists of a set I with a partial order $\prec$ such that for all pairs of elements $x, y \in I$, at least one of the conditions $x \prec y$ or $y \prec x$ holds. Such a set is said to be **well ordered** if every subset of I contains a smallest element. The most familiar example of a well-ordered set is the natural numbers. For the purposes of our example below, we need a well-ordered set that is uncountable.

Lemma 18.23. *There exists an uncountable well-ordered set $(\omega_1, \leqslant)$ such that for every $x \in \omega_1$, at most countably many elements $y \in \omega_1$ satisfy $y \leqslant x$.*

Understanding this lemma requires some knowledge of the **ordinal numbers** (*Ordinalzahlen*), which we do not have time to describe here in detail, but the intuitive idea is to think of any well-ordered set as a "number," call two such numbers equivalent if there exists an order-preserving bijection from one to the other, and write $x \leqslant y$ whenever there exists an order-preserving injection from x into y. Informally, an ordinal number can be regarded as an equivalence class of well-ordered sets under this notion of equivalence. We can then think of each natural number $n \in \mathbb{N}$ as an ordinal number by identifying it with the set $\{1, \dots, n\}$, and this identification obviously produces the correct ordering relation for the natural numbers. But there are also infinite ordinal numbers, e.g. the set $\mathbb{N}$ itself. Informally again, the set ω_1 in the above lemma is defined to be the "smallest uncountable ordinal".

To see what this really means, we need a slightly more formal definition of the ordinal numbers—the informal description above is a bit hard to make precise in formal set-theoretic terms. A more concrete description of the ordinal numbers was introduced by Johann von Neumann, and the idea is to regard each ordinal number as a set whose elements are also sets, namely each ordinal is the set of all ordinals that precede it. In particular, we label the empty set $\varnothing$ as 0, identify the natural number 1 with the set $\{0\} = \{\varnothing\}$, identify 2 with the set $\{0,1\} = \{\varnothing, \{\varnothing\}\}$, identify

$$3 = \{0,1,2\} = \{\varnothing, \{\varnothing\}, \{\{\varnothing\}\}\}$$

and so forth. Although the notation quickly becomes confusing, one can make sense of von Neumann's general definition:

Definition 18.24. A set S is an ordinal number if and only if S is well ordered with respect to set membership and every element of S is also a subset of S.

If this definition makes your head spin, rest assured that I have the same reaction, but the concept of the ordinal numbers does not rely on anything other than the standard axioms of set theory. With this definition in place, one can define ω_1 as the union of all countable ordinals, which is necessarily uncountable since it would otherwise contain itself.

We now use this to construct a Hausdorff space that is path-connected and locally homeomorphic to $\mathbb{R}$ but is not second countable. This space and various related constructions are sometimes referred to as the **long line**. Let

$$L = \omega_1 \times [0,1),$$

and define a total order on L such that $(x,s) \leqslant (y,t)$ whenever either $x \leqslant y$ or both $x = y$ and $s \leqslant t$ hold. Writing $x < y$ to mean $x \leqslant y$ and $x \neq y$ for $x, y \in L$, the total order determines a natural topology on L, called the **order topology**, whose base is the collection of all "open" intervals

$$(a,b) := \{x \in L \mid a < x < b\}$$

for arbitrary values $a, b \in L$. The proof of the following statement is an amusing exercise for a rainy day.

Proposition 18.25. *Every point of L has a neighborhood homeomorphic to either $\mathbb{R}$ or (in the case of $(0,0) \in L$) the half-interval $[0,\infty)$. Moreover, L is Hausdorff and is sequentially compact, but not compact; in particular the set $\{(x,1/2) \mid x \in \omega_1\} \subset L$ is an uncountable discrete subset of L, implying that L cannot be second countable.* □

I'm guessing you find it especially surprising that this enormous space L is sequentially compact, but that has to do with a peculiar property built into the definition of the set ω_1: every sequence in ω_1 has an upper bound. This is almost immediate from the definition of the ordinal numbers, as for any given sequence $x_n \in \omega_1$, the elements x_n are also (necessarily countable) sets of ordinal numbers, hence their union $\bigcup_n x_n$ is another ordinal number and is countable, meaning it is an element of ω_1, and it clearly bounds the sequence from above.

In dimensions $n \geqslant 2$, there are further constructions of non-second countable but locally Euclidean Hausdorff spaces which do not rely on anything so exotic as the ordinal numbers. An example is the *Prüfer surface*; see the exercise below. But I'm only talking about these things now in order to explain why I will never mention them again.

Exercise 18.26. The **Prüfer surface** is an example of a space that would be a connected 2-dimensional manifold if we did not require manifolds to be second countable. It is defined as follows: let $\mathbb{H} = \{(x, y) \in \mathbb{R}^2 \mid y > 0\}$, and associate to each $a \in \mathbb{R}$ a copy of the plane $X_a := \mathbb{R}^2$. The Prüfer surface is then

$$\Sigma := \mathbb{H} \amalg \left(\coprod_{a \in \mathbb{R}} X_a \right) \Big/ \sim$$

where the equivalence relation identifies each point $(x, y) \in X_a$ for $y > 0$ with the point $(a + yx, y) \in \mathbb{H}$. Notice that $\mathbb{H}$ and X_a for each $a \in \mathbb{R}$ can be regarded naturally as subspaces of Σ.

(a) Prove that Σ is Hausdorff.
(b) Prove that Σ is path-connected.
(c) Prove that every point in Σ has a neighborhood homeomorphic to $\mathbb{R}^2$.
(d) Prove that a second countable space can never contain an uncountable discrete subset. Then find an uncountable discrete subset of Σ.

19. Surfaces and triangulations (June 22, 2023)

As far as I'm aware, dimension one is the only case in which the problem of classifying *arbitrary* (compact or noncompact) manifolds up to homeomorphism has a reasonable solution. In this lecture we will do the next best thing in dimension two: we will classify all *compact* surfaces. We will focus in particular on closed and connected surfaces. The classification of compact connected surfaces with boundary can easily be derived from this (see Exercise 20.13), and of course compact disconnected surfaces are all just disjoint unions of finitely many connected surfaces, so we lose no generality by restricting to the connected case.

Let us first enumerate the closed connected surfaces that we are already familiar with.

Examples 19.1. The sphere $S^2 = \Sigma_0$ and torus $\mathbb{T}^2 = \Sigma_1$ are both examples of "oriented surfaces of genus g," which can be defined for any nonnegative integer $g \geqslant 0$ and denoted by Σ_g. In particular, we've seen that for each $g \geqslant 1$, Σ_g is homeomorphic to the g-fold connected sum of copies of $\mathbb{T}^2$, and we have also computed its fundamental group

$$\pi_1(\Sigma_g) \cong \left\{ a_1, b_1, \ldots, a_g, b_g \;\middle|\; \prod_{i=1}^{g} [a_i, b_i] = e \right\},$$

whose abelianization is isomorphic to $\mathbb{Z}^{2g}$.

Examples 19.2. An analogous sequence of surfaces can be defined by taking repeated connected sums of copies of $\mathbb{RP}^2$, e.g. $\mathbb{RP}^2 \# \mathbb{RP}^2$ is homeomorphic to the Klein bottle. By the same trick that we used in Lecture 13 to understand Σ_g, the g-fold connected sum $\#_{i=1}^g \mathbb{RP}^2$ is homeomorphic to a space obtained from a polygon with $2g$ edges by identifying them in pairs according to the sequence $a_1, a_1, \ldots, a_g, a_g$, thus

$$\pi_1\left(\#_{i=1}^g \mathbb{RP}^2\right) \cong \left\{ a_1, \ldots, a_g \;\middle|\; a_1^2 \ldots a_g^2 = e \right\}.$$

Exercise 19.3. For $i = 1, \ldots, g-1$, let $e_i \in \mathbb{Z}^{g-1}$ denote the ith standard basis vector. Show that there is a well-defined homomorphism $G := \{a_1, \ldots, a_g \mid a_1^2 \ldots a_g^2 = e\} \to \mathbb{Z}^{g-1} \oplus \mathbb{Z}_2$ such that

$$a_i \mapsto \begin{cases} (e_i, 0) & \text{for } i = 1, \ldots, g-1, \\ (-1, \ldots, -1, 1) & \text{for } i = g, \end{cases}$$

and that it descends to an isomorphism of the abelianization of G to $\mathbb{Z}^{g-1} \oplus \mathbb{Z}_2$.

Appealing to the standard classification of finitely generated abelian groups, we deduce from the above exercise that all of our examples so far are topologically distinct:

LEMMA 19.4. *No two of the closed surfaces listed in Examples 19.1 and 19.2 are homeomorphic.* □

You might now be wondering whether new examples can be constructed by taking the connected sum of a surface from Example 19.1 with some surface from Example 19.2. The answer is no:

PROPOSITION 19.5. $\mathbb{RP}^2 \# \mathbb{T}^2$ *is homeomorphic to the connected sum of* $\mathbb{RP}^2$ *with the Klein bottle.*[26]

PROOF. Given any surface Σ with two disjoint disks removed, one can construct a new surface by attaching a "handle" of the form $[-1,1] \times S^1$:

$$\Sigma' := \left(\Sigma \backslash (\mathring{\mathbb{D}}^2 \amalg \mathring{\mathbb{D}}^2)\right) \cup_{S^1 \amalg S^1} \left([-1,1] \times S^1\right).$$

This operation is essentially the same as the connected sum, except we allow the two disks to be embedded (disjointly) into a single surface Σ rather than two separate surfaces; we sometimes call this a "self-connected sum". As with the connected sum, it depends on a choice of embedding

$$i_1 \amalg i_2 : \mathbb{D}^2 \amalg \mathbb{D}^2 \hookrightarrow \Sigma,$$

but only up to homotopy through embeddings, i.e. modifying the embedding through a continuous 1-parameter family of embeddings will change Σ' into something homeomorphic to the original Σ'.

Let us now shift our perspective on the operation that changes Σ into Σ'. For this it would be helpful to have some pictures, and I do not have time to draw them, but I recommend having a look at Figure 1 in [**FW99**]. Suppose the two holes you're drilling in Σ are right next to each other, but before you drill them, you push the surface up a bit from underneath, creating a disk-shaped lump. Now pick two smaller disk-shaped areas within that lump and push those up even further. Then drill the holes in those two places and attach the handle. We haven't changed any of the topology in creating these "lumps," but we have changed the picture, and if you're imagining it the way that I intended, it now looks like instead of cutting out two holes and attaching a handle, you cut out *one* hole (the base of the original lump) and attached $\Sigma_{1,1}$, the torus with a disk removed. In other words, you performed the connected sum of Σ with $\mathbb{T}^2$:

$$\Sigma' \cong \Sigma \# \mathbb{T}^2.$$

So far so good. . . now let's modify the procedure once more. Viewing $\mathbb{D}^2$ as the unit disk in $\mathbb{C}$, let's replace one of our embeddings $i_1 : \mathbb{D}^2 \to \Sigma$ with another one that has the same image but changes the parametrization by complex conjugation:

$$i_1' : \mathbb{D}^2 \hookrightarrow \Sigma : z \mapsto i_1(\bar{z}).$$

While we will now be cutting out the same two holes in Σ, the way that we attach the handle at the first hole needs to change because $i_1'|_{\partial \mathbb{D}^2}$ parametrizes the circle in the opposite direction from $i_1|_{\partial \mathbb{D}^2}$. The effect is the same as if you were to cut open Σ' along the circle at the boundary of the first hole, flip it's orientation and then glue it back together. Unfortunately you cannot do this in 3-dimensional space—for the same reasons that you cannot embed a Klein bottle into $\mathbb{R}^3$—but it's easy to define the topological space that results from this modification. The effect is precisely to replace the torus in the above description of a connected sum with the Klein bottle; if we call Σ'' the space that results from attaching the handle along this modified gluing map, we have

$$\Sigma'' \cong \Sigma \# K^2,$$

where K^2 denotes the Klein bottle.

[26]This proposition has its very own Youtube video, see https://www.youtube.com/watch?v=aBbDvKq4JqE&t=20s. Maybe you'll find it helpful. . . I'm not entirely sure if I did.

Finally, let's specify this to the case $\Sigma = \mathbb{RP}^2$. The projective plane has a special property that many surfaces don't: it contains an embedded Möbius band, call it $\mathbb{M}$. Now suppose we construct $\mathbb{RP}^2 \# \mathbb{T}^2$ by embedding two small disks disjointly into $\mathbb{M} \subset \mathbb{RP}^2$, then cutting both out and gluing in a handle. By the previous remarks, the homeomorphism type of the resulting surface will not change if we now move the first hole continuously along a circle traversing $\mathbb{M}$, and the orientation reversal as we traverse $\mathbb{M}$ thus allows us to deform $i_1 : \mathbb{D}^2 \hookrightarrow \mathbb{RP}^2$ to $i_1' : \mathbb{D}^2 \hookrightarrow \mathbb{RP}^2$ through a continuous family of embeddings disjoint from the second disk. This proves that if $\Sigma = \mathbb{RP}^2$, then the two surfaces Σ' and Σ'' described above are homeomorphic. $\square$

It is sometimes useful to make a distinction between two types of handle attachment that were described in the above proof. In one case, the two holes $\mathbb{D}^2 \hookrightarrow \Sigma$ are embedded "right next to each other" and with opposite orientations—in precise terms, this means we focus on the domain of a single chart on Σ, assume both holes are in this domain, define i_1' by translating the image of i_2 in some direction to make it disjoint, and then define $i_1(z) = i_1'(\bar{z})$. The handle attachment that results is straightforward to draw, see e.g. Figure 1 in [**FW99**]. If we then leave the positions of the two holes the same but reverse an orientation by replacing i_1 with i_1', the handle attachment can no longer be embedded in $\mathbb{R}^3$, though this does not stop some authors from trying to draw pictures of it anyway (see Figure 2 in [**FW99**]). This type of handle attachment is sometimes referred to as a **cross-handle**. One should not take this terminology too seriously since the main point of the above prove was that in certain cases such as $\Sigma = \mathbb{RP}^2$, there is no globally meaningful distinction between ordinary handles and cross-handles, i.e. if the two holes do not lie in the same chart, it is not always possible to say that we are dealing with one type of handle and not the other. The distinction does make sense however if both holes are in the same chart, so we will occasionally also use the term "cross-handle" in this situation.

Proposition 19.5 told us that the most obvious way to produce new examples of closed connected surfaces out of the inventory in Examples 19.1 and 19.2 does not actually give anything new. The reason for this turns out to be that there are no others:

THEOREM 19.6. *Every closed connected surface is homeomorphic to either Σ_g for some $g \geqslant 0$ or $\#_{i=1}^g \mathbb{RP}^2$ for some $g \geqslant 1$, where the integer g is in each case unique.*

The uniqueness in this statement already follows from the computations of fundamental groups explained above, so in light of Proposition 19.5, we only still need to show that every closed connected surface other than the sphere is homeomorphic to something constructed out of copies of $\mathbb{T}^2$ and $\mathbb{RP}^2$ by connected sums. (Note that whenever both $\mathbb{T}^2$ and $\mathbb{RP}^2$ appear in this collection, Prop. 19.5 allows us to replace $\mathbb{T}^2$ with two copies of $\mathbb{RP}^2$, as $\mathbb{RP}^2 \# \mathbb{RP}^2$ is the Klein bottle.) We will sketch a proof of this below that is due to John Conway and known colloquially as Conway's "ZIP proof". Another readable account of it is given in [**FW99**].

To frame the problem properly, let us say that for Σ a compact (but not necessarily closed or connected) surface, Σ is *ordinary* if there is a finite sequence of compact surfaces

$$\Sigma^{(0)}, \Sigma^{(1)}, \dots, \Sigma^{(m)} = \Sigma$$

such that $\Sigma^{(0)}$ is a finite disjoint union of spheres $\coprod_{i=1}^N S^2$, and each $\Sigma^{(j+1)}$ is homeomorphic to something obtained from $\Sigma^{(j)}$ by performing one of the following operations:

(1) Removing an open disk from the interior, i.e.

$$\Sigma^{(j+1)} \cong \Sigma^{(j)} \backslash \mathring{\mathbb{D}}^2$$

for some embedding $\mathbb{D}^2 \hookrightarrow \Sigma^{(j)} \backslash \partial \Sigma^{(j)}$;

(2) Attaching a handle (or "cross-handle") to connect two separate boundary components $\ell_1, \ell_2 \subset \partial\Sigma^{(j)}$, i.e.
$$\Sigma^{(j+1)} \cong \Sigma^{(j)} \cup_{\ell_1 \amalg \ell_2} ([-1,1] \times S^1)$$
for some choice of homeomorphism $\partial([-1,1] \times S^1) = S^1 \amalg S^1 \to \ell_1 \amalg \ell^2$;
(3) Attaching a disk (called a **cap**) to a boundary component $\ell \subset \partial\Sigma^{(j)}$, i.e.
$$\Sigma^{(j+1)} \cong \Sigma^{(j)} \cup_\ell \mathbb{D}^2$$
for some choice of homeomorphism $\partial\mathbb{D}^2 = S^1 \to \ell$;
(4) Attaching a Möbius band (called a **cross-cap**) $\mathbb{M}$ to a boundary component $\ell \subset \partial\Sigma^{(j)}$, i.e.
$$\Sigma^{(j+1)} \cong \Sigma^{(j)} \cup_\ell \mathbb{M}$$
for some choice of homeomorphism $\partial\mathbb{M} \cong S^1 \to \ell$.

The classification of 1-manifolds is implicitly in the background of the last three operations: since $\Sigma^{(j)}$ is a compact 2-manifold, $\partial\Sigma^{(j)}$ is a closed 1-manifold and is therefore always a finite disjoint union of circles. Observe now that each of the operations can be reinterpreted in terms of connected sums, e.g. cutting out two holes and then attaching a handle or cross-handle is equivalent to taking the connected sum with $\mathbb{T}^2$ or $\mathbb{RP}^2 \# \mathbb{RP}^2$, while attaching a cap or cross-cap gives connected sums with S^2 or $\mathbb{RP}^2$ respectively. It follows that any ordinary surface that is also closed and connected necessarily belongs to our existing inventory of closed and connected surfaces, thus it will suffice to prove:

Lemma 19.7. *Every closed surface is ordinary.*

At this point in almost every topology class, it becomes necessary to cheat a bit and appeal to a fundamental result about surfaces that is believable and yet far harder to prove than we have time to discuss in any detail. I'm referring to the existence of *triangulations*. This is not only a useful tool in classifying surfaces, but also will play a large motivational role when we introduce homology. The following is thus simultaneously a necessary digression behind the proof of Lemma 19.7 and also a preview of things to come.

The idea of a triangulation is to decompose a topological n-manifold into many homeomorphic pieces that we think of as "n-dimensional triangles". More precisely, the **standard n-simplex** is defined as the set
$$\Delta^n := \left\{ (t_0, \ldots, t_n) \in I^{n+1} \;\middle|\; t_0 + \ldots + t_n = 1 \right\}$$
for each integer $n \geqslant 0$. This makes Δ^0 the one-point space $\{1\} \subset \mathbb{R}$, while Δ^1 is a compact line segment in $\mathbb{R}^2$ homeomorphic to the interval I, Δ^2 is the compact region in a plane bounded by a triangle, Δ^3 is the compact region in a 3-dimensional vector space bounded by a tetrahedron, and so forth. For a surface Σ, we would now like to view copies of Δ^2 as fundamental building blocks of Σ, arranged in such a way that the intersection between any two of those building blocks is either empty or is a copy of Δ^1 or Δ^0. One can express this condition in purely combinatorial terms by thinking of Δ^n as the convex hull of its $n+1$ *vertices*, which are the standard basis vectors of $\mathbb{R}^{n+1}$. In this way, an n-simplex is always determined by $n+1$ vertices, and this idea can be formalized via the notion of a *simplicial complex*.

Definition 19.8. A **simplicial complex** (*Simplizialkomplex*) K consists of two sets V and S, called the sets of **vertices** (*Eckpunkte*) and **simplices** (*Simplizes*) respectively, where the elements of S are nonempty finite subsets of V, and $\sigma \in S$ is called an **n-simplex** of K if it has $n+1$ elements. We require the following conditions:
(1) Every vertex $v \in V$ gives rise to a 0-simplex in K, i.e. $\{v\} \in S$;
(2) If $\sigma \in S$ then every subset $\sigma' \subset \sigma$ is also an element of S.

For any n-simplex $\sigma \in S$, its subsets are called its **faces** (*Seiten* or *Facetten*), and in particular the subsets that are $(n-1)$-simplices are called **boundary faces** (*Seitenflächen*) of σ. The second condition above thus says that for every simplex in the complex, all of its boundary faces also belong to the complex. With this condition in place, the first condition is then equivalent to the requirement that every vertex in the set V belongs to at least one simplex.

The complex K is said to be **finite** if V is finite, and it is n**-dimensional** if

$$\sup_{\sigma \in S} |\sigma| = n + 1,$$

i.e. n is the largest number for which K contains an n-simplex.

Though the definition above is purely combinatorial, there is a natural way to associate a topological space $|K|$ to any simplicial complex K. We shall describe it only in the case of a finite complex,[27] since that is what we need for our discussion of compact surfaces. Given $K = (V, S)$, choose a numbering of the vertices $V = \{v_1, \dots, v_N\}$ and associate to each k-simplex $\sigma = \{v_{i_0}, \dots, v_{i_k}\}$ the set

$$\Delta_\sigma := \left\{ (t_1, \dots, t_N) \in I^N \;\middle|\; t_{i_0} + \ldots + t_{i_k} = 1 \text{ and } t_j = 0 \text{ for all } v_j \notin \sigma \right\}.$$

Notice that Δ_σ is homeomorphic to the standard k-simplex Δ^k, but lives in the subspace of $\mathbb{R}^N$ spanned by the specific coordinates corresponding to its vertices. The **polyhedron** (*Polyeder*) of K is then the compact space

$$|K| := \bigcup_{\sigma \in S} \Delta_\sigma \subset \mathbb{R}^N.$$

While the definition above makes $|K|$ a subset of a Euclidean space that may have very large dimension in general, it is not so hard to picture $|K|$ in a few simple examples.

Example 19.9. Suppose $V = \{v_0, v_1, v_2\}$ and S is defined to consist of all subsets of V. Then $|K|$ is just the standard 2-simplex Δ^2.

Example 19.10. Suppose $V = \{v_0, v_1, v_2, v_3\}$ and S contains the subsets $A := \{v_0, v_1, v_2\}$ and $B := \{v_1, v_2, v_3\}$, plus all of their respective subsets. Then $|K|$ contains two copies of the triangle Δ^2, which we can label A and B, and they intersect each other along a single common edge connecting the vertices labeled v_1 and v_2. In particular, $|K|$ is homeomorphic to a 2-dimensional square I^2, formed by gluing two triangles together along one edge.

Definition 19.11. A **triangulation** (*Triangulierung*) of a compact topological n-manifold M is a homeomorphism of M to the polyhedron of a finite n-dimensional simplicial complex.

In particular, this makes precise the notion of decomposing a surface Σ into triangles (copies of Δ^2) whose intersections with each other are always simplices of lower dimension. Observe that in a triangulated surface Σ with $\partial\Sigma = \varnothing$, the fact that every point in one of the 1-simplices σ has a neighborhood homeomorphic to $\mathbb{R}^2$ implies that σ is a boundary face of *exactly two* 2-simplices in the triangulation. One can say the same about the $(n-1)$-simplices in any triangulation of a closed n-manifold. This is not a property that arbitrary simplicial complexes have, but it is a general property of the complexes that appear in triangulations of closed manifolds.

Theorem 19.12. *Every closed surface admits a triangulation.*

---

[27] The polyhedron of a finite simplicial complex has an obvious topology because it comes with an embedding into some finite-dimensional Euclidean space. For infinite complexes this is not true, and thus more thought is required to define the ***right*** topology on $|K|$. We would need to talk about this if we wanted to define triangulations of noncompact spaces, but since we don't want that right now, we will not. The correct topology on infinite complexes will be discussed next semester when we generalize all this to CW-complexes; see Lecture 36.

This theorem is old enough for the first proof to have been published in German [**Rad25**], and it was not the main result of the paper in which it appeared, yet it is in some sense far harder than it has any right to be—it seems to be one of the rare instances in mathematics where learning cleverer high-powered techniques does not really help. I can at least sketch what is involved. Since a closed surface Σ can be covered by finitely many charts, it can also be covered by a finite collection of regions homeomorphic to $\mathbb{D}^2$, which is homeomorphic to the standard 2-simplex Δ^2. Of course the interiors of these 2-simplices overlap, which is not allowed in a triangulation, but the idea is to examine each of the overlap regions and subdivide it further into simplices. By "overlap region," what I mean is the following: if $D_1, \ldots, D_N \subset \Sigma$ denote the finite collection of disks $D_i \cong \Delta^2$ covering Σ, whose boundaries are loops ∂D_i, then the closure of each connected component of $\Sigma \backslash \bigcup_i \partial D_i$ is a region that needs to be subdivided into triangles. After perturbing each of the disks D_i so that its boundary intersects the other boundaries only finitely many times, we can arrange for each of these overlap regions to be bounded by embedded circles, and notice that since each of the regions is contained in at least one of the disks D_i, we can view them as subsets of $\mathbb{R}^2$. Now, I don't know about you, but I find it not so hard to believe that regions in $\mathbb{R}^2$ bounded by embedded circles can be subdivided into triangles in a reasonable way—I would imagine that writing down a complete algorithm to do this is a pain in the neck, but it sounds plausible. It may surprise you however to know that it is very far from obvious what the region bounded by an embedded circle in $\mathbb{R}^2$ can look like in general. Actually the answer is simple and is what you would expect: the region is homeomorphic to a disk, but this is not at all easy to prove, it is an important theorem in classical topology known as the *Schönflies theorem*. With this result in hand, one can formulate an algorithm for triangulating surfaces as sketched above by triangulating the disk-like overlap regions. Complete accounts of this are given in [**Moi77**] and [**Tho92**].

Note that if Σ is not just a topological 2-manifold but also has a *smooth* structure, then one can avoid the Schönflies theorem by appealing to some basic facts from Riemannian geometry. Choosing a Riemannian metric allows us to define the notion of a "straight line" (geodesic) on the manifold, and one can arrange in this case for the disks D_i to be convex, so that the overlap regions are also convex and therefore obviously homeomorphic to disks. This trick actually works in arbitrary dimensions, leading to the result that *smooth* manifolds can be triangulated in any dimension. For topological manifolds this is not true in general: it is true in dimension three (see [**Moi77**]), but from dimension four upwards there are examples of topological manifolds that do not admit triangulations. The case of dimension five has only been understood since 2013—see [**Man14**] for a readable survey of this subject and its history.

But enough about triangulations: let's just assume that surfaces can be triangulated and use this to finish the classification theorem.

Proof of Lemma 19.7. Assume Σ is a closed surface homeomorphic to the polyhedron $|K|$ of a finite 2-dimensional simplicial complex $K = (V, S)$ with 2-simplices $\sigma_1, \ldots, \sigma_N$. By abuse of notation, we shall also denote by $\sigma_1, \ldots, \sigma_N$ the corresponding subsets of Σ homeomorphic to the standard 2-simplex Δ^2. The latter is homeomorphic to $\mathbb{D}^2 \cong S^2 \backslash \mathring{\mathbb{D}}^2$, thus

$$\Sigma^{(0)} := \sigma_1 \amalg \ldots \amalg \sigma_N$$

is ordinary. The idea now is to reconstruct Σ from this disjoint union by gluing pairs of 2-simplices together along corresponding boundary faces one at a time, producing a sequence of compact surfaces $\Sigma^{(j)}$, each of which may be disconnected and have nonempty boundary except for the last in the sequence, which is Σ. The operation changing $\Sigma^{(j)}$ to $\Sigma^{(j+1)}$ is performed by gluing together two arcs $\ell_1, \ell_2 \subset \partial \Sigma^{(j)}$, i.e. we can write

$$\Sigma^{(j+1)} = \Sigma^{(j)} / \sim \qquad \text{where} \qquad \sim \text{ identifies } \ell_1 \text{ with } \ell_2,$$

with ℓ_1 and ℓ_2 assumed to be individual boundary faces of two distinct 2-simplices. These boundary faces are each homeomorphic to the compact interval I, and their interiors are disjoint subsets of $\Sigma^{(j)}$, but they may have boundary points (vertices of the triangulation) in common if some neighboring pair of corresponding boundary faces has already been glued together in the process of turning $\Sigma^{(0)}$ into $\Sigma^{(j)}$. One can now imagine various scenarios, based on the knowledge (thanks to the classification of 1-manifolds) that every connected component of $\partial\Sigma^{(j)}$ is a circle:

Case 1: $\ell_1 \cup \ell_2$ forms a single connected component of $\partial\Sigma^{(j)}$. Gluing them together is then equivalent to attaching either a cap or a cross-cap to that boundary component, depending on the orientation of the homeomorphism that identifies them.

Case 2: ℓ_1 and ℓ_2 form part of a single connected component of $\partial\Sigma^{(j)}$, but not all of it, i.e. their boundary vertices are not exactly the same, so that there are either one or two gaps between them forming additional arcs on some circle in $\partial\Sigma^{(j)}$. Gluing them together then is equivalent to attaching a cap or cross-cap as in case 1, except that it leaves one or two holes where the gaps were, so we can realize this operation by attaching the cap/cross-cap and drilling holes afterward.

Case 3: ℓ_1 and ℓ_2 lie on different connected components of $\partial\Sigma^{(j)}$. Then neither can be the entirety of a boundary component since both are homeomorphic to I instead of S^1, though it's useful to imagine what would happen if both really were the entirety of a boundary component: gluing them together would then be equivalent to attaching a handle. The useful way to turn this picture into reality is to imagine both ℓ_1 and ℓ_2 as making up *most* of their respective boundary components, each leaving a very small gap where their end points fail to come together. Gluing ℓ_1 to ℓ_2 is then equivalent to attaching a handle but then drilling a small hole in it.

In all of these cases, the operation that converts $\Sigma^{(j)}$ into $\Sigma^{(j+1)}$ can be realized by a finite sequence of operations from our stated list, so carrying out this procedure as many times as necessary to convert $\Sigma^{(0)}$ into Σ produces a surface that is ordinary. $\square$

EXERCISE 19.13. Recall that if Σ is a surface with boundary, the **boundary** $\partial\Sigma$ is defined as the set of all points $p \in \Sigma$ such that some chart $\varphi : \mathcal{U} \xrightarrow{\cong} \Omega \subset \mathbb{H}^2$ defined on a neighborhood $\mathcal{U} \subset \Sigma$ of p satisfies $\varphi(p) \in \partial\mathbb{H}^2$. Here $\mathbb{H}^2 := [0,\infty) \times \mathbb{R} \subset \mathbb{R}^2$, $\partial\mathbb{H}^2 := \{0\} \times \mathbb{R} \subset \mathbb{H}^2$, and Ω is an open subset of $\mathbb{H}^2$. One can analogously define $p \in \Sigma$ to be an *interior point* of Σ of some chart maps it to $\mathbb{H}^2\backslash\partial\mathbb{H}^2$. Prove that no point on $\partial\Sigma$ is also an interior point of Σ.
Hint: If you have two charts defined near p such that one sends p to $\partial\mathbb{H}^2$ while the other sends it to $\mathbb{H}^2\backslash\partial\mathbb{H}^2$, then a transition map relating these two charts maps some neighborhood in $\mathbb{H}^2$ of a point $x \in \mathbb{H}^2\backslash\partial\mathbb{H}^2$ to a neighborhood in $\mathbb{H}^2$ of a point $y \in \partial\mathbb{H}^2$. What happens to this homeomorphism if you remove the points x and y? Think about the fundamental group.
Remark: A similar result is true for topological manifolds of arbitrary dimension, but you do not yet have enough tools at your disposal to prove this. A proof using singular homology will be possible before the end of the semester.

EXERCISE 19.14. This exercise concerns manifolds with smooth structures, which were discussed briefly in Lecture 18 (see especially Definition 18.10 and Theorem 18.11). We will need the following additional notions:

- For two smooth manifolds M and N, a map $f : M \to N$ is called **smooth** if for every pair of smooth charts ψ_β on N and φ_α on M, the map $f_{\beta\alpha} := \psi_\beta \circ f \circ \varphi_\alpha^{-1}$ is C^∞ wherever it is defined. (In other words, f is "C^∞ in local coordinates".)
- For $f : M \to N$ a smooth map between smooth manifolds, a point $q \in N$ is a **regular value** of f if for all charts φ_α on M and ψ_β on N such that q is in the domain of ψ_β, $\psi_\beta(q)$ is a regular value of $f_{\beta\alpha}$. (In other words, q is a "regular value of f in local coordinates".)

An easy corollary of the usual implicit function theorem (Theorem 18.11) then states that if M is a smooth m-manifold without boundary, N is a smooth n-manifold and $f : M \to N$ is a smooth map that has $q \in N$ as a regular value, the preimage $f^{-1}(q) \subset M$ is a smooth submanifold[28] of dimension $m-n$. If M has boundary, then one should assume additionally that q is a regular value of the restricted map $f|_{\partial M} : \partial M \to N$, and the conclusion is then that $Q := f^{-1}(q)$ is a smooth manifold of dimension $m-n$ with boundary $\partial Q = Q \cap \partial M$.

We will use the following perturbation lemma as a block box: if M and N are compact smooth manifolds, $q \in N$ and $f : M \to N$ is continuous, then every neighborhood of f in $C(M, N)$ with the compact-open topology (cf. Exercise 7.28) contains a smooth map $f_\epsilon : M \to N$ for which q is a regular value of both f_ϵ and $f_\epsilon|_{\partial M}$. Moreover, if $f|_{\partial M}$ is already smooth and has q as a regular value, then the perturbation can be chosen such that $f_\epsilon|_{\partial M} = f|_{\partial M}$. Proofs of these statements can be found in standard books on differential topology such as [**Hir94**].

If you take all of this as given, then you can use it to define something quite beautiful. Assume M and N are closed connected smooth manifolds of the same dimension n. Then for any smooth map $f : M \to N$ with regular value $q \in N$, the implicit function theorem implies that $f^{-1}(q)$ is a compact 0-manifold, i.e. a finite set of points. Define the **mod 2 mapping degree** $\deg_2(f) \in \mathbb{Z}_2$ of f by

$$\deg_2(f) := |f^{-1}(q)| \pmod 2,$$

i.e. $\deg_2(f)$ is $0 \in \mathbb{Z}_2$ if the number of points in $f^{-1}(q)$ is even, and $1 \in \mathbb{Z}_2$ if it is odd.

(a) Prove that for any given choice of the point $q \in N$, the degree $\deg_2(f) \in \mathbb{Z}_2$ depends only on the homotopy class of the map $f : M \to N$.
Hint: If you have a homotopy $H : I \times M \to N$ between two maps, perturb it as necessary and look at $H^{-1}(q)$. Use the classification of compact 1-manifolds.
Remark: One can show with a little more effort that $\deg_2(f)$ also does not depend on the choice of the point q, and moreover, it has a well-defined extension to continuous (but not necessarily smooth) maps $f : M \to N$, defined by setting $\deg_2(f) := \deg_2(f_\epsilon)$ for any sufficiently close smooth perturbation f_ϵ that has q as a regular value.

(b) Prove that every continuous map $f : S^2 \to S^2$ homotopic to the identity is surjective.

(c) What goes wrong with this discussion of we allow M to be a noncompact manifold? Describe two homotopic maps $f, g : \mathbb{R} \to S^1$ for which $\deg_2(f)$ and $\deg_2(g)$ can be defined in the manner described above but are not equal.

(d) Prove that if $n > m$, every continuous map $S^m \to S^n$ is homotopic to a constant map.
Hint: What does it mean for a point $q \in S^n$ to be a regular value of $f : S^m \to S^n$ if $n > m$?

20. Orientations (June 27, 2023)

This lecture is in part an addendum to the classification of surfaces, though it will also introduce some concepts that will be useful to have in mind when we discuss homology.

I have used the word "orientation" many times in this course without giving any precise explanation of what it means. I want to do that now, at least for manifolds of dimensions one and two. The canonical example to have in mind is the Klein bottle:

---

[28]A subset $Y \subset M$ of a smooth m-manifold M is called a **smooth submanifold** (*glatte Untermannigfaltigkeit*) of dimension k if every point $p \in Y$ has a neighborhood $\mathcal{U} \subset M$ admitting a so-called **slice chart** (*Bügelkarte*), meaning a smooth chart $\varphi : \mathcal{U} \to \mathbb{R}^n$ with the property that $Y \cap \mathcal{U} = \varphi^{-1}(\mathbb{R}^k \times \{0\})$. Covering Y with slice charts then gives Y the structure of a smooth k-manifold for which the inclusion $Y \hookrightarrow M$ is a smooth map. As an important special case: the boundary $\partial M \subset M$ of a smooth m-manifold is always a smooth $(m-1)$-dimensional submanifold.

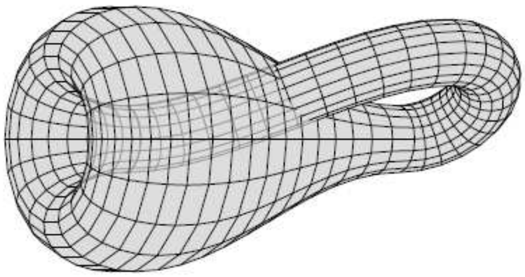

This standard picture of the Klein bottle is unfortunately the image of a *non-injective* map $i : K^2 \to \mathbb{R}^3$ into 3-dimensional Euclidean space from a certain closed 2-manifold K^2: in differential geometry, one would call $i : K^2 \to \mathbb{R}^3$ an *immersion*, which fails to be an *embedding* (and its image is therefore not a *submanifold* of $\mathbb{R}^3$) because one can see a pair of disjoint circles $C_1, C_2 \subset K^2$ such that $i(C_1) = i(C_2)$. For the following informal discussion, however, let us ignore this detail and pretend that $i : K^2 \to \mathbb{R}^3$ is an embedding, with no self-intersections.[29] Now, aside from the fact that it cannot be embedded into $\mathbb{R}^3$, what most of us really find strange about the Klein bottle is that we cannot make a meaningful distinction between the "inside" and the "outside" of the surface. If, for instance, you were an insect and somebody tried to trap you inside a glass Klein bottle, then you could just walk along the surface until you are standing on the opposite side of the glass, and you are free. In mathematical terms, this means that the Klein bottle $K^2 \subset \mathbb{R}^3$ admits an embedded loop $\gamma : I \to K^2$ along which a continuous family of nonzero vectors $V(t) \in \mathbb{R}^3$ can be found which are orthogonal to the surface at each $\gamma(t)$ and satisfy $V(1) = -V(0)$. By contrast, if you take any embedded loop $\gamma : I \to \mathbb{T}^2 \subset \mathbb{R}^3$ on the torus in its standard representation as a tube-like subset of $\mathbb{R}^3$, and choose a normal vector field $V(t)$ along this loop, $V(1)$ will always need to be a positive multiple of $V(0)$. That's because there *is* a meaningful distinction between the outside and inside of the torus $\mathbb{T}^2 \subset \mathbb{R}^3$.[30]

But this discussion of "inside" vs. "outside" is not really satisfactory, because whenever we talk about normal vectors, we are referring to a piece of data that is not intrinsic to the spaces $\mathbb{T}^2$ or K^2. It depends rather on how we choose to embed or immerse them in $\mathbb{R}^3$. So how can we talk about orientations without mentioning normal vectors?

To answer this, imagine again that you are an insect standing on the surface of the Klein bottle, and while standing in place, you turn around in a circle, rotating 360 degrees to your left. An observer from the outside will see you turn, but the *direction* of the turn that observer sees will depend on which side of the glass you are standing on. In particular, if you turn around like this and then follow the aforementioned path to come back to the same point but on the other side of the glass, then when you turn again 360 degrees to the left, the outside observer will see you turning the other way. We can use this turning idea to formulate a precise notion of orientation without mentioning normal vectors.

Informally, let us agree that an orientation of a surface should mean a choice of which kinds of rotations at each point are to be labeled "clockwise" as opposed "counterclockwise". This is still not a precise mathematical definition, but now we are making progress. The term "counterclockwise rotation" has a precise and canonical definition in $\mathbb{R}^2$, for instance, thus we can agree that $\mathbb{R}^2$ has a canonical orientation. The natural thing to do is then to use charts to define orientations

[29] Notice that if we were willing to map K^2 into $\mathbb{R}^4$ instead of $\mathbb{R}^3$, then we could easily turn i into an injective map $K^2 \hookrightarrow \mathbb{R}^4$ just by slightly perturbing the fourth coordinate along C_1 but not along C_2.

[30] The fancy way of saying this in differential-geometric language is that the *normal bundle* of the standard immersion $K^2 \looparrowright \mathbb{R}^3$ is nontrivial, whereas the standard embedding $\mathbb{T}^2 \hookrightarrow \mathbb{R}^3$ has trivial normal bundle. If you don't know what that means, don't worry about it for now.

on a surface Σ via their local identifications with $\mathbb{R}^2$. There's just one obvious problem with this idea: if all charts are allowed, then the definition of an orientation at some point might depend on our choice of chart to use near that point, because the transition map relating two charts might interchange counterclockwise and clockwise rotations. It therefore becomes important to restrict the class of allowed charts so that transition maps do not change orientations, i.e. so that they are *orientation preserving*. Our main task is to give the latter term a precise definition, and this can be done in terms of winding numbers.

Recall the following notion from Exercise 10.27. For $z \in \mathbb{C}$ and $\epsilon > 0$, define a counterclockwise loop about z by

$$\gamma_{z,\epsilon} : S^1 \hookrightarrow \mathbb{C} : e^{i\theta} \mapsto z + \epsilon e^{i\theta}.$$

Note that for fixed $z \in \mathbb{C}$, varying the value of $\epsilon > 0$ does not change the homotopy class of this loop in $\mathbb{C}\backslash\{z\}$, and for a suitable choice of base point it is always a generator of $\pi_1(\mathbb{C}\backslash\{z\}) \cong \mathbb{Z}$. For $k \in \mathbb{Z}$, define also the loop

$$\gamma^k_{z,\epsilon} : S^1 \to \mathbb{C} : e^{i\theta} \mapsto z + \epsilon e^{ki\theta},$$

which covers $\gamma_{z,\epsilon}$ exactly k times if $k > 0$, covers it $|k|$ times with reversed orientation if $k < 0$, and is constant if $k = 0$. Now for any other loop $\alpha : S^1 \to \mathbb{C}\backslash\{z\}$, the **winding number** (*Windungszahl*) of α about z is an integer characterized uniquely by the condition

$$\operatorname{wind}(\alpha; z) = k \qquad \Longleftrightarrow \qquad \alpha \underset{h}{\sim} \gamma^k_{z,\epsilon} \quad \text{in} \quad \mathbb{C}\backslash\{z\}.$$

If $\mathcal{U}, \mathcal{V} \subset \mathbb{C}$ are open subsets and $f : \mathcal{U} \to \mathcal{V}$ is a homeomorphism, then for any $z \in \mathcal{U}$ with $f(z) = w \in \mathcal{V}$, we can assume the loop $\gamma_{z,\epsilon}$ lies in $\mathcal{U}$ for all $\epsilon > 0$ sufficiently small, and the fact that f is bijective makes $f \circ \gamma_{z,\epsilon}$ a loop in $\mathbb{C}\backslash\{w\}$. It follows that there is a well-defined winding number $\operatorname{wind}(f \circ \gamma_{z,\epsilon}; w) \in \mathbb{Z}$, and shrinking $\epsilon > 0$ to a smaller number $\epsilon' > 0$ obviously will not change it since $\gamma_{z,\epsilon}$ and $\gamma_{z,\epsilon'}$ are homotopic in $\mathcal{U}\backslash\{z\}$, so that $f \circ \gamma_{z,\epsilon}$ and $f \circ \gamma_{z,\epsilon'}$ are homotopic in $\mathbb{C}\backslash\{w\}$.

LEMMA 20.1. *In the situation described above,* $\operatorname{wind}(f \circ \gamma_{z,\epsilon}; w)$ *is always either* 1 *or* -1.

PROOF. Choose $\epsilon > 0$ small enough so that the image of $f \circ \gamma_{z,\epsilon}$ lies in a ball $B_r(w)$ about w with radius $r > 0$ sufficiently small such that $B_r(w) \subset \mathcal{V}$. Then for $\delta \in (0, r)$, the homotopy class of $\gamma_{w,\delta}$ generates $\pi_1(B_r(w)\backslash\{w\}) \cong \pi_1(\mathbb{C}\backslash\{w\}) \cong \mathbb{Z}$, and $k := \operatorname{wind}(f \circ \gamma_{z,\epsilon}; w)$ is the unique integer such that $f \circ \gamma_{z,\epsilon}$ is homotopic in $B_r(w)\backslash\{w\}$ to $\gamma^k_{w,\delta}$. Since $\gamma_{z,\epsilon}$ generates $\pi_1(\mathbb{C}\backslash\{z\})$, there is also a unique integer $\ell \in \mathbb{Z}$ such that $f^{-1} \circ \gamma_{w,\delta}$ is homotopic in $\mathbb{C}\backslash\{z\}$ to $\gamma^\ell_{z,\epsilon}$. This implies

$$\gamma_{z,\epsilon} = f^{-1} \circ f \circ \gamma_{z,\epsilon} \underset{h}{\sim} f^{-1} \circ \gamma^k_{w,\delta} \underset{h}{\sim} \gamma^{k\ell}_{z,\epsilon} \quad \text{in} \quad \mathbb{C}\backslash\{z\},$$

hence $k\ell = 1$. Since k and ℓ are both integers, we conclude both are ± 1. $\square$

EXERCISE 20.2. Show that in the setting of Lemma 20.1, the subsets $\mathcal{U}_\pm = \{z \in \mathcal{U} \mid \operatorname{wind}(f \circ \gamma_{z,\epsilon}; f(z)) = \pm 1\}$ are each both open and closed, so in particular, the sign of this winding number is constant on each connected component of $\mathcal{U}$.
Hint: Since the two sets are complementary, it suffices to prove both are open. What happens to $\operatorname{wind}(f \circ \gamma_{z,\epsilon}; w)$ *if you perturb* z *and* w *independently of each other by very small amounts?*

One can define winding numbers just as well for loops in $\mathbb{R}^2$ by identifying $\mathbb{R}^2$ with $\mathbb{C}$ via $(x, y) \leftrightarrow x + iy$. We have been using complex numbers purely for notational convenience, but in the following we will refer instead to domains in $\mathbb{R}^2$ or the half-plane $\mathbb{H}^2$. The discussion also makes sense for homeomorphisms between open subsets of $\mathbb{H}^2$ as long as we only consider points z in the interior $\mathbb{H}^2\backslash\partial\mathbb{H}^2$, since the loop $\gamma_{z,\epsilon}$ is then contained in $\mathbb{H}^2$ for ϵ sufficiently small. Note that by Exercise 19.13, a homeomorphism between open subsets of $\mathbb{H}^2$ always maps points in $\partial\mathbb{H}^2$ to $\partial\mathbb{H}^2$ and points in $\mathbb{H}^2\backslash\partial\mathbb{H}^2$ to $\mathbb{H}^2\backslash\partial\mathbb{H}^2$.

Definition 20.3. Given open subsets $\mathcal{U}, \mathcal{V} \subset \mathbb{H}^2$, a homeomorphism $f : \mathcal{U} \to \mathcal{V}$ is called **orientation preserving** (*orientierungserhaltend*) if $\text{wind}(f \circ \gamma_{z,\epsilon}; f(z)) = 1$ for all $z \in \mathbb{H}^2 \backslash \partial \mathbb{H}^2$ and $\epsilon > 0$ sufficiently small. It is called **orientation reversing** (*orientierungsumkehrend*) if $\text{wind}(f \circ \gamma_{z,\epsilon}; f(z)) = -1$ for all $z \in \mathbb{H}^2 \backslash \partial \mathbb{H}^2$ and $\epsilon > 0$ sufficiently small.

Lemma 20.1 and Exercise 20.2 together imply that a homeomorphism is always either orientation preserving or orientation reversing on each individual connected component. Similar notions can also be defined in all positive dimensions, not only dimension two, though one needs to replace winding numbers with a different way of measuring the local behavior of a homeomorphism in higher dimensions. In dimension one, the proper definition is fairly obvious:

Definition 20.4. Given open subsets $\mathcal{U}, \mathcal{V}$ in $\mathbb{R}$ or $\mathbb{H} := [0, \infty)$, a homeomorphism $f : \mathcal{U} \to \mathcal{V}$ is called **orientation preserving** if it is an increasing function, and **orientation reversing** if it is a decreasing function.

I will refrain for now from stating the definition for dimensions $n \geqslant 3$, since it requires a certain amount of language (involving degrees of maps between spheres) that we have not yet adequately defined. A more straightforward definition is available however if you are willing to restrict from homeomorphisms to *diffeomorphisms*, i.e. bijections that are C^∞ and have C^∞ inverses. Actually, C^1 is good enough: the point is that the derivative $df(x) : \mathbb{R}^n \to \mathbb{R}^n$ of such a map at any point x is guaranteed to be an *invertible* linear map, so it has a nonzero determinant. One then calls the map orientation preserving if the determinant of its derivative is everywhere positive, and orientation reversing if that determinant is everywhere negative. We will not worry about this in the following since we will almost exclusively talk about orientations for manifolds of dimension at most two. Nonetheless, there is no harm in stating a definition of orientation that is valid for topological manifolds of arbitrary dimension, and the definition will look slightly familiar if you recall our discussion of smooth structures in Lecture 18.

Definition 20.5. An **orientation** (*Orientierung*) of an n-manifold M for $n \geqslant 1$ is a maximal collection of charts $\{\varphi_\alpha : \mathcal{U}_\alpha \to \Omega_\alpha\}_{\alpha \in J}$ such that $M = \bigcup_{\alpha \in J} \mathcal{U}_\alpha$ and all transition maps $\varphi_\beta \circ \varphi_\alpha^{-1}$ are orientation preserving. If M is a 0-manifold, we define an orientation on M to be a function $\epsilon : M \to \{1, -1\}$, which partitions M into sets of positively/negatively oriented points $M_\pm := \epsilon^{-1}(\pm 1)$.

We say that M is **orientable** (*orientierbar*) if it admits an orientation, and refer to any manifold endowed with the extra structure of an orientation as an **oriented manifold** (*orientierte Mannigfaltigkeit*).

Specializing again to dimension 2, an orientation of M allows you to draw small loops around arbitrary points in M and label them "counterclockwise" or "clockwise" in a consistent way, where consistency means in effect that you can never deform a counterclockwise loop continuously through small loops around other points and end up with a clockwise loop. The actual definition of counterclockwise comes from the special collection of charts that an orientation provides: we call these **oriented charts**, and define a small loop about a point in M to be counterclockwise if and only if it looks counterclockwise in an oriented chart.

If M is a 1-manifold, then instead of talking about loops or rotations, we can simply label orientations with arrows: the orientation defines which paths in M can be called "increasing" as opposed to "decreasing".

Remark 20.6. One can show that any orientation-preserving homeomorphism between open subsets of $\mathbb{H}^2$ restricts to the boundary as an orientation-preserving homeomorphism between open subsets of $\partial \mathbb{H}^2 \cong \mathbb{R}$. It follows that there is a natural notion of induced **boundary orientation**, i.e. on any orientable surface Σ with boundary, a choice of orientation on Σ induces a natural orientation on $\partial\Sigma$ by taking the oriented charts on the latter to be restrictions of the oriented charts on Σ. An analogous statement is true for manifolds with boundary in all dimensions. For $\dim M = 1$, one defines the boundary orientation of ∂M by setting $\epsilon(p) = 1$ whenever the "increasing" direction of M points from the interior of M toward the boundary point $p \in \partial M$, and $\epsilon(p) = -1$ whenever this direction points from $p \in \partial M$ toward the interior. (Different authors may define this in slightly different ways, but it usually doesn't matter: the point is just to choose a convention and be consistent about it.)

Let us specialize this discussion to manifolds with triangulations, i.e. manifolds that are homeomorphic to the polyhedron of a simplicial complex. The latter is an essentially combinatorial notion, so orientations of such objects can also be defined in combinatorial terms. Recall that if J is any finite set, any bijection $\pi : J \to J$ is a permutation of its elements, that is, one can identify π with some element of the symmetric S_N group on N objects after choosing a numbering $v_1, \ldots, v_N$ for the elements in J. The symmetric group S_N is generated by *flips*, meaning permutations that interchange two elements of J while leaving the rest fixed, and we say that $\pi \in S_N$ is an **even** permutation if it can be written as a composition of evenly many flips; otherwise it is an **odd** permutation. If we represent π by an N-by-N matrix permuting the N standard basis vectors of $\mathbb{R}^N$, then we can recognize the even/odd permutations as those for which this matrix has positive/negative determinant respectively; in fact, the matrices of even permutations always have determinant $+1$, and those of odd permutations have determinant -1. To motivate the next definition, recall the definition of the standard n-simplex $\Delta^n = \{(t_0, \ldots, t_n) \mid t_0 + \ldots + t_n = 1\}$. Any element of the symmetric group on $n+1$ objects can be regarded as a permutation of the vertices of Δ^n numbered from 0 to n, and the matrix representation of this permutation then defines a linear map on $\mathbb{R}^{n+1}$ that permutes the standard basis vectors accordingly. That linear map preserves the subset $\Delta^n \subset \mathbb{R}^{n+1}$, and it is an orientation-preserving transformation on $\mathbb{R}^{n+1}$ if and only if its determinant is positive, which is equivalent to requiring the permutation to be even.

DEFINITION 20.7. For a simplicial complex $K = (V, S)$, an **orientation** of an n-simplex $\sigma \in S$ for $n \geqslant 1$ is an equivalence class of orderings of the vertices $v \in \sigma$, where two orderings are defined to be equivalent if and only if they are related to each other by an even permutation. An orientation of a 0-simplex is defined simply as an assignment of the number $+1$ or -1 to that vertex.

For simplices of dimension 1 or 2 there are easy ways to illustrate in pictures what this definition means; see Figure 11. The figure shows the six possible ways of ordering the three vertices of a 2-simplex, where the individual choices in each row are related to each other by even permutations and thus define equivalent orientations, whereas each choice is related to the one directly underneath it by a single flip, which is an odd permutation. We can represent the orientation itself by drawing a circular arrow that follows the direction of the sequence of vertices labeled $0, 1, 2$, and this arrow depends *only* on the orientation since even permutations of three objects are also *cyclical* permutations.

Another intuitive fact you can infer from Figure 11 is that an orientation of a 2-simplex induces a natural **boundary orientation** for each of its 1-dimensional boundary faces. The latter orientations are represented in the picture by arrows pointing from one vertex to another, meant to indicate the ordering of the two vertices, and the visual recipe is simply that the arrows of all three edges together should describe the same kind of rotation as the circular arrow on the 2-simplex. This can also be reduced to a purely combinatorial algorithm, and it makes sense in every dimension. For an n-simplex $\sigma = \{v_0, \ldots, v_n\}$, the k**th boundary face** $\partial_{(k)}\sigma$ of σ is the $(n-1)$-simplex whose vertices include all the $v_0, \ldots, v_n$ except v_k. Clearly if the vertices $v_0, \ldots, v_n$ come with an ordering, then the vertices of $\partial_{(k)}\sigma$ inherit an ordering from this, though here we

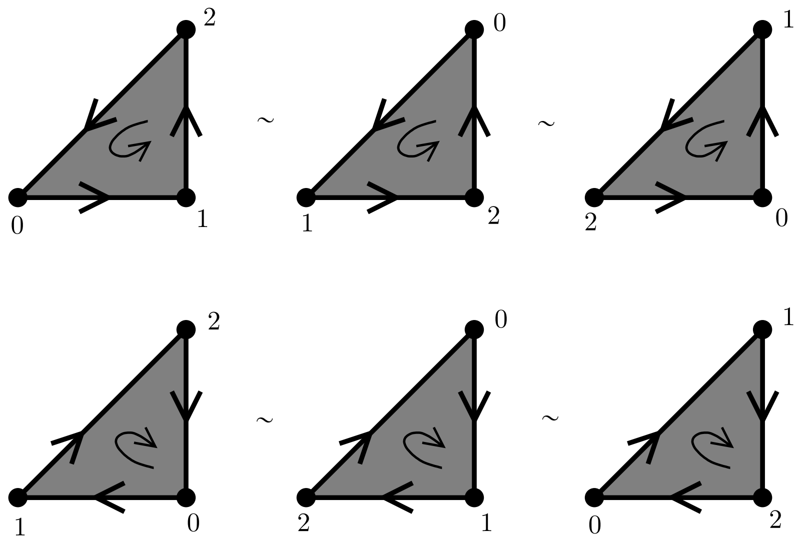

FIGURE 11. The six distinct orderings that define the two possible orientations of a 2-simplex.

have to be a bit careful because applying an even permutation to $v_0, \dots, v_n$ and then eliminating v_k may produce a sequence that differs from $v_0, \dots, v_{k-1}, v_{k+1}, \dots, v_n$ by an *odd* permutation. To get a well-defined orientation on $\partial_{(k)}\sigma$, one can instead do the following: notice that the sequence $v_0, \dots, v_k$ can be reordered as $v_k, v_0, \dots, v_{k-1}, v_{k+1}, \dots, v_n$ by a sequence of k flips. Permutations of this new sequence that fix the first object v_k are then equivalent to permutations of the vertices of $\partial_{(k)}\sigma$, so the even/odd parity of the permutation does not change if we remove v_k from the list. We must not forget however that in order to produce the list with v_k at the front, we performed k flips, meaning a permutation that is even if and only if k is even. This discussion implies that the following notion of boundary orientation is well defined.

DEFINITION 20.8. Given an oriented n-simplex for $n \geqslant 2$ with vertices $v_0, \dots, v_n$ ordered accordingly, the induced **boundary orientation** of its kth boundary face $\partial_{(k)}\sigma$ is defined as the same ordering of its vertices (with v_k removed) if k is even, and otherwise it is defined by any odd permutation of this ordering. For $n = 1$, the boundary orientations are defined by assigning the sign $+1$ to $\partial_{(0)}\sigma = \{v_1\}$ and -1 to $\partial_{(1)}\sigma = \{v_0\}$.

You should now take a moment to stare again at Figure 11 and assure yourself that the boundary orientations indicated there are consistent with this definition.

DEFINITION 20.9. An **oriented triangulation** of a closed surface Σ is a triangulation $\Sigma \cong |K|$ together with a choice of orientation for each 2-simplex in the complex K such that for every 1-simplex σ in K, the two induced boundary orientations that it inherits as a boundary face of two distinct 2-simplices are opposite.

The point of the condition on 1-simplices is to ensure that the orientations of any two neighboring 2-simplices are "compatible" in the sense that each of the circular arrows can be pushed continuously into the other. Figure 12 (left) shows an example of an oriented triangulation of $\mathbb{T}^2$. The arrows on 1-simplices in this picture are not meant to represent boundary orientations, but are just the usual indications of which 1-simplices on the boundary of the square should be glued

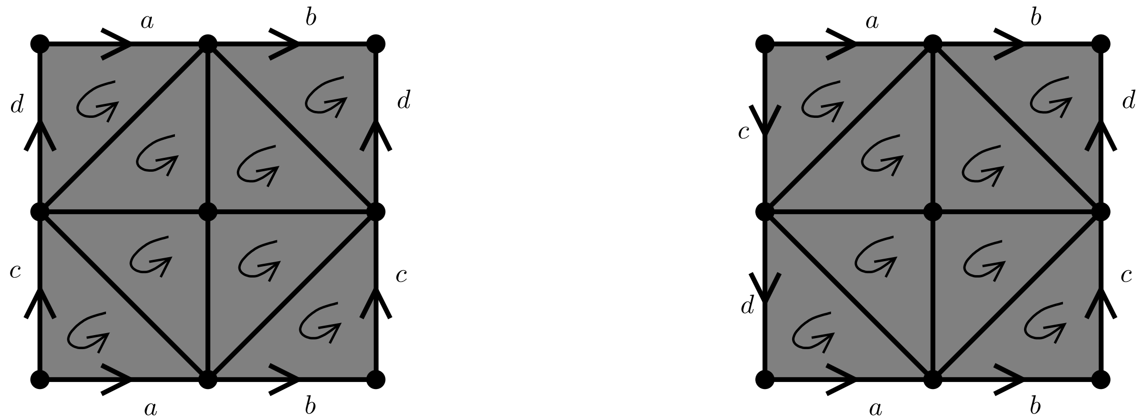

FIGURE 12. An oriented triangulation of the 2-torus (left) and a failed attempt to orient a triangulation of the Klein bottle (right).

together and how. We see in particular that the orientations indicated by these arrows on simplices c and d are the right boundary orientation on the right hand side but the wrong one on the left hand side. According to Definition 20.9, this is exactly what we want. Figure 12 (right) then shows what goes wrong if we try to do the same thing with a Klein bottle. If we imagine that this triangulation admits an orientation, then it will be represented by either clockwise or counterclockwise loops in each 2-simplex in the picture, all of them the same because they must induce opposite orientations on all the 1-dimensional boundary faces between them. In the picture they are all drawn counterclockwise. But notice that in both copies of each of the 1-simplices c and d, the arrow matches the induced boundary orientation, so this picture does not define a valid oriented triangulation. The next theorem implies in fact that no triangulation of the Klein bottle can be oriented.

THEOREM 20.10. *The following conditions are equivalent for any closed connected surface Σ.*

(1) Σ is orientable.
(2) Σ admits an oriented triangulation.
(3) Σ does not contain any subset homeomorphic to the Möbius band.

COROLLARY 20.11. *Every closed, connected and orientable surface is homeomorphic to Σ_g for some $g \geqslant 0$.* □

All of the ideas required for proving Theorem 20.10 have been discussed already, so let us merely sketch how they need to be put together. The equivalence of (1) and (2) is easy to understand by drawing small loops: clearly a choice of "counterclockwise loops" around points in the interior of any 2-simplex $\sigma \subset \Sigma$ determines a cyclic ordering of the vertices of that simplex, and conversely. Notice that this correspondence has a slightly non-obvious corollary: if some triangulation of Σ can be oriented, then so can all others. It should also be intuitively clear why (1) implies (3): if Σ contains a Möbius band, then no globally consistent notion of counterclockwise loops can be defined, since deforming it continuously along certain closed paths around the Möbius band would reverse it. For the converse, we can appeal to the classification of surfaces and observe that any surface Σ satisfying the third condition is homeomorphic to one of the surfaces Σ_g, which can be represented by a polygon with $4g$ sides. In the polygon picture, it is an easy exercise to construct an oriented triangulation for Σ_g. Alternatively, one can understand the relationship between (2) and (3) in terms of the presence of cross-caps or cross-handles in our proof of the classification of surfaces: the orientable surfaces are precisely those which can be constructed without any cross-caps or cross-handles, which turns out to work if and only if the 2-simplices can be assigned orientations for which the gluing maps between matching 1-simplices are orientation reversing.

Exercise 20.12. Construct an explicit oriented triangulation of Σ_g for each $g \geqslant 0$. Then, just for fun, count how many k-simplices it has for each $k = 0, 1, 2$. You will find that the number of 0-simplices minus the number of 1-simplices plus the number of 2-simplices is $2 - 2g$. (Someday next semester we'll discuss the Euler characteristic, and then you'll see why this is true.)

Exercise 20.13. In Exercise 14.13 we considered the space $\Sigma_{g,m}$, defined by cutting the interiors of $m \geqslant 0$ disjoint disks out of the oriented surface Σ_g of genus $g \geqslant 0$.

(a) Prove that every compact, orientable, connected surface with boundary is homeomorphic to $\Sigma_{g,m}$ for some values of $g, m \geqslant 0$.
Hint: If Σ is a compact 2-manifold, then $\partial\Sigma$ is a closed 1-manifold, and we classified all of the latter. With this knowledge, there is a cheap trick by which you can turn any compact surface with boundary into a closed surface, and then apply what you have learned about the classification of closed surfaces. Don't forget to keep track of orientations.
(b) Prove that $\Sigma_{g,m}$ is homeomorphic to $\Sigma_{h,n}$ if and only if $g = h$ and $m = n$.

This concludes our discussion of surfaces.

21. Higher homotopy, bordism, and simplicial homology (June 29, 2023)

The rest of this semester's course will be about homology, but before defining it, I want to discuss some related ideas that should help motivate the definition. In some sense, all of the algebraic topological invariants we discuss in this course can be viewed as methods for "detecting holes" in a topological space. Let me start by describing a few concrete examples in which the fundamental group either does or does not succeed in this task.

Example 21.1. If we replace $\mathbb{R}^2$ with $\mathbb{R}^2\backslash\mathring{\mathbb{D}}^2$, then the fundamental group changes from 0 to $\mathbb{Z}$, with the boundary of $\mathbb{D}^2$ representing a generator of $\pi_1(\mathbb{R}^2\backslash\mathring{\mathbb{D}}^2)$, so this is one type of hole that π_1 detects very well.

Example 21.2. A 3-dimensional generalization of Example 21.1 is to replace $\mathbb{R}^3$ by $(\mathbb{R}^2\backslash\mathring{\mathbb{D}}^2) \times \mathbb{R}$, which amounts to cutting the neighborhood of a line $\{0\} \times \mathbb{R} \subset \mathbb{R}^2 \times \mathbb{R}$ out of $\mathbb{R}^3$. Since the extra factor $\mathbb{R}$ is contractible, this example essentially admits a deformation retraction to the previous one, so we still find a generator of $\pi_1((\mathbb{R}^2\backslash\mathring{\mathbb{D}}^2) \times \mathbb{R}) \cong \pi_1(\mathbb{R}^2\backslash\mathring{\mathbb{D}}^2) \cong \mathbb{Z}$ which detects the removal of the tube $\mathring{\mathbb{D}}^2 \times \mathbb{R}$.

Example 21.3. A different type of generalization of Example 21.1 is to remove a 3-dimensional ball from $\mathbb{R}^3$, and here the fundamental group performs less well: $\pi_1(\mathbb{R}^3)$ is 0, and $\pi_1(\mathbb{R}^3\backslash\mathring{\mathbb{D}}^3)$ is still zero since $\mathbb{R}^3\backslash\mathring{\mathbb{D}}^3$ is homotopy equivalent to S^2 and the latter is simply connected. There clearly is a "hole" here, but π_1 does not see it.

Example 21.4. There are also examples in which π_1 seems to detect something other than a hole. Let $\Sigma_{g,m}$ denote the surface of genus g with m holes cut out, so Σ_2 is homeomorphic to a surface constructed by gluing together two copies of $\Sigma_{1,1}$ along their common boundary:

$$\Sigma_2 \cong \Sigma_{1,1} \cup_{\partial\Sigma_{1,1}} \Sigma_{1,1}.$$

Let $\gamma : S^1 \to \Sigma_2$ denote a loop parametrizing the common boundary of these copies of $\Sigma_{1,1}$. As we saw in Exercise 14.13, γ represents a nontrivial element in $\pi_1(\Sigma_2)$, though it is in the kernel of the natural homomorphism of $\pi_1(\Sigma_2)$ to its abelianization. The latter will turn out to be related to the following geometric observation: while γ cannot be extended to any map $\mathbb{D}^2 \to \Sigma_2$, it can be extended to a map on *some* surface with boundary S^1, e.g. it admits an extension to the inclusion $\Sigma_{1,1} \hookrightarrow \Sigma_2$. In this sense, there is no actual hole there for γ to detect; it is instead detecting a different phenomenon that has to do with the distinction between "disk-shaped" holes and "holes with genus".

I'm now going to start suggesting possible remedies for the drawbacks encountered in the last two examples. We will have to try a few times before we can point to the "right" remedy, but all of the objects we discuss along the way are also interesting and worthy of study.

Remedy 1: Higher homotopy groups. For any integer $k \geqslant 0$, fix a base point $t_0 \in S^k$ and associate to any pointed space (X, x_0) the set

$$\pi_k(X, x_0) = \left\{ f : (S^k, t_0) \to (X, x_0) \right\} / \underset{h+}{\sim},$$

where the equivalence relation $\underset{h+}{\sim}$ here means base-point preserving homotopy. This clearly reproduces the fundamental group when $k = 1$. When $k = 0$, $S^0 = \partial \mathbb{D}^1 = \{1, -1\}$ is a discrete space with two points, one of which must be the base point and is thus constrained to map to x_0, but the other can move freely within each path-component of X, so $\pi_0(X, x_0)$ is in bijective correspondence with the set of path-components of X. This set does not naturally have any group structure, though it does naturally have a "neutral" element, represented by the map that sends both points in S^0 to the base point x_0. It turns out that for $k \geqslant 2$, $\pi_k(X, x_0)$ can always be given the structure of an *abelian* group whose identity element is represented by the constant map

$$0 := [(S^k, t_0) \to (X, x_0) : t \mapsto x_0].$$

The precise definition of the group operation is a bit less obvious than for $k = 1$, so I will not go into it in this brief sketch. As with the fundamental group, one can show that $\pi_k(X, x_0)$ is independent of the base point up to isomorphism whenever X is path-connected, and it is also isomorphic for any two spaces that are homotopy equivalent. We will prove these statements next semester in *Topologie II*, but feel free to have a look at [**Hat02**, §4.1] if you can't bear to wait.

Here are a couple of things that can be proved about the higher homotopy groups using something resembling our present state of knowledge in this course:

Example 21.5. The identity map $S^k \to S^k$ represents a nontrivial element of $\pi_k(S^k)$ for every $k \geqslant 1$. This follows from Exercise 19.14, which sketches the notion of the mod 2 mapping degree in order to show that every map $S^k \to S^k$ homotopic to the identity is surjective (and therefore nonconstant). More generally, one can use the integer-valued mapping degree for maps $S^k \to S^k$ to prove that $\pi_k(S^k) \cong \mathbb{Z}$, just like the case $k = 1$. A very nice account of this is given in [**Mil97**].

Example 21.6. For every pair of integers $k, n \in \mathbb{N}$ with $n > k$, $\pi_k(S^n) = 0$. This follows easily from a general result in differential topology that allows us to approximate any continuous map between smooth manifolds by a smooth map for which any given point in the target space can be assumed to be a regular value. When $n > k$, the latter means that for any given $q \in S^n$ and a continuous map $f : S^k \to S^n$, we can approximate f with a map whose image does not contain q and is thus contained in $S^n \backslash \{q\} \cong \mathbb{R}^n$. The latter admits a deformation retraction to any point it contains, so composing the perturbed map $S^k \to S^n \backslash \{q\}$ with a deformation retraction of $S^n \backslash \{q\}$ to the base point gives a homotopy of f to the constant map.

Now here is the first piece of bad news about π_k: in general it is rather hard to compute. So hard, in fact, that the answers to certain basic questions about π_k remain unknown, e.g. one of the most popular open questions in modern topology is how to compute $\pi_k(S^n)$ in general when $k > n$. Various special cases are known, but the as-yet incomplete effort to extend these special cases to a general theorem has played a large role in motivating the development of modern homotopy theory.

We will need to have more and easier techniques at our disposal before we can discuss such things in earnest.

Remedy 2: Bordism groups. The higher homotopy groups do remedy one of the drawbacks of π_1 that I pointed out above: e.g. π_2 can be used to detect the hole in $\mathbb{R}^3 \backslash \mathring{\mathbb{D}}^3$ since, by homotopy invariance,

$$\pi_2(\mathbb{R}^3 \backslash \mathring{\mathbb{D}}^3) \cong \pi_2(S^2) \cong \mathbb{Z},$$

with the inclusion $S^2 \hookrightarrow \mathbb{R}^3 \backslash \mathring{\mathbb{D}}^3$ representing a generator. But there's another drawback here: while π_k can detect higher-dimensional holes, they are still holes of a fairly specific type which one might call "sphere-shaped" holes. What kind of hole is not sphere-shaped, you ask? Is there such a thing as a "torus-shaped" hole? How about this one:

EXAMPLE 21.7. Let $X = S^1 \times \mathbb{R}^2$ and $X_0 = S^1 \times \mathring{\mathbb{D}}^2$, so $X \backslash X_0 = S^1 \times (\mathbb{R}^2 \backslash \mathring{\mathbb{D}}^2)$ admits a deformation retraction to $\partial \bar{X}_0 = S^1 \times S^1 = \mathbb{T}^2$. By homotopy invariance, we have $\pi_1(X) \cong \pi_1(S^1) \cong \mathbb{Z}$ and $\pi_1(X \backslash X_0) \cong \pi_1(\mathbb{T}^2) \cong \mathbb{Z}^2$, so π_1 does at least partly detect the removal of X_0 from X. But since $X \backslash X_0$ is homotopy equivalent to a surface, there is also an intrinsically 2-dimensional phenonomenon going on in this picture, and it seems natural to ask: does $X \backslash X_0$ contain any surface detecting the fact that X_0 has been removed from X? We can almost immediately give the following answer: if such a surface exists, it is *not* a sphere, in fact $\pi_2(X) = \pi_2(X \backslash X_0) = 0$. To see this, we can use the homotopy invariance of π_2: the spaces X and $X \backslash X_0$ are homotopy equivalent to S^1 and $\mathbb{T}^2$ respectively, so it suffices to prove $\pi_2(S^1) = \pi_2(\mathbb{T}^2) = 0$. Now observe that both S^1 and $\mathbb{T}^2$ are spaces whose universal covers ($\mathbb{R}$ and $\mathbb{R}^2$ respectively) happen to be contractible. In general, suppose $p : \widetilde{Y} \to Y$ denotes the universal cover of some reasonable space Y, and $\widetilde{Y}$ is contractible. Since S^2 is simply connected, any map $f : S^2 \to Y$ can be lifted to $\tilde{f} : S^2 \to \widetilde{Y}$, but the contractibility of $\widetilde{Y}$ then implies that $\tilde{f}$ is homotopic to a constant map. Composing that homotopy with $p : \widetilde{Y} \to Y$ gives a corresponding homotopy of $f = p \circ \tilde{f} : S^2 \to Y$ to a constant map, proving $\pi_2(Y) = 0$.

The preceding example is meant to provide motivation for a new invariant that might be able to detect holes that are not "sphere-shaped". The idea is to forget about the special roll played by spheres in the definition of π_k, but remember the fact that S^k is a closed k-dimensional manifold. Similarly, if M is a k-manifold, the homotopy relation for maps defined on M is defined in terms of maps on $I \times M$, which gives a special status to a very particular class of $(k+1)$-manifolds with boundary. Since we are now allowing arbitrary closed k-manifolds in place of spheres, it also seems natural to allow arbitrary compact $(k+1)$-manifolds with boundary for defining equivalence, instead of just manifolds of the form $I \times M$. Following this train of thought to its logical conclusion leads to *bordism theory.*[31]

For any space X and each integer $k \geqslant 0$, let

$$\Omega_k(X) := \{(M, f)\} / \sim,$$

[31] In the older literature, "bordism theory" was usually called "cobordism theory," and it is still common in most subfields of geometry and topology to refer to manifolds whose boundaries are disjoint unions of a given pair of closed manifolds as "cobordisms" instead of "bordisms". The elimination of the "co-" in "cobordism" is presumably motivated by the fact that bordism groups define a covariant functor instead of a contravariant functor, which makes it more analogous to *homology* than to *cohomology*. I promise you this footnote will make more sense after *Topologie II*.

where M is any closed (but not necessarily connected or nonempty)[32] k-manifold, $f : M \to X$ is a continuous map, and we write $(M_+, f_+) \sim (M_-, f_-)$ if and only if there exists a compact $(k+1)$-manifold W with $\partial W \cong M_- \amalg M_+$ and a map $F : W \to X$ such that $F|_{M_\pm} = f_\pm$. You should take a moment to think about why $\sim$ defines an equivalence relation. Any two pairs that are equivalent in this sense are said to be **bordant**, and the pair (W, F) is called a **bordism** between them.

EXAMPLE 21.8. $(M, f) \sim (M, g)$ whenever f and g are homotopic maps $M \to X$, as the homotopy $H : I \times M \to X$ defines a bordism $(I \times M, H)$.

EXAMPLE 21.9. Recall from Example 21.4 the loop $\gamma : S^1 \to \Sigma_2$ whose image separates Σ_2 into two pieces both homeomorphic to $\Sigma_{1,1}$. Either of the two inclusions $\Sigma_{1,1} \hookrightarrow \Sigma_2$ in this picture can be viewed as a bordism between (S^1, γ) and $(\varnothing, \cdot)$, where $\cdot$ denotes the unique map $\varnothing \to X$. Hence $[(S^1, \gamma)] = [(\varnothing, \cdot)] \in \Omega_1(\Sigma_2)$.

Since the manifolds representing elements of $\Omega_k(X)$ need not be connected, the disjoint union provides an obvious definition for a group operation on $\Omega_k(X)$. This operation is necessarily commutative since $X \amalg Y$ has a natural identification with $Y \amalg X$ for any two spaces X and Y. Now would be a good moment to mention the following notational convention: whenever a group G is known a priori to be abelian, we shall from now on denote the group operation in G as *addition* (with a "+" sign) rather than multiplication.

DEFINITION 21.10. We give $\Omega_k(X)$ the structure of an abelian group by defining

$$[(M_1, f_1)] + [(M_2, f_2)] := [(M_1 \amalg M_2, f_1 \amalg f_2)],$$

where $f_1 \amalg f_2 : M_1 \amalg M_2 \to X$ denotes the unique map whose restriction to $M_i \subset M_1 \amalg M_2$ is f_i for $i = 1, 2$. The identity element is

$$0 := [(\varnothing, \cdot)],$$

with $\cdot : \varnothing \to X$ denoting the unique map. The group $\Omega_k(X)$ is called the **k-dimensional unoriented bordism group** of X. We say that a pair (M, f) is **null-bordant** whenever $[(M, f)] = 0$, meaning there exists a compact $(k+1)$-manifold W with $\partial W \cong M$ and a map $F : W \to X$ with $F|_M = f$.

Referring back to Example 21.7, one can now show that the bordism class represented by the inclusion $\mathbb{T}^2 = \partial \bar{X}_0 \hookrightarrow X \backslash X_0$ is nontrivial in $\Omega_2(X \backslash X_0)$. One way to prove this uses the mod 2 mapping degree (cf. Exercise 19.14) for maps $f : \mathbb{T}^2 \to \mathbb{T}^2$: by an argument similar to the proof that $\deg_2(f)$ depends only on the homotopy class of f, one can show that $\deg(f) = 0$ whenever $(\mathbb{T}^2, f)$ is null-bordant. It follows that $[(\mathbb{T}^2, \mathrm{Id})] \neq 0 \in \Omega_2(\mathbb{T}^2)$ since $\deg_2(\mathrm{Id}) = 1$, and this element of $\Omega_2(\mathbb{T}^2)$ can be identified with the aforementioned inclusion using the homotopy equivalence between $\mathbb{T}^2$ and $X \backslash X_0$. In summary, Ω_2 does indeed detect "$\mathbb{T}^2$-shaped" holes.

The algebraic structure of $\Omega_k(X)$ is also extremely simple, one might even say too simple, in light of the following result saying that every element in $\Omega_k(X)$ is its own inverse:

PROPOSITION 21.11. *For every $[(M, f)] \in \Omega_k(X)$, $[(M, f)] + [(M, f)] = 0$.*

PROOF. Let $W = I \times M$ and $F : W \to X : (s, x) \mapsto f(x)$. Then $\partial W \cong \varnothing \amalg (M \amalg M)$ and $F|_{M \amalg M} = f \amalg f$, hence (W, F) is a bordism between $(M \amalg M, f \amalg f)$ and $(\varnothing, \cdot)$.[33] $\square$

---

[32]Note that the empty set is a k-manifold for every $k \in \mathbb{Z}$. Look again at the definition of manifolds, and you will see that this is true.

[33]One of the slightly confusing things about $\Omega_k(X)$ is that there is always some ambiguity about how to split up the various connected components of ∂W into M_- and M_+. For the bordism in the proof of Prop. 21.11, one can equally well view it as a bordism between (M, f) and (M, f), but we are ignoring this because it does not give us any information beyond the fact that the bordism relation is reflexive.

One obtains a slightly more interesting algebraic structure by restricting to orientable manifolds and keeping track of orientations. Recall from the previous lecture that a manifold endowed with the extra structure of an orientation is called an *oriented manifold*; we will continue to denote such objects by single letters such as M, but you should keep in mind that they include slightly more data than just a set with its topology. If M is an oriented manifold, we shall denote by $-M$ the same manifold with its orientation reversed: this can always be defined by replacing each of the oriented charts on M by their compositions with an orientation-reversing homeomorphism $\mathbb{H}^n \to \mathbb{H}^n$ such as $(x_1, \dots, x_{n-1}, x_n) \mapsto (x_1, \dots, x_{n-1}, -x_n)$. Recall also from Remark 20.6 that any oriented manifold W with boundary determines a natural *boundary orientation* on ∂W. Whenever we write expressions like $\partial W \cong M$ in the context of oriented manifolds, we will always mean there is a homeomorphism $\partial W \to M$ that matches the given orientation of M to the boundary orientation of ∂W induced by the given orientation of W.

Definition 21.12. The **k-dimensional oriented bordism group** of X is[34]

$$\Omega_k^{\mathrm{SO}}(X) := \{(M, f)\} / \sim,$$

where M is a closed (but not necessarily connected or nonempty) oriented k-manifold, $f : M \to X$ is continuous, and the oriented bordism relation $(M_+, f_+) \sim (M_-, f_-)$ means that there exists a compact oriented $(k+1)$-manifold W and a map $F : W \to X$ such that

$$\partial W \cong -M_- \amalg M_+$$

and $F|_{M_\pm} = f_\pm$. The group operation on $\Omega_k^{\mathrm{SO}}(X)$ is defined via disjoint union as with $\Omega_k(X)$.

Proposition 21.11 is not true for oriented bordism groups: its proof fails due to the fact that the oriented boundary of $I \times M$ is $-M \amalg M$, not $M \amalg M$.

Let us compare both groups in the case $k = 0$. We claim that

$$\Omega_0(X) \cong \bigoplus_{\pi_0(X)} \mathbb{Z}_2,$$

while

$$\Omega_0^{\mathrm{SO}}(X) \cong \bigoplus_{\pi_0(X)} \mathbb{Z},$$

where $\pi_0(X)$ is an abbreviation for the set of path-components of X. For concreteness, consider a case where X has exactly three path-components $X_1, X_2, X_3 \subset X$, so the claim is that $\Omega_0(X) \cong \mathbb{Z}_2^3$ and $\Omega_0^{\mathrm{SO}}(X) \cong \mathbb{Z}^3$. An element of $\Omega_0(X)$ is an equivalence class of pairs (M, f), where M is a closed 0-manifold, i.e. a finite discrete set, and $f : M \to X$. Let us number the elements of M as $x_1, \dots, x_N$, and suppose there are two elements that are mapped by f to the same path-component, say $f(x_1), f(x_2) \in X_1$. Then there exists a path $\gamma : I_{12} \to X$, where $I_{12} := I$, satisfying $\gamma(0) = f(x_1)$ and $\gamma(1) = f(x_2)$. Now define $W := I_{12} \amalg I_3 \amalg \dots \amalg I_N$ where each I_j for $j = 3, \dots, N$ is another copy of I, and decompose the boundary $\partial W = M_- \amalg M_+$ so that M_+ contains ∂I_{12} and $1 \in \partial I_j$ for every $j = 3, \dots, N$, while M_- contains $0 \in \partial I_j$ for every $j = 3, \dots, N$. Defining $F : W \to X$ such that $F|_{I_{12}} := \gamma$ and F sends I_j to the constant $f(x_j)$ for each $j = 3, \dots, N$, we now have a bordism between (M, f) and (M', f') where $M' := M \backslash \{x_1, x_2\}$ and f' is the restriction of f. One can do this for any pair of points in M that are mapped to the same path-component, so that whenever (M, f) and (N, g) have the same number of points (mod 2) mapped into each path-component, there exists a bordism between them. Conversely, any bordism between two pairs (M, f) and (N, g) is of the form (W, F) where W is a compact 1-manifold with boundary,

[34]The "SO" in the notation $\Omega_k^{\mathrm{SO}}(X)$ stands for the group $\mathrm{SO}(k)$, the special orthogonal group. This has to do with the fact that $\mathrm{SO}(k)$ is precisely the subgroup of $\mathrm{O}(k)$ consisting of orthogonal transformations that are *orientation preserving*.

and by the classification of 1-manifolds, this can only mean a finite disjoint union of circles and compact intervals. Since each of these components individually can only be mapped into one of the path-components X_1, X_2, X_3 and each has either zero or two boundary points, it follows that for each $i = 1, 2, 3$, the number of points of M or N that are mapped into X_i can only differ by an even number. We have just proved the following: given $[(M, f)] \in \Omega_0(X)$, let $f_i \in \mathbb{Z}_2$ for $i = 1, 2, 3$ denote the number (mod 2) of points in M that f maps into X_i. Then

$$\Omega_0(X) \to \mathbb{Z}_2^3 : [(M, f)] \mapsto (f_1, f_2, f_3)$$

is an isomorphism.

To understand $\Omega_0^{\mathrm{SO}}(X)$, we need to keep in mind that an oriented 0-manifold M is not just a finite set of points, but it also comes with a map $\epsilon : M \to \{1, -1\}$ telling us which points are to be regarded as "positively oriented" as opposed to "negatively oriented" (cf. Definition 20.5). It is now no longer possible to cancel arbitrary pairs as in the unoriented case, but suppose $M = \{x_1, \ldots, x_N\}$ and f sends both x_1 and x_2 into X_1, and also that $\epsilon(x_1) = -1$ while $\epsilon(x_2) = +1$. We can again choose a path $\gamma : I_{12} \to X_1$ with $\gamma(0) = f(x_1)$ and $\gamma(1) = f(x_2)$, and define $W = I_{12} \amalg I_3 \amalg \ldots \amalg I_N$ and $F : W \to X$ as before. Before we can call (W, F) an oriented bordism, we need to specify the orientation of W. Let us assume I_{12} is oriented so that $\epsilon(1) = +1$ and $\epsilon(0) = -1$, while for $j = 3, \ldots, N$, orient I_j such that $\epsilon(1) = \epsilon(x_j)$ and $\epsilon(0) = -\epsilon(x_j)$. We now have $\partial W = -M' \amalg M$ where $M' = M \backslash \{x_1, x_2\}$ with the same orientations on the points $x_3, \ldots, x_N$, hence (W, F) is an oriented bordism between (M, f) and (M', f'). It is possible to construct such a bordism to eliminate any pair of points in M that have opposite signs and are mapped to the same path-component of X. Thus if we define $f_i \in \mathbb{Z}$ for each $i = 1, 2, 3$ by

$$f_i := \sum_{x \in f^{-1}(X_i)} \epsilon(x),$$

it follows that any two pairs (M, f) and (N, g) for which $f_i = g_i$ for every i must admit an oriented bordism. Conversely, the classification of 1-manifolds again implies that an arbitrary oriented bordism (W, F) between two pairs (M, f) and (N, g) is a map defined on a finite disjoint union of oriented intervals and circles, and since the two boundary points of an oriented interval I are always oriented with opposite signs, any component of W whose boundary lies entirely in one of M or $-N$ contributes zero to the counts defining the numbers f_i and g_i, while components that have one boundary point in M and one in $-N$ make the same contribution ± 1 to f_i and g_i. This proves that the map

$$\Omega_0^{\mathrm{SO}}(X) \to \mathbb{Z}^3 : [(M, f)] \mapsto (f_1, f_2, f_3)$$

is well defined and is also an isomorphism.

While computing the 0-dimensional bordism groups is not hard, we run into a serious (though interesting!) difficulty with the higher-dimensional bordism groups: they can be nontrivial even if X is only a one-point space. When $X = \{\mathrm{pt}\}$, we abbreviate

$$\Omega_k := \Omega_k(\{\mathrm{pt}\}), \qquad \Omega_k^{\mathrm{SO}} := \Omega_k^{\mathrm{SO}}(\{\mathrm{pt}\}),$$

and notice that since there is only one map from each manifold to $\{\mathrm{pt}\}$, the elements of Ω_k^{SO} are equivalence classes of oriented closed manifolds M where $M \sim N$ whenever $\partial W \cong -M \amalg N$ for some compact oriented manifold W; elements of Ω_k can be described in the same way after deleting the word "oriented" everywhere. In particular, we have $[M] = 0 \in \Omega_k$ if and only if M is homeomorphic to the boundary of some compact $(k+1)$-manifold. The question of whether a given manifold can be the boundary of another compact manifold is interesting, and the answer is often not obvious. For $k = 1$ it is not so hard: the classification of 1-manifolds implies that every bordism class $[M]$ in Ω_1 or Ω_1^{SO} is represented by a finite disjoint union of circles, and since

$S^1 = \partial\mathbb{D}^2$, all of these are (oriented) boundaries, hence

$$\Omega_1 = \Omega_1^{\mathrm{SO}} = 0.$$

It is similarly easy to see that all closed oriented surfaces are boundaries of compact oriented 3-manifolds: just take your favorite embedding of Σ_g into $\mathbb{R}^3$ and consider the region bounded by that embedded surface. For the oriented 3-dimensional case, we do not have any simple classification result to rely upon, but one can instead appeal to a standard (though not so trivial) result from low-dimensional topology known as the Dehn-Lickorish theorem, which can be interpreted as presenting arbitrary closed oriented 3-manifolds as boundaries of compact oriented 4-manifolds obtained by attaching "2-handles" to $\mathbb{D}^4$. We can therefore say

$$\Omega_2^{\mathrm{SO}} = \Omega_3^{\mathrm{SO}} = 0.$$

However, in the unoriented case there is already trouble in dimension two: it is known that there does not exist any compact 3-manifold whose boundary is homeomorphic to $\mathbb{RP}^2$. This can be proved using methods that we will cover in *Topologie II*, notably the Poincaré duality isomorphism between the homology and cohomology groups of closed manifolds. A similar argument implies that the complex counterpart of $\mathbb{RP}^2$, the complex projective space $\mathbb{CP}^2$, is a closed oriented 4-manifold that never occurs as the boundary of any compact oriented 5-manifold. This implies

$$[\mathbb{RP}^2] \neq 0 \in \Omega_2, \qquad \text{and} \qquad [\mathbb{CP}^2] \neq 0 \in \Omega_4^{\mathrm{SO}}.$$

This reveals that in general, the k-dimensional bordism groups of a one-point space contain a lot more information than one might expect: instead of just telling us something about the rather boring space $\{\mathrm{pt}\}$, they tell us something about the classification of closed k-manifolds, namely which ones can appear as boundaries of other compact manifolds and which ones cannot. That is an interesting question, and one that is very much worth studying at some point, but as with the higher homotopy groups, we will need to have a much wider range of simpler techniques at our disposal before we are equipped to tackle it.

Remedy 3: Simplicial homology (AKA "triangulated bordism"). The first version of homology theory that we will now discuss can be regarded as an attempt to capture much of the same information about X that is seen by the bordism groups $\Omega_n(X)$ and $\Omega_n^{\mathrm{SO}}(X)$, but without requiring us to know anything about the (generally quite hard) problem of classifying closed n-manifolds. The first idea is that instead of allowing arbitrary closed manifolds as domains, we consider manifolds with triangulations, so that all the data can be expressed in terms of simplices. The followup idea is that now that everything is expressed in terms of simplices, there is no need to mention manifolds at all.

Consider a simplicial complex $K = (V, S)$ with associated polyhedron $X := |K|$, and for each integer $n \geqslant 0$, let $S_{(n)} \subset S$ denote the set of n-simplices. As auxiliary data, we also fix an abelian group G, which in principle can be arbitrary, but for reasons related to the distinction between oriented and unoriented bordism, we will typically want to choose G to be either $\mathbb{Z}$ or $\mathbb{Z}_2$.

Definition 21.13. The group of **n-chains** in K (with **coefficients** in G) is the abelian group

$$C_n(K; G) := \bigoplus_{\sigma \in S_{(n)}} G,$$

whose elements can be written as finite sums $\sum_i a_i\sigma_i$ with $a_i \in G$ and $\sigma_i \in S_{(n)}$, with the group operation defined by

$$\sum_i a_i\sigma_i + \sum_i b_i\sigma_i = \sum_i (a_i + b_i)\sigma_i.$$

An n-chain is in some sense an abstract algebraic object, but if we choose $G = \mathbb{Z}$ and consider an n-chain $\sum_i a_i \sigma_i$ whose coefficients are all $a_i = \pm 1$, then you can picture the chain geometrically as the union of the n-simplices in X corresponding to each σ_i in the sum, with orientations determined by the signs a_i. These subsets are always compact, and if the particular set of n-simplices is chosen appropriately, then they will sometimes look like n-dimensional manifolds embedded in X. Our goal is now to single out a special class of n-chains that are analogous to *closed* n-dimensional manifolds embedded in X, i.e. the n-chains that have "empty boundary". This can be done by writing down an algebraic operation that describes the boundary of each individual simplex. To define this properly, we need to choose an orientation for every simplex in S; note that this has nothing intrinsically to do with oriented triangulations, as it is a completely arbitrary choice with no compatibility conditions required, so it can always be done. With this choice in place, for each $\sigma = \{v_0, \dots, v_n\} \in S_{(n)}$, set

$$\partial \sigma := \sum_{k=0}^{n} \epsilon_k \partial_{(k)} \sigma \in C_{n-1}(K; \mathbb{Z}),$$

where as usual $\partial_{(k)} \sigma = \{v_0, \dots, v_{k-1}, v_{k+1}, \dots, v_n\}$ denotes the kth boundary face of σ, and $\epsilon_k \in \{1, -1\}$ is defined to be $+1$ if the chosen orientation of the $(n-1)$-simplex $\partial_{(k)}\sigma$ matches the boundary orientation it inherits from σ (see Definition 20.8), and -1 if these two orientations are opposite. There is now a uniquely determined group homomorphism

$$\partial_n : C_n(K; G) \to C_{n-1}(K; G) : \sum_i a_i \sigma_i \mapsto \sum_i a_i (\partial \sigma_i),$$

where the multiplication of each coefficient $a_i \in G$ by a sign $\epsilon_k = \pm 1$ is defined in the obvious way as an element of G. (Notice that if $G = \mathbb{Z}_2$, the signs ϵ_k become irrelevant because every coefficient a_i then satisfies $a_i = -a_i$.) Strictly speaking, the definition above only makes sense for $n \geqslant 1$ since there are no (-1)-simplices; in light of this, we set

$$\partial_0 := 0.$$

We call the subgroup $\ker \partial_n \subset C_n(K; G)$ the group of n-**cycles**, or equivalently, the **closed** n-chains. The elements of the subgroup $\operatorname{im} \partial_{n+1} \subset C_n(K; G)$ are called **boundaries**.

LEMMA 21.14. *$\partial_{n-1} \circ \partial_n = 0$ for all $n \in \mathbb{N}$.*

PROOF. You should think of this as an algebraic or combinatorial expression of the geometric fact that the boundary of any n-manifold with boundary is always an $(n-1)$-manifold with *empty* boundary. On a more mundane level, the result holds due to cancelations, e.g. suppose A is an oriented 2-simplex whose oriented 1-dimensional boundary faces are denoted by a, b, c, giving

$$\partial_2 A = a + b + c.$$

Suppose further that the vertices of A are denoted by α, β, γ, all oriented with positive signs, but the arrow determined by the orientation of a points toward α and away from γ, while b points toward β and away from α, and c points toward γ but away from β. This gives the three relations

$$\partial_1 a = \alpha - \gamma, \quad \partial_1 b = \beta - \alpha, \quad \partial_1 c = \gamma - \beta,$$

thus $\partial_1 \circ \partial_2 A = \partial_1(a + b + c) = (\alpha - \gamma) + (\beta - \alpha) + (\gamma - \beta) = 0$. Similar cancelations occur in every dimension. □

Lemma 21.14 is often abbreviated with the formula

$$\partial^2 = 0,$$

and we will sometimes abbreviate $\partial := \partial_n$ when there is no chance of confusion. The formula implies in particular that $\operatorname{im}\partial_{n+1}$ is a subgroup of ∂_n for every $n \geqslant 0$. Since all these groups are abelian and subgroups are therefore normal, we can now consider quotients:

DEFINITION 21.15. The nth **simplicial homology** group of the complex K (with coefficients in G) is

$$H_n^{\Delta}(K;G) := \ker\partial_n \big/ \operatorname{im}\partial_{n+1}.$$

It is worth comparing this definition to the bordism groups $\Omega_n(X)$ and $\Omega_n^{\mathrm{SO}}(X)$, as the extra layer of algebra involved in the definition of homology obscures a fairly direct analogy. Instead of closed n-manifolds M with maps $f : M \to X$, homology considers n-cycles, meaning formal linear combinations of n-simplices $c := \sum_i a_i\sigma_i$ with $\partial c = 0$. The bordism relation $(M_+, f_+) \sim (M_-, f_-)$ is now replaced by the conditition that two cycles $c, c' \in \ker\partial_n$ represent the same homology class $[c] = [c'] \in H_n^{\Delta}(K;G)$ if $c - c' \in \operatorname{im}\partial_{n+1}$, i.e. their difference is the boundary of an $(n+1)$-chain (analogous to a map defined on a compact $(n+1)$-manifold with boundary). When this holds, we say that the cycles c and c' are **homologous**. Finally, we will see that the distinction between $\Omega_n^{\mathrm{SO}}(X)$ and $\Omega_n(X)$ now corresponds to the distinction between $H_n^{\Delta}(K;\mathbb{Z})$ and $H_n^{\Delta}(K;\mathbb{Z}_2)$.

Let's compute an example. Figure 13 shows an oriented triangulation of $\mathbb{T}^2$ with eight 2-simplices, twelve 1-simplices and four vertices labeled as follows:

$$\begin{aligned}
S_2 &= \{A, B, C, D, E, F, G, H\},\\
S_1 &= \{a, b, c, d, e, f, g, h, i, j, k, \ell\},\\
S_0 &= \{\alpha, \beta, \gamma, \delta\}.
\end{aligned}$$

In addition to the orientations of the 2-simplices that come from this being an oriented triangulation, the figure shows (via arrows) an arbitrary choice of orientations for all 1-simplices, and we shall assume all the 0-simplices are oriented with a positive sign. One can now begin writing down relations such as

$$\partial A = a - h - c, \quad \partial B = i - k + h, \quad \partial a = \beta - \alpha$$

and so forth, but writing down all such relations would be rather tedious, so let us instead try to reason more geometrically. The computation of $H_0^{\Delta}(K;\mathbb{Z})$ is not hard in any case: all 0-chains are cycles since $\partial_0 = 0$, including the four generators α, β, γ and δ, but all four of them are also homologous to each other since any pair of them can be connected by an oriented 1-simplex pointing from one to the other, e.g. $\partial a = \beta - \alpha$ implies $[\alpha] = [\beta]$, and $\partial i = \delta - \beta$ implies $[\beta] = [\delta]$. The result is

$$H_0^{\Delta}(K;\mathbb{Z}) \cong \mathbb{Z},$$

with a canonical generator represented by any of the vertices in the complex. Notice that this matches the oriented bordism group $\Omega_0^{\mathrm{SO}}(\mathbb{T}^2)$ since $\mathbb{T}^2$ is path-connected.

Let's look at the 1-cycles. There is a 1-cycle for every continuous loop we can find that follows a path through 1-simplices—we just have to insert minus signs wherever there is an arrow pointing the wrong way in order to ensure the necessary cancelation of 0-simplices. For example, traversing the boundary of the lower-right square gives

$$\partial(i + \ell - c - b) = 0,$$

so $i + \ell - c - b$ is a 1-cycle, but not a very interesting one since it is also the boundary of the region filled by the 2-simplices C and D: in particular,

$$\partial(-C - D) = i + \ell - c - b,$$

hence $[i + \ell - c - b] = 0 \in H_1^{\Delta}(K;\mathbb{Z})$. To find more interesting 1-cycles, it helps to remember what we already know about $\pi_1(\mathbb{T}^2) \cong \mathbb{Z}^2$. We can easily find two loops through 1-simplices that

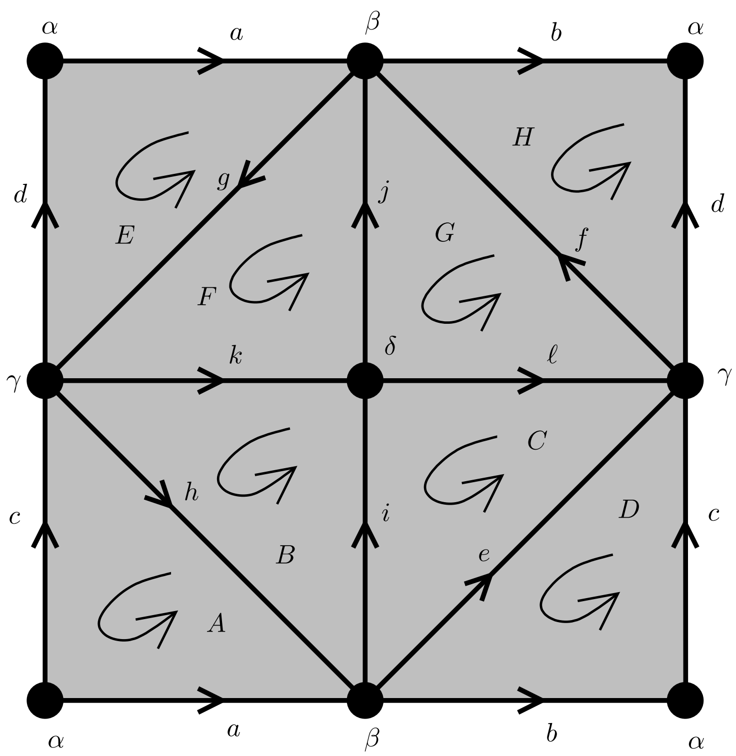

FIGURE 13. A simplicial complex with $|K| = \mathbb{T}^2$.

represent the two distinct generators of this fundamental group: one of them is $i+j$, and we easily see that

$$\partial(i+j) = (\delta - \beta) + (\beta - \delta) = 0.$$

Another is $c+d$, but notice that the loops corresponding to these two 1-cycles are homotopic in $\mathbb{T}^2$, and relatedly, they form the boundary of the region filled by the 2-simplices C, D, G and H, so

$$\partial(C + D + G + H) = c + d - (i + j),$$

implying $[c+d] = [i+j] \in H_1^\Delta(K;\mathbb{Z})$. One can show however that this homology class really is nontrivial, and it is not the only one: the other generator of $\pi_1(\mathbb{T}^2)$ corresponds to either of the two homologous 1-cycles $a+b$ or $k+\ell$. The end result is

$$H_1^\Delta(K;\mathbb{Z}) \cong \mathbb{Z}^2,$$

the same as the fundamental group.

As observed at the beginning of this lecture, the fact that $\mathbb{T}^2$ has a contractible universal cover implies that $\pi_2(\mathbb{T}^2) = 0$, so if there are any interesting 2-cycles in $\mathbb{T}^2$, they will not look like spheres. But if you think that $H_2(K;\mathbb{Z})$ should have something to do with the oriented bordism group $\Omega_2^{\mathrm{SO}}(\mathbb{T}^2)$, then there is a fairly obvious candidate for a 2-cycle in this picture: $\mathbb{T}^2$ itself is a closed oriented manifold, and the oriented triangulation we have chosen turns it into a 2-cycle:

$$\partial(A + B + C + D + E + F + G + H) = 0.$$

The point is that since the triangulation is oriented, writing down each individual term in this sum would produce a linear combination of 1-simplicies in which every 1-simplex in the complex appears exactly twice, but with opposite signs, thus adding up to 0. It should be easy to convince yourself that no nontrivial 2-chain that does not include all eight of the 2-simplices can ever be a cycle, as its boundary will have to include some 1-simplices that have nothing to cancel with. It follows easily that all 2-cycles in this complex are integer multiples of the one found above, and none of them are boundaries since there are no 3-simplices, thus

$$H_2^{\Delta}(K;\mathbb{Z}) \cong \mathbb{Z}.$$

I can now state a theorem that is really rather amazing, though I'm sorry to say that we will not be able to prove it until next semester:

THEOREM 21.16. *For any simplicial complex K, the simplicial homology groups $H_n^{\Delta}(K;G)$ depend (up to isomorphism) on the topological space $X = |K|$, i.e. the polyhedron of K, but not on the complex K itself.*

This theorem seems to have been known for quite a while before the reasons behind it were properly understood. At the dawn of homology theory, the subject had a very combinatorial flavor, and the use of triangulations as a tool for understanding manifolds proved to be very successful. A fairly natural strategy for proving Theorem 21.16 was formulated near the beginning of the 20th century and was based on a conjecture called the **Hauptvermutung**:[35] it claims essentially that any two triangulations of the same topological space can be turned into the same triangulation by a process of subdivision. Subdivision replaces each individual simplex σ with a triangulation by smaller simplices, so it makes the chain groups $C_n(K;G)$ much larger, but it is not too hard to show that the homology resulting from these enlarged chain groups is isomorphic to the original, hence if the Hauptvermutung is true, Theorem 21.16 follows. The only trouble is that the Hauptvermutung is false, as was discovered in the 1960's; moreover, we now also know examples of closed topological manifolds that cannot be triangulated at all, so that simplicial complexes do not provide the ideal framework for understanding manifolds in general. But in the mean time, the mathematical community discovered much better ways of proving Theorem 21.16, namely by defining another invariant for arbitrary topological spaces X that manifestly only depends on the topology of X without any auxiliary structure, but also can be shown to match simplicial homology whenever X is a polyhedron. That invariant is singular homology, and it will be our topic for the rest of this semester.

22. Singular homology (July 4, 2023)

So here's the challenge: how do we define a topological invariant that captures the same information as simplicial homology, but without ever referring to a simplicial complex? The answer to this turns out to be fairly simple, but speaking for myself, the first time I heard it, I thought it sounded crazy. There seemed to be no way that one could ever compute such a thing, or if one could, then it was hard to imagine what geometric insight would be gained from the computation. I've been leading up to this definition gradually over the last few lectures in order to give you some intuition about what kind of invariant we are looking for and why. The hope is that, equipped with this intuition, your first reaction to seeing the definition of singular homology might be that it has a fighting chance of answering some question you actually care about.

It will be convenient to first establish some basic principles of the subject known as *homological algebra*. We have already seen an example of the first definition in our discussion of simplicial homology.

[35] This is what the conjecture was called in English—one does not translate the word *Hauptvermutung*.

Definition 22.1. A ($\mathbb{Z}$-graded) **chain complex** (*Kettenkomplex*) of abelian groups (C_*, ∂) consists of a sequence $\{C_n\}_{n\in\mathbb{Z}}$ of abelian groups together with homomorphisms $\partial_n : C_n \to C_{n-1}$ for each $n \in \mathbb{Z}$ such that $\partial_{n-1} \circ \partial_n : C_n \to C_{n-2}$ is the trivial homomorphism for every n.

We sometimes denote the direct sum of all the chain groups C_n in a chain complex by

$$C_* := \bigoplus_{n\in\mathbb{Z}} C_n,$$

whose elements can all be written as finite sums $\sum_i a_i$ with $a_i \in C_{n_i}$ for some integers $n_i \in \mathbb{Z}$. An element $x \in C_*$ is said to have **degree** (*Grad*) n if $x \in C_n$. The individual homomorphisms $\partial_n : C_n \to C_{n-1}$ extend uniquely to a homomorphism $\partial : C_* \to C_*$ which has **degree** -1, meaning it maps elements of any given degree to elements of one degree less. We sometimes indicate this by abusing notation and writing

$$\partial : C_* \to C_{*-1}.$$

The collection of relations $\partial_{n-1} \circ \partial_n = 0$ for all n can now be abbreviated by the single relation

$$\partial^2 = 0,$$

which is equivalent to the condition that $\operatorname{im} \partial_{n+1} \subset \ker \partial_n$ for every n. We call ∂ the **boundary map** (*Randoperator*) in the complex. Elements in $\ker \partial \subset C_*$ are called **cycles** (*Zykel*), while elements in $\operatorname{im} \partial \subset C_*$ are called **boundaries** (*Ränder*).

Definition 22.2. The **homology** (*Homologie*) of a chain complex (C_*, ∂) is the sequence of abelian groups

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

for $n \in \mathbb{Z}$. We sometimes denote

$$H_*(C_*, \partial) := \bigoplus_{n\in\mathbb{Z}} H_n(C_*, \partial),$$

which makes $H_*(C_*, \partial)$ a $\mathbb{Z}$-graded abelian group.

Every element of $H_n(C_*, \partial)$ can be written as an equivalence class $[c]$ for some n-cycle $c \in \ker \partial_n$, and we call $[c]$ the **homology class** (*Homologieklasse*) represented by c. Two cycles $a, b \in \ker \partial_n$ are called **homologous** (*homolog*) if $[a] = [b] \in H_n(C_*, \partial)$, meaning $a - b \in \operatorname{im} \partial_{n+1}$.

Remark 22.3. For the examples of chain complexes (C_*, ∂) we consider in this course, C_n is always the trivial group for $n < 0$, mainly because the degree n typically corresponds to a geometric dimension and dimensions cannot be negative. But there is no need to assume this in the general algebraic definitions. In other settings, there are plenty of interesting examples of chain complexes that have nontrivial elements of negative degree.

The next definition will be needed when we want to show that continuous maps between topological spaces induce homomorphisms of their singular homology groups.

Definition 22.4. Given two chain complexes (A_*, ∂^A) and (B_*, ∂^B), a **chain map** (*Kettenabbildung*) from (A_*, ∂^A) to (B_*, ∂^B) is a sequence of homomorphisms $f_n : A_n \to B_n$ for $n \in \mathbb{Z}$ such that the following diagram commutes:

$$\begin{tikzcd}
\ldots \arrow[r] & A_{n+1} \arrow[r, "\partial^A_{n+1}"] \arrow[d, "f_{n+1}"] & A_n \arrow[r, "\partial^A_n"] \arrow[d, "f_n"] & A_{n-1} \arrow[r, "\partial^A_{n-1}"] \arrow[d, "f_{n-1}"] & \ldots \\
\ldots \arrow[r] & B_{n+1} \arrow[r, "\partial^B_{n+1}"] & B_n \arrow[r, "\partial^B_n"] & B_{n-1} \arrow[r, "\partial^B_{n-1}"] & \ldots
\end{tikzcd} \tag{22.1}$$

In other words, a chain map is a homomorphism $f : A_* \to B_*$ of degree zero satisfying $\partial^B \circ f = f \circ \partial^A$.

PROPOSITION 22.5. *Any chain map $f : (A_*, \partial^A) \to (B_*, \partial^B)$ determines homomorphisms $f_* : H_n(A_*, \partial^A) \to H_n(B_*, \partial^B)$ for every $n \in \mathbb{Z}$ via the formula*

$$f_*[a] := [f(a)].$$

PROOF. There are two things to prove: first, that whenever $a \in A_n$ is a cycle, so is $f(a) \in B_n$. This is clear since $\partial^A a = 0$ implies $\partial^B(f(a)) = f(\partial^A a) = 0$ by the chain map condition. Second, we need to know that f maps boundaries to boundaries, so that it descends to a well-defined homomorphism $\ker \partial_n^A / \operatorname{im} \partial_{n+1}^A \to \ker \partial_n^B / \operatorname{im} \partial_{n+1}^B$. This is equally clear, since $a = \partial^A x$ implies $f(a) = f(\partial^A x) = \partial^B f(x)$. □

With these algebraic preliminaries out of the way, we now proceed to define the chain complex of singular homology. As in simplicial homology, we fix an arbitrary abelian group G as auxiliary data, called the **coefficient group**; in practice it will usually be either $\mathbb{Z}$ or $\mathbb{Z}_2$, occasionally $\mathbb{Q}$. Recall that for integers $n \geqslant 0$, the **standard n-simplex** is the set

$$\Delta^n = \{(t_0, \ldots, t_n) \in I^{n+1} \mid t_0 + \ldots + t_n = 1\}.$$

For each $k = 0, \ldots, n$, the **kth boundary face** of Δ^n is the subset

$$\partial_{(k)}\Delta^n := \{t_k = 0\} \subset \Delta^n,$$

which is canonically homeomorphic to Δ^{n-1} via the map

$$\partial_{(k)}\Delta^n \to \Delta^{n-1} : (t_0, \ldots, t_{k-1}, 0, t_{k+1}, \ldots, t_n) \mapsto (t_0, \ldots, t_{k-1}, t_{k+1}, \ldots, t_n). \tag{22.2}$$

DEFINITION 22.6. Given a topological space X, a **singular n-simplex** in X is a continuous map $\sigma : \Delta^n \to X$.

Let $\mathcal{K}_n(X)$ denote the set of all singular n-simplices in X, and define the **singular n-chain group** with coefficients in G by

$$C_n(X; G) = \bigoplus_{\sigma \in \mathcal{K}_n(X)} G.$$

Note that this definition also makes sense for $n < 0$ if we agree that $\mathcal{K}_n(X)$ is then empty since there is no such thing as a simplex of negative dimension, hence the groups $C_n(X; G)$ are trivial in these cases. In general, elements in $C_n(X; G)$ can be written as finite sums $\Sigma_i a_i \sigma_i$ where $a_i \in G$ and $\sigma_i \in \mathcal{K}_n(X)$. This clearly looks similar to the simplicial chain groups, but if you're paying attention properly, you may at this point be feeling nervous about the fact that $C_n(X; G)$ is a *bloody enormous* group: algebraically it is very simple, but the set $\mathcal{K}_n(X)$ that generates it is usually uncountably infinite. It's probably even larger than you are imagining, because a singular n-simplex is not just a "simplex-shaped" subset of X, but it is also the parametrization of that subset, so any two distinct parametrizations $\sigma : \Delta^n \to X$, even if they have exactly the same image, define different elements of $\mathcal{K}_n(X)$ and thus different generators of $C_n(X; G)$.[36] If this makes you nervous, then you are right to feel nervous: it is a minor miracle that we will eventually be able to extract useful and computable information from groups as large as $C_n(X; G)$. You will see.

The next step is to define a boundary map $C_n(X; G) \to C_{n-1}(X; G)$. As in simplicial homology, this is done by writing a formula for $\partial\sigma$ for each generator $\sigma \in \mathcal{K}_n(X)$, and the formula follows the same orientation convention that we saw in our discussion of oriented triangulations, cf. Definition 20.8: set

$$\partial\sigma := \sum_{k=0}^{n} (-1)^k \left(\sigma|_{\partial_{(k)}\Delta^n}\right) \in C_{n-1}(X; \mathbb{Z}),$$

[36]The word "singular" in this context refers to the fact that there is no condition beyond continuity required for the maps $\sigma : \Delta^n \to X$, i.e. they need not be injective, nor differentiable (even if X happens to be a smooth manifold), and so their images might not look "simplex-shaped" at all, but could instead be full of singularities.

where each $\sigma|_{\partial_{(k)}\Delta^n}$ is regarded as a singular $(n-1)$-simplex using the identification $\partial_{(k)}\Delta^n = \Delta^{n-1}$ from (22.2).

This uniquely determines a homomorphism

$$\partial : C_n(X;G) \to C_{n-1}(X;G) : \sum_i a_i\sigma_i \mapsto \sum_i a_i\ \partial\sigma_i,$$

and the usual cancelation phenomenon implies:

LEMMA 22.7. $\partial^2 = 0$. □

The n**th singular homology group** (*singuläre Homologiegruppe*) with coefficients in G is now defined by

$$H_n(X;G) := H_n\left(C_*(X;G),\partial\right).$$

In the case $G = \mathbb{Z}$, this is often abbreviated by

$$H_n(X) := H_n(X;\mathbb{Z}).$$

The direct sum of these groups for all n is denoted by $H_*(X;G)$, though informally, this notation is also sometimes used with the symbol "$*$" acting as an integer-valued variable just like n.

I encourage you to compare the following result with our computation of the bordism groups $\Omega_0(X)$ and $\Omega_0^{\mathrm{SO}}(X)$ in Lecture 21.

PROPOSITION 22.8. *For any space X and any coefficient group G, $H_0(X;G) \cong \bigoplus_{\pi_0(X)} G$, i.e. it is a direct sum of copies of G for every path-component of X.*

PROOF. Since Δ^0 is a one-point space, the set $\mathcal{K}_0(X)$ of singular 0-simplices $\sigma : \Delta^0 \to X$ can be identified naturally with X, and we shall write 0-chains accordingly as finite sums $\sum_i a_i x_i$ with $a_i \in G$ and $x_i \in X$. Similarly, Δ^1 is homeomorphic to the unit interval $I = [0,1]$, and if we choose a homeomorphism $[0,1] \to \Delta^1$ sending 1 to $\partial_{(0)}\Delta^1$ and 0 to $\partial_{(1)}\Delta^1$, we can think of each $\sigma \in \mathcal{K}_1(X)$ as a path $\sigma : I \to X$ and write the boundary operator as

$$\partial\sigma = \sigma(1) - \sigma(0) \in C_0(X;\mathbb{Z}).$$

Since there are no (-1)-chains, every $a \in G$ and $x \in X$ then define a 0-cycle $ax \in C_0(X;G)$, but ax and ay are homologous whenever x and y belong to the same path-component since then any path $\sigma : I \to X$ from x to y gives $\partial(a\sigma) = ay - ax$. Choosing a point x_α in each path-component X_α, we can now say that every 0-cycle is homologous to a unique 0-cycle of the form $\sum_\alpha c_\alpha x_\alpha$, where the sum ranges over all the path-components of X but only finitely many of the coefficients $c_\alpha \in G$ are nonzero. If two cycles of this form are homologous, then they differ by the boundary of a 1-chain, which is a finite linear combination of paths, and since each path is confined to a single path-component and has two end points with opposite orientations, the conclusion is that both 0-cycles have the same coefficients. □

The next result is a straightforward exercise based on the definitions, and you should also compare it with our previous discussion of the bordism groups of a point, if only to observe that the result is very different: while bordism groups require some information about the classification of manifolds which has nothing to do with the one-point space, the singular homology of $\{\mathrm{pt}\}$ is much simpler.

EXERCISE 22.9. Show that for the 1-point space $\{\mathrm{pt}\}$ and any coefficient group G, singular homology satisfies

$$H_n(\{\mathrm{pt}\};G) \cong \begin{cases} G & \text{for } n = 0, \\ 0 & \text{for } n \neq 0. \end{cases}$$

Hint: For each integer $n \geqslant 0$, there is exactly one singular n-simplex $\Delta^n \to \{\mathrm{pt}\}$, so the chain groups $C_n(\{\mathrm{pt}\}; G)$ are all naturally isomorphic to G. What is $\partial : C_n(\{\mathrm{pt}\}; G) \to C_{n-1}(\{\mathrm{pt}\}; G)$?

Let us discuss the group $H_1(X; \mathbb{Z})$ for an arbitrary space X. As noted above in our proof of Proposition 22.8, Δ^1 is homeomorphic to the interval I, thus there is a bijection

$$\{\text{paths } I \to X\} \leftrightarrow \mathcal{K}_1(X) \tag{22.3}$$

which identifies each path γ with a singular 1-simplex (denoted by the same symbol) such that, under the canonical identification of $\mathcal{K}_0(X)$ with X,

$$\partial\gamma = \gamma(1) - \gamma(0).$$

Notice in particular that if γ is a loop, then it also defines a 1-cycle. More generally, let us write elements of $C_1(X; \mathbb{Z})$ as finite sums $\sum_i m_i\gamma_i$ where $m_i \in \mathbb{Z}$ and the γ_i are understood as singular 1-simplices via the above bijection, so

$$\partial \sum_i m_i\gamma_i = \sum_i m_i\,(\gamma_i(1) - \gamma_i(0)) \in C_0(X; \mathbb{Z}).$$

Now observe that since the coefficients m_i are integers, we are free to assume they are all ± 1 at the cost of allowing repeats in the finite list of paths γ_i. It will then be convenient to think of $-\gamma_i$ as the reversed path γ_i^{-1}, which makes sense if you look at the boundary formula since

$$\partial(-\gamma_i) = -(\gamma_i(1) - \gamma_i(0)) = \gamma_i(0) - \gamma_i(1) = \gamma_i^{-1}(1) - \gamma_i^{-1}(0) = \partial(\gamma_i^{-1}).$$

Thinking in these terms and continuing to assume $m_i = \pm 1$, $\sum_i m_i\gamma_i$ will now be a cycle if and only if the finite set of paths $\gamma_i^{m_i}$ can be arranged in some order so that they form a loop, i.e. each can be concatenated with the next in the list, and the last can be concatenated with the first. This is precisely what is needed in order to ensure that every 0-simplex in $\partial\sum_i m_i\gamma_i$ cancels out. This suggests a relationship between $H_1(X; \mathbb{Z})$ and $\pi_1(X)$, but notice that there is some ambiguity in the correspondence: in general there may be multiple ways that the paths $\gamma_i^{m_i}$ can be ordered to produce a loop, and different loops produced in this way need not always be homotopic to each other. In fact, one should not expect $H_1(X; \mathbb{Z})$ and $\pi_1(X)$ to be the same, since $H_1(X; \mathbb{Z})$ is abelian by definition, but $\pi_1(X)$ usually is not. It turns out that the next best thing is true.

Theorem 22.10. *For any path-connected space X with base point $x_0 \in X$, the bijection (22.3) determines a group homomorphism*

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

which descends to an isomorphism of the abelianization $\pi_1(X, x_0)/[\pi_1(X, x_0), \pi_1(X, x_0)]$ to $H_1(X; \mathbb{Z})$.

We say that a cycle $c \in C_*(X; G)$ is **nullhomologous** if $[c] = 0 \in H_*(X; G)$, or equivalently, c is a boundary. According to the discussion above, every loop $\gamma : I \to X$ with $\gamma(0) = \gamma(1) = x_0$ can be viewed as a 1-cycle, and that cycle is nullhomologous if and only if $[\gamma]$ belongs to the commutator subgroup of $\pi_1(X, x_0)$.

Example 22.11. Recall from Exercise 14.13 the embedded loop $\gamma : S^1 \to \Sigma_g$ for $g \geqslant 2$ whose image separates Σ_g into two surfaces of genus $h \geqslant 1$ and $k \geqslant 1$ respectively with one boundary component each:

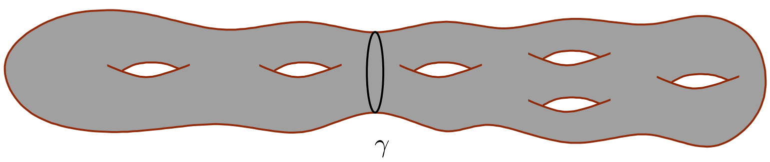

We computed in that exercise that $[\gamma]$ is a nontrivial element of the commutator subgroup of $\pi_1(\Sigma_g)$, thus by Theorem 22.10, γ represents the trivial class in $H_1(\Sigma_g;\mathbb{Z})$. This should not be surprising, since γ also parametrizes the boundary of a compact oriented submanifold of Σ_g, e.g. for this same reason, γ also represents the trivial bordism class in $\Omega_1^{\mathrm{SO}}(\Sigma_g)$. One can find an explicit 2-chain whose boundary is γ by decomposing the surface $\Sigma_{h,1}$ into 2-simplices so as to reinterpret the inclusion $\Sigma_{h,1} \hookrightarrow \Sigma_g$ as a linear combination of singular 2-simplices in Σ_g.

The proof of Theorem 22.10 is not trivial, but it is simple enough to leave as a guided homework problem (see Exercise 22.12 below). The homomorphism $h : \pi_1(X) \to H_1(X;\mathbb{Z})$ is called the **Hurewicz map**. There exists a similar Hurewicz homomorphism $\pi_k(X) \to H_k(X;\mathbb{Z})$ for every $k \geqslant 1$, which we will discuss near the end of *Topologie II* if time permits. Note that for $k \geqslant 2$, $\pi_k(X)$ is always abelian, so it is reasonable in those cases to hope that the Hurewicz map might be an honest isomorphism. A result called Hurewicz's theorem gives conditions under which this turns out to hold, thus providing a nice way to compute higher homotopy groups in some cases since, as we will see, computing homology is generally easier. But there are also simple examples in which $\pi_k(X)$ and $H_k(X;\mathbb{Z})$ are totally different. We saw for instance in the previous lecture that $\pi_2(\mathbb{T}^2) = 0$ due to the lifting theorem, but one can use any oriented triangulation of $\mathbb{T}^2$ to produce a singular 2-cycle that can be shown to be nontrivial in $H_2(\mathbb{T}^2;\mathbb{Z})$. Homology classes in the image of the Hurewicz map are sometimes called *spherical* homology classes. The example of $\mathbb{T}^2$ shows that for $n \geqslant 2$, one cannot generally expect all classes in $H_n(X;\mathbb{Z})$ to be spherical.

EXERCISE 22.12. Let us prove Theorem 22.10. Assume X is a path-connected space, fix $x_0 \in X$ and abbreviate $\pi_1(X) := \pi_1(X, x_0)$, so elements of $\pi_1(X)$ are represented by paths $\gamma : I \to X$ with $\gamma(0) = \gamma(1) = x_0$. Identifying the standard 1-simplex

$$\Delta^1 := \left\{ (t_0, t_1) \in \mathbb{R}^2 \;\middle|\; t_0 + t_1 = 1,\ t_0, t_1 \geqslant 0 \right\}$$

with $I := [0, 1]$ via the homeomorphism $\Delta^1 \to I : (t_0, t_1) \mapsto t_1$, every path $\gamma : I \to X$ corresponds to a singular 1-simplex $\Delta^1 \to X$, which we shall denote by $\tilde{h}(\gamma)$ and regard as an element of the singular 1-chain group $C_1(X;\mathbb{Z})$. Show that $\tilde{h}$ has each of the following properties:

(a) If $\gamma : I \to X$ satisfies $\gamma(0) = \gamma(1)$, then $\partial \tilde{h}(\gamma) = 0$.
(b) For any constant path $e : I \to X$, $\tilde{h}(e) = \partial\sigma$ for some singular 2-simplex $\sigma : \Delta^2 \to X$.
(c) For any paths $\alpha, \beta : I \to X$ with $\alpha(1) = \beta(0)$, the concatenated path $\alpha \cdot \beta : I \to X$ satisfies $\tilde{h}(\alpha) + \tilde{h}(\beta) - \tilde{h}(\alpha \cdot \beta) = \partial\sigma$ for some singular 2-simplex $\sigma : \Delta^2 \to X$.
Hint: Imagine a triangle whose three edges are mapped to X via the paths α, β and $\alpha \cdot \beta$. Can you extend this map continuously over the rest of the triangle?
(d) If $\alpha, \beta : I \to X$ are two paths that are homotopic with fixed end points, then $\tilde{h}(\alpha) - \tilde{h}(\beta) = \partial f$ for some singular 2-chain $f \in C_2(X;\mathbb{Z})$.
Hint: If you draw a square representing a homotopy between α and β, you can decompose this square into two triangles.
(e) Applying $\tilde{h}$ to paths that begin and end at the base point x_0, deduce that $\tilde{h}$ determines a group homomorphism $h : \pi_1(X) \to H_1(X;\mathbb{Z}) : [\gamma] \mapsto [\tilde{h}(\gamma)]$.

We call $h : \pi_1(X) \to H_1(X;\mathbb{Z})$ the **Hurewicz homomorphism**. Notice that since $H_1(X;\mathbb{Z})$ is abelian, $\ker h$ automatically contains the commutator subgroup $[\pi_1(X), \pi_1(X)] \subset \pi(X)$ (see Exercise 12.21), thus h descends to a homomorphism on the abelianization of $\pi_1(X)$,

$$\Phi : \pi_1(X) \big/ [\pi_1(X), \pi_1(X)] \to H_1(X;\mathbb{Z}).$$

We will now show that this is an isomorphism by writing down its inverse. For each point $p \in X$, choose arbitrarily a path $\omega_p : I \to X$ from x_0 to p, and choose ω_{x_0} in particular to be the constant path. Regarding singular 1-simplices $\sigma : \Delta^1 \to X$ as paths $\sigma : I \to X$ under the usual identification of I with Δ^1, we can then associate to every singular 1-simplex $\sigma \in C_1(X;\mathbb{Z})$ a concatenated path

$$\widetilde{\Psi}(\sigma) := \omega_{\sigma(0)} \cdot \sigma \cdot \omega_{\sigma(1)}^{-1} : I \to X$$

which begins and ends at the base point x_0, hence $\widetilde{\Psi}(\sigma)$ represents an element of $\pi_1(X)$. Let $\Psi(\sigma)$ denote the equivalence class represented by $\widetilde{\Psi}(\sigma)$ in the abelianization $\pi_1(X)/[\pi_1(X), \pi_1(X)]$. This uniquely determines a homomorphism[37]

$$\Psi : C_1(X;\mathbb{Z}) \to \pi_1(X)/[\pi_1(X), \pi_1(X)] : \sum_i m_i\sigma_i \mapsto \sum_i m_i\Psi(\sigma_i).$$

(f) Show that $\Psi(\partial\sigma) = 0$ for every singular 2-simplex $\sigma : \Delta^2 \to X$, and deduce that Ψ descends to a homomorphism $\Psi : H_1(X;\mathbb{Z}) \to \pi_1(X)/[\pi_1(X), \pi_1(X)]$.
(g) Show that $\Psi \circ \Phi$ and $\Phi \circ \Psi$ are both the identity map.
(h) For a closed surface Σ_g of genus $g \geqslant 2$, find an example of a nontrivial element in the kernel of the Hurewicz homomorphism $\pi_1(\Sigma_g) \to H_1(\Sigma_g)$. *Hint: See Exercise 14.13.*

23. Relative homology and long exact sequences (July 6, 2023)

The above results for $H_0(X;G)$ and $H_1(X;\mathbb{Z})$ provide some evidence that in spite of being defined as quotients of groups with uncountably many generators, the singular homology groups $H_n(X;G)$ might turn out to be computable more often than we'd expect. In this lecture we'll introduce a powerful computational tool that is also a fundamental concept in homological algebra. But before that, let us clarify in what sense singular homology is a topological invariant.

LEMMA 23.1. *Every continuous map $f : X \to Y$ determines a chain map $f_* : C_*(X;G) \to C_*(Y;G)$ via the formula $f_*\sigma := f \circ \sigma$ for singular n-simplices $\sigma : \Delta^n \to X$.*

PROOF. It is straightforward to check that $\partial(f_*\sigma) = f_*(\partial\sigma) \in C_{n-1}(Y;\mathbb{Z})$ for all $\sigma : \Delta^n \to X$, thus the uniquely determined homomorphism

$$f_* : C_n(X;G) \to C_n(Y;G) : \sum_i a_i\sigma_i \mapsto \sum_i a_i(f \circ \sigma_i)$$

defines a chain map. □

Notice that the chain maps in the above lemma also satisfy $(f \circ g)_* = f_* \circ g_*$ whenever f and g are composable continuous maps, and the chain map induced by the identity map on X is simply the identity homomorphism on $C_*(X;G)$. Applying Proposition 22.5 thus gives the following result, which implies that homeomorphic spaces always have isomorphic singular homology groups:

COROLLARY 23.2. *Continuous maps $f : X \to Y$ determine group homomorphisms $f_* : H_n(X;G) \to H_n(Y;G)$ for every n and G such that $(f \circ g)_* = f_* \circ g_*$ whenever f and g can be composed, and the identity map satisfies* $(\mathrm{Id})_* = \mathbb{1}$. □

REMARK 23.3. Recall that in the analogue of Corollary 23.2 for the fundamental group, the map $f : X \to Y$ is required to be base-point preserving, due to the fact that the definitions of $\pi_1(X)$ and $\pi_1(Y)$ require choices of base points in X and Y respectively. In most applications, base points are an extra piece of data that one doesn't actually care about but needs to keep track of anyway. One of the advantages of singular homology in comparison with the fundamental group is that its definition does not require any choice of base point, and Corollary 23.2 thus holds for *arbitrary* continuous maps $f : X \to Y$.

[37] Since $\pi_1(X)/[\pi_1(X), \pi_1(X)]$ is abelian, we are adopting the convention of writing its group operation as addition, so the multiplication of an integer $m \in \mathbb{Z}$ by an element $\Psi(\sigma) \in \pi_1(X)/[\pi_1(X), \pi_1(X)]$ is defined accordingly.

We will show in the next lecture that the homomorphisms f_* induced by continuous maps f only depend on f up to homotopy, which has the easy consequence that $H_*(X;G)$ only depends on the homotopy type of X.

But first, let us generalize the discussion somewhat. Algebraic gadgets often have the feature that they become easier to compute if you add more structure to them, sometimes at the cost of making the basic definitions slightly more elaborate. We will now do that with singular homology by introducing the *relative homology* groups of pairs. A **pair of spaces** (X, A), often abbreviated as simply a "pair," (*topologisches Paar*) consists of a topological space X and a subset $A \subset X$. Given two pairs (X, A) and (Y, B), a map $f : X \to Y$ is called a **map of pairs** if $f(A) \subset B$, and in this case we write

$$f : (X, A) \to (Y, B).$$

This is an obvious generalization of the definition of a pointed map, where arbitrary subsets have now replaced base points. Similarly, two maps of pairs $f, g : (X, A) \to (Y, B)$ are **homotopic** if there exists a homotopy $H : I \times X \to Y$ between f and g such that $H(s, \cdot) : (X, A) \to (Y, B)$ is a map of pairs for every $s \in I$, or equivalently,

$$H(I \times A) \subset B.$$

Two pairs (X, A) and (Y, B) are **homeomorphic** if there exist maps of pairs $f : (X, A) \to (Y, B)$ and $g : (Y, B) \to (X, A)$ such that $g \circ f$ and $f \circ g$ are the identity maps on (X, A) and (Y, B) respectively, and f and g are in this case called **homeomorphisms of pairs**. If $g \circ f$ and $f \circ g$ are not necessarily equal but are homotopic (as maps of pairs) to the respective identity maps, then we call each of them a **homotopy equivalence of pairs** and say that (X, A) and (Y, B) are homotopy equivalent, written

$$(X, A) \underset{h.e.}{\simeq} (Y, B).$$

One can regard every individual space X as a pair by identifying it with $(X, \varnothing)$, in which case the above definitions reproduce the usual ones for maps between ordinary spaces.

The relative homology of a pair (X, A) is based on the trivial observation that since every singular simplex in A is also a singular simplex in X whose boundary faces are all contained in A, $C_n(A; G)$ is naturally a subgroup of $C_n(X; G)$ for each n, and the boundary map $\partial : C_n(X; G) \to C_{n-1}(X; G)$ sends $C_n(A; G)$ to $C_{n-1}(A; G)$. It follows that ∂ descends to a sequence of well-defined homomorphisms on the quotients

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

and since ∂^2 is still zero, $(C_*(X, A; G), \partial)$ is a chain complex, called the **relative singular chain complex** of the pair (X, A) with coefficients in G. Its homology groups are the **relative singular homology** (*relative singuläre Homologie*),

$$H_n(X, A; G) := H_n\left(C_*(X, A; G), \partial\right).$$

The case $A = \varnothing$ reproduces $H_n(X; G)$ as we defined it in the previous lecture, and these are sometimes called the **absolute** homology groups of X so as to distinguish them from relative homology groups. As in absolute homology, we may sometimes abbreviate the case of integer coefficients by

$$H_n(X, A) := H_n(X, A; \mathbb{Z}).$$

Lemma 23.1 extends in an obvious way to the relative chain complex: if $f : (X, A) \to (Y, B)$ is a map of pairs, then the absolute chain map $f_* : C_*(X; G) \to C_*(Y; G)$ sends the subgroup $C_*(A; G)$ into $C_*(B; G)$ and thus descends to a chain map

$$f_* : C_*(X, A; G) \to C_*(Y, B; G),$$

implying the relative version of Corollary 23.2:

THEOREM 23.4. *Maps of pairs $f : (X, A) \to (Y, B)$ determine group homomorphisms $f_* : H_n(X, A; G) \to H_n(Y, B; G)$ for every n and G such that $(f \circ g)_* = f_* \circ g_*$ whenever f and g can be composed, and the identity map on (X, A) induces the identity homomorphism on $H_n(X, A; G)$.* $\square$

Since $C_n(X, A; G)$ is a quotient, its elements are technically equivalence classes, but in order to avoid having too many equivalence relations floating around in the same discussion, let us instead think of them as ordinary n-chains $c \in C_n(X; G)$, keeping in mind that two such n-chains $a, b \in C_n(X; G)$ define the same element of $C_n(X, A; G)$ whenever $a - b \in C_n(A; G)$, meaning a and b differ by a linear combination of simplices that are all contained in A. A chain $c \in C_n(X; G)$ can then be called a **relative cycle** if the element of $C_n(X, A; G)$ it determines is a cycle, which means ∂c belongs to $C_{n-1}(A; G)$. Notice that a relative cycle need not be an **absolute cycle** in general (meaning $\partial c = 0$), though absolute cycles also define relative cycles. Relative cycles $c \in C_n(X; G)$ define relative homology classes $[c] \in C_n(X, A; G)$, and two relative cycles $b, c \in C_n(X; G)$ are homologous (meaning $[b] = [c] \in H_n(X, A; G)$) if and only if

$$b - c = a + \partial x \quad \text{for some } a \in C_n(A; G),\ x \in C_{n+1}(X; G).$$

In particular, a relative cycle is nullhomologous if and only if it is the sum of a boundary plus a chain contained in A. If you find these algebraic relations overly abstract and would like some advice on how to actually *visualize* relative cycles, see the extended digression at the end of this lecture.

The reason for introducing the relative homology groups $H_*(X, A; G)$ was *not* that we wanted a tool for distinguishing non-homeomorphic pairs—the relative homology is such a tool, but our primary interest remains the space X on its own, rather than the pair (X, A). The usefulness of relative homology lies in the fact that there is a relation between the three groups $H_*(X; G)$, $H_*(A; G)$ and $H_*(X, A; G)$ for any pair (X, A), and indeed, one might hope to encounter situations in which two out of these three groups are easy to compute, so that a computation of the third one then comes for free. Let's make this idea more precise.

We begin with a seemingly trivial observation: let $i : A \hookrightarrow X$ and $j : X = (X, \varnothing) \hookrightarrow (X, A)$ denote the natural inclusions,[38] and consider the sequence of chain maps

$$0 \longrightarrow C_*(A; G) \xrightarrow{i_*} C_*(X; G) \xrightarrow{j_*} C_*(X, A; G) \to 0, \tag{23.1}$$

where the first and last maps are each trivial. The map j_* is obviously surjective, as it is actually just the quotient projection

$$C_*(X; G) \to C_*(X, G)/C_*(A; G) = C_*(X, A; G).$$

The map i_* is similarly the inclusion $C_*(A; G) \hookrightarrow C_*(X; G)$ and is thus injective, and its image is precisely the kernel of j_*. This means that every term in this sequence has the property that the image of the preceding map equals the kernel of the next one. In general, a sequence of abelian groups with homomorphisms

$$\ldots \longrightarrow A_{n-2} \xrightarrow{f_{n-2}} A_{n-1} \xrightarrow{f_{n-1}} A_n \xrightarrow{f_n} A_{n+1} \xrightarrow{f_{n+1}} A_{n+2} \longrightarrow \ldots$$

is called **exact** (*exakt*) if $\ker f_n = \operatorname{im} f_{n-1}$ for every $n \in \mathbb{Z}$. If all the groups except for two neighboring groups in the sequence are trivial, then it suffices to look at a sequence of four groups with only one nontrivial homomorphism

$$0 \longrightarrow A_1 \xrightarrow{f} A_2 \longrightarrow 0,$$

[38] Strictly speaking, j in this context is just the identity map on X, but we cannot call it that since we are viewing it as a map between two non-identical pairs of spaces. It is a map of pairs due to the trivial fact that $\varnothing \subset A$.

and the exactness of the sequence then simply means that $f : A_1 \to A_2$ is both injective and surjective, i.e. it is an isomorphism. In this sense, one can think of an exact sequence as a generalization of the notion of an isomorphism between two abelian groups. The next simplest case is what is called a **short exact sequence** (*kurze exakte Sequenz*), in which all except three of the groups and two of the homomorphisms are trivial,

$$0 \longrightarrow A_1 \xrightarrow{f_1} A_2 \xrightarrow{f_2} A_3 \longrightarrow 0.$$

Exactness in this case means three things: f_1 is injective, f_2 is surjective, and $\operatorname{im} f_1 = \ker f_2$. The sequence in (23.1) is what we call a **short exact sequence of chain maps**, because the abelian groups in each term are also chain complexes and the homomorphisms between them are chain maps. One can now wonder what happens if we replace these chain complexes with their homology groups and the chain maps with the induced homomorphisms on homology: will the resulting sequence be exact? The answer is no, but what is actually true is much better and more useful than this:

Theorem 23.5. *Suppose (A_*, ∂^A), (B_*, ∂^B) and (C_*, ∂^C) are chain complexes and*

$$0 \longrightarrow A_* \xrightarrow{f} B_* \xrightarrow{g} C_* \longrightarrow 0$$

is a short exact sequence of chain maps. Then there exists a natural homomorphism $\partial_* : H_n(C_*, \partial^C) \to H_{n-1}(A_*, \partial^A)$ for each $n \in \mathbb{Z}$ such that the sequence

$$\begin{aligned} \ldots \xrightarrow{\partial_*} H_{n+1}(A_*, \partial^A) \xrightarrow{f_*} H_{n+1}(B_*, \partial^B) \xrightarrow{g_*} H_{n+1}(C_*, \partial^C) \\ \xrightarrow{\partial_*} H_n(A_*, \partial^A) \xrightarrow{f_*} H_n(B_*, \partial^B) \xrightarrow{g_*} H_n(C_*, \partial^C) \\ \xrightarrow{\partial_*} H_{n-1}(A_*, \partial^A) \xrightarrow{f_*} H_{n-1}(B_*, \partial^B) \xrightarrow{g_*} H_{n-1}(C_*, \partial^C) \xrightarrow{\partial_*} \ldots \end{aligned} \tag{23.2}$$

is exact.

The sequence of homology groups in this theorem is called a **long exact sequence** (*lange exakte Sequenz*), and the maps $\partial_* : H_n(C_*, \partial^C) \to H_{n-1}(A_*, \partial^A)$ are called the **connecting homomorphisms** in this sequence. In particular, this result turns (23.1) into the so-called **long exact sequence of the pair** (X, A),

$$\ldots \to H_{n+1}(X, A; G) \xrightarrow{\partial_*} H_n(A; G) \xrightarrow{i_*} H_n(X; G) \xrightarrow{j_*} H_n(X, A; G) \xrightarrow{\partial_*} H_{n-1}(A; G) \to \ldots \tag{23.3}$$

To see why this might be useful, notice what it implies if we happen to know for some reason that one of the three groups $H_n(X; G)$, $H_n(A; G)$ or $H_n(X, A; G)$ is trivial for every n; for concreteness, let's suppose it is known that $H_*(X, A; G) = 0$. This knowledge turns the long exact sequence (23.3) into an infinite collection of two-term exact sequences

$$0 \longrightarrow H_n(A; G) \xrightarrow{i_*} H_n(X; G) \longrightarrow 0,$$

implying that for every n, the map $i_* : H_n(A; G) \to H_n(X; G)$ is an isomorphism. If we are also lucky enough to know already what $H_*(A; G)$ is, then the computation of $H_*(X; G)$ is thus complete. An argument of this type will be used in Lecture 25 as the final step in computing $H_*(S^n; \mathbb{Z})$ for every $n \geqslant 1$.

Theorem 23.5 is a purely algebraic statement, and it is proved by a straightforward but nonetheless slightly surprising procedure known as "diagram chasing". I will not give the full argument here, because that would bore you to tears, but I will explain the first couple of steps, and I highly recommend that you work through the rest yourself the next time you are half-asleep and in need of amusement on an airplane, or recovering from surgery on heavy pain medication, as the case may be.[39] The basic idea is to write down a great big commutative diagram, examine at each step exactly what information you can deduce from exactness and commutativity, and then let the diagram tell you what to do.

Here is the diagram we need—it commutes because f and g are chain maps, and each of its rows is an exact sequence of abelian groups:

$$\begin{tikzcd}
 & \vdots \arrow[d, "\partial^A"] & \vdots \arrow[d, "\partial^B"] & \vdots \arrow[d, "\partial^C"] & \\
0 \arrow[r] & A_{n+1} \arrow[r, "f"] \arrow[d, "\partial^A"] & B_{n+1} \arrow[r, "g"] \arrow[d, "\partial^B"] & C_{n+1} \arrow[r] \arrow[d, "\partial^C"] & 0 \\
0 \arrow[r] & A_n \arrow[r, "f"] \arrow[d, "\partial^A"] & B_n \arrow[r, "g"] \arrow[d, "\partial^B"] & C_n \arrow[r] \arrow[d, "\partial^C"] & 0 \\
0 \arrow[r] & A_{n-1} \arrow[r, "f"] \arrow[d, "\partial^A"] & B_{n-1} \arrow[r, "g"] \arrow[d, "\partial^B"] & C_{n-1} \arrow[r] \arrow[d, "\partial^C"] & 0 \\
0 \arrow[r] & A_{n-2} \arrow[r, "f"] \arrow[d, "\partial^A"] & B_{n-2} \arrow[r, "g"] \arrow[d, "\partial^B"] & C_{n-2} \arrow[r] \arrow[d, "\partial^C"] & 0 \\
 & \vdots & \vdots & \vdots &
\end{tikzcd}$$

We start by writing down a reasonable candidate for the map $\partial_* : H_n(C_*, \partial^C) \to H_{n-1}(A_*, \partial^A)$. Given $[c] \in H_n(C_*, \partial^C)$, $c \in C_n$ is necessarily a cycle, and exactness tells us that $g : B_n \to C_n$ is surjective, hence $c = g(b)$ for some $b \in B_n$. Then using commutativity,

$$0 = \partial^C c = \partial^C g(b) = g(\partial^B b),$$

so $\partial^B b \in \ker g \subset B_{n-1}$, and using exactness again, this implies $\partial^B b = f(a)$ for some $a \in A_{n-1}$. Notice that a is uniquely determined by b since (using exactness again) f is injective. Applying commutativity again, we now observe that

$$f(\partial^A a) = \partial^B(f(a)) = \partial^B \partial^B b = 0$$

since $(\partial^B)^2 = 0$, and the injectivity of f then implies $\partial^A a = 0$. So just by chasing the diagram from C_n to A_{n-1}, we found a cycle $a \in A_{n-1}$, and it seems reasonable to define

$$\partial_*[c] := [a] \in H_{n-1}(A, \partial^A).$$

We need to check that this is well defined, as two arbitrary choices were made in the procedure going from $[c]$ to $[a]$. One was the choice of an element $b \in B_n$ with $g(b) = c$, so we could get a different cycle $a' \in A_{n-1}$ by choosing a different element $b' \in g^{-1}(c)$ and requiring $f(a') = \partial^B b'$. But then $b' - b$ belongs to $\ker g = \operatorname{im} f$, hence we can write $b' - b = f(x)$ for some $x \in A_n$, implying

$$f(a' - a) = f(a') - f(a) = \partial^B(b' - b) = \partial^B(f(x)) = f(\partial^A(x)),$$

and since f is injective, $a' - a = \partial^A x$, implying that a and a' are homologous cycles. The other choice we made was the cycle $c \in C_n$, which in principle we are free to replace by any homologous cycle $c' \in C_n$ and then follow the same procedure to produce a different cycle $a' \in A_{n-1}$. If we do

[39] I first learned about exact sequences around the same time that I had all four of my wisdom teeth removed in a complicated procedure that left me drowsily dependent on prescription pain medication for about three weeks afterward. It turns out that that was exactly the right frame of mind in which to work through diagram chasing arguments without getting bored.

this, then $c' - c = \partial^C z$ for some $z \in C_{n+1}$, and since g is surjective, $z = g(y)$ for some $y \in B_{n+1}$. We then have

$$c' - c = \partial^C(g(y)) = g(\partial^B(y)),$$

and since we now know that we are free to choose any $b \in g^{-1}(c)$ and $b' \in g^{-1}(c')$, we can set

$$b' := b + \partial^B(y).$$

This implies $\partial^B b' = \partial^B b$, thus the condition $f(a') = \partial^B b'$ produces $a' = a$, and we have finished the proof that ∂_* is well defined.

It remains to prove that ∂_* really is a homomorphism, and that the long exact sequence really is exact, i.e. that $\ker \partial_* = \operatorname{im} g_*$, $\ker g_* = \operatorname{im} f_*$ and $\ker f_* = \operatorname{im} \partial_*$. This can all be done by the same kinds of straightforward arguments as above, but I'm sure you can see now why I'm not going to write down the complete details here.

I have one final remark however about the long exact sequence of a pair (X, A). If you redo the diagram chase above for the particular short exact sequence (23.1), you end up with a precise and very natural formula for the connecting homomorphisms

$$\partial_* : H_n(X, A; G) \to H_{n-1}(A; G).$$

The procedure starts with a relative n-cycle $c \in C_n(X, A; G)$, from which we need to pick $b \in j_*^{-1}(c) \subset C_n(X; G)$, but if we apply the usual convention of regarding relative cycles in (X, A) as chains in X, then c is already in $C_n(X; G)$ and we can pick b to be exactly the same chain c. Next we look at $\partial c \in C_{n-1}(X; G)$ and find the unique cycle $a \in C_{n-1}(A; G)$ that is sent to ∂c under the inclusion $C_{n-1}(A; G) \hookrightarrow C_{n-1}(X; G)$. In other words, $a = \partial c$, so the "obvious" formula is the right one:

$$\partial_*[c] = [\partial c]. \tag{23.4}$$

This looks more trivial than it is, e.g. you might think that $[\partial c]$ should automatically be 0 because ∂c is a boundary, but the point is that c is a chain in X, it might not be confined to A, so ∂c is certainly a cycle in A (as a consequence of the fact that c is a relative chain in (X, A)) but it need not be the boundary of any chain in A, and $[\partial c]$ may very well be a nontrivial homology class in $H_{n-1}(A; G)$.

Exercise 23.6. Use the formula (23.4) to give a direct proof that the sequence (23.3) is exact.

Remark 23.7. Exercise 23.6 is straightforward and doable in a much shorter time than the proof of Theorem 23.5, so we could have skipped the abstract homological algebra discussion without losing anything that is essential for the current semester. However, I wanted to make the point that the long exact sequence of a pair is not just an isolated topological phenomenon—it is a special case of a much more general algebraic principle, and that principle reappears in many other contexts in various branches of mathematics. We will see it again several times in *Topologie II*.

The following **extended digression** is not logically necessary for our development of basic homology theory, but you might still appreciate some intuition on the following question: what do relative n-cycles actually *look like*? Actually, that's also a valid question when applied to absolute n-cycles, and we've only really addressed it so far for $n = 0$ and $n = 1$. The best way I know for visualizing absolute cycles is via the analogy with bordism theory. Recall that elements of $\Omega_n^{\mathrm{SO}}(X)$ are equivalence classes of maps $f : M \to X$ where M is a closed oriented n-manifold. If M admits an oriented triangulation, then after choosing an ordering for all the vertices in this triangulation and assigning orientations accordingly to each simplex in the triangulation, one can identify each k-simplex $\sigma \subset M$ with a map $\Delta^k \to M$ that parametrizes it, thus defining a singular k-simplex in M. For $k = n$ in particular, the condition in Definition 20.9 relating the orientations of neighboring n-simplices implies that the sum $\sum_i \epsilon_i \sigma_i$ of all the singular n-simplices in the triangulation—with appropriate signs $\epsilon_i = \pm 1$ attached in order to describe their orientations in the triangulation—is a cycle in $C_n(M;\mathbb{Z})$. This is true because in $\partial \sum_i \epsilon_i \sigma_i$, every $(n-1)$-simplex of the triangulation appears exactly twice, but the orientation condition requires these two instances to appear with opposite signs. The resulting singular homology class is denoted by

$$[M] := \left[\sum_i \epsilon_i \sigma_i \right] \in H_n(M;\mathbb{Z})$$

and called the **fundamental class** (*Fundamentalklasse*) of M. We cannot prove it right now, but we will see in *Topologie II* that $[M]$ does not depend on the choice of triangulation, and it can even be defined for arbitrary closed and oriented topological manifolds, which need not admit triangulations. The map $f : M \to X$ then determines a corresponding cycle $\sum_i \epsilon_i (f \circ \sigma_i) \in C_n(X;\mathbb{Z})$ and an n-dimensional homology class $f_*[M] \in H_n(X;\mathbb{Z})$.

How can we recognize when two n-cycles in X defined in this way are homologous, or equivalently, when $\sum_i \epsilon_i (f \circ \sigma_i)$ is nullhomologous? A nice answer can again be extracted from bordism theory. If $[(M,f)] = 0 \in \Omega_n^{\mathrm{SO}}(X)$, it means there exists a compact oriented $(n+1)$-manifold W with $\partial W \cong M$ and a map $F : W \to X$ with $F|_M = f$. Suppose W admits an oriented triangulation that restricts to ∂W as an oriented triangulation of M. Identifying the $(n+1)$-simplices τ_j in this triangulation with singular $(n+1)$-simplices in W and then adding them up with suitable signs $\epsilon_j = \pm 1$ as in the previous paragraph produces an $(n+1)$-chain in X of the form $\sum_j \epsilon_j (F \circ \tau_j)$, whose boundary is the n-cycle representing $f_*[M]$. Thus if oriented triangulations can always be assumed to exist, then $f_*[M] = 0 \in H_n(X;\mathbb{Z})$ whenever (M,f) is nullbordant, and similarly, $f_*[M] = g_*[N] \in H_n(X;\mathbb{Z})$ will hold whenever (M,f) and (N,g) are related by an oriented bordism. We will also see in *Topologie II* that these statements remain true without mentioning triangulations.

You may be wondering how general this discussion really is, i.e. does *every* integral homology class in X arise from a map of a closed manifold into X? The answer is in general no, but if X is a nice enough space like the polyhedron of a finite simplicial complex, then something almost as good is true. The proof of the following famous result of Thom would be far beyond the scope of this course, and we will not make use of it, but it is nice to know that it exists.

THEOREM 23.8 (R. Thom [Tho54]). *If X is a compact polyhedron, then for every $n \geqslant 0$ and $A \in H_n(X;\mathbb{Z})$, there exists a closed n-manifold M, a map $f : M \to X$ and a number $k \in \mathbb{N}$ such that $kA = f_*[M]$.* □

To talk about relative homology classes, we could now allow M to be a compact oriented n-manifold with boundary and assume that its oriented triangulation also defines an oriented triangulation of ∂M. The chain $\sum_i \epsilon_i \sigma_i \in C_n(M;\mathbb{Z})$ is then no longer a cycle, because $(n-1)$-simplices on ∂M are not canceled, they each appear exactly once. Instead, $\partial \sum_i \epsilon_i \sigma_i$ is an $(n-1)$-cycle representing the fundamental class of ∂M, and $\sum_i \epsilon_i \sigma_i$ is therefore a relative cycle in $(M, \partial M)$, defining a **relative fundmental class**

$$[M] \in H_n(M, \partial M;\mathbb{Z}).$$

Given a pair (X, A), any map $f : (M, \partial M) \to (X, A)$ now determines a relative cycle $\sum_i \epsilon_i (f \circ \sigma_i) \in C_n(X, A;\mathbb{Z})$ and relative homology class $f_*[M] \in H_n(X, A;\mathbb{Z})$. For intuition, it is usually helpful to assume that f is an embedding, so a relative n-cycle in (X, A) then looks like an oriented and triangulated compact n-dimensional submanifold in X whose boundary lies in A.

Finally, note that one can drop the orientations from this entire discussion at the cost of replacing $\mathbb{Z}$ coefficients with $\mathbb{Z}_2$. Indeed, if M is closed and has a triangulation but not one that is orientable, then the n-chain defined by adding up the n-simplices may not be a cycle because its boundary may include some $(n-1)$-simplex that appears twice without canceling. But since $2 = 0 \in \mathbb{Z}_2$, this sum still defines a cycle in $C_n(M;\mathbb{Z}_2)$ and therefore also a fundamental class

$$[M] \in H_n(M;\mathbb{Z}_2).$$

This reveals that unoriented bordism classes in $\Omega_n(X)$ determine homology classes in $H_n(X;\mathbb{Z}_2)$, and the analogue of Theorem 23.8 remains true in this case without any need for the multiplicative factor $k \in \mathbb{N}$.

24. Homotopy invariance and excision (July 11, 2023)

We need to prove two more theorems about singular homology before it becomes a truly useful tool. Both will require a bit of work, but the almost immediate payoff will be that we can then compute the homology of spheres in every dimension. This has several important applications, including the general case of the Brouwer fixed point theorem, and the basic fact that open sets in $\mathbb{R}^n$ are never homeomorphic to open sets in $\mathbb{R}^m$ unless $n = m$. It is also the first step in developing an algorithm to compute the singular homology of any CW-complex, a general class of "reasonable" spaces that includes all smooth manifolds and all simplicial complexes.

Our first task for today is homotopy invariance.

THEOREM 24.1. *The map $f_* : H_n(X,A;G) \to H_n(Y,B;G)$ induced for each $n \in \mathbb{Z}$ by a map of pairs $f : (X,A) \to (Y,B)$ depends only on the homotopy class of f (as a map of pairs).*

The obvious corollary about homotopy equivalent spaces is a result of tremendous theoretical importance, and I would like to point out how much simpler its proof is than that of the corresponding statement about fundamental groups (Theorem 10.23). The complication in the case of π_1 was that its definition depends on a choice of base point, but the notion of homotopy equivalence does not—as a result, we had to find a workaround to cope with the fact that homotopy inverses need not be base-point preserving. In homology, one can also allow for base points by considering pairs (X,A) where $A \subset X$ is a single point, but homotopies between maps of pairs are required to respect this extra data, which makes the proofs easier. And unlike the fundamental group, homology also makes sense for pairs (X,A) with $A = \varnothing$, in which case the terms "homotopy" and "homotopy equivalence" mean the same thing that they always did.

COROLLARY 24.2. *If $f : (X,A) \to (Y,B)$ is a homotopy equivalence of pairs, then the induced maps $f_* : H_n(X,A;G) \to H_n(Y,B;G)$ are isomorphisms.*

PROOF. Suppose $f : (X,A) \to (Y,B)$ is a homotopy equivalence, so it has a homotopy inverse $g : (Y,B) \to (X,A)$. Then $f \circ g$ and $g \circ f$ are homotopic to the identity maps on (Y,B) and (X,A) respectively, so that Theorem 24.1 gives $f_* \circ g_* = \mathbb{1}$ and $g_* \circ f_* = \mathbb{1}$ for the induced maps on homology, implying that both are isomorphisms. □

The proof of Theorem 24.1 requires another fundamental notion from homological algebra. It should be clear that if $f, g : X \to Y$ are two non-identical maps, then the induced chain maps $f_*, g_* : C_*(X;G) \to C_*(Y;G)$ will not be identical, even if f and g are homotopic. It is still possible however for two distinct chain maps to descend to exactly the same map between homology groups. What we need for Theorem 24.1 is an algebraic mechanism to recognize when this happens, and that mechanism is called *chain homotopy*.

DEFINITION 24.3. A **chain homotopy** (*Kettenhomotopie*) between two chain maps $f, g : (A_*, \partial^A) \to (B_*, \partial^B)$ is a sequence of homomorphisms $h_n : A_n \to B_{n+1}$ such that for every $n \in \mathbb{Z}$,

$$f_n - g_n = \partial^B_{n+1} \circ h_n + h_{n-1} \circ \partial^A_n.$$

In other words, a chain homotopy between f and g is a homomorphism $h : A_* \to B_*$ of degree $+1$ such that $f - g = \partial^B \circ h + h \circ \partial^A$. We sometimes abuse notation and write

$$h : A_* \to B_{*+1}$$

to emphasize that a chain homotopy is a homomorphism of degree 1.

Two chain maps that admit a chain homotopy between them are called **chain homotopic** (*kettenhomotop*), and it is not hard to show that this defines an equivalence relation on chain maps. You can picture a chain homotopy as a sequence of down-left diagonal arrows in the diagram (22.1), though you need to be a little careful with that diagram since a chain homotopy does not make it commute. The main importance of chain homotopies comes from the following result.

Proposition 24.4. *If there exists a chain homotopy between two chain maps f and g from (A_*, ∂^A) to (B_*, ∂^B), then they induce the same homomorphisms*

$$f_* = g_* : H_n(A_*, \partial^A) \to H_n(B_*, \partial^B)$$

for all $n \in \mathbb{Z}$.

Proof. If $h : A_* \to B_{*+1}$ is a chain homotopy, then given any $[a] \in H_n(A_*, \partial^A)$, we have $\partial^A a = 0$ and thus

$$f(a) - g(a) = \partial^B h(a) + h(\partial^A a) = \partial^B (h(a)),$$

hence $f(a)$ and $g(a)$ are homologous cycles. □

If you're seeing the notion of chain homotopies for the first time, you might think that the definition above looks a bit unmotivated—it is not obvious for instance whether this is the *only* reasonable algebraic condition that makes two chain maps induce the same map on homology. However, the following lemma and its proof provide convincing evidence that this definition is the right one: it turns out that chain homotopies are the *natural* algebraic structure that arises in the singular chain complex from a homotopy between continuous maps. We will see that they arise naturally in many other contexts as well.

Lemma 24.5. *If there exists a homotopy between the maps of pairs $f, g : (X, A) \to (Y, B)$, then there also exists a chain homotopy between the induced chain maps $f_*, g_* : C_*(X, A; G) \to C_*(Y, B; G)$.*

Theorem 24.1 is an immediate consequence of this lemma and Proposition 24.4, so our remaining task is to prove the lemma. For notational simplicity, let us start under the assumption

$$A = B = \varnothing,$$

as the general case will only require a few extra remarks beyond this. Suppose $H : I \times X \to Y$ is a homotopy between $f = H(0, \cdot)$ and $g = H(1, \cdot)$. Associate to each singular n-simplex $\sigma : \Delta^n \to X$ the map

$$h_\sigma : I \times \Delta^n \to Y : (s, t) \mapsto H(s, \sigma(t)),$$

so $h_\sigma(0, \cdot) = f \circ \sigma$ and $h_\sigma(1, \cdot) = g \circ \sigma$. If we pretend for a moment that the maps in this picture are all embeddings, then we can picture h_σ as tracing out a "prism-shaped" region in Y whose boundary consists of three pieces, two of which are the n-simplices traced about by $f_*\sigma$ and $g_*\sigma$. If we pay proper attention to orientations, then $f_*\sigma$ will get a negative orientation because the boundary orientation for $\partial(I \times \Delta^n)$ induces opposite orientations on $\{0\} \times \Delta^n$ and $\{1\} \times \Delta^n$. But there is a third piece of $\partial(I \times \Delta^n)$ that we haven't mentioned yet, namely $I \times \partial\Delta^n$. If we regard

$I \times \Delta^n$ as a compact oriented $(n+1)$-manifold with boundary, then its oriented boundary turns out to be[40]

$$\partial(I \times \Delta^n) = (-\{0\} \times \Delta^n) \cup (\{1\} \times \Delta^n) \cup (-I \times \partial\Delta^n). \tag{24.1}$$

This relation will be the geometric motivation behind the chain homotopy formula.

The idea now is to define a chain homotopy $h : C_*(X;G) \to C_{*+1}(Y;G)$ by associating to each singular n-simplex $\sigma : \Delta^n \to X$ a linear combination of singular $(n+1)$-simplices in Y determined by the prism map $h_\sigma : I \times \Delta^n \to Y$. Unfortunately, $I \times \Delta^n$ is not a simplex, but there are various natural ways to decompose it into simplices, i.e. to triangulate it. In principle, the result should not depend on how this is done, so long as the triangulation has reasonable properties, thus we will not explain the details here except to state what properties are needed:

LEMMA 24.6. *There exists a sequence of oriented triangulations of the sequence of spaces $I \times \Delta^n$ for $n = 0, 1, 2, \dots$ satisfying the following properties:*

(1) *$\{0\} \times \Delta^n$ and $\{1\} \times \Delta^n$ are boundary faces of $(n+1)$-simplices in the triangulation of $I \times \Delta^n$;*
(2) *Under the natural identification of each boundary face $\partial_{(k)}\Delta^n$ with Δ^{n-1}, the triangulation of $I \times \Delta^n$ restricts to $I \times \partial_{(k)}\Delta^n$ as the triangulation of $I \times \Delta^{n-1}$.*

A precise algorithm to produce such triangulations of $I \times \Delta^n$ is described in [**Hat02**, p. 112]. I recommend taking a moment to draw pictures of how it might be done for $n = 1$ and $n = 2$. In the following, we will assume that parametrizations $\tau_i : \Delta^{n+1} \to I \times \Delta^n$ of the finite set of $(n+1)$-simplices in these triangulations have also been chosen such that for a suitable choice of signs $\epsilon_i = \pm 1$ determined by their orientations,

$$\sum_i \epsilon_i \tau_i \in C_{n+1}(I \times \Delta^n; \mathbb{Z})$$

defines a relative cycle in $(I \times \Delta^n, \partial(I \times \Delta^n))$; in other words, all interior n-simplices in the triangulation of $I \times \Delta^n$ appear twice with opposite signs in $\partial \sum_i \epsilon_i \tau_i$, so that what remains is an n-chain in the boundary. The stated conditions on the triangulation guarantee in fact that $\partial \sum_i \epsilon_i \tau_i$ will consist of the following terms:

(1) A single term for the obvious parametrization $\Delta^n \to \{1\} \times \Delta^n$, whose attached coefficient we can assume without loss of generality is $+1$;
(2) Another term for the obvious parametrization $\Delta^n \to \{0\} \times \Delta^n$, whose attached coefficient must now be -1 for orientation reasons;
(3) Linear combinations (with coefficients ± 1) of the n-simplices triangulating $I \times \partial_{(k)}\Delta^n = I \times \Delta^{n-1}$ for each boundary face of Δ^n.

With this in hand, there is a unique homomorphism $h : C_n(X;G) \to C_{n+1}(Y;G)$ defined on each singular n-simplex $\sigma : \Delta^n \to X$ by the formula

$$h(\sigma) := \sum_i \epsilon_i (h_\sigma \circ \tau_i) \in C_{n+1}(Y;\mathbb{Z}),$$

where the sum is over all the parametrized $(n+1)$-simplices $\tau_i : \Delta^{n+1} \to I \times \Delta^n$ in our triangulation from Lemma 24.6, and the $\epsilon_i = \pm 1$ are determined by their orientations as outlined above. In light of (24.1), we then have

$$\partial h(\sigma) = g_*\sigma - f_*\sigma - h(\partial\sigma),$$

[40] One can deduce the signs in (24.1) from things that were said in Lecture 20, though it's a bit tedious, and for now I would encourage you to just believe me that the signs are correct. There is an easier way to see it using the notion of orientation for *smooth* manifolds and their tangent spaces, which we do not have space to talk about here, but you'll likely see things like this again in differential geometry at some point.

where the third term comes from the restriction of h_σ to the triangulated subset $-I \times \partial\Delta^n$ in the oriented boundary of $I \times \Delta^n$. It follows that $h : C_*(X;G) \to C_{*+1}(Y;G)$ satisfies $\partial \circ h + h \circ \partial = g_* - f_*$, i.e. h is a chain homotopy.

This concludes the proof of Lemma 24.5 in the case $A = B = \emptyset$. In the general case, the given homotopy satisfies the additional assumption

$$H(I \times A) \subset B,$$

thus following through with the above construction, h_σ has image contained in B whenever σ has image in A. It follows that the chain homotopy we constructed sends $C_n(A;G)$ into $C_{n+1}(B;G)$ and thus descends to the quotients as a chain homotopy

$$h_* : C_*(X,A;G) \to C_{*+1}(Y,B;G)$$

between the relative chain maps $f_*, g_* : C_*(X,A;G) \to C_*(Y,B;G)$. The proof of the lemma is now complete, and with it, the proof of the homotopy invariance of singular homology.

Let us pick some low-hanging fruit from this result.

Corollary 24.7 (via Exercise 22.9). *For any contractible space X and any coefficient group G, $H_n(X;G)$ is isomorphic to G for $n = 0$ and vanishes for $n \neq 0$.* □

Corollary 24.8 (via Theorem 22.10). *If X is homotopy equivalent to S^1, then $H_1(X;\mathbb{Z}) \cong \mathbb{Z}$.* □

The second big theorem for today is called the *excision* property. It is based on the intuition that since $H_*(X,A;G)$ is supposed to ignore anything that happens entirely inside the subset A, removing smaller subsets $B \subset A$ should not change the relative homology, i.e. we expect

$$H_*(X\backslash B, A\backslash B;G) \cong H_*(X,A;G).$$

This works under a mild assumption on what it means for a subset B to be "smaller" than A.

Theorem 24.9 (excision). *For any pair (X,A), if $B \subset A$ is a subset with closure contained in the interior of A, then the inclusion of pairs $i : (X\backslash B, A\backslash B) \hookrightarrow (X,A)$ induces isomorphisms*

$$i_* : H_n(X\backslash B, A\backslash B;G) \xrightarrow{\cong} H_n(X,A;G)$$

for all n and G.

The assumption $B \subset \bar{B} \subset \mathring{A} \subset A \subset X$ means essentially that the two open subsets $\mathring{A}$ and $X\backslash\bar{B}$ cover X. In this setting, let us say that a chain $c \in C_n(X;G)$ is *decomposable* if c can be written as a sum of a chain in A plus a chain in $X\backslash B$, i.e. c belongs to the subgroup $C_n(A;G) + C_n(X\backslash B;G) \subset C_n(X;G)$. The excision theorem is closely related to the observation that every relative n-cycle in (X,A) is homologous to one that is decomposable. Indeed, if this is true and every $[c] \in H_n(X,A;G)$ can be written without loss of generality as $c = c_A + c_{X\backslash B}$ for some $c_A \in C_n(A;G)$ and $C_{X\backslash B} \in C_n(X\backslash B;G)$, then since c is a relative cycle, $\partial c \in C_{n-1}(A;G)$, implying $\partial c_{X\backslash B}$ is also in $C_{n-1}(A;G)$ since ∂c_A must be as well, thus $\partial c_{X\backslash B} \in C_{n-1}(A\backslash B;G)$. This proves that $c_{X\backslash B}$ is a relative n-cycle for the pair $(X\backslash B, A\backslash B)$, so it represents a homology class in $H_n(X\backslash B, A\backslash B;G)$, and obviously

$$i_*[c_{X\backslash B}] = [c]$$

since $c_A \in C_n(A;G)$ represents the trivial element of $C_n(X,A;G)$. This proves surjectivity in Theorem 24.9, modulo the detail about why we are allowed to restrict our attention to decomposable chains. The latter is where most of the hard work is hidden.

Let us reframe the discussion slightly and suppose $\mathcal{U}, \mathcal{V} \subset X$ are two subsets whose interiors form an open cover of X,

$$X = \mathring{\mathcal{U}} \cup \mathring{\mathcal{V}}.$$

We would like to develop a procedure for replacing any given chain $c \in C_n(X;G)$ with one that is in the subgroup $C_n(\mathcal{U};G) + C_n(\mathcal{V};G) \subset C_n(X;G)$ but represents the same homology class in cases where c is a (relative) cycle. If you followed the extended digression on how to visualize n-cycles at the end of the previous lecture, then you can imagine an intuitive reason why this should be possible: consider a homology class that is presented in the form $f_*[M] \in H_n(X;\mathbb{Z})$ for some triangulated oriented n-manifold M and a map $f : M \to X$. In this case, the definition of a cycle representing $f_*[M]$ depends on a choice of oriented triangulation for M, but we do not really expect the homology class $f_*[M]$ to depend on this triangulation, and in particular, we should be free to replace the triangulation by a *finer* one, which has more simplices but each one small enough to be contained in either $\mathcal{U}$ or $\mathcal{V}$ (or both). It is not hard to imagine that one could achieve this simply by triangulating each individual simplex in M to decompose it into strictly smaller simplices, and the process could then be repeated finitely many times to make the simplices as small as we like. This process is called *subdivision*. We shall now describe an inductive algorithm that makes the idea precise.

The **barycentric subdivision** of the standard n-simplex Δ^n is an oriented triangulation of Δ^n defined as follows. If $n = 0$, then Δ^0 is only a single point, so it cannot be subdivided any further and our triangulation of Δ^0 will consist only of that single 0-simplex. Now by induction, assume the desired triangulation of Δ^m has already been defined for all $m \leqslant n-1$. Under the natural identification of each boundary face $\partial_{(k)}\Delta^n$ with Δ^{n-1}, this means in particular that a triangulation of $\partial_{(k)}\Delta^n$ has been chosen for each $k = 0, \dots, n$. Now for each $(n-1)$-simplex σ in that triangulation, define σ' to be the n-simplex in Δ^n that is linearly spanned by the n vertices of σ plus one extra vertex that is in the interior of Δ^n, the so-called **barycenter**

$$b_n := \left(\frac{1}{n+1}, \dots, \frac{1}{n+1}\right) \in \Delta^n.$$

It is straightforward to check that the collection of all n-simplices σ' defined in this way from $(n-1)$-simplices σ in boundary faces $\partial_{(k)}\Delta^n$ forms a triangulation of Δ^n, and one can also assign it an orientation based on the orientations of the triangulations of $\partial_{(k)}\Delta^n$. Some pictures for $n = 1, 2, 3$ are shown in [**Hat02**, p. 120].

As usual with triangulations of manifolds, one can assign to each n-simplex $\sigma' \subset \Delta^n$ in the barycentric subdivision of Δ^n a parametrization $\tau : \Delta^n \xrightarrow{\cong} \sigma' \subset \Delta^n$ such that the sum over all such parametrized simplices τ_i with attached signs $\epsilon_i = \pm 1$ determined by their orientations in the triangulation produces a relative n-cycle in $(\Delta^n, \partial\Delta^n)$,

$$\sum_i \epsilon_i \tau_i \in C_n(\Delta^n;\mathbb{Z}), \qquad \partial \sum_i \epsilon_i \tau_i \in C_{n-1}(\partial\Delta^n;\mathbb{Z}),$$

where $(n-1)$-simplices in the interior of Δ^n do not appear in $\partial \sum_i \epsilon_i \tau_i$ because each is a boundary face of two n-simplices whose induced boundary orientations cancel. We can then use this to define a homomorphism

$$S : C_n(X;G) \to C_n(X;G)$$

via the formula

$$S(\sigma) := \sum_i \epsilon_i (\sigma \circ \tau_i)$$

for each $n \geqslant 0$ and $\sigma : \Delta^n \to X$. Essentially, S replaces each singular n-simplex σ by a linear combination (with coefficients ± 1) of the restrictions of σ to the subdivided pieces of its domain.

Lemma 24.10. *$S : C_*(X;G) \to C_*(X;G)$ is a chain map.*

Proof. This follows from the relation $\partial S(\sigma) = S(\partial\sigma)$ for each $\sigma : \Delta^n \to X$, which is a direct consequence of the inductive nature of the subdivision algorithm: boundary faces of the smaller simplices in the subdivision are also the simplices in a subdivision of the original boundary faces. □

LEMMA 24.11. *$S : C_*(X;G) \to C_*(X;G)$ is chain homotopic to the identity map.*

PROOF. As in the proof of Lemma 24.5, the chain homotopy here comes from a particular choice of oriented triangulation of the prism $I \times \Delta^n$. A picture of this triangulation and a precise algorithm to construct it are given in [**Hat02**, p. 122]. We want it in particular to have the following properties:

(1) Its restriction to $\{1\} \times \Delta^n$ is the barycentric subdivision of Δ^n;
(2) Its restriction to $\{0\} \times \Delta^n$ consists only of that one n-simplex, with no subdivision;
(3) Its restriction to each $I \times \partial_{(k)}\Delta^n$ matches the chosen triangulation of $I \times \Delta^{n-1}$.

The third property means that the construction is again inductive: we start with $n = 0$ by choosing the trivial triangulation of $I \times \Delta^0 = I$, and then increase the dimension one at a time such that the triangulation already defined for $I \times \Delta^{n-1}$ determines the triangulation of $I \times \Delta^n$. Since it is an oriented triangulation, one can now define a relative $(n+1)$-cycle in $(I \times \Delta^n, \partial(I \times \Delta^n))$ of the form

$$\sum_i \epsilon_i \tau_i \in C_{n+1}(I \times \Delta^n; \mathbb{Z}),$$

where $\tau_i : \Delta^{n+1} \to I \times \Delta^n$ are parametrizations of the simplices in the triangulation and the signs $\epsilon_i = \pm 1$ are determined by their orientations. Let

$$\pi : I \times \Delta^n \to \Delta^n$$

denote the obvious projection map. The desired chain homotopy $h : C_n(X;G) \to C_{n+1}(X;G)$ is then determined by the formula

$$h(\sigma) = \sum_i \epsilon_i \left(\sigma \circ \pi \circ \tau_i\right).$$

In computing $\partial h(\sigma)$, n-simplices in the interior of $I \times \Delta^n$ make no contribution due to the usual cancelations, but there are contributions from the induced triangulation of $\partial(I \times \Delta^n)$, and the chain homotopy relation again follows from the geometric formula (24.1) for the oriented boundary of $I \times \Delta^n$. Namely, restricting to $\{1\} \times \Delta^n$ gives the barycentric subdivision $S(\sigma)$, restricting to $-\{0\} \times \Delta^n$ gives $-\sigma$, and restricting to $-I \times \partial\Delta^n$ gives the same operator applied to $\partial\sigma$, hence

$$\partial h(\sigma) = S(\sigma) - \sigma - h(\partial\sigma),$$

proving $S - \mathbb{1} = \partial h + h\partial$. □

The chain homotopy result implies that our subdivision map $S : C_*(X;G) \to C_*(X;G)$ has the main property we want, namely it induces the identity homomorphism $H_*(X;G) \to H_*(X;G)$, and since S clearly also preserves $C_*(A;G)$ for any $A \subset X$, the same is also true for the relative homology groups of (X, A). It then remains true if we replace S by any iteration S^m for integers $m \geqslant 1$, thus we can apply S repeatedly in order to make the individual simplices in a chain as small as we like. In particular, for any $c \in C_*(X;G)$, we will have $S^m c \in C_*(\mathcal{U};G) + C_*(\mathcal{V};G)$ for m sufficiently large. This is enough information to prove the excision theorem, so let's go ahead and do that.

PROOF OF THEOREM 24.9. The hypotheses of the theorem imply that X is the union of the interiors of $X \backslash B$ and A, so given any class $[c] \in H_n(X, A;G)$ with a relative n-cycle $c \in C_n(X;G)$ representing it, c can be replaced by an iterated subdivision $S^m c$ for large $m \in \mathbb{N}$ that represents the same relative homology class $[S^m c] = [c] \in H_n(X, A;G)$ but is also decomposable, meaning it is the sum of a chain in $X\backslash B$ with a chain in A. Let's assume that c has already been replaced with $S^m c$ in this way, so that without loss of generality,

$$c = c_A + c_{X\backslash B} \qquad \text{for some} \qquad c_A \in C_n(A;G),\ c_{X\backslash B} \in C_n(X\backslash B;G).$$

Having made this assumption, the reason why $i_* : H_n(X\backslash B, A\backslash B; G) \to H_n(X, A; G)$ is surjective was explained already in the paragraph after the statement of the theorem: the fact that $c \in C_n(X, A; G)$ is a relative n-cycle means $\partial c \in C_n(A;G)$ and therefore also $\partial c_{X\backslash B} \in C_n(A;G)$, so that $c_{X\backslash B}$ is a relative n-cycle in $(X\backslash B, A\backslash B)$, thus representing a class $[c_{X\backslash B}] \in H_n(X\backslash B, A\backslash B; G)$ that satisfies

$$i_*[c_{X\backslash B}] = [c].$$

The proof that $i_* : H_n(X\backslash B, A\backslash B; G) \to H_n(X, A; G)$ is injective uses subdivision in a slightly different way. Suppose $c \in C_n(X\backslash B; G)$ is a relative n-cycle representing a homology class $[c] \in H_n(X\backslash B, A\backslash B; G)$ with $i_*[c] = 0 \in H_n(X, A; G)$. Since i is just an inclusion map, $i_*[c] = 0$ means that after reinterpreting c as an n-chain in X instead of just in $X\backslash B$, c is a boundary of some $(n+1)$-chain in X, modulo one that is contained in A, i.e. we have

$$c = \partial b + a \qquad \text{for some } b \in C_{n+1}(X;G) \text{ and } a \in C_n(A;G).$$

Applying ∂ to both sides of this equation gives $\partial c = \partial a$, which implies since c is a relative n-cycle in $(X\backslash B, A\backslash B)$ that $\partial a \in C_n(A\backslash B; G)$, i.e. none of the singular simplices that make up the $(n-1)$-cycle ∂a intersect B. If we happened to know that the chains $b \in C_{n+1}(X;G)$ and $a \in C_n(A;G)$ also have that property, i.e. that they are made up only of singular simplices that do not intersect B, then we would be done: indeed, we could then interpret b as an $(n+1)$-chain in $X\backslash B$ and a as an n-chain in $A\backslash B$, so that the relation $c = \partial b + a$ also implies $[c] = 0 \in H_n(X\backslash B, A\backslash B; G)$. As it stands, each of b and a might very well intersect B, but we can now use subdivision to replace them with chains that do not. Indeed, the homology class $[c] \in H_n(X\backslash B, A\backslash B; G)$ does not change if we replace c with $S^m c$ for any $m \geqslant 1$, and since S is a chain map, the relation $c = \partial b + a$ then implies $S^m c = S^m(\partial b) + S^m a = \partial(S^m b) + S^m a$. Choosing m sufficiently large and replacing each of a, b, c with their m-fold subdivisions, we can now assume without loss of generality that all three are decomposable; for $c \in C_n(X\backslash B; G)$ and $a \in C_n(A;G)$ this is not new information since we already assumed them to be contained in $X\backslash B$ or A respectively, but for $b \in C_{n+1}(X;G)$ we can now write

$$b = b_A + b_{X\backslash B} \qquad \text{for some} \qquad b_A \in C_{n+1}(A;G),\ b_{X\backslash B} \in C_{n+1}(X\backslash B;G).$$

The relation $c = \partial b + a$ thus becomes

$$c = \partial b_{X\backslash B} + (\partial b_A + a)\,,$$

and we observe that since c and $\partial b_{X\backslash B}$ are both n-chains in $X\backslash B$, the same must therefore be true for $\partial b_A + a$, meaning it is actually contained in $A\backslash B$. This proves $[c] = 0 \in H_n(X\backslash B, A\backslash B; G)$. ☐

The remainder of this lecture should be considered optional for now, as it is not needed for the purposes of this semester's course. However, when we study cohomology next semester, we will need a slightly better version of the excision result than Theorem 24.9. One thing you've probably gathered by now is that a chain homotopy is always a useful thing to have, so when one exists, we should take note of it. Theorem 24.9 can be seen as a consequence of the stronger result that the inclusion $i : (X\backslash B, A\backslash B) \hookrightarrow (X, A)$ induces a **chain homotopy equivalence** (*Kettenhomotopieäquivalenz*)

$$i_* : C_*(X\backslash B, A\backslash B; G) \to C_*(X, A; G).$$

In case the meaning of this terminology is not obvious, this means there exists a chain map $\psi : C_*(X, A; G) \to C_*(X \backslash B, A \backslash B; G)$ such that $\psi \circ i_*$ and $i_* \circ \psi$ are each chain homotopic to the identity; we call ψ a **chain homotopy inverse** of i_*.

The following statement turns our previous discussion of subdivision into an actual chain homotopy equivalence that has several applications in the further development of the theory, e.g. we will use it again next semester when we discuss the homology analogue of the Seifert-van Kampen theorem, known as the *Mayer-Vietoris* exact sequence. To understand the statement, it is important to be aware that for any subsets $\mathcal{U}, \mathcal{V} \subset X$, the subgroup $C_*(\mathcal{U}; G) + C_*(\mathcal{V}; G) \subset C_*(X; G)$ is also a chain complex in a natural way. Indeed, the boundary operator on $C_*(X; G)$ maps each of $C_*(\mathcal{U}; G)$ and $C_*(\mathcal{V}; G)$ to themselves, thus it also preserves their sum.

LEMMA 24.12. *For any subsets $\mathcal{U}, \mathcal{V} \subset X$ with $X = \mathring{\mathcal{U}} \cup \mathring{\mathcal{V}}$, the inclusion map*

$$j : C_*(\mathcal{U}; G) + C_*(\mathcal{V}; G) \hookrightarrow C_*(X; G)$$

admits a chain homotopy inverse

$$\rho : C_*(X; G) \to C_*(\mathcal{U}; G) + C_*(\mathcal{V}; G)$$

such that $\rho \circ j = \mathbb{1}$, and moreover, there is a chain homotopy $h : C_*(X; G) \to C_{*+1}(X; G)$ of $j \circ \rho$ to the identity such that h vanishes on $C_*(\mathcal{U}; G) + C_*(\mathcal{V}; G)$.

PROOF. Let me first point out how one would intuitively wish to prove this, and why it will not work. As observed above, any chain $c \in C_*(X; G)$ can be mapped into $C_*(\mathcal{U}; G) + C_*(\mathcal{V}; G)$ via S^m if the integer m is sufficiently large, so S^m seems like a good candidate for the chain homotopy inverse ρ. The problem however is that we don't know in general how large m needs to be, and in fact the answer depends on the chain c: for any fixed integer m, one can always find a singular n-simplex $\sigma : \Delta^n \to X$ whose boundary is close enough to the boundary of $\mathcal{U}$ or $\mathcal{V}$ so that the m-fold subdivision $S^m(\sigma)$ includes some simplex that is not fully contained in either one. This means that regardless of how large we make m, S^m can never map *all* of $C_*(X; G)$ into $C_*(\mathcal{U}; G) + C_*(\mathcal{V}; G)$, and it will require a bit more cleverness to come up with a candidate for a map ρ that does this. Our approach will be somewhat indirect: instead of writing down ρ, we will first write down a (somewhat naive) candidate for the chain homotopy h in terms of the chain homotopies between S^m and $\mathbb{1}$ for varying values of m. We will then be able to verify that h really is a chain homotopy between $\mathbb{1}$ and something; that so-called "something" will be defined to be ρ, whose further properties we can then verify.

Let $h_1 : C_*(X; G) \to C_{*+1}(X; G)$ denote the chain homotopy provided by Lemma 24.11 for the barycentric subdivision chain map $S : C_*(X; G) \to C_*(X; G)$, i.e. it satisfies $S - \mathbb{1} = \partial h_1 + h_1 \partial$. We claim that for all integers $m \geqslant 0$, the map

$$h_m := h_1 \sum_{k=0}^{m-1} S^k : C_*(X; G) \to C_{*+1}(X; G)$$

then satisfies

$$S^m - \mathbb{1} = \partial h_m + h_m \partial, \tag{24.2}$$

so h_m is a chain homotopy between S^m and the identity. Note that the case $m = 0$ is included here, with $S^0 = \mathbb{1}$ and $h_0 = 0$, so the claim is trivial in that case, and the definition of h_1 establishes it for $m = 1$. If we now use induction and assume that the claim holds for powers of S up to

$m-1 \geqslant 1$, then since S commutes with ∂,

$$\begin{aligned} S^m - \mathbb{1} &= (S^{m-1} - \mathbb{1})S + (S - \mathbb{1}) = (\partial h_{m-1} + h_{m-1}\partial)S + \partial h_1 + h_1\partial \\ &= \left(\partial h_1 \sum_{k=0}^{m-2} S^k + h_1 \sum_{k=0}^{m-2} S^k \partial\right) S + \partial h_1 + h_1 \partial = \partial h_1 \sum_{k=1}^{m-1} S^k + h_1 \sum_{k=1}^{m-1} S^k \partial + \partial h_1 + h_1 \partial \\ &= \partial h_1 \sum_{k=0}^{m-1} S^k + h_1 \sum_{k=0}^{m-1} S^k \partial = \partial h_m + h_m \partial. \end{aligned}$$

For any given $\sigma : \Delta^n \to X$, the iterated subdivision maps S^m can be assumed to satisfy

$$S^m(\sigma) \in C_*(\mathcal{U}; G) + C_*(\mathcal{V}; G) \tag{24.3}$$

if m is large enough, so for each each $n \geqslant 0$ and $\sigma : \Delta^n \to X$, let $m_\sigma \geqslant 0$ denote the smallest integer for which (24.3) holds with $m = m_\sigma$. We can then define a homomorphism $h : C_n(X; G) \to C_{n+1}(X; G)$ for each $n \geqslant 0$ via

$$h(\sigma) := h_{m_\sigma}(\sigma).$$

Let us see whether this is a chain homotopy. We have

$$\begin{aligned} (\partial h + h\partial)(\sigma) &= \partial h_{m_\sigma}(\sigma) + h_{m_\sigma}(\partial\sigma) + (h - h_{m_\sigma})(\partial\sigma) \\ &= (S^{m_\sigma} - \mathbb{1})(\sigma) + (h - h_{m_\sigma})(\partial\sigma) = \left([S^{m_\sigma} + (h - h_{m_\sigma})\partial] - \mathbb{1}\right)(\sigma). \end{aligned}$$

Use this to define $\rho : C_*(X; G) \to C_*(X; G)$ by

$$\rho(\sigma) := S^{m_\sigma}(\sigma) + (h - h_{m_\sigma})(\partial\sigma),$$

so the relation

$$\partial h + h\partial = \rho - \mathbb{1} \tag{24.4}$$

is satisfied. The latter implies that ρ is a chain map since applying ∂ from either the left or right on the left hand side of (24.4) gives $\partial h \partial$, thus on the right hand side we obtain $(\rho - \mathbb{1})\partial = \partial(\rho - \mathbb{1})$. To understand ρ better, we need to observe that each boundary face τ appearing in $\partial\sigma$ satisfies $m_\tau \leqslant m_\sigma$ since m_σ is clearly enough (but need not be the minimal number of) iterations of S to put σ (and therefore also τ) in $C_*(\mathcal{U}; G) + C_*(\mathcal{V}; G)$. Now if $\sigma \in C_*(\mathcal{U}; G) + C_*(\mathcal{V}; G)$, then $S^{m_\sigma}(\sigma) = \sigma$ since $m_\sigma = 0$, and the above remarks imply $h(\partial\sigma) = h_0(\partial\sigma) = 0$ as well, thus $\rho(\sigma) = \sigma$ and we conclude

$$\rho \circ j = \mathbb{1}.$$

It remains to show that for all $\sigma : \Delta^n \to X$, $\rho(\sigma)$ is a linear combination of simplices that are each contained in either $\mathcal{U}$ or $\mathcal{V}$. We have $S^{m_\sigma}(\sigma) \in C_*(\mathcal{U}; G) + C_*(\mathcal{V}; G)$ by the definition of m_σ, so it suffices to inspect the other term $(h - h_{m_\sigma})(\partial\sigma)$. Here again we observe that $\partial\sigma$ is a sum of singular $(n-1)$-simplices τ for which $m_\tau \leqslant m_\sigma$, and

$$(h - h_{m_\sigma})\tau = (h_{m_\tau} - h_{m_\sigma})\tau = -h_1 \sum_{k=m_\tau}^{m_\sigma - 1} S^k(\tau) \in C_n(\mathcal{U}; G) + C_n(\mathcal{V}; G).$$

This last conclusion requires you to recall how h_1 was constructed in the proof of Lemma 24.11: in particular, it maps any simplex that is contained in either $\mathcal{U}$ or $\mathcal{V}$ to a linear combination of simplices that have this same property.

One last detail: the chain homotopy $h : C_*(X; G) \to C_{*+1}(X; G)$ vanishes on $C_*(\mathcal{U}; G) + C_*(\mathcal{V}; G)$ since every singular n-simplex $\sigma : \Delta^n \to X$ with image in either $\mathcal{U}$ or $\mathcal{V}$ satisfies $m_\sigma = 0$, thus $h(\sigma) = h_{m_\sigma}(\sigma) = h_0(\sigma) = 0$. □

Now we can prove the "chain level" result that implies Theorem 24.9.

LEMMA 24.13. *If $A, B \subset X$ are subsets with $\bar{B} \subset \mathring{A}$, then the inclusion $i : (X\backslash B, A\backslash B) \hookrightarrow (X, A)$ induces a chain homotopy equivalence $i_* : C_*(X\backslash B, A\backslash B; G) \to C_*(X, A; G)$.*

PROOF. Consider the quotient chain complex $(C_*(X\backslash B; G) + C_*(A; G))\,/C_*(A; G)$, which has a natural identification with the group of all finite sums $\sum_i a_i\sigma_i$ with coefficients $a_i \in G$ and singular simplices $\sigma_i : \Delta^n \to X$ that have image in $X\backslash B$ but not contained in A. The point here is that while simplices with $\sigma(\Delta^n) \subset A$ are also generators of $C_*(X\backslash B; G) + C_*(A; G)$, they are all equivalent to zero in the quotient. As it happens, the quotient complex $C_*(X\backslash B, A\backslash B; G) = C_*(X\backslash B; G)/C_*(A\backslash B; G)$ can be described in exactly the same way, with the same set of generators: singular simplices that are contained in $X\backslash B$ but not contained in A. Since the obvious inclusion $C_*(X\backslash B; G) \hookrightarrow C_*(X\backslash B; G) + C_*(A; G)$ sends $C_*(A\backslash B; G)$ into $C_*(A; G)$, it follows that this inclusion descends to a chain map of quotient complexes

$$C_*(X\backslash B, A\backslash B; G) \to (C_*(X\backslash B; G) + C_*(A; G))\,/C_*(A; G)$$

which is in fact an *isomorphism* of chain complexes, i.e. it has an inverse, which is also a chain map. This is a trivial observation; we have not done anything interesting yet.

But in light of this identification of two quotient chain complexes, it will suffice to prove that the chain map

$$(C_*(X\backslash B; G) + C_*(A; G))\,/C_*(A; G) \xrightarrow{j} C_*(X; G)/C_*(A; G) = C_*(X, A; G) \tag{24.5}$$

induced on these quotients by the obvious inclusion

$$C_*(X\backslash B; G) + C_*(A; G) \overset{j}{\hookrightarrow} C_*(X; G)$$

is a chain homotopy equivalence. Since $X\backslash\bar{B}$ and $\mathring{A}$ form an open cover of X, Lemma 24.12 provides a chain homotopy inverse for j, namely the map $\rho : C_*(X; G) \to C_*(X\backslash B; G) + C_*(A; G)$, defined in terms of subdivision. That map satisfies $\rho \circ j = \mathbb{1}$, thus ρ restricts to the identity on the subgroup $C_*(A; G) \subset C_*(X; G)$ and therefore descends to a map on quotients going the opposite direction to j in (24.5). It also satisfies $j \circ \rho - \mathbb{1} = \partial h + h\partial$ for a chain homotopy $h : C_*(X; G) \to C_{*+1}(X; G)$ that vanishes on $C_*(A; G)$, thus h also descends to the quotient $C_*(X; G)/C_*(A; G)$ as a chain homotopy $h : C_*(X, A; G) \to C_{*+1}(X, A; G)$ satisfying $j \circ \rho - \mathbb{1} = \partial h + h\partial$ on the quotient complexes. □

REMARK 24.14. We will not need it this semester, but since the notions of chain maps and chain homotopies did not appear in our discussion of simplicial homology, you might wonder if they nonetheless have some role to play in that context. Chain maps arise for instance from *simplicial maps*: given two simplicial complexes $K = (V, S)$ and $K' = (V', S')$, a map $f : V \to V'$ is called a simplicial map if for every simplex σ of K, the images under f of the vertices of σ form the vertices (possibly with repetition) of a simplex of K'. A simplicial map naturally determines a continuous map of the associated polyhedra $|K| \to |K'|$ which maps each n-simplex in $|K|$ linearly to a k-simplex in $|K'|$ for some $k \leqslant n$. It is not hard to show that f also naturally induces a chain map $f_* : C_*(K; G) \to C_*(K'; G)$, defined by sending each n-simplex σ in K to its image k-simplex in K' if $k = n$ and otherwise sending σ to 0. In light of this, Proposition 22.5 implies (unsurprisingly) that any *bijective* simplicial map from K to K' induces an isomorphism of the simplicial homology groups $H_*^\Delta(K; G) \to H_*^\Delta(K'; G)$. Chain homotopies play an important role when one considers subdivisions of a simplicial complex, e.g. one can adapt the notion of barycentric subdivision so that it naturally associates to any simplicial complex K a larger complex K' with a homeomorphism of $|K'|$ to $|K|$ such that the simplices in K' triangulate the individual simplices of K into smaller pieces. This defines a chain map $S : C_*(K; G) \to C_*(K'; G)$ sending each simplex of K to the linear combination of simplices of K' that triangulate it, and importantly, S turns out to be a chain homotopy equivalence, so it follows from Proposition 24.4 that the induced homomorphism

$S_* : H_*^\Delta(K;G) \to H_*^\Delta(K';G)$ is an isomorphism. This was historically considered one of the major motivations to believe that simplicial homology depends only on the underlying space $|K|$ and not on the simplicial complex itself (cf. Theorem 21.16). We saw a closely analogous phenomenon in our proof of the excision property above, though in the simplicial context, one usually has to consult some of the older textbooks (e.g. [**Spa95**] is quite nice) to find adequate discussions of such topics.

25. The homology of the spheres, and applications (July 13, 2023)

It is time to put the results of the last few lectures together and compute $H_*(S^n;\mathbb{Z})$. The computation proceeds by induction on the dimension n, making use of the convenient fact that the suspension of S^n is homeomorphic to S^{n+1}. Suspensions, in fact, provide us with our first interesting example of a homotopy equivalence of pairs.

EXAMPLE 25.1. Recall from Lecture 11 that the **suspension** (*Einhängung*) SX of a space X is defined by gluing together two copies of its cone,

$$SX = C_+X \cup_X C_-X, \tag{25.1}$$

where $C_+X := ([0,1] \times X)/(\{1\} \times X)$, $C_-X := ([-1,0] \times X)/(\{-1\} \times X)$, and we identify X with the subset $\{0\} \times X$ in each. Let $p_\pm \in SX$ denote the points at the tips of the two cones, defined by collapsing $\{\pm 1\} \times X$. Then the inclusion

$$(C_+X, X) \hookrightarrow (SX\backslash\{p_-\}, C_-X\backslash\{p_-\})$$

is a homotopy equivalence of pairs. Indeed, one can define a deformation retraction $H : I \times (SX\backslash\{p_-\}) \to SX\backslash\{p_-\}$ by pushing points in $C_-X\backslash\{p_-\}$ continuously upward toward X while leaving C_+X fixed, so that $H(1,\cdot)$ is the identity while $H(0,\cdot)$ retracts $SX\backslash\{p_-\}$ to C_+X and $H(s,\cdot)$ preserves $C_-X\backslash\{p_-\}$ for every $s \in I$. The resulting retraction of pairs $(SX\backslash\{p_-\}, C_-X\backslash\{p_-\}) \to (C_+X, X)$ is a homotopy inverse for the inclusion. Let us spell this out more explicitly in the special case where $X = S^{n-1}$, so SX is then homeomorphic to S^n. The decomposition (25.1) then becomes a splitting of S^n into two hemispheres $\mathbb{D}^n_+ \cong \mathbb{D}^n \cong \mathbb{D}^n_-$ glued along an "equator" homeomorphic to S^{n-1},

$$S^n \cong \mathbb{D}^n_+ \cup_{S^{n-1}} \mathbb{D}^n_-,$$

and our homotopy equivalence of pairs is now the resulting inclusion map

$$(\mathbb{D}^n_+, S^{n-1}) \hookrightarrow (S^n\backslash\{p_-\}, \mathbb{D}^n_-\backslash\{p_-\}),$$

where p_- is now the "south pole," i.e. the center of $\mathbb{D}^n_-$.

The homotopy equivalence in Example 25.1 gives rise to an interesting relationship between $H_*(X;G)$ and $H_*(SX;G)$ for any space X. Ponder the following diagram:

$$\begin{tikzcd} H_k(X;G) & & H_{k+1}(SX;G) \arrow[d, "\varphi_*"] \\ H_{k+1}(C_+X, X; G) \arrow[u, "\partial_*"] \arrow[r, "i_*"] & H_{k+1}(SX\backslash\{p_-\}, C_-X\backslash\{p_-\}; G) \arrow[r, "j_*"] & H_{k+1}(SX, C_-X; G) \end{tikzcd} \tag{25.2}$$

Here ∂_* denotes the connecting homomorphism from the long exact sequence of the pair (C_+X, X), while the maps j_* and φ_* are induced by the obvious inclusions of pairs

$$(SX\backslash\{p_-\}, C_-X\backslash\{p_-\}) \xhookrightarrow{j} (SX, C_-X),$$
$$(SX, \varnothing) \xhookrightarrow{\varphi} (SX, C_-X).$$

Since $\{p_-\} \subset C_-X$ is a closed subset in the interior of C_-X, excision (Theorem 24.9) implies that j_* is an isomorphism. We claim that if $k \geqslant 1$, then ∂_* and φ_* are both also isomorphisms. For the first, consider the long exact sequence of the pair (C_+X, X):

$$\ldots \longrightarrow H_{k+1}(C_+X; G) \longrightarrow H_{k+1}(C_+X, X; G) \xrightarrow{\partial_*} H_k(X; G) \longrightarrow H_k(C_+X; G) \longrightarrow \ldots$$

Since C_+X is contractible, homotopy invariance implies that the first and last of these four terms vanish, as $H_n(\{\text{pt}\}; G) = 0$ for all $n > 0$. The sequence thus becomes

$$0 \longrightarrow H_{k+1}(C_+X, X; G) \xrightarrow{\partial_*} H_k(X; G) \longrightarrow 0$$

for each $k \geqslant 1$, so exactness implies that ∂_* is an isomorphism. For φ_*, we instead take an exerpt from the long exact sequence of (SX, C_-X):

$$\ldots \longrightarrow H_{k+1}(C_-X; G) \longrightarrow H_{k+1}(SX; G) \xrightarrow{\varphi_*} H_{k+1}(SX, C_-X; G) \longrightarrow H_k(C_-X; G) \longrightarrow \ldots$$

The contractibility of C_-X again makes the first and last terms vanish if $k \geqslant 1$, leaving

$$0 \longrightarrow H_{k+1}(SX; G) \xrightarrow{\varphi_*} H_{k+1}(SX, C_-X; G) \longrightarrow 0,$$

so that φ_* is also an isomorphism. We have proved:

THEOREM 25.2. *For all spaces X, abelian groups G and integers $k \geqslant 1$, the diagram (25.2) defines an isomorphism*

$$S_* = \varphi_*^{-1} \circ j_* \circ i_* \circ \partial_*^{-1} : H_k(X; G) \to H_{k+1}(SX; G).$$

□

EXERCISE 25.3. Show that for any k-cycle $b \in C_k(X; G) \subset C_k(SX; G)$, there exists a pair of $(k+1)$-chains $c_\pm \in C_{k+1}(C_\pm X; G) \subset C_{k+1}(SX; G)$ satisfying

$$\partial c_+ = -\partial c_- = b \tag{25.3}$$

and

$$S_*[b] = [c_+ + c_-]. \tag{25.4}$$

Note that $c_+ + c_- \in C_{n+1}(SX; G)$ is automatically a cycle since $\partial c_+ = -\partial c_-$. Show moreover that (25.4) is satisfied for any pair of chains $c_\pm$ satisfying (25.3).

For the spheres S^n with $n \geqslant 1$, we already know $H_0(S^n; G)$ and $H_1(S^n; \mathbb{Z})$; the former is G because S^n is path-connected (Proposition 22.8), and the latter is the abelianization of $\pi_1(S^n)$ by Theorem 22.10. Since $SS^n \cong S^{n+1}$, we can now compute $H_*(S^n; \mathbb{Z})$ inductively for every $n \geqslant 1$:

THEOREM 25.4. *For every $n \in \mathbb{N}$,*

$$H_k(S^n; \mathbb{Z}) \cong \begin{cases} \mathbb{Z} & \text{for } k = 0, n, \\ 0 & \text{for all other } k. \end{cases}$$

PROOF. Proposition 22.8 gives $H_0(S^n; \mathbb{Z}) \cong \mathbb{Z}$. For $k = n$, $H_n(S^n; \mathbb{Z}) \cong \mathbb{Z}$ follows by an inductive argument starting from $H_1(S^1; \mathbb{Z}) \cong \pi_1(S^1) \cong \mathbb{Z}$ and applying Theorem 25.2. For any $k = 1, \ldots, n-1$, a similar inductive argument starting from $H_1(S^{n-k+1}; \mathbb{Z}) = \pi_1(S^{n-k+1}) = 0$ gives $H_k(S^n; \mathbb{Z}) = 0$. For $k > n$, repeatedly applying Theorem 25.2 identifies $H_k(S^n; \mathbb{Z})$ with $H_{k-n}(S^0; \mathbb{Z})$, where $k - n > 0$ and S^0 is a discrete space of two points. But one can easily adapt Exercise 22.9 to prove by direct computation that $H_m(X; G) = 0$ for any $m > 0$ whenever X is a discrete space. □

We can now extend our proof of the Brouwer fixed point theorem to all dimensions. The basic ingredients are the same as before: first, if a map $f : \mathbb{D}^n \to \mathbb{D}^n$ has no fixed point, then we can use it to define a retraction $g : \mathbb{D}^n \to S^{n-1} = \partial\mathbb{D}^n$. In Lecture 10, we used the fundamental group to prove that no such retraction exists when $n = 2$. The argument for this did not require many specific properties of the fundamental group: the key point was just the fact that continuous maps $X \to Y$ induce homomorphisms $\pi_1(X) \to \pi_1(Y)$ in a way that is compatible with composition of maps, and the homology groups have this same property. In particular:

Exercise 25.5. Show that if $f : X \to A$ is a retraction to a subset $A \subset X$ with inclusion $i : A \hookrightarrow X$, then for all $n \in \mathbb{Z}$ and abelian groups G, $f_* : H_n(X; G) \to H_n(A; G)$ is surjective, while $i_* : H_n(A; G) \to H_n(X; G)$ is injective.

Proof of the Brouwer fixed point theorem. Arguing by contradiction, assume a map $f : \mathbb{D}^n \to \mathbb{D}^n$ without fixed points exists, and therefore also a retraction $g : \mathbb{D}^n \to S^{n-1}$. We may assume $n \geqslant 2$ since the case $n = 1$ follows already from the intermediate value theorem for continuous functions on $[-1, 1]$. By Exercise 25.5, g induces a surjective homomorphism

$$g_* : H_{n-1}(\mathbb{D}^n; \mathbb{Z}) \to H_{n-1}(S^{n-1}; \mathbb{Z}).$$

But this is impossible since $H_{n-1}(\mathbb{D}^n; \mathbb{Z}) \cong H_{n-1}(\{\text{pt}\}; \mathbb{Z}) = 0$ and $H_{n-1}(S^{n-1}; \mathbb{Z}) \cong \mathbb{Z}$. □

Here is another easy application.

Theorem 25.6. *A topological manifold of dimension n is not also a topological manifold of dimension $m \neq n$.*

Proof. Let us assume m and n are both at least 2, as the result can otherwise be proved via easier methods. (Hint: removing a point from $\mathbb{R}$ makes it disconnected.) We argue by contradiction and assume M is a manifold with an interior point admitting a neighborhood homeomorphic to $\mathbb{R}^n$ and also a neighborhood homeomorphic to $\mathbb{R}^m$ for $m \neq n$. By choosing a suitable pair of charts and writing down their transition maps, we can produce from this a pair of open neighborhoods of the origin $\Omega_n \subset \mathbb{R}^n$ and $\Omega_m \subset \mathbb{R}^m$ admitting a homeomorphism $f : \Omega_n \to \Omega_m$ with $f(0) = 0$. Choose $\epsilon > 0$ small enough so that f maps the ϵ-ball $B^n_\epsilon(0) \subset \Omega_n$ about the origin into the δ-ball $B^m_\delta(0) \subset \mathbb{R}^m$ for some $\delta > 0$, where the latter is also small enough so that $B^m_\delta(0) \subset \Omega_m$. Now pick a generator

$$A \in H_{n-1}(B^n_\epsilon(0)\backslash\{0\}; \mathbb{Z}) \cong H_{n-1}(S^{n-1}; \mathbb{Z}) \cong \mathbb{Z}.$$

Since $m \neq n$,

$$H_{n-1}(B^m_\delta(0)\backslash\{0\}; \mathbb{Z}) \cong H_{n-1}(S^{m-1}; \mathbb{Z}) = 0,$$

so restricting f to a map $B^n_\epsilon(0)\backslash\{0\} \to B^m_\delta(0)\backslash\{0\}$ gives $f_*A = 0 \in H_{n-1}(B^m_\delta(0)\backslash\{0\}; \mathbb{Z})$. But f^{-1} is also defined on $B^m_\delta(0)$, and restricting both f and f^{-1} to maps on punctured neighborhoods with the origin removed, we deduce

$$A = (f^{-1} \circ f)_* A = f^{-1}_* f_* A = 0,$$

which is a contradiction since A was assumed to generate $H_{n-1}(B^n_\epsilon(0)\backslash\{0\}; \mathbb{Z}) \neq 0$. □

26. Axioms, cells, and the Euler characteristic (July 18, 2023)

At this point, I believe I've proved everything that I promised to prove in earlier lectures, so the course *Topologie I* is officially over. Since we nonetheless have a bit of time left, the present lecture is included partly just for fun: none of what it contains should be considered examinable in the current semester, though some of it may provide a useful wider perspective on the material we've previously covered. All of it will also be treated in much more detail in next semester's *Topologie II* course.

The Eilenberg-Steenrod axioms. First a bit of good news: while the proofs of homotopy invariance and excision in Lecture 24 may have seemed somewhat unpleasant, we will hardly ever need to engage in such hands-on constructions via subdivision of simplices in the future. That is because almost everything one actually needs to know in order to use homology in applications follows from a small set of results that we've spent the last few lectures proving. These results form an axiomatic description of general "homology theories," which was first codified by Eilenberg-Steenrod [ES52] and Milnor [Mil62] around the middle of the 20th century. An **axiomatic homology theory** can be thought of as a function

$$(X, A) \mapsto h_*(X, A)$$

that associates to each pair of spaces a sequence of abelian groups $\{h_n(X, A)\}_{n\in\mathbb{Z}}$, and has some additional properties that make it computable for nice spaces and useful for applications in the same way that singular homology is. Identifying each single space X with the pair $(X, \varnothing)$ as usual, one abbreviates

$$h_n(X) := h_n(X, \varnothing).$$

Besides the actual groups $h_n(X, A)$, the theory h_* comes with some additional data: first, it should also associate to each map of pairs $f : (X, A) \to (Y, B)$ a sequence of homomorphisms

$$f_* : h_n(X, A) \to h_n(Y, B), \qquad n \in \mathbb{Z}$$

with the properties that $(f \circ g)_* = f_* \circ g_*$ whenever the composition of f and g makes sense, and the identity map $\mathrm{Id} : (X, A) \to (X, A)$ gives rise to the identity homomorphism $\mathrm{Id}_* = \mathbb{1} : h_n(X, A) \to h_n(X, A)$. Category theory has a technical term for things like this: we call h_* a **functor** from the category of pairs of topological spaces to the category of $\mathbb{Z}$-graded abelian groups. There is one additional piece of data: since the long exact sequences of pairs in singular homology were very useful in the computation of $H_*(S^n)$, we would like to have similar exact sequences for h_*, and one of the ingredients required for this is a sequence of *connecting homomorphisms*

$$\partial_* : h_n(X, A) \to h_{n-1}(A), \qquad n \in \mathbb{Z}.$$

Aside from fitting into an exact sequence as described below, we want these maps to be compatible with the homomorphisms induced on h_* by maps of pairs, in the following sense: any map of pairs $f : (X, A) \to (Y, B)$ restricts to a continuous map $A \to B$, so it induces homomorphisms $f_* : h_n(X, A) \to h_n(Y, B)$ and $f_* : h_n(A) \to h_n(B)$, which we would like to fit together with ∂_* into the following commutative diagram for each n:

$$\begin{tikzcd}
h_n(X, A) \arrow[r, "\partial_*"] \arrow[d, "f_*"] & h_{n-1}(A) \arrow[d, "f_*"] \\
h_n(Y, B) \arrow[r, "\partial_*"] & h_{n-1}(B)
\end{tikzcd}$$

The fancy category-theoretic term for this condition is "naturality": more specifically, ∂_* defines for each $n \in \mathbb{Z}$ a so-called **natural transformation** from the functor $(X, A) \mapsto h_n(X, A)$ to the functor $(X, A) \mapsto h_n(A) := h_n(A, \varnothing)$. The precise meanings of these terms from category theory will be discussed in the first lecture of next semester's course.

The original list of axioms stated in [ES52] included the properties described above, but they are usually not regarded as actual axioms in modern treatments, since they can instead be summarized with category-theoretic terminology such as "h_* is a functor and ∂_* is a natural transformation". The further conditions we want these things to satisfy are then the following:

- (HOMOTOPY) $f_* : h_*(X, A) \to h_*(Y, B)$ depends only on the homotopy class of $f : (X, A) \to (Y, B)$.

- (EXACTNESS) For the inclusions $i : A \hookrightarrow X$ and $j : (X, \varnothing) \hookrightarrow (X, A)$, the sequence
$$\ldots \longrightarrow h_{n+1}(X, A) \xrightarrow{\partial_*} h_n(A) \xrightarrow{i_*} h_n(X) \xrightarrow{j_*} h_n(X, A) \xrightarrow{\partial_*} h_{n-1}(A) \longrightarrow \ldots$$
is exact.
- (EXCISION) If $B \subset \bar{B} \subset \mathring{A} \subset A \subset X$, then the inclusion $(X\backslash B, A\backslash B) \hookrightarrow (X, A)$ induces an isomorphism $h_*(X\backslash B, A\backslash B) \to h_*(X, A)$.
- (DIMENSION) $h_n(\{\text{pt}\}) = 0$ for all $n \neq 0$. The potentially nontrivial abelian group
$$G := h_0(\{\text{pt}\})$$
is then called the **coefficient group** of h_*.
- (ADDITIVITY) For any collection of spaces $\{X_\alpha\}_{\alpha\in J}$ with inclusion maps $i^\alpha : X_\alpha \hookrightarrow \coprod_{\beta\in J} X_\beta$, the homomorphisms $i^\alpha_* : h_*(X_\alpha) \to h_*(\coprod_\beta X_\beta)$ determine an isomorphism
$$\bigoplus_{\alpha\in J} h_*(X_\alpha) \to h_*\Big(\coprod_{\alpha\in J} X_\alpha\Big).$$

Put together, these properties of an axiomatic homology theory h_* are known as the **Eilenberg-Steenrod axioms**, and they were first written down in [ES52] with the exception of the additivity axiom, which was added later by Milnor [Mil62].[41] We have already done most of the work of proving that for any given abelian group G, the singular homology $H_*(\cdot; G)$ defines an axiomatic homology theory with coefficient group G. The next two exercises fill the remaining gaps in proving this.

EXERCISE 26.1. Assume G is any abelian group and abbreviate the singular homology of a pair (X, A) with coefficients in G by $H_*(X, A) := H_*(X, A; G)$.

(a) Show that the connecting homomorphisms $\partial_* : H_n(X, A) \to H_{n-1}(A)$ in singular homology satisfy naturality, i.e. for any map $f : (X, A) \to (Y, B)$ and every $n \in \mathbb{Z}$, the diagram
$$\begin{tikzcd} H_n(X,A) \arrow[r, "\partial_*"] \arrow[d, "f_*"] & H_{n-1}(A) \arrow[d, "f_*"] \\ H_n(Y,B) \arrow[r, "\partial_*"] & H_{n-1}(B) \end{tikzcd}$$
commutes.

(b) Deduce that for any map $f : (X, A) \to (Y, B)$, the long exact sequences of (X, A) and (Y, B) in singular homology form the rows of a commutative diagram
$$\begin{tikzcd} \ldots \arrow[r] & H_n(A) \arrow[r] \arrow[d, "f_*"] & H_n(X) \arrow[r] \arrow[d, "f_*"] & H_n(X,A) \arrow[r] \arrow[d, "f_*"] & H_{n-1}(A) \arrow[r] \arrow[d, "f_*"] & \ldots \\ \ldots \arrow[r] & H_n(B) \arrow[r] & H_n(Y) \arrow[r] & H_n(Y,B) \arrow[r] & H_{n-1}(B) \arrow[r] & \ldots \end{tikzcd}$$

EXERCISE 26.2. Prove directly from the definition of singular homology $H_*(\cdot; G)$ with any coefficient group G that it satisfies the additivity axiom.

If you look again at our computation of $H_*(S^n; \mathbb{Z})$, you'll see that it mostly only used the axioms listed above—I say "mostly" because we did cheat slightly in using the isomorphism $H_1(S^n; \mathbb{Z}) \cong \pi_1(S^n)$, the proof of which is a fairly hands-on argument with singular simplices

[41]One can show that for *finite* disjoint unions, the additivity axiom follows from the others—it was thus unnecessary from the perspective of Eilenberg and Steenrod because they were mainly interested in compact spaces, in particular the polyhedra of finite simplicial complexes. The extra axiom becomes important however as soon as the discussion is extended to include noncompact spaces with infinitely many connected components.

and does not follow from the axioms. But actually, we could have gotten around this with a little more effort, and it is even possible to compute $H_1(S^n; G)$ for arbitrary coefficient groups G without knowing anything about the fundamental group. The reason we had to appeal to the fundamental group was that Theorem 25.2 is not true for $k = 0$, and it fails for a very specific reason: since H_0 of a contractible space does not vanish, the exact sequences do not always give isomorphisms when this term appears. But there is a formal trick to avoid this problem, called **reduced homology**: it is a variant $\widetilde{H}_*$ of the usual singular homology H_* that fits into all the same exact sequences, but is defined in a slightly more elaborate way so that $\widetilde{H}_n(\{\mathrm{pt}\}) = 0$ for all n, not just for $n \neq 0$. If we had used this, we could have done an inductive argument reducing the homology of every sphere S^n to the homology of S^0, which is the disjoint union of two one-point spaces, so the dimension and additivity axioms then provide the answer. This version of the argument eliminates any need for specifying the coefficients $G = \mathbb{Z}$, and it also works for any axiomatic homology theory, thus giving:

THEOREM. *For every $n \in \mathbb{N}$ and any theory h_* satisfying the Eilenberg-Steenrod axioms with coefficient group G,*

$$h_k(S^n) \cong \begin{cases} G & \text{for } k = 0, n, \\ 0 & \text{for all other } k. \end{cases}$$

Now a word of caution: in the last few lectures, we proved two things about singular homology that cannot be deduced merely from the formal properties codified in the Eilenberg-Steenrod axioms, and they are in fact *not true* for arbitrary axiomatic homology theories. One of these was Proposition 22.8, which related H_0 of an arbitrary space X to the set $\pi_0(X)$ of path-components of X via the formula

$$H_0(X; G) \cong \bigoplus_{\pi_0(X)} G. \tag{26.1}$$

This looks at first like it should be related to the additivity axiom: if X is homeomorphic to the disjoint union of its path-components $X_\alpha \subset X$, then additivity gives $H_0(X; G) \cong \bigoplus_\alpha H_0(X_\alpha; G)$, but there is unfortunately nothing in the axioms to imply $H_0(X_\alpha; G) \cong G$ for an arbitrary path-connected space X_α, unless X_α happens to be contractible. There is also a more serious problem, though you may have forgotten about it since we started focusing only on "nice" spaces after Lecture 7: not every space is homeomorphic to the disjoint union of its path-components. Manifolds have this property, and so do locally path-connected spaces in general—the latter follows from a combination of Exercise 7.12, Proposition 7.18 and Theorem 7.19. But not every space is locally path-connected, and no such assumption was imposed on X when we computed $H_0(X; G)$.

Another important result that does not follow from the axioms is Theorem 22.10, on the natural homomorphism

$$\pi_1(X) \to H_1(X; \mathbb{Z}) \tag{26.2}$$

for any path-connected space X, and the isomorphism it induces between $H_1(X; \mathbb{Z})$ and the abelianization of $\pi_1(X)$. Its proof (carried out in Exercise 22.12) similarly required a hands-on examination of the chain complex $C_*(X; \mathbb{Z})$ that underlies the definition of $H_*(X; \mathbb{Z})$. In this context, allow me to point out an odd detail that you may or may not have noticed about the Eilenberg-Steenrod axioms: they never mention any chain complex at all. Homology theories in the sense of Eilenberg-Steenrod need not generally come from chain complexes—in practice, most of them do, though often in less direct ways than singular homology, and one cannot derive from the axioms any direct intuition about the geometric meaning of elements in the groups $h_0(X)$ and $h_1(X)$. Part of the point of the axioms is that for most of the interesting applications of homology, it should suffice to know that a homology theory *exists* and satisfies the right formal properties, because if those properties hold, then one can typically carry out the applications one wants without even knowing how the theory itself is defined. This "highbrow" perspective does not suffice however for computations like (26.1) and (26.2), which are unique to singular homology and its underlying chain complex.

A sketch of Čech homology. Singular homology is not the only theory that satisfies the Eilenberg-Steenrod axioms, though it has been the standard one that people use for over half a century. While the alternatives have gone out of fashion, a few of them do still occasionally resurface in research articles. I would like to give a quick sketch of one of them, if only to demonstrate how two completely different ideas can sometimes lead to invariants that detect more-or-less the same information.

While singular homology tries to understand spaces by viewing singular n-simplices as basic building blocks of n-dimensional objects, the Čech homology theory studies them instead via the combinatorial properties of their open coverings. Suppose in particular that $\mathcal{O} := \{\mathcal{U}_\alpha \subset X\}_{\alpha \in J}$ is an open covering of a space X. One can associate to any such covering an abstract simplicial complex $K_{\mathcal{O}} = (V, S)$, called the **nerve** of the covering: its set of vertices V is the index set J, or equivalently the set of open sets that belong to the covering, and a subset $\sigma := \{\alpha_0, \ldots, \alpha_n\} \subset V$ is defined to be an n-simplex $\sigma \in S$ of the complex $K_{\mathcal{O}}$ if and only if

$$\mathcal{U}_{\alpha_0} \cap \ldots \cap \mathcal{U}_{\alpha_n} \neq \varnothing.$$

This easily satisfies the required conditions for a simplicial complex: each vertex $\alpha \in V$ defines a 0-simplex $\{\alpha\} \in S$ since $\mathcal{U}_\alpha \neq \varnothing$, and each face of $\sigma = \{\alpha_0, \ldots, \alpha_n\} \in S$ is also a simplex in the complex since every nontrivial subcollection of the sets $\mathcal{U}_{\alpha_0}, \ldots, \mathcal{U}_{\alpha_n}$ must still have nonempty intersection. As with all simplicial complexes, $K_{\mathcal{O}}$ gives rise to a topological space, its polyhedron $|K_{\mathcal{O}}|$, but that space need not look at all similar to X: for example, if X is something as simple as S^1, then even if the open covering $\{\mathcal{U}_\alpha\}_{\alpha \in J}$ is finite, the simplicial complex $K_{\mathcal{O}}$ may have arbitrarily large dimension, namely the largest number $n \geqslant 0$ such that $n+1$ of the sets in the covering have a nonempty intersection.

The example $X = S^1$ is quite instructive, however, if one compares what $K_{\mathcal{O}}$ looks like for a few simple choices of open coverings. Figure 14 shows three such choices, two of which give rise to 1-dimensional simplicial complexes, and in the third case, the simplicial complex is 2-dimensional. The polyhedra of these three simplicial complexes are all different spaces, none homeomorphic to any of the others, but you may notice that the last two have something in common: they are homotopy equivalent, and not just to each other, but also to the original space, $X = S^1$. The polyhedron in the first example is not homotopy equivalent to S^1, but the other two open coverings also happen to have a nice property that this one does not: in the other two, the intersection sets $\mathcal{U}_{\alpha_0} \cap \ldots \cap \mathcal{U}_{\alpha_n}$ are always contractible, whereas in the first covering, $\mathcal{U}_1 \cap \mathcal{U}_2$ is a disconnected set. Open coverings in which the sets $\mathcal{U}_{\alpha_0} \cap \ldots \cap \mathcal{U}_{\alpha_n}$ are always contractible have a special status: they are called *good* covers, and for sufficiently nice spaces such as smooth manifolds, one can show that every open covering has a refinement that is a good cover. Figure 14 hints at an intriguing general phenomenon: for sufficiently nice open coverings of sufficiently nice spaces X, the nerve of the cover can be viewed as a simplicial model for X itself, up to homotopy type. This suggests that the *simplicial* homology $H_*^\Delta(K_{\mathcal{O}}; G)$ of the nerve should encode interesting topological information about X, and that is how Čech homology is defined: for sufficiently nice open coverings $\mathcal{O}$ of X, the **Čech homology** of X with coefficient group G is

$$\check{H}_*(X; G) := H_*^\Delta(K_{\mathcal{O}}; G).$$

I am being deliberately vague now, because making this definition more precise would require a discussion of inverse limits and chain homotopy equivalences which we do not have time for right

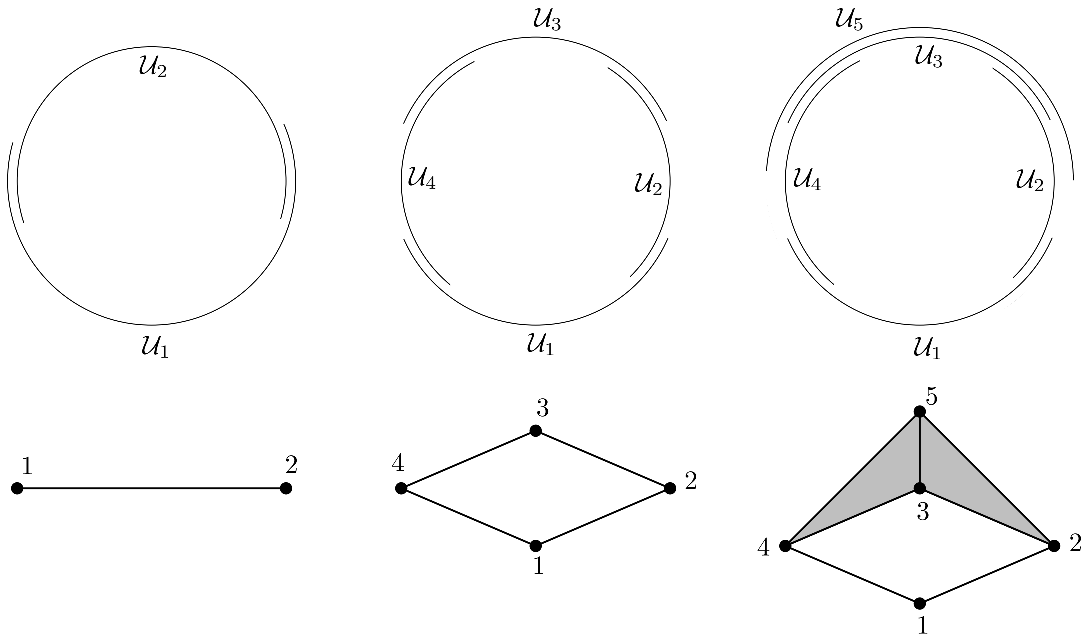

FIGURE 14. Three examples of open coverings of S^1 and their nerves, with vertices labeled $k \in \{1, 2, 3, 4, 5\}$ in correspondence with the open sets $\mathcal{U}_k \subset S^1$. The rightmost example includes two 2-simplices in addition to vertices and 1-simplices.

now: in particular, some serious work would be required in order to show that $H_*^\Delta(K_\mathcal{O}; G)$ up to isomorphism is independent of the choice of (sufficiently nice!) open covering $\mathcal{O}$. The examples on the circle in Figure 14 are intended to convince you that this idea might not be completely outlandish.

Since the definitions of $H_*(X; G)$ and $\check{H}_*(X; G)$ seem very different, it is somewhat remarkable that for a wide class of spaces that includes all compact manifolds, they are isomorphic. One way to explain this is by ignoring the definitions of these two invariants and concentrating instead on their formal properties: after extending Čech homology to an invariant of pairs (X, A) rather than just individual spaces X, one can show (under one or two extra assumptions) that it satisfies the Eilenberg-Steenrod axioms, just like singular homology. As a consequence, any computation that relies only on the formal properties of homology theories—homotopy invariance, excision, long exact sequences and so forth—applies equally well to $H_*(X; G)$ and $\check{H}_*(X; G)$.

It is not true that $H_*(X; G)$ and $\check{H}_*(X; G)$ are always isomorphic, but one has to consider fairly ugly spaces in order to see the difference. A hint of where to look comes from our computation $H_0(X; G) \cong \bigoplus_{\pi_0(X)} G$: as mentioned above, this result does not follow from the axioms. As it turns out, $\check{H}_0(X; G)$ does not care whether the space X is *path*-connected, but cares instead whether it is connected:

EXERCISE 26.3. Show that if X is a connected space, then for any open cover $\mathcal{O}$ of X, the polyhedron $|K_\mathcal{O}|$ of its nerve is path-connected.

Way back in Lecture 7, we saw examples of spaces that are connected but not path-connected. One can deduce from Exercise 26.3 that whenever X is such a space, $\check{H}_0(X; G) \cong G$, but according to (26.1), $H_0(X;G)$ is larger. Using suspensions, one can also derive from this examples of path-connected spaces X for which $\check{H}_1(X;\mathbb{Z})$ is not isomorphic to the abelianization of $\pi_1(X)$. But again: spaces like this are ugly, they are not the kinds of spaces that arise naturally in most applications.

REMARK 26.4. In the discussion above, I have swept an uncomfortable fact about $\check{H}_*(X;G)$ under the rug: most versions of Čech homology satisfy *most* of the Eilenberg-Steenrod axioms, but not quite all of them. For technical reasons having to do with the formal properties of inverse limits in homological algebra, $\check{H}_*(X;G)$ does not generally satisfy the exactness axiom unless one restricts to *compact* pairs (X,A) and a restrictive class of coefficient groups G, e.g. any *finite* abelian group or finite-dimensional vector space over a field will do. This shortcoming is one reason why Čech homology has not been used very much in the past half-century. On the other hand, another major topic for next semester's course will be *cohomology*, which is a kind of dualization of homology that has its own closely related set of axioms. The most popular cohomology theory is singular cohomology, but there is also a Čech cohomology theory, which has strictly better formal properties than its undualized counterpart, i.e. it satisfies all of the conditions required for an axiomatic cohomology theory, and even has one or two desirable properties that singular cohomology does not. The ability of Čech cohomology to relate local and global properties of spaces via the combinatorics of their open coverings makes it an essential and frequently used tool in certain branches of mathematics, especially in algebraic geometry.

Cell complexes. We've seen that all axiomatic homology theories are isomorphic on the spaces S^n, though they need not be isomorphic in peculiar examples such as connected spaces that are not path-connected. It is natural to wonder: how large is the class of spaces X for which the Eilenberg-Steenrod axioms completely determine their homologies $h_*(X)$? The spaces with this property happen to be the spaces for which most of the more advanced techniques of algebraic topology have something interesting to say, so they play a starring role in the subject from this point forward.

A plausible first guess for the class of spaces we want to consider would be *polyhedra*: the topological spaces associated to abstract simplicial complexes. But there is a larger class of spaces called, *cell complexes* (or the fancier term "CW-complexes"), which are actually easier to work with and much more general. It is known that all smooth manifolds or simplicial complexes are also cell complexes, and all topological manifolds are at least homotopy equivalent to cell complexes. We saw one concrete example in Lecture 14: when we proved that every finitely presented group occurs as the fundamental group of some compact Hausdorff space (Theorem 14.20), the space we constructed was a wedge of circles with a finite set of disks attached. The general idea of a cell complex is to build up a space inductively as a nested sequence of "skeleta" of various dimensions, where the n-skeleton is always constructed by attaching n-disks to the $(n-1)$-skeleton. In this language, the space constructed in the proof of Theorem 14.20 was a 2-dimensional cell complex, because it had a 1-skeleton (the wedge of circles) and a 2-skeleton (the attached disks). Here is the general definition in the case where there are only finitely many cells.

DEFINITION 26.5. A space X is called a (finite) **cell complex** (*Zellenkomplex*) of dimension n if it contains a nested sequence of subspaces $X^0 \subset X^1 \subset \ldots \subset X^{n-1} \subset X^n = X$ such that:

(1) X^0 is a finite discrete set;
(2) For each $m = 1, \ldots, n$, X^m is homeomorphic to a space constructed by attaching finitely many m-disks $\mathbb{D}^m$ to X^{m-1} along maps $\partial\mathbb{D}^m \to X^{m-1}$.

In general, the collection of m-disks attached to X^{m-1} at each step need not be nonempty; if it is empty, then $X^m = X^{m-1}$, but we implicitly assume $X^n \neq X^{n-1}$ when we call X "n-dimensional".

We call $X^m \subset X$ the **m-skeleton** of X. The definition implies that for each $m = 1, \dots, n$, there is a finite set $\mathcal{K}_m(X)$ and a so-called **attaching map** $\varphi_\alpha : S^{m-1} \to X^{m-1}$ associated to each $\alpha \in \mathcal{K}_m(X)$ such that

$$X^m \cong \left(\coprod_{\alpha \in \mathcal{K}_m(X)} \mathbb{D}^m \right) \cup_{\varphi_m} X^{m-1},$$

where $\varphi_m : \coprod_{\alpha \in \mathcal{K}_m(X)} \partial \mathbb{D}^m \to X^{m-1}$ denotes the disjoint union of the maps $\varphi_\alpha : S^{m-1} \to X^{m-1}$, each defined on the boundary of the disk indexed by α. As a set, X^m is the union of X^{m-1} with a disjoint union of open disks

$$e_\alpha^m \cong \mathring{\mathbb{D}}^m \quad \text{for each} \quad \alpha \in \mathcal{K}_m(X),$$

called the **m-cells** of the complex. For $m = 0$, we call the discrete points of the 0-skeleton X^0 the **0-cells** and denote this set by $\mathcal{K}_0(X)$.

Since $\Delta^n \cong \mathbb{D}^n$, it is easy to see that polyhedra are also cell complexes: the n-cells are the interiors of the n-simplices, while the n-skeleton is the union of all simplices of dimension at most n and the attaching maps $S^{n-1} \cong \partial \Delta^n \to X^{n-1}$ are each homeomorphisms onto their images. In general, the attaching maps in a cell complex do not need to be injective, they only must be continuous, so while the m-cells e_α^m look like open m-disks, their closures in X might not be homeomorphic to closed disks. For instance, here is an example with an n-cell whose boundary is collapsed to a point, so its closure is not a disk, but a sphere:

EXAMPLE 26.6. Consider a cell complex that has one 0-cell and no cells of dimensions $1, \dots, n-1$, so its m-skeleton for every $m < n$ is a one-point space, but there is one n-cell e_α^n attached via the unique map $\varphi_\alpha : S^{n-1} \to \{\mathrm{pt}\}$. The resulting space $X = X^n$ is homeomorphic to S^n.

The **cellular homology** of a cell complex $X = \bigcup_{n \geqslant 0} X^n$ is now defined as follows. Given an abelian coefficient group G, let

$$C_n^{\mathrm{CW}}(X; G) := \bigoplus_{\alpha \in \mathcal{K}_n(X)} G = \left\{ \text{finite sums } \sum_i c_i e_{\alpha_i}^n \;\middle|\; c_i \in G,\ \alpha_i \in \mathcal{K}_n(X) \right\}$$

denote the abelian group of finite linear combinations of generators e_α^n corresponding to the n-cells in the complex, with coefficients in G. A boundary map $\partial : C_n^{\mathrm{CW}}(X; G) \to C_{n-1}^{\mathrm{CW}}(X; G)$ is determined by the formula

$$\partial e_\alpha^n = \sum_{\beta \in \mathcal{K}_{n-1}(X)} [e_\beta^{n-1} : e_\alpha^n] e_\beta^{n-1},$$

where the **incidence numbers** $[e_\beta^{n-1} : e_\alpha^n] \in \mathbb{Z}$ are determined as follows. For each $\alpha \in \mathcal{K}_n(X)$ and $\beta \in \mathcal{K}_{n-1}(X)$, let

$$X_\beta := X^{n-1} / (X^{n-1} \backslash e_\beta^{n-1}),$$

i.e. it is a space obtained by collapsing everything in the $(n-1)$-skeleton except for the individual cell e_β^{n-1} to a point. Since e_β^{n-1} is an open $(n-1)$-disk with a canonical homeomorphism to $\mathring{\mathbb{D}}^{n-1}$, there is a canonical homeomorphism

$$X_\beta = \mathbb{D}^{n-1} / \partial \mathbb{D}^{n-1} \cong S^{n-1}.$$

There is also a quotient projection $q : X^{n-1} \to X_\beta$, so composing this with the attaching map $\varphi_\alpha : S^{n-1} \to X^{n-1}$ gives a map between two $(n-1)$-dimensional spheres

$$q \circ \varphi_\alpha : S^{n-1} \to X_\beta \cong S^{n-1}.$$

This induces a homomorphism

$$\mathbb{Z} \cong H_{n-1}(S^{n-1};\mathbb{Z}) \xrightarrow{(q\circ\varphi_\alpha)_*} H_{n-1}(X_\beta;\mathbb{Z}) \cong \mathbb{Z},$$

and all homomorphisms $\mathbb{Z} \to \mathbb{Z}$ are of the form $x \mapsto dx$ for some $d \in \mathbb{Z}$. The integer d appearing here is called the **degree** of $q \circ \varphi_\alpha$, and that is how we define the incidence number:

$$[e^{n-1}_\beta : e^n_\alpha] := \deg(q \circ \varphi_\alpha).$$

Strictly speaking, this definition only makes sense for $n \geqslant 2$ since our computation of the homology of spheres does not apply to S^0, but this is a minor headache that can easily be fixed with an extra definition, as in simplicial homology.

It would take a lot more time than we have right now to explain why this definition of ∂ is the right one, and why it implies $\partial^2 = 0$ in particular. But if you are willing to accept that for now, then we can define the **cellular homology** (*zelluläre Homologie*) groups

$$H^{\mathrm{CW}}_n(X;G) := H_n\left(C^{\mathrm{CW}}_*(X;G),\partial\right),$$

and we can almost immediately carry out a surprisingly easy computation:

Example 26.7. The cell decomposition of S^n in Example 26.6 gives

$$H^{\mathrm{CW}}_k(S^n;G) \cong \begin{cases} G & \text{for } k = 0, n, \\ 0 & \text{for all other } k. \end{cases}$$

Indeed, for $n \geqslant 2$ we can see this without doing any work, because $C^{\mathrm{CW}}_0(S^n;G) \cong C^{\mathrm{CW}}_n(S^n;G) \cong G$ are the only nontrivial chain groups, so ∂ simply vanishes and the homology groups are the chain groups. For $n = 1$ you need a little bit more information that I haven't given you, but one can show also in this case that $\partial = 0$, so the result is the same.

In reality, cellular homology is not a *new* homology theory as such, it is just an extremely efficient way of computing any axiomatic homology theory for spaces that are nice enough to have cell decompositions. The following result has been the main tool used for computations of singular homology for most of its history, and it implies in particular the fact that *simplicial* homology is a topological invariant (cf. Theorem 21.16). We will work through a complete proof next semester, and the first step in that proof will be the computation of $h_*(S^n)$.

Theorem. *For any cell complex X and any axiomatic homology theory h_* with coefficient group G, $H^{\mathrm{CW}}_*(X;G) \cong h_*(X)$.*

This theorem is the real reason why homology is considered one of the "easier" invariants to work with in algebraic topology: for most of the spaces that arise in practice, and all compact manifolds in particular, $H_*(X)$ can be computed after replacing the unmageably large singular chain complex with the cellular chain complex, which is *finitely* generated. Having only finitely many generators means that in principle, one can always just feed all the information from the chain complex into a computer program, then press a button and get an answer.

The Euler characteristic. Here is a remarkable application of cellular homology. To make our lives algebraically a bit easier, let's choose the coefficient group G to be a field $\mathbb{K}$, e.g. $\mathbb{Q}$ or $\mathbb{R}$ will do. This has the advantage of making our chain complexes naturally into vector spaces over $\mathbb{K}$, and the boundary maps are $\mathbb{K}$-linear, so the homology groups are also $\mathbb{K}$-vector spaces. Whenever $H_*(X;\mathbb{K})$ is finite dimensional, we then define the **Euler characteristic** of X as the integer

$$\chi(X) := \sum_{n=0}^{\infty} (-1)^n \dim_{\mathbb{K}} H_n(X;\mathbb{K}) \in \mathbb{Z}.$$

Although each individual term $\dim_{\mathbb{K}} H_n(X;\mathbb{K})$ may in general depend on the choice of field $\mathbb{K}$, one can show that their alternating sum does not.[42] This fact admits a purely algebraic proof, but if X is a finite cell complex, then it also follows from the following much more surprising observation. It is not difficult to prove that whenever (C_*, ∂) is a finite-dimensional chain complex of $\mathbb{K}$-vector spaces, the alternating sum of the dimensions of its homology groups can be computed without computing the homology at all: in fact,

$$\sum_{n\in\mathbb{Z}} (-1)^n \dim_{\mathbb{K}} H_n(C_*, \partial) = \sum_{n\in\mathbb{Z}} (-1)^n \dim_{\mathbb{K}} C_n. \tag{26.3}$$

This follows essentially from the fact that for each $n \in \mathbb{Z}$, writing $Z_n := \ker \partial_n \subset C_n$ and $B_n := \operatorname{im} \partial_{n+1} \subset C_n$, the map $\partial_n : C_n \to C_{n-1}$ descends to an isomorphism $C_n/Z_n \to B_{n-1}$, implying

$$\dim_{\mathbb{K}} C_n - \dim_{\mathbb{K}} Z_n = \dim_{\mathbb{K}} B_{n-1}.$$

Since $H_n(C_*, \partial) = Z_n/B_n$, we also have $\dim_{\mathbb{K}} H_n(C_*, \partial) = \dim_{\mathbb{K}} Z_n - \dim_{\mathbb{K}} B_n$, so combining these two relations and adding things up with alternating signs produces lots of cancelations leading to (26.3). Now apply this to the cellular chain complex, in which each $C_n^{\mathrm{CW}}(X;\mathbb{K})$ is a $\mathbb{K}$-vector space whose dimension is the number of n-cells in the complex. What we learn is that we don't need to know anything about homology in order to compute $\chi(X)$—all we have to do is count cells and add up the counts with signs. The isomorphism $H_*(X;\mathbb{K}) \cong H_*^{\mathrm{CW}}(X;\mathbb{K})$ now implies that the result of this counting game only depends on the space, and not on our choice of how to decompose it into cells:

Theorem. *For any finite cell complex X,*

$$\chi(X) = \sum_{n=0}^{\infty} (-1)^n \, (\textit{the number of } n\textit{-cells}) .$$

In particular this applies to simplicial complexes, e.g. if you build a 2-sphere by gluing together triangles along common edges, then no matter how you do it or how many triangles are involved, the number of triangles minus the number of glued edges plus the number of glued vertices will always be

$$\chi(S^2) = \dim_{\mathbb{R}} H_0(S^2;\mathbb{R}) - \dim_{\mathbb{R}} H_1(S^2;\mathbb{R}) + \dim_{\mathbb{R}} H_2(S^2;\mathbb{R}) = 1 - 0 + 1 = 2.$$

It is not much harder to work out the result for Σ_g with any $g \geqslant 0$: the answer is

$$\chi(\Sigma_g) = 2 - 2g,$$

and off the top of my head, I can think of two completely different ways to prove this by decomposing Σ_g into cells and counting them with signs: regardless of the choices in the decomposition, the answer will always be the same. Go ahead. Try it.

[42] One can also define $\chi(X)$ using integer coefficients in terms of the *ranks* of the abelian groups $H_n(X;\mathbb{Z})$. This is one of the algebraic details I wanted to avoid by using field coefficients.

Second semester (Topologie II)

27. Categories and functors (October 17, 2023)

Basics of category theory. The general approach of algebraic topology is to associate to each topological space some algebraic object that can be used to tell "different" spaces apart. One important example we saw last semester was the fundamental group (π_1), which assigns to every pair (X, p) consisting of a topological space X with a choice of "base" point $p \in X$ a group $\pi_1(X, p)$. Another—which we shall reintroduce at the end of this lecture and subsequently have a lot more to say about—is singular homology (H_*), which assigns to each space X a whole sequence of abelian groups $H_n(X)$ indexed by the nonnegative integers $n \geqslant 0$. It is reasonable to think of these in some sense as "functions" with domains consisting of the collection of all topological spaces (possibly with extra data such as a base point), and targets consisting of the collection of all groups (or in the case of homology, all abelian groups). We have not yet developed the right language to make this notion of a "function" precise, so it is time to do so now.

One reason why π_1 and H_* cannot actually be called "functions" is that their domains, strictly speaking, are not sets (*Mengen*). I encourage you to skip the rest of this paragraph if you are not interested in the finer points of axiomatic set theory or the classic set-theoretic paradoxes... but for those who are still reading, let us agree that there is no such thing as the "set of all topological spaces". Indeed, every set can be made into a topological space by assigning it the discrete topology, so if one can talk about the set of all topological spaces, then one must also be able to talk about the set of all sets, and it is a short step from there to the "set of all sets that do not contain themselves"—at which point we may find ourselves asking whether that particular set contains itself, and promptly jumping off the nearest bridge. The architects of abstract set theory solved this dilema by coming up with a set of axioms that tell you how to construct new sets from old ones, together with a very short list of sets (e.g. the empty set) whose existence clearly needs to be assumed, and insisting that *only* collections of objects that arise from these axioms should be called *sets*. Of course, we do sometimes also need to discuss collections of objects that do not arise from the axioms of set theory, and the collection of all topological spaces is an example. Such collections are generally called (proper) **classes** (*Klassen*), but since I do not wish to go any further into the subtleties of set theory in this course, I shall continue to refer to them via the informal word **collections**. You should just keep in mind that while such things can be defined, they are not considered equivalent to sets and cannot be used for all the same purposes that sets can—in particular, an arbitrary "collection" cannot serve as the domain of a function according to the traditional definitions.

Leaving set theory aside, it also must be observed that π_1 and H_* are not just arbitrary "functions" that associate algebraic objects to topological spaces, but they do so in ways that make the algebraic objects into *topological invariants*. In both cases, this results mainly from the fact that continuous maps of spaces induce homomorphisms between the corresponding fundamental groups or homology groups, implying in particular that homeomorphisms induce group isomorphisms. The notion of a *functor* is meant as a form of abstract packaging for this idea.

DEFINITION 27.1. A **category** (*Kategorie*) $\mathscr{C}$ consists of the following data:

- A collection (i.e. class) $\mathrm{Ob}_{\mathscr{C}}$, whose elements are called the **objects** (*Objekte*) of $\mathscr{C}$;
- For each $X, Y \in \mathrm{Ob}_{\mathscr{C}}$ a set $\mathrm{Mor}(X,Y)$, whose elements are called the **morphisms from** X **to** Y (*Morphismen von* X *nach* Y), such that for each $X \in \mathrm{Ob}_{\mathscr{C}}$ there is a distinguished[43] element $\mathrm{Id}_X \in \mathrm{Mor}(X,X)$;
- For each $X, Y, Z \in \mathrm{Ob}_{\mathscr{C}}$, a function

$$\mathrm{Mor}(X,Y) \times \mathrm{Mor}(Y,Z) \to \mathrm{Mor}(X,Z) : (f,g) \mapsto g \circ f \tag{27.1}$$

such that $(f \circ g) \circ h = f \circ (g \circ h)$, and whenever two of the objects match and Id denotes the corresponding distinguished morphism, $f \circ \mathrm{Id} = f = \mathrm{Id} \circ f$.

EXAMPLE 27.2. The category Top has $\mathrm{Ob}_{\mathsf{Top}} = \{\text{topological spaces}\}$ and $\mathrm{Mor}(X,Y) = \{f : X \to Y \mid f \text{ a continuous map}\}$, with Id_X defined for each space X as the identity map and the function (27.1) defined as the usual composition of maps. This defines a category since the identity map is always continuous and the composition of two continuous maps is also continuous.

EXAMPLE 27.3. The category Set has $\mathrm{Ob}_{\mathsf{Set}} = \{\text{sets}\}$ and $\mathrm{Mor}(X,Y) = \{f : X \to Y\}$, with no requirement on continuity of maps since there is no topology.

EXAMPLE 27.4. The objects of Diff are the smooth finite dimensional manifolds, and its morphisms are smooth maps. (As in Example 27.2, the identity is always smooth and the composition of two smooth maps is smooth.)

EXAMPLE 27.5. The category Grp has $\mathrm{Ob}_{\mathsf{Grp}} = \{\text{groups}\}$ and $\mathrm{Mor}(G,H) = \{\text{homomorphisms } G \to H\}$ for $G, H \in \mathrm{Ob}_{\mathsf{Grp}}$.

EXAMPLE 27.6. There is a **subcategory** (*Unterkategorie*) Ab of Grp whose objects consist of all *abelian* groups, with morphisms defined the same way as in Grp.

The examples above might give you the impression that in every category, a morphism is just a map that may be required to satisfy some specific properties. But nothing in Definition 27.1 says either that an object must be a kind of set or that a morphism is a map. Here is an example in which the objects are still sets, but the morphisms are *equivalence classes* of maps.

EXAMPLE 27.7. Let Top^h denote the category whose objects are the same as in Top, but with $\mathrm{Mor}(X,Y)$ defined as the set of homotopy classes of continuous maps $X \to Y$ and $\mathrm{Id}_X \in \mathrm{Mor}(X,X)$ as the homotopy class of the identity map. The function (27.1) is defined in terms of the usual composition of continuous maps $f : X \to Y$ and $g : Y \to Z$ by

$$[g] \circ [f] := [g \circ f].$$

(Exercise: check that this is well defined!)

For some interesting examples in which objects are not sets and the function (27.1) has nothing to do with composition of maps, see Exercises 27.16 and 27.17.

DEFINITION 27.8. In any category, a morphism $f \in \mathrm{Mor}(X,Y)$ is called an **isomorphism** (*Isomorphismus*) if there exists a morphism $f^{-1} \in \mathrm{Mor}(Y,X)$ such that $f^{-1} \circ f = \mathrm{Id}_X$ and $f \circ f^{-1} = \mathrm{Id}_Y$. If an isomorphism exists in $\mathrm{Mor}(X,Y)$, we say that the objects X and Y are **isomorphic** (*isomorph*).

[43] The word "distinguished" appears here because part of the structure of the category $\mathscr{C}$ is the knowledge of which morphism should be called "Id_X" for each object X. If we simply required the existence of a morphism that satisfies the conditions stated in the third bullet point, then there might be more than one such element and we would not know which one to call Id_X. But the structure of $\mathscr{C}$ requires each set $\mathrm{Mor}(X,X)$ to contain a specific element that carries that name; there might in theory exist additional morphisms that have the same properties, but only one is called Id_X.

According to this definition, the word "isomorphism" no longer has a strictly algebraic meaning, but will mean whatever is considered to be the notion of "equivalence" in whichever category we are working with. Let's run through the list: an isomorphism in Top is a homeomorphism, in Set it is simply a bijection, in Diff a diffeomorphism, and in Grp or Ab it is the usual notion of group isomorphism. The most interesting case so far is Top^h: two objects in Top^h are isomorphic if and only if they are homotopy equivalent!

Definition 27.9. Given two categories $\mathscr{C}$ and $\mathscr{D}$, a **functor** (*Funktor*) $\mathcal{F} : \mathscr{C} \to \mathscr{D}$ assigns to each $X \in \mathrm{Ob}_{\mathscr{C}}$ an object $\mathcal{F}(X) \in \mathrm{Ob}_{\mathscr{D}}$ and to each $f \in \mathrm{Mor}(X, Y)$ for $X, Y \in \mathrm{Ob}_{\mathscr{C}}$ a morphism $\mathcal{F}(f) \in \mathrm{Mor}(\mathcal{F}(X), \mathcal{F}(Y))$ such that:

(1) $\mathcal{F}(\mathrm{Id}_X) = \mathrm{Id}_{\mathcal{F}(X)}$ for all $X \in \mathrm{Ob}_{\mathscr{C}}$;
(2) $\mathcal{F}(f \circ g) = \mathcal{F}(f) \circ \mathcal{F}(g)$ for all $g \in \mathrm{Mor}(X, Y)$ and $f \in \mathrm{Mor}(Y, Z)$, $X, Y, Z \in \mathrm{Ob}_{\mathscr{C}}$.

Example 27.10. Denote by Top_* the category whose objects are the pointed spaces (X, p), i.e. a topological space X together with a point $p \in X$, and

$$\mathrm{Mor}((X, p), (Y, q)) := \left\{ f : X \to Y \;\middle|\; f \text{ continuous and } f(p) = q \right\}.$$

Then the fundamental group defines a functor $\pi_1 : \mathsf{Top}_* \to \mathsf{Grp}$; indeed, it associates to each pointed space (X, p) the group $\pi_1(X, p)$ and to each pointed map $f : (X, p) \to (Y, q)$ the group homomorphism

$$\pi_1(f) := f_* : \pi_1(X, p) \to \pi_1(Y, q)$$

such that Id_* is the identity homomorphism and $(f \circ g)_* = f_* \circ g_*$.

Example 27.11. The fundamental group also defines a functor $\pi_1 : \mathsf{Top}_*^h \to \mathsf{Grp}$ where Top_*^h is defined to have the same objects as Top_*, but with $\mathrm{Mor}((X, p), (Y, q))$ defined as the set of pointed homotopy classes of maps $(X, p) \to (Y, q)$. (See Theorem 8.11 in Lecture 8 from last semester.)

Example 27.12. As we will review within the next few lectures, the singular homology group $H_n(X; G)$ for each integer $n \geqslant 0$ and any fixed (abelian) coefficient group G defines functors

$$H_n(\cdot; G) : \mathsf{Top} \to \mathsf{Ab} \quad \text{and} \quad \mathsf{Top}^h \to \mathsf{Ab}.$$

The latter makes sense due to the fact that the homomorphism $f_* : H_n(X; G) \to H_n(Y; G)$ induced by a continuous map $f : X \to Y$ depends only on the homotopy class of f.

We will later encounter several algebraic constructions and related topological invariants that satisfy most of the conditions of a functor but differ in one crucial respect: the morphisms they induce go *the other way*. In practice, this phenomenon often arises from the algebraic notion of dualization, and we'll give an example of this kind immediately after the definition.

Definition 27.13. Given two categories $\mathscr{C}$ and $\mathscr{D}$, a **contravariant functor** (*kontravarianter Funktor*) $\mathcal{F} : \mathscr{C} \to \mathscr{D}$ assigns to each $X \in \mathrm{Ob}_{\mathscr{C}}$ some $\mathcal{F}(X) \in \mathrm{Ob}_{\mathscr{D}}$ and to each $f \in \mathrm{Mor}(X, Y)$ for $X, Y \in \mathrm{Ob}_{\mathscr{C}}$ a morphism $\mathcal{F}(f) \in \mathrm{Mor}(\mathcal{F}(Y), \mathcal{F}(X))$ such that

(1) $\mathcal{F}(\mathrm{Id}_X) = \mathrm{Id}_{\mathcal{F}(X)}$ for all $X \in \mathrm{Ob}_{\mathscr{C}}$;
(2) $\mathcal{F}(f \circ g) = \mathcal{F}(g) \circ \mathcal{F}(f)$ for all $g \in \mathrm{Mor}(X, Y)$ and $f \in \mathrm{Mor}(Y, Z)$, $X, Y, Z \in \mathrm{Ob}_{\mathscr{C}}$.

A functor that satisfies the original Definition 27.9 instead of Definition 27.13 can be called **covariant** (*kovariant*) when we want to emphasize that it is not contravariant.

Example 27.14. Let $\mathsf{Vec}_{\mathbb{K}}$ denote the category of vector spaces over a fixed field $\mathbb{K}$, so $\mathrm{Mor}(V, W) := \mathrm{Hom}_{\mathbb{K}}(V, W)$ is the space of $\mathbb{K}$-linear maps $V \to W$. There is a contravariant functor $\boldsymbol{\Delta} : \mathsf{Vec}_{\mathbb{K}} \to \mathsf{Vec}_{\mathbb{K}}$ which sends each vector space V to its dual space $\boldsymbol{\Delta}(V) := V^* := \mathrm{Hom}_{\mathbb{K}}(V, \mathbb{K})$ and sends each morphism $A : V \to W$ to its transpose $\boldsymbol{\Delta}(A) := A^* : W^* \to V^*$, defined by $A^*(\lambda)v = \lambda(Av)$ for $\lambda \in W^*$ and $v \in V$. It satisfies the conditions of a functor since $(AB)^* = B^* A^*$ and the transpose of the identity $V \to V$ is the identity $V^* \to V^*$.

EXERCISE 27.15. One can speak of "functors of multiple variables" in much the same way as with functions. Show for instance that on the category Ab of abelian groups and homomorphisms,

$$\mathrm{Hom} : \mathsf{Ab} \times \mathsf{Ab} \to \mathsf{Ab}$$

defines a functor that is contravariant in the first variable and covariant in the second, assigning to each pair of abelian groups (G, H) the group $\mathrm{Hom}(G, H)$ of homomorphisms $G \to H$.

EXERCISE 27.16. Suppose $\mathscr{A}$ is a category whose objects form a set X, such that for each pair $x, y \in X$, the set of morphisms $\mathrm{Mor}(x, y)$ contains either exactly one element or none. We can turn this into a binary relation by writing $x \bowtie y$ for every pair such that $\mathrm{Mor}(x, y) \neq \emptyset$.

(a) What properties does the relation $\bowtie$ need to have in order for it to define a category in the way indicated above?
(b) If $\mathscr{B}$ is another category whose objects form a set Y with morphisms determined by a binary relation $\bowtie$ as indicated above, what properties does a map $f : X \to Y$ need to have in order for it to define a functor from $\mathscr{A}$ to $\mathscr{B}$?

EXERCISE 27.17. In any category $\mathscr{C}$, each object X has an **automorphism group** (also called **isotropy** group) $\mathrm{Aut}(X)$, consisting of all the isomorphisms in $\mathrm{Mor}(X, X)$. A **groupoid** is a category in which all morphisms are also isomorphisms.

(a) Show that if $\mathscr{G}$ is a groupoid and Grp denotes the usual category of groups with homomorphisms, there exists a contravariant functor from $\mathscr{G}$ to Grp that assigns to each object X of $\mathscr{G}$ its automorphism group $\mathrm{Aut}(X)$. How does this functor act on morphisms $X \to Y$? Could you alternatively define it as a *covariant* functor? Conclude either way that whenever X and Y are isomorphic objects in $\mathscr{G}$ (meaning there exists an isomorphism in $\mathrm{Mor}(X, Y)$), the groups $\mathrm{Aut}(X)$ and $\mathrm{Aut}(Y)$ are isomorphic.
(b) Given a topological space X and two points x, y, let $\mathrm{Mor}(x, y)$ denote the set of homotopy classes (with fixed end points) of paths $[0, 1] \to X$ from x to y, and define a composition function $\mathrm{Mor}(x, y) \times \mathrm{Mor}(y, z) \to \mathrm{Mor}(x, z) : (\alpha, \beta) \mapsto \alpha \cdot \beta$ by the usual notion of concatenation of paths. Show that this notion of morphisms defines a groupoid whose objects are the points in X.[44] In this case, what are the automorphism groups $\mathrm{Aut}(x)$ and the isomorphisms $\mathrm{Aut}(y) \to \mathrm{Aut}(x)$ given by the functor in part (a)?

We have one more piece of abstract language to add to this story before we can get back to studying topology. You've often seen the words "natural" or "naturally" appearing in statements of theorems in order to emphasize that something does not depend on any arbitrary choices. In category theory, these words can be given a precise definition.

DEFINITION 27.18. Given two covariant functors $\mathcal{F}, \mathcal{G} : \mathscr{C} \to \mathscr{D}$, a **natural transformation** (*natürliche Transformation*) T from $\mathcal{F}$ to $\mathcal{G}$ associates to each $X \in \mathrm{Ob}_{\mathscr{C}}$ a morphism $T_X \in \mathrm{Mor}(\mathcal{F}(X), \mathcal{G}(X))$ such that for all $X, Y \in \mathrm{Ob}_{\mathscr{C}}$ and $f \in \mathrm{Mor}(X, Y)$, the following diagram commutes:

$$\begin{tikzcd} \mathcal{F}(X) \arrow[r, "T_X"] \arrow[d, "\mathcal{F}(f)"] & \mathcal{G}(X) \arrow[d, "\mathcal{G}(f)"] \\ \mathcal{F}(Y) \arrow[r, "T_Y"] & \mathcal{G}(Y) \end{tikzcd}$$

A natural transformation of contravariant functors can be defined analogously.

A nice topological example of a natural transformation arises from the Hurewicz homomorphism, which sends the fundamental group of a space to its first singular homology group; see Exercise 28.9 in the next lecture. Here is an algebraic example.

[44]It is called the **fundamental groupoid** of X.

EXERCISE 27.19. Consider again the category $\mathsf{Vec}_{\mathbb{K}}$ of vector spaces over a fixed field $\mathbb{K}$ as in Example 27.14.

(a) Show that there is a covariant functor $\mathbf{\Delta}^2$ from $\mathsf{Vec}_{\mathbb{K}}$ to itself, assigning to each $V \in \mathsf{Vec}_{\mathbb{K}}$ the dual of its dual space $(V^*)^*$. Describe how this functor acts on morphisms.
(b) Let Id denote the identity functor on $\mathsf{Vec}_{\mathbb{K}}$, which sends each object and morphism to itself. Construct a natural transformation from Id to $\mathbf{\Delta}^2$ that assigns to every $V \in \mathsf{Vec}_{\mathbb{K}}$ a vector space isomorphism $V \to (V^*)^*$.

REMARK 27.20. Whenever a vector space V is finite dimensional, the map $V \to (V^*)^*$ given by the natural transformation in Exercise 27.19(b) is an isomorphism, and a large part of the reason why it turns out to define a *natural* transformation is that the definition of this map does not depend on any choices. By contrast, every finite-dimensional vector space is isomorphic to its dual space V^*, but there is no *canonical* way to define such isomorphisms for all vector spaces at once. Notice that since $\mathsf{Id} : \mathsf{Vec}_{\mathbb{K}} \to \mathsf{Vec}_{\mathbb{K}}$ is a covariant functor while the dualization functor $\mathbf{\Delta} : \mathsf{Vec}_{\mathbb{K}} \to \mathsf{Vec}_{\mathbb{K}}$ from Example 27.14 is contravariant, there is no sensible notion of natural transformations from Id to $\mathbf{\Delta}$.

EXERCISE 27.21. The **conjugate** $\bar{V}$ of a complex vector space V is defined as the same set $\bar{V} := V$ with the same notion of vector addition but with multiplication by scalars $\lambda = a + ib \in \mathbb{C}$ defined as multiplication by the complex conjugate $\bar{\lambda} = a - ib$. In other words, if $V \to \bar{V} : v \mapsto \bar{v}$ denotes the identity map, then scalar multiplication on $\bar{V}$ is defined to make this map complex antilinear, giving the formula

$$\lambda\bar{v} := \overline{\bar{\lambda}v} \in \bar{V} \qquad \text{for } \lambda \in \mathbb{C},\ v \in V.$$

(a) Show that there is a covariant functor $\boldsymbol{\kappa} : \mathsf{Vec}_{\mathbb{C}} \to \mathsf{Vec}_{\mathbb{C}}$ that sends each $V \in \mathsf{Vec}_{\mathbb{C}}$ to its conjugate $\bar{V}$, and describe how this functor acts on morphisms.
(b) Show that if T is a natural transformation from $\mathsf{Id} : \mathsf{Vec}_{\mathbb{C}} \to \mathsf{Vec}_{\mathbb{C}}$ to $\boldsymbol{\kappa} : \mathsf{Vec}_{\mathbb{C}} \to \mathsf{Vec}_{\mathbb{C}}$, then T assigns to each $V \in \mathsf{Vec}_{\mathbb{C}}$ the zero map $V \to \bar{V}$.
Hint: What does T imply about the specific morphism $V \to V : v \mapsto iv$?

Comment: The map $V \to \bar{V} : v \mapsto \bar{v}$ is always a real-linear isomorphism, but it is not complex linear and is thus not a morphism in $\mathsf{Vec}_{\mathbb{C}}$. Every finite-dimensional complex vector space is of course complex-linearly isomorphic to its conjugate, simply because both spaces have the same dimension, but the lack of any nontrivial natural transformation $\mathsf{Id} \to \boldsymbol{\kappa}$ demonstrates that there is generally no canonical way to define such isomorphisms.

The singular homology functor. Let us now give a definition of singular homology in this category-theoretic context. For now we will consider only homology with integer coefficients and thus leave the coefficient group out of the notation, writing e.g. $H_*(X)$ instead of $H_*(X; \mathbb{Z})$. It is natural to view singular homology as the composition of *two* covariant functors: one that transforms topological information into algebra (in the form of a chain complex), and another that performs the purely algebraic step of replacing a chain complex with something that is less unwieldy and (hopefully) more computable, namely its homology groups. Let us define the algebraic functor first.

Definition 27.22. A **$\mathbb{Z}$-graded abelian group** (*$\mathbb{Z}$-graduierte abelsche Gruppe*)[45] G_* is an abelian group that is equipped with a direct sum splitting

$$G_* = \bigoplus_{n\in\mathbb{Z}} G_n,$$

i.e. there is a subgroup $G_n \subset G_*$ for each $n \in \mathbb{Z}$ such that every $g \in G_*$ can be written as $g = \sum_{n\in\mathbb{Z}} g_n$ for uniquely determined elements $g_n \in G_n$, at most finitely many of which are nonzero. Another way to say this is that the canonical homomorphism $\bigoplus_{n\in\mathbb{Z}} G_n \to G_*$ determined by the inclusions $G_n \hookrightarrow G_*$ of the subgroups is an isomorphism. An element $g \in G_*$ such that $g \in G_n$ for some $n \in \mathbb{Z}$ is called a **homogeneous** (*homogen*) element of **degree** (*Grad*) n. Let $\mathsf{Ab}_{\mathbb{Z}}$ denote the category whose objects are $\mathbb{Z}$-graded abelian groups, with morphisms from G_* to H_* defined as group homomorphisms that send G_n into H_n for every $n \in \mathbb{Z}$.

Definition 27.23. A **chain complex** (*Kettenkomplex*) of abelian groups is a $\mathbb{Z}$-graded abelian group C_* equipped with the additional structure of a homormophism $\partial : C_* \to C_*$ that satisfies $\partial(C_n) \subset C_{n-1}$ for every $n \in \mathbb{Z}$ and $\partial^2 = 0$. Given two chain complexes (A_*, ∂^A) and (B_*, ∂^B), a **chain map** (*Kettenabbildung*) from (A_*, ∂^A) to (B_*, ∂^B) is a morphism $\Phi : A_* \to B_*$ in the sense of Definition 27.22 such that $\Phi \circ \partial^A = \partial^B \circ \Phi$. Let Chain denote the category that has chain complexes as objects and chain maps as morphisms. (Notice in particular that the identity map is always a chain map, and the composition of two chain maps is also a chain map.)

In a chain complex C_*, the homogeneous elements $c \in C_n$ of a given degree n are sometimes called **n-chains** (*n-Ketten*), the homomorphism $\partial : C_* \to C_*$ is called the **boundary operator** (*Randoperator*), elements in its image are called **boundaries** (*Ränder*), and the elements of its kernel are called **cycles** (*Zykel*). In situations where it's important to keep track of the degree, the n-chains that are also boundaries or cycles can also be called **n-boundaries** or **n-cycles** respectively. The condition $\partial^2 = 0$ thus means that for every $n \in \mathbb{Z}$, the n-boundaries form a subgroup of the n-cycles, and the definition below makes the "nth homology group" of a chain complex the quotient of the latter by the former; we also call a pair of cycles **homologous** (*homolog*) if their difference is a boundary. There is some geometric intuition behind the idea that one might obtain useful topological information by measuring "n-dimensional cycles modulo n-dimensional boundaries," and for instance in the case $n = 1$, we will be able to view it as a variation on the definition of the fundamental group, in which loops are replaced by 1-cycles and we consider them equivalent when they are homologous instead of homotopic. The extension of this intuition to arbitrary $n \geqslant 0$ will become more apparent when we talk about the homology of triangulated manifolds in a few lectures. But for now, it should also suffice to regard these as purely algebraic notions.[46]

Definition 27.24. The **homology** (*Homologie*) of a chain complex (C_*, ∂) is the graded abelian group $H_*(C_*, \partial) = \bigoplus_{n\in\mathbb{Z}} H_n(C_*, \partial)$ where

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

with ∂_n denoting the restriction of $\partial : C_* \to C_*$ to $C_n \to C_{n-1}$.

[45] We follow the common algebraic convention that whenever a group G is known *a priori* to be abelian, its group operation is denoted by "+", its identity element by $0 \in G$, and the inverse of each $g \in G$ by $-g \in G$, with subtraction "$g - h$" then being an abbreviation for "$g + (-h)$". This applies to all chain complexes and all homology groups, though it does not apply e.g. to cases in which $\pi_1(X)$ happens to be abelian, since fundamental groups are not *always* abelian.

[46] If you haven't seen the intuition behind homology before, it might also help to take a look at last semester's Lecture 21 on bordism groups and simplicial homology.

PROPOSITION 27.25. *There is a functor* $H_* : \mathsf{Chain} \to \mathsf{Ab}_{\mathbb{Z}}$ *that assigns to each chain complex* (C_*, ∂) *its homology* $H_*(C_*, \partial)$ *and assigns to each chain map* $\Phi : (A_*, \partial^A) \to (B_*, \partial^B)$ *the homomorphism*

$$H_*(A_*, \partial^A) \to H_*(B_*, \partial^B) : [a] \mapsto [\Phi(a)].$$

PROOF. A straightforward exercise. □

A **chain homotopy** (*Kettenhomotopie*) between two chain maps $\Phi, \Psi : (A_*, \partial^A) \to (B_*, \partial^B)$ is a homomorphism $h : A_* \to B_*$ such that $h(A_n) \subset B_{n+1}$ for each $n \in \mathbb{Z}$ and

$$\partial^B \circ h + h \circ \partial^A = \Phi - \Psi,$$

and we say that Φ and Ψ are **chain homotopic** (*kettenhomotop*) whenever a chain map exists. It is not too hard to show that chain homotopy defines an equivalence relation for chain maps, and the notion of composition for chain maps descends to a well-defined composition of *chain homotopy classes* of chain maps, which is why the following definition makes sense:

DEFINITION 27.26. Let Chain^h denote the category whose objects are chain complexes and whose morphisms are chain homotopy classes of chain maps.

It is an easy algebraic exercise to show that whenever two chain maps $\Phi, \Psi : (A_*, \partial^A) \to (B_*, \partial^B)$ are chain homotopic, the homomorphisms $H_*(A_*, \partial^A) \to H_*(B_*, \partial^B)$ that they induce are the same. Proposition 27.25 thus extends as follows:

PROPOSITION 27.27. *The prescription of Proposition 27.25 also defines a functor* $\mathsf{Chain}^h \to \mathsf{Ab}_{\mathbb{Z}}$. □

So much for algebra; now back to topology. Since we will frequently find ourselves talking about paths and homotopies, a very useful abbreviation to have at our disposal is

$$I := [0, 1],$$

and if we think of I as a compact 1-dimensional manifold with boundary, then its set of boundary points can likewise be denoted by

$$\partial I = \{0, 1\}.$$

This is not the only possible meaning for the symbol I when it is used in these notes, but it is by far the most common, and it will hopefully always be clear from context when I am allowing I to mean something else. Needless to say, $I = [0, 1]$ should always be assumed to carry the standard topology that it inherits from $\mathbb{R}$.

DEFINITION 27.28. A **pair of spaces** (*topologisches Paar*) (X, A) consists of a topological space X together with a subset $A \subset X$. A **map of pairs** (*Abbildung von Paaren*) $f : (X, A) \to (Y, B)$ is a continuous map $f : X \to Y$ such that $f(A) \subset B$. Two such maps $f, g : (X, A) \to (Y, B)$ are **homotopic** (*homotop*) if there exists a continuous map $h : I \times X \to Y$ such that $h(0, \cdot) = f$, $h(1, \cdot) = g$, and $h(I \times A) \subset B$. Let $\mathsf{Top}_{\mathrm{rel}}$ denote the category whose objects are pairs of spaces and whose morphisms are maps of pairs. Similarly, $\mathsf{Top}^h_{\mathrm{rel}}$ will denote the category with the same objects, but whose morphisms are homotopy classes of maps of pairs.

DEFINITION 27.29. For any integer $n \geqslant 0$, define the **standard n-simplex** (*Standard n-Simplex*) as the set

$$\Delta^n = \left\{ (t_0, \ldots, t_n) \in I^{n+1} \;\middle|\; \sum_j t_j = 1 \right\},$$

with the standard topology that it inherits as a subset of $\mathbb{R}^{n+1}$. For $n \geqslant 1$ and $k = 0, \dots, n$, the **kth boundary face** (*k-te Seitenfläche*) of Δ^n is the subset

$$\partial_{(k)}\Delta^n = \{t_k = 0\} \subset \Delta^n,$$

which we will sometimes identify with Δ^{n-1} via the obvious homeomorphism

$$\partial_{(k)}\Delta^n \to \Delta^{n-1} : (t_0, \dots, t_{k-1}, 0, t_{k+1}, \dots, t_n) \mapsto (t_0, \dots, t_{k-1}, t_{k+1}, \dots, t_n). \tag{27.2}$$

A **singular n-simplex** (*singulärer n-Simplex*) in a space X is defined to be a continuous map $\sigma : \Delta^n \to X$. Let

$$\mathcal{K}_n(X) := \left\{\sigma : \Delta^n \to X \;\middle|\; \sigma \text{ is continuous}\right\}$$

denote the set of all singular n-simplices in X.

We can now state the prescription for turning a pair of spaces (X, A) into its **singular chain complex** (*singulärer Kettenkomplex*) $C_*(X, A)$. First, define $C_n(X) = 0$ for all $n < 0$, by which we mean $C_n(X)$ is the trivial abelian group. For $n \geqslant 0$, define $C_n(X)$ to be the free abelian group generated by the set $\mathcal{K}_n(X)$, i.e.

$$C_n(X) = \bigoplus_{\sigma \in \mathcal{K}_n(X)} \mathbb{Z}.$$

Elements of $C_n(X)$ are called **singular n-chains** (*singuläre n-Ketten*) and can be written as finite sums $\sum_j m_j \sigma_j$, where $m_j \in \mathbb{Z}$ and $\sigma_j \in \mathcal{K}_n(X)$, with the understanding that $(k + \ell)\sigma = k\sigma + \ell\sigma$ for any individual $\sigma \in \mathcal{K}_n(X)$ and $k, \ell \in \mathbb{Z}$. The boundary operator $\partial : C_n(X) \to C_{n-1}(X)$ is necessarily trivial for $n \leqslant 0$ since $C_{n-1}(X)$ is trivial, whereas for $n \geqslant 1$, we can exploit the fact that $C_n(X)$ is freely generated and thus specify $\partial : C_n(X) \to C_{n-1}(X)$ by saying what it does to each of the generators $\sigma \in \mathcal{K}_n(X)$: the prescription is

$$\partial\sigma = \sum_{k=0}^{n} (-1)^k \left(\sigma|_{\partial_{(k)}\Delta^n}\right),$$

where the identification (27.2) is used in order to view each term in the summation as a singular $(n-1)$-simplex $\sigma|_{\partial_{(k)}\Delta^n} : \Delta^{n-1} \to X$, making the linear combination an element of $C_{n-1}(X)$. It is now straightforward to verify that $\partial^2 = 0$, so ∂ equips the graded abelian group $C_*(X) = \bigoplus_{n \in \mathbb{Z}} C_n(X)$ with the structure of a chain complex.

Finally, for any pair (X, A), one can assign the subspace topology to A and define $C_*(A)$ as above, with the consequence that $C_*(A)$ becomes a subgroup of $C_*(X)$ that is preserved by ∂, hence the latter descends to the quotients

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

and thus endows $C_*(X, A) := \bigoplus_{n \in \mathbb{Z}} C_n(X, A)$ with the structure of a chain complex. Algebraically, $C_n(X, A)$ is still quite simple: one can identify it with the free abelian group generated by the set of all singular n-simplices in X that are not fully contained in A, so that all elements of $C_n(X, A)$ can still be written as singular n-chains $\sum_j m_j \sigma_j$, with the understanding that terms having $\sigma_j(\Delta^n) \subset A$ should be ignored. Such an n-chain is then called a **relative n-cycle** if

$$\partial\left(\sum_j m_j \sigma_j\right) \subset C_{n-1}(A),$$

which just means that its equivalence class in $C_n(X, A) = C_n(X)/C_n(A)$ is in the kernel of the map $\partial : C_n(X, A) \to C_{n-1}(X, A)$, and two relative n-cycles are **homologous** if they differ by something that is the sum of an n-boundary with an arbitrary element of $C_n(A)$.

REMARK 27.30. If $A = \varnothing$, then $C_*(X, A)$ is the same thing as $C_*(X)$. It is often convenient to think of Top as the subcategory of $\mathsf{Top}_{\text{rel}}$ whose objects are all of the form $(X, \varnothing)$.

In order to view C_* as a functor from $\mathsf{Top}_{\text{rel}}$ to Chain, we need to explain what it does to morphisms. The answer is again straightforward: for any continuous map $f : X \to Y$, there is a chain map $f_* : C_*(X) \to C_*(Y)$ whose action on the generators $\sigma \in \mathcal{K}_n(X)$ is

$$f_*\sigma := f \circ \sigma \in \mathcal{K}_n(Y).$$

If f is also a map of pairs $(X, A) \to (Y, B)$, then the chain map $C_*(X) \to C_*(Y)$ sends $C_*(A)$ into $C_*(B)$ and thus descends to a relative chain map $f_* : C_*(X, A) \to C_*(Y, B)$. The subtler part of the story involves homotopies: as we will show within the next two lectures, any homotopy h between maps of pairs $f, g : (X, A) \to (Y, B)$ induces a chain homotopy h_* between the corresponding chain maps $f_*, g_* : C_*(X, A) \to C_*(Y, B)$. Putting all of this together gives the following:

PROPOSITION 27.31. *There exist functors* $\mathsf{Top}_{\text{rel}} \to \mathsf{Chain}$ *and* $\mathsf{Top}_{\text{rel}}^h \to \mathsf{Chain}^h$ *that assign to each pair* (X, A) *its relative singular chain complex* $C_*(X, A)$ *and associate to each map of pairs* $f : (X, A) \to (Y, B)$ *(or the homotopy class thereof) the induced chain map* $f_* : C_*(X, A) \to C_*(Y, B)$ *(or the chain homotopy class thereof).*

The **relative singular homology** (*relative singuläre Homologie*) of the pair (X, A) is defined as the homology of the chain complex $C_*(X, A)$ and is denoted by

$$H_*(X, A) = \bigoplus_{n \in \mathbb{Z}} H_n(X, A).$$

If $A = \varnothing$ as in Remark 27.30, then we abbreviate it as $H_*(X)$ and call it the **absolute** singular homology (*absolute singuläre Homologie*) of X. Composing the functors of Propositions 27.25 or 27.27 with those in Proposition 27.31 now gives:

THEOREM 27.32. H_* *defines functors* $\mathsf{Top}_{\text{rel}} \to \mathsf{Ab}_{\mathbb{Z}}$ *and* $\mathsf{Top}_{\text{rel}}^h \to \mathsf{Ab}_{\mathbb{Z}}$. □

28. Properties of singular homology (October 20, 2023)

In this lecture, we shall begin a quick survey of the most important properties of singular homology, including its computation for spheres of arbitrary dimension, modulo some technical details that will need to wait for the next lecture. A portion of this material was covered in the last few weeks of last semester's *Topologie I* course, but not quite at the level of generality that we will need going forward, so a few important things in this lecture (e.g. the *reduced* homology groups) will be new even if you took last semester's course.

Coefficients. The singular homology groups $H_n(X)$ we defined in Lecture 27 are sometimes also denoted by $H_n(X; \mathbb{Z})$ and called *homology with integer coefficients*, due to the fact that elements of the chain complex $C_*(X)$ can be written as finite linear combinations $\sum_j m_j \sigma_j$ of generators $\sigma_j \in \mathcal{K}_n(X)$ with coefficients $m_j \in \mathbb{Z}$. There is a more general version of this chain complex, in which the coefficients m_j are allowed to belong to an arbitrary abelian group G, which is sometimes useful in applications because e.g. taking coefficients in $\mathbb{Z}_2$ or $\mathbb{Q}$ may detect a slightly different range of topological information than coefficients in $\mathbb{Z}$. We will especially see a distinction between coefficients in $\mathbb{Z}_2$ as opposed to $\mathbb{Z}$ when we study the homology of non-orientable vs. orientable manifolds.

In order to put the notion of more general coefficient groups on a firm footing, we begin with an algebraic remark: every abelian group G is, in a canonical way, also a module over $\mathbb{Z}$. Indeed, the obvious multiplication operation $\mathbb{Z} \times G \to G$ can be defined for positive integers m by

$$mg := \underbrace{g + \ldots + g}_{m}.$$

Combining this with $0g := 0$ and $(-1)g := -g$ gives a definition of mg that makes sense for every $m \in \mathbb{Z}$. With this understood, we can now generalize the definition of the chain complex $C_*(X, A)$ in the previous lecture to

$$C_n(X, A; G) = C_n(X; G)/C_n(A; G), \quad \text{where} \quad C_n(X; G) = \bigoplus_{\sigma \in \mathcal{K}_n(X)} G.$$

Elements of $C_n(X, A; G)$ can thus be written as finite sums $\sum_i m_i \sigma_i$ where the coefficients m_i now belong to the group G, and addition in $C_n(X, A; G)$ is defined via the obvious relation $k\sigma + m\sigma := (k + m)\sigma$. The natural generalization of $\partial : C_n(X) \to C_{n-1}(X)$ to a map $C_n(X, A; G) \to C_{n-1}(X, A; G)$ is determined by the formula

$$\partial(g\sigma) := \sum_{k=0}^{n} (-1)^k g \left(\sigma|_{\partial_{(k)}\Delta^n}\right) \in C_{n-1}(X, A; G), \quad \text{for} \quad g \in G,\ \sigma \in \mathcal{K}_n(X),$$

which (as before) must be understood with the caveat that any term for which $\sigma|_{\partial_{(k)}\Delta^n}$ has image contained in A gets dropped from the sum, as it vanishes in the quotient $C_{n-1}(X; G)/C_{n-1}(A; G)$. For $G = \mathbb{Z}$, we recover the original chain complex $C_*(X, A; \mathbb{Z}) = C_*(X, A)$.

EXAMPLE 28.1. Since every element of $\mathbb{Z}_2$ is its own inverse, the formula for $\partial\sigma$ on $C_n(X; \mathbb{Z}_2)$ simplifies to $\partial\sigma = \sum_{k=0}^n \sigma|_{\partial_{(k)}\Delta^n}$.

The transformation of $C_*(X, A)$ into $C_*(X, A; G)$ can be reframed in terms of another purely algebraic functor, and this will be a useful perspective going forward. We first need to recall a few standard notions from the theory of abelian groups.

Given a set S, the **free abelian group** (*freie abelsche Gruppe*) on S is defined as a direct sum of copies of $\mathbb{Z}$, one for each element of S:

$$F^{\text{ab}}(S) := \bigoplus_{s \in S} \mathbb{Z}.$$

We can write elements of $F^{\text{ab}}(S)$ as finite sums $\sum_i m_i s_i$ for $m_i \in \mathbb{Z}$ and the **generators** (*Erzeuger*) $s_i \in S$, with the addition operation determined by $ks + ms := (k + m)s$ for any $k, m \in \mathbb{Z}$ and $s \in S$.

EXERCISE 28.2.

(a) Show that for any abelian group H, set S, and map $f : S \to H$, there exists a unique homomorphism $\Phi : F^{\text{ab}}(S) \to H$ such that $\Phi(s) = f(s)$ for each of the generators $s \in S$.
(b) Show that there is a natural isomorphism between $F^{\text{ab}}(S)$ and the abelianization of the free (non-abelian) group $F(S)$.

Given abelian groups G, H, K, a map $\Phi : G \oplus H \to K$ is called **bilinear** if for every fixed $g_0 \in G$ and $h_0 \in H$, the maps $G \to K : g \mapsto \Phi(g, h_0)$ and $H \to K : h \mapsto \Phi(g_0, h)$ are both homomorphisms.

The **tensor product** (*Tensorprodukt*) of two abelian groups G and H can be defined as the abelian group

$$G \otimes H := F^{\text{ab}}(G \times H)/N$$

where $N \subset F^{\text{ab}}(G \times H)$ is the smallest subgroup containing all elements of the form $(g + g', h) - (g, h) - (g', h)$ and $(g, h + h') - (g, h) - (g, h')$ for $g, g' \in G$ and $h, h' \in H$. We denote the equivalence class represented by $(g, h) \in F^{\text{ab}}(G \times H)$ in the quotient by

$$g \otimes h \in G \otimes H.$$

EXERCISE 28.3.

(a) Show that the map $G \oplus H \to G \otimes H : (g,h) \mapsto g \otimes h$ is bilinear, and deduce from this that for any $g \in G$ and $h \in H$, $0 \otimes h = g \otimes 0 = 0 \in G \otimes H$.
(b) Show that for any bilinear map $\Phi : G \oplus H \to K$ of abelian groups, there exists a unique homomorphism $\Psi : G \otimes H \to K$ such that $\Phi(g,h) = \Psi(g \otimes h)$ for all $(g,h) \in G \oplus H$.
(c) Show that for any abelian group G, the map $G \to G \otimes \mathbb{Z} : g \mapsto g \otimes 1$ is a group isomorphism. Write down its inverse.
Hint: Use part (b) to write down homomorphisms in terms of bilinear maps.
(d) Find a natural isomorphism from $(G \oplus H) \otimes K$ to $(G \otimes K) \oplus (H \otimes K)$.
(e) Given two sets S and T, find a natural isomorphism from $F^{\mathrm{ab}}(S) \otimes F^{\mathrm{ab}}(T)$ to $F^{\mathrm{ab}}(S \times T)$.
(f) Let $\mathbb{K}$ be a field, regarded as an abelian group with respect to its addition operation. Show that the abelian group $G \otimes \mathbb{K}$ naturally admits the structure of a vector space over $\mathbb{K}$ such that scalar multiplication takes the form

$$\lambda(g \otimes k) = g \otimes (\lambda k)$$

for every $\lambda, k \in \mathbb{K}$ and $g \in G$, and every group homomorphism $\Phi : G \to H$ determines a unique $\mathbb{K}$-linear map $\Psi : G \otimes \mathbb{K} \to H \otimes \mathbb{K}$ such that $\Psi(g \otimes k) = \Phi(g) \otimes k$ for $g \in G$, $k \in \mathbb{K}$.
(g) For any abelian groups A, B, C, D and homomorphisms $f : A \to B$, $g : C \to D$, show that there exists a homomorphism

$$f \otimes g : A \otimes C \to B \otimes D$$

defined uniquely by the condition $(f \otimes g)(a \otimes c) = f(a) \otimes g(c)$ for all $a \in A$ and $c \in C$.
(h) An element $a \in G$ is said to be **torsion** if $ma = 0$ for some $m \in \mathbb{Z}$. Show that if every element of G is torsion and $\mathbb{K}$ is a field (regarded as an abelian group with respect to addition), then $G \otimes \mathbb{K} = 0$.

The proof of the following result should now be an easy exercise.

PROPOSITION 28.4. *For any fixed abelian group G, there is a covariant functor*

$$\otimes G : \mathsf{Ab} \to \mathsf{Ab}$$

that sends each abelian group A to $A \otimes G$ and sends each group homomorphism $\Phi : A \to B$ to $\Phi \otimes \mathbb{1} : A \otimes G \to B \otimes G$.

Similarly, $\otimes G$ defines functors

$$\mathsf{Ab}_{\mathbb{Z}} \to \mathsf{Ab}_{\mathbb{Z}}, \qquad \mathsf{Chain} \to \mathsf{Chain}, \qquad \textit{and} \qquad \mathsf{Chain}^h \to \mathsf{Chain}^h$$

which send $\mathbb{Z}$-graded abelian groups $C_* = \bigoplus_{n\in\mathbb{Z}} C_n$ to $C_* \otimes G = \bigoplus_{n\in\mathbb{Z}} C_n \otimes G$ and morphisms $\Phi : C_* \to C_*$ to $\Phi \otimes \mathbb{1} : C_* \otimes G \to C_* \otimes G$. For a chain complex (C_*, ∂), the boundary map on $C_* \otimes G$ is defined as $\partial \otimes \mathbb{1}$.

EXERCISE 28.5. Fill in the details of the proof of Proposition 28.4. In particular, check that for any chain map $\Phi : (A_*, \partial^A) \to (B_*, \partial^B)$, the chain homotopy class of $\Phi \otimes \mathbb{1} : (A_* \otimes G, \partial^A \otimes \mathbb{1}) \to (B_* \otimes G, \partial^B \otimes \mathbb{1})$ depends only on the chain homotopy class of Φ, so that the functor $\otimes G : \mathsf{Chain}^h \to \mathsf{Chain}^h$ is well defined.
Hint: If $h_* : A_* \to B_*$ is a chain homotopy, what can you say about $h_* \otimes \mathbb{1} : A_* \otimes G \to B_* \otimes G$?

EXERCISE 28.6. Show that there is a canonical identification between $C_*(X, A; G)$ as defined above and the tensor product chain complex $C_*(X, A) \otimes G$ arising from Proposition 28.4.

In light of this discussion, the **singular homology with coefficients in** G

$$H_*(X, A; G) := H_*(C_*(X, A; G))$$

can now be understood as a composition of *three* covariant functors: first from $\mathsf{Top}_{\mathrm{rel}}$ to Chain to construct the singular chain complex with $\mathbb{Z}$-coefficients, then $\otimes G : \mathsf{Chain} \to \mathsf{Chain}$ to introduce the coefficient group G, and finally the homology functor $H_* : \mathsf{Chain} \to \mathsf{Ab}_{\mathbb{Z}}$. If we prefer to emphasize homotopy invariance, we can also view this as a composition of three functors $\mathsf{Top}^h_{\mathrm{rel}} \to \mathsf{Chain}^h \to \mathsf{Chain}^h \to \mathsf{Ab}_{\mathbb{Z}}$. The case $A = \varnothing$ will be abbreviated as usual by

$$H_*(X;G) := H_*(X,\varnothing;G).$$

Convention. From now on in these notes, we will often drop the coefficient group G from the notation and simply write

$$H_*(X) := H_*(X;G), \qquad C_*(X,A) := C_*(X,A;G) \qquad \text{etc.}$$

in situations where nothing important depends on the choice of the group G. When it is important to assume $G = \mathbb{Z}$, we will usually be explicit about this by writing $H_*(X;\mathbb{Z})$ instead of simply $H_*(X)$. This is a commonly used convention in the literature, but one should also be aware that there is another commonly used convention in which $H_*(X)$ is understood by default to be a synonym for $H_*(X;\mathbb{Z})$, and we will *not* be following that convention from now on.

When there is a coefficient group other than $\mathbb{Z}$ in the picture, we will usually continue to denote elements of the chain complex $C_*(X,A)$ by $\sum_i m_i\sigma_i$ instead of $\sum_i \sigma_i \otimes m_i \in C_*(X,A;\mathbb{Z}) \otimes G$, but the tensor product perspective will also serve us well, as it can often be used to turn results about homology with integer coefficients into results for general coefficient groups with almost no extra effort.

Absolute homology in degrees 0 and 1. Our survey of the properties of H_* begins with two fairly general computations that are easy to carry out, though for reasons that we'll get into later, they are often considered less important than the more "formal" properties that will be discussed subsequently.

Theorem 28.7. *For any space X and any coefficient group G, there is a canonical isomorphism*

$$H_0(X;G) = \bigoplus_{\pi_0(X)} G,$$

where $\pi_0(X)$ is an abbreviation for the set of path-components of X.

The isomorphism in this theorem arises from a pair of convenient coincidences: first, since the standard 0-simplex Δ^0 contains only one point, there is a natural bijection between the set $\mathcal{K}_0(X)$ of singular 0-simplices in X and the set X itself, allowing us to write singular 0-chains as finite linear combinations

$$\sum_i m_i x_i \in C_0(X)$$

of generators $x_i \in X$ with coefficients $m_i \in G$. The second coincidence is that the unit interval $I = [0,1]$, which we normally use for parametrizing paths in X, is homeomorphic to the standard 1-simplex $\Delta^1 \subset I^2$, e.g. via the map

$$I \xrightarrow{\cong} \Delta^1 : t \mapsto (1-t,t). \tag{28.1}$$

This is of course not the only possible choice of such a homeomorphism, but we will use it consistently in this course, for the following reason. The map (28.1) matches boundary points via the correspondence

$$\partial I \ni 0 \mapsto \partial_{(1)}\Delta^1 \subset \partial\Delta^1, \qquad \partial I \ni 1 \mapsto \partial_{(0)}\Delta^1 \subset \partial\Delta^1,$$

which may seem backwards when you see it for the first time, but if you recall the way in which signs were associated to the various boundary faces of Δ^n in our definition of the boundary operator $\partial : C_n(X) \to C_{n-1}(X)$, you might recognize that this particular correspondence is consistent with certain orientation conventions in differential geometry, where the standard orientation of the 1-manifold $I \subset \mathbb{R}$ induces a positive boundary orientation on $1 \in \partial I$ and a negative boundary orientation on $0 \in \partial I$. This detail is unimportant for our present purposes, but what matters is that if we use (28.1) to identify singular 1-simplices in X with paths $\gamma : I \to X$ and likewise identify singular 0-simplices with points $x \in X$ in the canonical way, then the operator $\partial : C_1(X) \to C_0(X)$ is now determined by the formula

$$\partial\gamma = \gamma(1) - \gamma(0). \tag{28.2}$$

This tells you why two 0-cycles of the form $mx, my \in C_0(X)$ for $m \in G$ and $x, y \in X$ will always be homologous if x and y lie in the same path-component, and from there it is not a difficult exercise to find an explicit isomorphism $H_0(X) \cong \bigoplus_{\pi_0(X)} G$.

For any choice of base point $p \in X$, the identification (28.1) between I and Δ^1 also gives rise to a natural homomorphism

$$h : \pi_1(X, p) \to H_1(X; \mathbb{Z}) \tag{28.3}$$

sending the homotopy class of the loop $\gamma : I \to X$ to the homology class that it represents when regarded as a singular 1-chain with integer coefficients; note that by (28.2), this 1-chain is a cycle because $\gamma : I \to X$ has the same start and end point. The map (28.3) is called the **Hurewicz homomorphism**, and the proof that it is well defined (see e.g. Exercise 22.12 from last semester's *Topologie I* course) relies on several straightforward lemmas, showing for instance that any two homotopic loops based at p give rise to homologous 1-cycles, and the 1-cycle arising from a concatenation of two loops is homologous to the sum of the two corresponding 1-cycles. Since $H_1(X; \mathbb{Z})$ is abelian, the Hurewicz map automatically vanishes on the commutator subgroup of $\pi_1(X, p)$, so it descends to a map of the abelianization of $\pi_1(X, p)$ to $H_1(X; \mathbb{Z})$.

THEOREM 28.8. *If X is path-connected, then the Hurewicz map (28.3) descends to the abelianization of $\pi_1(X) := \pi_1(X, p)$ as an isomorphism*

$$\pi_1(X) \big/ [\pi_1(X), \pi_1(X)] \xrightarrow{\cong} H_1(X; \mathbb{Z}).$$

One can prove Theorem 28.8 by writing down an inverse map that transforms any singular 1-cycle (viewed as a formal sum of paths whose end points must satisfy some matching conditions in order to produce a cycle) into a loop based at p by concatenating the associated paths. There are typically many ways that this can be done, but the ambiguity turns out to lie in the commutator subgroup $[\pi_1(X, p), \pi_1(X, p)]$; see last semester's Exercise 22.12 for further hints.

EXERCISE 28.9. Let Top_* denote the category of pointed spaces with base-point preserving continuous maps, so that we can regard both π_1 and $H_1(\cdot; \mathbb{Z})$ as functors from Top_* to the category Grp of groups with homomorphisms. (Note that the base point is irrelevant for the definition of $H_1(\cdot, \mathbb{Z})$, which actually takes values in the smaller subcategory of *abelian* groups, but these details are unimportant for now.) In this context, show that the Hurewicz homomorphism (28.3) defines a natural transformation from π_1 to $H_1(\cdot; \mathbb{Z})$.

I've put Theorems 28.7 and 28.8 first in our survey because, among the standard properties of singular homology, they are somewhat special: their proofs depend heavily on the precise definition of $H_*(X)$ in terms of singular simplices, and while both results are sometimes useful in concrete situations, it will turn out that the most important topological applications of homology do not depend on them. We will later discuss some other functors that are defined very differently from singular homology, and yet deserve nonetheless to be called "homology theories" due to a certain set of formal properties that they share with H_*. Those other theories do not generally have the two properties described above, but they will have the rest of the properties that we discuss below.

One-point spaces and disjoint unions. Here are the simplest two of the so-called "formal properties" of H_*. When we later start writing down axiomatic conditions that homology theories should be required to satisfy, these two will be known as the "dimension" and "additivity" axioms respectively.

THEOREM 28.10. *For a space* $\{\mathrm{pt}\}$ *consisting of only one point, we have*

$$H_n(\{\mathrm{pt}\};G) \cong \begin{cases} G & \text{if } n = 0, \\ 0 & \text{if } n \neq 0. \end{cases}$$

THEOREM 28.11. *For any collection of spaces* $\{X_\alpha\}_{\alpha \in J}$ *with inclusion maps* $i^\alpha : X_\alpha \hookrightarrow \coprod_{\beta \in J} X_\beta$*, the induced homomorphisms*

$$i^\alpha_* : H_*(X_\alpha) \to H_*\left(\coprod_{\alpha \in J} X_\beta\right)$$

determine an isomorphism

$$\bigoplus_{\alpha \in J} i^\alpha_* : \bigoplus_{\alpha \in J} H_*(X_\alpha) \to H_*\left(\coprod_{\alpha \in J} X_\beta\right).$$

Both theorems are straightforward consequences of the definitions, e.g. Theorem 28.10 is based on the observation that for each $n \geqslant 0$, the set $\mathcal{K}_n(\{\mathrm{pt}\})$ of singular n-simplices in a ont-point space has only one element, hence $C_n(X;G)$ is naturally isomorphic to G, and computing the boundary operator becomes a simple matter of recognizing when the signs cause cancelations and when they don't.

Suspensions and spheres. The remaining three formal properties of H_* will take more effort to explain than the first two, so before doing so, I want to state a computational result that provides some clear motivation for the effort: namely, the computation of $H_*(S^n;G)$ for every $n \geqslant 0$ and every coefficient group G. It is based on the observation that for each $n \geqslant 0$, the suspension of an n-sphere is homeomorphic to an $(n+1)$-sphere,

$$SS^n \cong S^{n+1},$$

where we recall (cf. Lecture 11 from last semester) that for an arbitrary space X, the **suspension** (*Einhängung*) of X is a space SX formed by gluing together two cones $C_+X := CX := (X \times [0,1])/(X \times \{1\})$ and $C_-X := (X \times [-1,0])/(X \times \{-1\})$ along $X = X \times \{0\} \subset C_\pm X$, in short,

$$SX := C_+X \cup_X C_-X.$$

The idea is then to find a general relation between $H_*(X;G)$ and $H_*(SX;G)$ for every space X, so that applying it repeatedly to the exceedingly simple space $X := S^0 \cong \{\mathrm{pt}\} \amalg \{\mathrm{pt}\}$ gives a formula for $H_*(S^n;G)$ by induction on n. This idea works, but there is a catch: the most useful relation one can prove between the homologies of X and SX is not actually a statement about H_*, but is instead about a slightly different object, called the *reduced* singular homology $\widetilde{H}_*$.

THEOREM 28.12. *One can associate to every space* X *and abelian group* G *a* $\mathbb{Z}$*-graded abelian group*

$$\widetilde{H}_*(X;G) = \bigoplus_{n \in \mathbb{Z}} \widetilde{H}_n(X;G),$$

abbreviated in the following as $\widetilde{H}_*(X) = \bigoplus_{n \in \mathbb{Z}} \widetilde{H}_n(X)$ *and called the* reduced *singular homology of* X*, which has the following properties:*

(1) $H_n(X) = \widetilde{H}_n(X)$ *for all* $n \neq 0$*, while* $H_0(X) \cong \widetilde{H}_0(X) \oplus G$*.*

(2) $\widetilde{H}_0(S^0) \cong G$.
(3) $\widetilde{H}_n(X) \cong \widetilde{H}_{n+1}(SX)$ *for every* $n \in \mathbb{Z}$ *and every space* X.

We will see by the end of this lecture where the suspension isomorphism $H_n(X) \cong H_{n+1}(SX)$ comes from when $n \geqslant 1$, and the next lecture will introduce the reduced homology $\widetilde{H}_*(X)$ in order to extend this result for $n \leqslant 0$. But first, let's clarify why we want this theorem to be true.

COROLLARY 28.13. *For every* $n \in \mathbb{N}$ *and* $k \in \mathbb{Z}$ *and every abelian group* G*,*

$$H_k(S^n; G) \cong \begin{cases} G & \text{if } k = 0 \text{ or } k = n, \\ 0 & \text{otherwise.} \end{cases}$$

PROOF. In light of the first property listed in Theorem 28.12, the formula is equivalent to the statement that $\widetilde{H}_n(S^n) \cong G$ for all $n \geqslant 1$ and $\widetilde{H}_k(S^n) = 0$ whenever $k \neq n$. If $k < n$, the latter follows by applying the suspension isomorphism of Theorem 28.12 repeatedly to

$$\widetilde{H}_{-1}(S^{n-k-1}) = H_{-1}(S^{n-k-1}) = 0.$$

If $k > n$, one can instead apply the suspension isomorphism to replace S^n with $S^0 \cong \{\text{pt}\} \amalg \{\text{pt}\}$ and then appeal to Theorems 28.10 and 28.11, giving

$$\widetilde{H}_{k-n}(S^0) = H_{k-n}(S^0) \cong H_{k-n}(\{\text{pt}\}) \oplus H_{k-n}(\{\text{pt}\}) = 0.$$

In the same manner, the computation of $\widetilde{H}_n(S^n)$ for each $n \in \mathbb{N}$ follows from $\widetilde{H}_0(S^0) \cong G$. $\square$

The computation of $H_*(S^n)$ is the first major step in a long series of computational results that make singular homology into an indispensable tool for understanding topological spaces. We will spend the next several weeks pursuing that story further, but at this stage it's worth noting that $H_*(S^n)$ on its own already has some impressive applications. One of those is the *Brouwer fixed point theorem*, which is an easy consequence of Corollary 28.13 in combination with an even simpler computation that will be sketched in a moment: for the closed unit disk $\mathbb{D}^n \subset \mathbb{R}^n$ of each dimension $n \geqslant 0$,

$$H_k(\mathbb{D}^n; G) \cong \begin{cases} G & \text{if } k = 0, \\ 0 & \text{otherwise.} \end{cases} \tag{28.4}$$

COROLLARY 28.14 (Brouwer). *For each* $n \geqslant 0$*, every continuous map* $f : \mathbb{D}^n \to \mathbb{D}^n$ *has a fixed point.*

PROOF. The result is trivial for $n = 0$, and the case $n = 1$ is an easy application of the intermediate value theorem, so let us assume $n \geqslant 2$. If $f : \mathbb{D}^n \to \mathbb{D}^n$ has no fixed point, then by following the unique line segment from $f(x)$ through x to the boundary $\partial\mathbb{D}^n = S^{n-1}$ for each $x \in \mathbb{D}^n$, one obtains a retraction

$$r : \mathbb{D}^n \to S^{n-1}.$$

Being a retraction, this map satisfies $r \circ i = \mathrm{Id}_{S^{n-1}}$ for the inclusion $i : S^{n-1} \hookrightarrow \mathbb{D}^n$, thus $r_* i_* : H_{n-1}(S^{n-1}) \to H_{n-1}(S^{n-1})$ is an isomorphism and $r_* : H_{n-1}(\mathbb{D}^n) \to H_{n-1}(S^{n-1})$ must therefore be surjective. But the latter is impossible for any choice of nontrivial coefficient group G, since $H_{n-1}(\mathbb{D}^n) = 0$ and $H_{n-1}(S^{n-1}) \cong G$. $\square$

Homotopy invariance and excision. Three further formal properties of H_* must be in place before we can prove Theorem 28.12. The first of these has already been mentioned a few times:

THEOREM 28.15 (homotopy invariance). *The homomorphism* $f_* : H_*(X, A) \to H_*(Y, B)$ *induced by a map of pairs* $f : (X, A) \to (Y, B)$ *depends only on its homotopy class (as a map of pairs).*

COROLLARY 28.16. *The homology groups of homotopy-equivalent pairs are isomorphic.* $\square$

In particular, since $\mathbb{D}^n$ is a contractible space, the corollary reduces the computation of $H_*(\mathbb{D}^n)$ stated in (28.4) to Theorem 28.10 on the homology of $\{\mathrm{pt}\}$.

The next of the formal properties expresses the intuition that for any pair (X, A) with $A \neq \varnothing$, $H_*(X, A)$ should ignore everything that happens entirely within the interior of A. Here we denote by

$$\mathring{A} \subset X$$

the interior of a subset $A \subset X$, and let

$$\bar{A} \subset X$$

denote the closure of A.

THEOREM 28.17 (excision). *For any subsets* $B \subset A \subset X$ *such that* $\bar{B} \subset \mathring{A}$, *the map*

$$H_*(X \backslash B, A \backslash B) \to H_*(X, A)$$

induced by the inclusion $(X \backslash B, A \backslash B) \hookrightarrow (X, A)$ *is an isomorphism.*

There's no getting around it: Theorems 28.15 and 28.17 are both somewhat tricky to prove, though in both cases, it is not so hard to point to intuitive reasons why they should be expected to hold. The ideas required are in any case interesting in their own right, and they deserve a more detailed discussion, so we shall make these two of the main topics for the next lecture, and take both theorems for now as black boxes.

The long exact sequence of a pair. The usefulness of being able to define homology for pairs (X, A) rather than just individual spaces X and A is that the three groups $H_*(X, A)$, $H_*(X)$ and $H_*(A)$ have a natural algebraic relation to each other. A sequence of abelian groups and homomorphisms

$$\ldots \longrightarrow G_{n-1} \xrightarrow{f_{n-1}} G_n \xrightarrow{f_n} G_{n+1} \longrightarrow \ldots$$

is called **exact** if $\operatorname{im} f_{n-1} = \ker f_n$ for every n. An exact sequence of the form

$$0 \longrightarrow G_1 \longrightarrow G_2 \longrightarrow G_3 \longrightarrow 0$$

is called a **short exact sequence** (*kurze exakte Sequenz*), and one can equally well consider an exact sequence of chain complexes, in which the homomorphisms are all assumed to be chain maps. An obvious example is

$$0 \longrightarrow C_*(A) \xrightarrow{i_*} C_*(X) \xrightarrow{j_*} C_*(X, A) \longrightarrow 0 \tag{28.5}$$

for any pair of spaces (X, A), where i_* and j_* are the chain maps induced by the natural inclusions $i : A \hookrightarrow X$ and $j : (X, \varnothing) \hookrightarrow (X, A)$. Notice that, in purely algebraic terms, i_* is also the inclusion of the subgroup $C_*(A) \hookrightarrow C_*(X)$, and j_* is the quotient projection $C_*(X) \to C_*(X, A) = C_*(X)/C_*(A)$. This short exact sequence can then be plugged into the following purely algebraic result:

PROPOSITION 28.18 (cf. Theorem 23.5). *Suppose* $0 \to A_* \xrightarrow{f} B_* \xrightarrow{g} C_* \to 0$ *is a short exact sequence of chain complexes. Then for each* $n \in \mathbb{Z}$ *there exists a so-called* ***connecting***

homomorphism *$\partial_* : H_n(C_*) \to H_{n-1}(A_*)$ such that the sequence*

$$\begin{aligned} \ldots \xrightarrow{\partial_*} H_{n+1}(A_*) \xrightarrow{f_*} H_{n+1}(B_*) &\xrightarrow{g_*} H_{n+1}(C_*) \\ &\xrightarrow{\partial_*} H_n(A_*) \xrightarrow{f_*} H_n(B_*) \xrightarrow{g_*} H_n(C_*) \\ &\xrightarrow{\partial_*} H_{n-1}(A_*) \xrightarrow{f_*} H_{n-1}(B_*) \xrightarrow{g_*} H_{n-1}(C_*) \xrightarrow{\partial_*} \ldots \end{aligned}$$

is exact. Moreover, this result is functorial in the following sense: suppose we are given another triple of chain complexes A'_*, B'_* and C'_*, with a commuting diagram

$$\begin{tikzcd} 0 \arrow[r] & A_* \arrow[r, "f"] \arrow[d, "\alpha"] & B_* \arrow[r, "g"] \arrow[d, "\beta"] & C_* \arrow[r] \arrow[d, "\gamma"] & 0 \\ 0 \arrow[r] & A'_* \arrow[r, "f'"] & B'_* \arrow[r, "g'"] & C'_* \arrow[r] & 0 \end{tikzcd}$$

in which all maps are chain maps and the bottom row is also exact, and we denote the resulting connecting homomorphisms by $\partial'_* : H_n(C'_*) \to H_{n-1}(A'_*)$. Then the diagram

$$\begin{tikzcd} \ldots \arrow[r] & H_{n+1}(C_*) \arrow[r, "\partial_*"] \arrow[d, "\gamma_*"] & H_n(A_*) \arrow[r, "f_*"] \arrow[d, "\alpha_*"] & H_n(B_*) \arrow[r, "g_*"] \arrow[d, "\beta_*"] & H_n(C_*) \arrow[r, "\partial_*"] \arrow[d, "\gamma_*"] & H_{n-1}(A_*) \arrow[r] \arrow[d, "\alpha_*"] & \ldots \\ \ldots \arrow[r] & H_{n+1}(C'_*) \arrow[r, "\partial'_*"] & H_n(A'_*) \arrow[r, "f'_*"] & H_n(B'_*) \arrow[r, "g'_*"] & H_n(C'_*) \arrow[r, "\partial'_*"] & H_{n-1}(A'_*) \arrow[r] & \ldots \end{tikzcd}$$

also commutes.

The proof of this result uses the standard method known as "diagram chasing". Let's do the first step, which is to write down a reasonable candidate for the map $\partial_* : H_n(C_*) \to H_{n-1}(A_*)$. We are given a commuting diagram of the form

$$\begin{tikzcd} & \vdots \arrow[d] & \vdots \arrow[d] & \vdots \arrow[d] & \\ 0 \arrow[r] & A_n \arrow[r, "f"] \arrow[d, "\partial"] & B_n \arrow[r, "g"] \arrow[d, "\partial"] & C_n \arrow[r] \arrow[d, "\partial"] & 0 \\ 0 \arrow[r] & A_{n-1} \arrow[r, "f"] \arrow[d, "\partial"] & B_{n-1} \arrow[r, "g"] \arrow[d, "\partial"] & C_{n-1} \arrow[r] \arrow[d, "\partial"] & 0 \\ 0 \arrow[r] & A_{n-2} \arrow[r, "f"] \arrow[d, "\partial"] & B_{n-2} \arrow[r, "g"] \arrow[d, "\partial"] & C_{n-2} \arrow[r] \arrow[d, "\partial"] & 0 \\ & \vdots & \vdots & \vdots & \end{tikzcd}$$

in which every column is a chain complex and every row is exact. Given $[c] \in H_n(C_*)$, choose a representative $c \in C_n$, which necessarily satisfies $\partial c = 0$. We would like to find some element $a \in A_{n-1}$ that satisfies $\partial a = 0$ so that we can set $\partial_*[c] := [a]$. The idea is to use whatever information the diagram gives us to forge a path from C_n to A_{n-1}. To start with, the exactness of the top row implies that g is surjective, so choose $b \in B_n$ with $g(b) = c$. Since $\partial c = 0$ and the diagram commutes, we also know $\partial g(b) = g(\partial b) = 0$, and exactness of the middle row then implies $\partial b = f(a)$ for some $a \in A_{n-1}$. To see that a is a cycle, we use commutativity again and observe $f(\partial a) = \partial f(a) = \partial\partial b = 0$, and since the bottom row is exact, f is injective, so this implies $\partial a = 0$. We can therefore sensibly set $\partial_*[c] = [a]$, and step 1 of the proof is complete.

There are still several things to check: steps 2 through 4000 consist of first verifying that the definition of $\partial_* : H_n(C_*) \to H_{n-1}(A_*)$ we just proposed does not depend on any of the choices we made (e.g. of the representative $c \in C_n$ and the element $b \in g^{-1}(c)$), and after that, we still need to show that the sequence of homology groups really is exact. All of this follows by the same style of diagram chasing—it becomes a bit tedious at some point, but it is not fundamentally difficult. If you haven't done it before, I recommend finding a quiet evening to do so once, so that you never have to do it again.

The long exact sequence in Proposition 28.18 appeared during our introduction to homology at the end of last semester's course, and while the "functoriality" aspect of the statement was not mentioned at the time, it will not be hard to see why it is true once you have understood the basic idea of diagram chasing. Functoriality in this situation amounts to the statement that there exist natural definitions of categories whose objects are short exact sequences of chain maps or long exact sequences of $\mathbb{Z}$-graded abelian groups, with morphisms defined in each case via commutative diagrams, such that Proposition 28.18 produces a functor from the former category to the latter. See Exercise 28.21 at the end of this lecture for a precise formulation in these terms.

Applying Proposition 28.18 to the short exact sequence (28.5) yields the last in our list of formal properties of singular homology:

THEOREM 28.19 (exactness). *For every pair of spaces (X, A) and $n \in \mathbb{Z}$, there exists a natural transformation ∂_* from the functor $(X, A) \mapsto H_n(X, A)$ to the functor $(X, A) \mapsto H_{n-1}(A)$, both regarded as functors $\mathsf{Top}_{\mathrm{rel}} \to \mathsf{Ab}$, such that the sequence*

$$\ldots \longrightarrow H_{n+1}(X, A) \xrightarrow{\partial_*} H_n(A) \xrightarrow{i_*} H_n(X) \xrightarrow{j_*} H_n(X, A) \xrightarrow{\partial_*} H_{n-1}(A) \longrightarrow \ldots$$

is exact, where i_* and j_* are induced by the inclusions $i : A \hookrightarrow X$ and $j : (X, \varnothing) \hookrightarrow (X, A)$. □

The fact that ∂_* is a natural transformation concretely means the following: if $f : (X, A) \to (Y, B)$ is any map of pairs, then the connecting homomorphisms for both pairs fit into the commutative diagram

$$\begin{tikzcd} H_n(X, A) \arrow[r, "\partial_*"] \arrow[d, "f_*"] & H_{n-1}(A) \arrow[d, "f_*"] \\ H_n(Y, B) \arrow[r, "\partial_*"] & H_{n-1}(B) \end{tikzcd}$$

This follows from the functoriality in Proposition 28.18, as f_* also induces a commutative diagram of chain maps

$$\begin{tikzcd} 0 \arrow[r] & C_*(A) \arrow[r, "i_*"] \arrow[d, "f_*"] & C_*(X) \arrow[r, "j_*"] \arrow[d, "f_*"] & C_*(X, A) \arrow[r] \arrow[d, "f_*"] & 0 \\ 0 \arrow[r] & C_*(B) \arrow[r, "i_*"] & C_*(Y) \arrow[r, "j_*"] & C_*(Y, B) \arrow[r] & 0 \end{tikzcd}$$

where the rows are simply the short exact sequences of (X, A) and (Y, B) respectively.

One last comment about the connecting homomorphisms $\partial_* : H_n(X, A) \to H_{n-1}(A)$. In the discussion above, we deduced their existence from an algebraic result, but it is also not hard to write them down with an explicit formula. To express it properly, recall that any relative homology class $[c] \in H_n(X, A)$ can be represented by some singular n-chain $c = \sum_i m_i \sigma_i \in C_n(X)$, i.e. this is a choice of representative for some element of the quotient $C_n(X, A) = C_n(X)/C_n(A)$, and the fact that that element is a cycle (i.e. is in $\ker \partial$) translates into the condition that ∂c must be an $(n-1)$-chain contained in A, in which case we call c a *relative* n-cycle. But ∂c is manifestly also a cycle in A since $\partial^2 = 0$, so it represents a homology class, and in this way we obtain the simplest possible formula for $\partial_* : H_n(X, A) \to H_{n-1}(A)$, namely

$$\partial_*[c] = [\partial c]. \tag{28.6}$$

If you rederive ∂_* from the diagram chase in Proposition 28.18 for the short exact sequence of the pair (X, A), you'll find that this is what it produces. The simplicity of the formula is deceptive: it looks like the right hand side should be trivial since it is the homology class of a boundary, but you need to keep in mind that while this is understood as a homology class in $H_{n-1}(A)$, c is not generally a chain in A, but in the larger space X.

The suspension isomorphism, almost. According to Theorem 28.12, there should be an isomorphism

$$H_n(X) \cong H_{n+1}(SX)$$

for every space X and every $n \geqslant 1$. The formal properties of H_* stated above now provide enough machinery to explain why this is true. Along the way, we will see why the argument does *not* work when $n \leqslant 0$, and the need to fill this gap is what will eventually motivate the definition of reduced homology.

For any space X, let

$$p_+ \in C_+X \subset SX \qquad \text{and} \qquad p_- \in C_-X \subset SX$$

denote the summits of the two cones that are glued together to form the suspension, e.g. if we write $C_+X = (X \times [0,1])/(X \times \{1\})$, then $p_+ \in C_+X$ is the point that results from collapsing $X \times \{1\}$. We then consider the diagram

$$\begin{tikzcd} H_k(X) & & H_{k+1}(SX) \arrow[d, "\varphi_*"] \\ H_{k+1}(C_+X, X) \arrow[u, "\partial_*"] \arrow[r, "i_*"] & H_{k+1}(SX\backslash\{p_-\}, C_-X\backslash\{p_-\}) \arrow[r, "j_*"] & H_{k+1}(SX, C_-X) \end{tikzcd} \tag{28.7}$$

in which three of the maps are determined by the obvious inclusions of pairs,

$$(C_+X, X) \overset{i}{\hookrightarrow} (SX\backslash\{p_-\}, C_-X\backslash\{p_-\}),$$
$$(SX\backslash\{p_-\}, C_-X\backslash\{p_-\}) \overset{j}{\hookrightarrow} (SX, C_-X),$$
$$(SX, \varnothing) \overset{\varphi}{\hookrightarrow} (SX, C_-X).$$

The first of these is a homotopy equivalence, as there exists a deformation retraction of the pair $(SX\backslash\{p_-\}, C_-X\backslash\{p_-\})$ to (C_+X, X), thus i_* is an isomorphism by Theorem 28.15. Since $\{p_-\}$ is a closed set contained in the interior of C_-X, excision (Theorem 28.17) implies that j_* is also an isomorphism. For the other two maps, we consider the long exact sequences of the pairs (SX, C_-X) and (C_+X, X), that is

$$\ldots \longrightarrow H_{k+1}(C_-X) \longrightarrow H_{k+1}(SX) \xrightarrow{\varphi_*} H_{k+1}(SX, C_-X) \longrightarrow H_k(C_-X) \longrightarrow \ldots$$

and

$$\ldots \longrightarrow H_{k+1}(C_+X) \longrightarrow H_{k+1}(C_+X, X) \xrightarrow{\partial_*} H_k(X) \longrightarrow H_k(C_+X) \longrightarrow \ldots$$

If $k \geqslant 1$, then the contractibility of $C_\pm X$ implies via homotopy invariance that

$$H_k(C_\pm X) \cong H_k(\{\mathrm{pt}\}) = 0, \qquad \text{and} \qquad H_{k+1}(C_\pm X) \cong H_{k+1}(\{\mathrm{pt}\}) = 0,$$

thus the exactness of these two sequences implies that φ_* and ∂_* are both isomorphisms. We've proved:

THEOREM 28.20. *For every space X and integer $k \geqslant 1$, the diagram (28.7) gives rise to an isomorphism*

$$S_* := \varphi_*^{-1} \circ j_* \circ i_* \circ \partial_*^{-1} : H_k(X) \to H_{k+1}(SX).$$

□

Theorem 28.20 is enough to achieve at least *part* of the inductive computation of $H_*(S^n)$ that we described in Corollary 28.13: in particular, if $k > n$, then applying this isomorphism n times gives

$$H_k(S^n) \cong H_{k-n}(S^0) \cong H_{k-n}(\{\text{pt}\}) \oplus H_{k-n}(\{\text{pt}\}) = 0$$

since $k - n > 0$. If $k \leqslant n$ and we are content to compute $H_k(S^n)$ only with integer coefficients, then Theorem 28.20 could instead be combined with the relation in Theorem 28.8 between $H_1(\cdot; \mathbb{Z})$ and π_1; as it happens, the fundamental groups of spheres are all abelian, so one then obtains

$$H_k(S^n; \mathbb{Z}) \cong H_1(S^{n-k+1}; \mathbb{Z}) \cong \pi_1(S^{n-k+1}) \cong \begin{cases} \mathbb{Z} & \text{if } k = n, \\ 0 & \text{if } 1 \leqslant k < n. \end{cases}$$

This computation suffices for basic applications such as the Brouwer fixed point theorem, and that is exactly what we did at the end of last semester's *Topologie I* course. Philosophically, however, this method is not really ideal, for a few reasons: one is that the isomorphism between $H_1(X)$ and the abelianization of $\pi_1(X)$ is only valid if one takes integer coefficients, and for further applications of homology to be explored in this course, we will sometimes need other choices of coefficient group. Another reason is that, as mentioned above, the relationship between $H_1(X; \mathbb{Z})$ and $\pi_1(X)$ is not counted among the "formal" properties of homology theories, and there are other homology theories besides singular homology that satisfy all of the formal properties but not that one—thus if we wanted to use one of those theories instead, an inductive argument based on $\pi_1(S^{n-k+1})$ would not work.

What's really needed is a way to modify Theorem 28.20 so that the condition $k \geqslant 1$ becomes unnecessary, as this condition is what prevents us from reducing the inductive computation of $H_k(S^n)$ for $k \leqslant n$ to $H_0(S^0) \cong H_0(\{\text{pt}\}) \oplus H_0(\{\text{pt}\})$ or $H_{-1}(S^{n-k-1}) = 0$. The reason why Theorem 28.20 doesn't work for $k = 0$ is that even though the cones $C_\pm X$ are contractible, the groups $H_0(C_\pm X)$ are not trivial, so that the two exact sequences we considered above fail to prove that φ_* and ∂_* are isomorphisms. The problem is thus caused by the fact that $H_*(X)$ for a contractible space does not completely vanish—it only *mostly* vanishes. The solution to this problem is reduced homology, so that's where we'll begin the next lecture.

EXERCISE 28.21. In this exercise we will prove the *functoriality* statement in Proposition 28.18 and flesh out its consequence for singular homology as a topological invariant. Consider the categories Short and Long, defined as follows. Objects in Short are short exact sequences of chain complexes $0 \to A_* \xrightarrow{f} B_* \xrightarrow{g} C_* \to 0$, with a morphism from this object to another object $0 \to A'_* \xrightarrow{f'} B'_* \xrightarrow{g'} C'_* \to 0$ defined as a triple of chain maps $A_* \xrightarrow{\alpha} A'_*$, $B_* \xrightarrow{\beta} B'_*$ and $C_* \xrightarrow{\gamma} C'_*$ such that the following diagram commutes:

$$\begin{tikzcd} 0 \arrow[r] & A_* \arrow[r, "f"] \arrow[d, "\alpha"] & B_* \arrow[r, "g"] \arrow[d, "\beta"] & C_* \arrow[r] \arrow[d, "\gamma"] & 0 \\ 0 \arrow[r] & A'_* \arrow[r, "f'"] & B'_* \arrow[r, "g'"] & C_* \arrow[r] & 0 \end{tikzcd}$$

The objects in Long are long exact sequences of $\mathbb{Z}$-graded abelian groups $\ldots \to C_{n+1} \xrightarrow{\delta} A_n \xrightarrow{F} B_n \xrightarrow{G} C_n \xrightarrow{\delta} A_{n-1} \to \ldots$, with morphisms from this to another object $\ldots \to C'_{n+1} \xrightarrow{\delta'} A'_n \xrightarrow{F'} B'_n \xrightarrow{G'}$

$C'_n \xrightarrow{\delta'} A'_{n-1} \to \ldots$ defined as triples of homomorphisms $A_* \xrightarrow{\alpha} A'_*$, $B_* \xrightarrow{\beta} B'_*$ and $C_* \xrightarrow{\gamma} C'_*$ that preserve the $\mathbb{Z}$-gradings and make the following diagram commute:

$$\begin{tikzcd}
\ldots \arrow[r] & C_{n+1} \arrow[r, "\delta"] \arrow[d, "\gamma"] & A_n \arrow[r, "F"] \arrow[d, "\alpha"] & B_n \arrow[r, "G"] \arrow[d, "\beta"] & C_n \arrow[r, "\delta"] \arrow[d, "\gamma"] & A_{n-1} \arrow[r] \arrow[d, "\alpha"] & \ldots \\
\ldots \arrow[r] & C'_{n+1} \arrow[r, "\delta'"] & A'_n \arrow[r, "F'"] & B'_n \arrow[r, "G'"] & C'_n \arrow[r, "\delta'"] & A'_{n-1} \arrow[r] & \ldots
\end{tikzcd}$$

(a) Show that there is a covariant functor $\mathsf{Top}_{\mathrm{rel}} \to \mathsf{Short}$ assigning to each pair (X, A) its short exact sequence of singular chain complexes.

(b) Show that there is also a covariant functor $\mathsf{Short} \to \mathsf{Long}$ assigning to each short exact sequence of chain complexes the corresponding long exact sequence of their homology groups. (Note that this can be composed with the functor in part (a) to define a functor $\mathsf{Top}_{\mathrm{rel}} \to \mathsf{Long}$.)

(c) Let Short^h denote a category with the same objects as in Short, but with morphisms consisting of triples of chain homotopy classes of chain maps. Assuming (as will be proved in the next lecture in order to establish Theorem 28.15) that the singular chain complex functor $C_* : \mathsf{Top}_{\mathrm{rel}} \to \mathsf{Chain}$ descends to a functor $\mathsf{Top}^h_{\mathrm{rel}} \to \mathsf{Chain}^h$, show that the functors in parts (a) and (b) also define functors $\mathsf{Top}^h_{\mathrm{rel}} \to \mathsf{Short}^h$ and $\mathsf{Short}^h \to \mathsf{Long}$, which then compose to define a functor $\mathsf{Top}^h_{\mathrm{rel}} \to \mathsf{Long}$.

29. Reduced homology, homotopy, and excision (October 24, 2023)

Our computation of $H_*(S^n)$ in the previous lecture had three gaps: two of them were the homotopy invariance and excision properties, which were used as black boxes but still need to be proved, and the other was the fact that the isomorphism $H_k(X) \cong H_{k+1}(SX)$ only works for $k \geqslant 1$ since we have not yet defined the *reduced* homology groups $\widetilde{H}_*(X)$. We'll deal with the latter issue first.

Reduced homology. Recall why the construction of the isomorphism $H_k(X) \cong H_{k+1}(SX)$ in the previous lecture didn't work without assuming $k \geqslant 1$. The argument required the terms $H_k(CX)$ and $H_{k+1}(CX)$ to vanish when they appear in certain exact sequences, which is fine when $k \geqslant 1$ since the cone CX is a contractible space, but $H_0(CX; G) \cong H_0(\{\mathrm{pt}\}; G) \cong G$ is unfortunately not a trivial group. Reduced homology is a clever remedy for this problem: it has many of the same formal properties as ordinary singular homology, and in particular fits into the same exact sequences, but with the added feature that it vanishes *completely* (not just mostly) on contractible spaces.

A brief algebraic digression is in order before we continue.

Exercise 29.1. Given a short exact sequence of abelian groups $0 \to A \xrightarrow{f} B \xrightarrow{g} C \to 0$, show that the following conditions are equivalent:

(i) There exists a homomorphism $\pi : B \to A$ such that $\pi \circ f = \mathbb{1}_A$;
(ii) There exists a homomorphism $i : C \to B$ such that $g \circ i = \mathbb{1}_C$;
(iii) There exists an isomorphism $\Phi : B \to A \oplus C$ satisfying the relations $\Phi \circ f(a) = (a, 0)$ and $g \circ \Phi^{-1}(a, c) = c$.

$$\begin{tikzcd}
 & & B \arrow[dd, "\Phi"] \arrow[dr, "g"] & & \\
0 \arrow[r] & A \arrow[ur, "f"] \arrow[dr] & & C \arrow[r] & 0 \\
 & & A \oplus C \arrow[ur] & &
\end{tikzcd}$$

Definition 29.2. We say that a short exact sequence **splits** whenever it satisfies any of the three equivalent properties listed in Exercise 29.1.

Exercise 29.3. Show that if the groups in Exercise 29.1 are all finite-dimensional vector spaces and the homomorphisms are linear maps, then the sequence always splits. Show also that this is true for any sequence of abelian groups $0 \to A \to B \to C \to 0$ if C is free.[47]
Hint: Use a basis of C to write down a right-inverse for $g : B \to C$.

Example 29.4. The sequence $0 \to 2\mathbb{Z} \to \mathbb{Z} \to \mathbb{Z}_2 \to 0$, with the first and last nontrivial maps defined as the natural inclusion and quotient projection respectively, is exact but does not split. Indeed, a splitting in this case would imply via Exercise 29.1 that $\mathbb{Z}$ is isomorphic to $2\mathbb{Z} \oplus \mathbb{Z}_2 \cong \mathbb{Z} \oplus \mathbb{Z}_2$, which it clearly is not. (In light of Exercise 29.3, we notice of course that $\mathbb{Z}_2$ is not free.)

End of digression; now we can explain the clever remedy. Fix a one-point space $\{\text{pt}\}$, and let

$$\epsilon : X \to \{\text{pt}\}$$

denote the unique map, which is (trivially) continuous.

Definition 29.5. The **reduced singular homology** (*reduzierte singuläre Homologie*) groups of X are defined for each $n \in \mathbb{Z}$ as the subgroup

$$\widetilde{H}_n(X) = \ker \epsilon_* \subset H_n(X),$$

where $\epsilon_* : H_n(X) \to H_n(\{\text{pt}\})$ is the homomorphism induced by the unique map $\epsilon : X \to \{\text{pt}\}$. All of these together form a $\mathbb{Z}$-graded abelian group

$$\widetilde{H}_*(X) = \bigoplus_{n \in \mathbb{Z}} \widetilde{H}_n(X),$$

which we will also denote by $\widetilde{H}_*(X; G) = \bigoplus_{n \in \mathbb{Z}} \widetilde{H}_n(X; G) \subset H_*(X; G)$ whenever the choice of coefficient group G needs to be made explicit.

Proposition 29.6. *If X is contractible, then $\widetilde{H}_*(X) = 0$.*

Proof. Contractibility implies that the map $\epsilon : X \to \{\text{pt}\}$ is a homotopy equivalence, thus $\epsilon_* : H_*(X) \to H_*(\{\text{pt}\})$ is an isomorphism, and its kernel $\widetilde{H}_*(X)$ is therefore trivial. □

Proposition 29.7. *For all spaces X, the map $\epsilon_* : H_*(X) \to H_*(\{\text{pt}\})$ is surjective, and the resulting short exact sequence*

$$0 \to \widetilde{H}_*(X) \hookrightarrow H_*(X) \xrightarrow{\epsilon_*} H_*(\{\text{pt}\}) \to 0$$

splits. In particular, $H_*(X)$ is isomorphic to $\widetilde{H}_*(X) \oplus H_*(\{\text{pt}\})$, thus

$$H_n(X; G) \cong \begin{cases} \widetilde{H}_n(X; G) \oplus G & \text{if } n = 0, \\ \widetilde{H}_n(X; G) & \text{if } n \neq 0. \end{cases}$$

Proof. Choose any map $i : \{\text{pt}\} \hookrightarrow X$ and notice that this is also trivially continuous, though usually not unique. Then $\epsilon \circ i$ is the identity map on $\{\text{pt}\}$, hence $\epsilon_* \circ i_* = \mathbb{1}$ on $H_*(\{\text{pt}\})$, implying that ϵ_* is surjective and (via Exercise 29.1) that the sequence splits. □

[47] An abelian group G is called **free** whenever it is isomorphic to the free abelian group $F^{\text{ab}}(S)$ on some set S. Equivalently, this means that G admits a **basis**, meaning a subset $S \subset G$ such that every element of G is uniquely representable as a linear combination $\sum_{s \in S} m_s s$ for some coefficients $m_s \in \mathbb{Z}$, only finitely many of which are nonzero.

PROPOSITION 29.8. *The homomorphism $f_* : H_*(X) \to H_*(Y)$ induced by any continuous map $f : X \to Y$ sends $\widetilde{H}_*(X)$ into $\widetilde{H}_*(Y)$. In particular, $\widetilde{H}_*$ defines functors $\mathsf{Top} \to \mathsf{Ab}_{\mathbb{Z}}$ and $\mathsf{Top}^h \to \mathsf{Ab}_{\mathbb{Z}}$ in the obvious way.*

PROOF. Denote $\epsilon^X : X \to \{\mathrm{pt}\}$ and $\epsilon^Y : Y \to \{\mathrm{pt}\}$ for the unique maps, and notice that $\epsilon^Y \circ f = \mathrm{Id} \circ \epsilon^X$, thus the following diagram commutes.

$$\begin{tikzcd} H_*(X) \arrow[r, "f_*"] \arrow[d, "\epsilon^X_*"] & H_*(Y) \arrow[d, "\epsilon^Y_*"] \\ H_*(\{\mathrm{pt}\}) \arrow[r, "\mathbb{1}"] & H_*(\{\mathrm{pt}\}) \end{tikzcd}$$

This implies that $f_*(\ker \epsilon^X_*) \subset \ker \epsilon^Y_*$. □

The relative version of reduced homology is defined in a trivial way: we set

$$\widetilde{H}_*(X, A) := H_*(X, A) \qquad \text{whenever} \qquad A \neq \varnothing.$$

This seemingly naive definition is justified by the following considerations. Note first that the functor $\widetilde{H}_* : \mathsf{Top} \to \mathsf{Ab}_{\mathbb{Z}}$ now extends to pairs as a functor $\mathsf{Top}_{\mathrm{rel}} \to \mathsf{Ab}_{\mathbb{Z}}$; here there is nothing to check since the existence of a map of pairs $(X, A) \to (Y, B)$ with $A \neq \varnothing$ implies $B \neq \varnothing$, so that both reduced relative homology groups match the unreduced case. Next, observe that for any space X, the relative homology groups $H_*(X, X)$ all vanish; one can prove this either directly from the definition of relative singular homology or indirectly via the long exact sequence of the pair. It follows that $\widetilde{H}_*(X, A)$ for $A \neq \varnothing$ is in fact the kernel of the map

$$H_*(X, A) \xrightarrow{\epsilon_*} H_*(\{\mathrm{pt}\}, \{\mathrm{pt}\}) = 0$$

induced by the unique map of pairs $\epsilon : (X, A) \to (\{\mathrm{pt}\}, \{\mathrm{pt}\})$. Moreover, the naturality of connecting homomorphisms gives a commutative diagram

$$\begin{tikzcd} H_{n+1}(X, A) \arrow[r, "\partial_*"] \arrow[d, "\epsilon_*"] & H_n(A) \arrow[d, "\epsilon_*"] \\ H_{n+1}(\{\mathrm{pt}\}, \{\mathrm{pt}\}) \arrow[r, "\partial_*"] & H_n(\{\mathrm{pt}\}) \end{tikzcd}$$

Since the term $H_{n+1}(\{\mathrm{pt}\}, \{\mathrm{pt}\})$ is trivial, this diagram proves that the image of $\partial_* : H_{n+1}(X, A) \to H_n(A)$ is always in the subgroup $\widetilde{H}_n(A)$. We can therefore write down a well-defined sequence of homomorphisms

$$\ldots \to \widetilde{H}_{n+1}(X, A) \stackrel{\partial_*}{\to} \widetilde{H}_n(A) \stackrel{i_*}{\to} \widetilde{H}_n(X) \stackrel{j_*}{\to} \widetilde{H}_n(X, A) \stackrel{\partial_*}{\to} \widetilde{H}_{n-1}(A) \to \ldots$$

using the usual inclusions $i : A \hookrightarrow X$ and $j : (X, \varnothing) \hookrightarrow (X, A)$. It is not immediately obvious, however, whether this sequence is exact. This is where the magic of diagram chasing again comes into play. Consider the commutative diagram

$$\begin{tikzcd}
 & 0 \arrow[d] & 0 \arrow[d] & 0 \arrow[d] & 0 \arrow[d] & \\
\ldots \arrow[r] & \widetilde{H}_n(A) \arrow[r, "i_*"] \arrow[d, hook] & \widetilde{H}_n(X) \arrow[r, "j_*"] \arrow[d, hook] & \widetilde{H}_n(X,A) \arrow[r, "\partial_*"] \arrow[d, equal] & \widetilde{H}_{n-1}(A) \arrow[r] \arrow[d, hook] & \ldots \\
\ldots \arrow[r] & H_n(A) \arrow[r, "i_*"] \arrow[d, "\epsilon_*"] & H_n(X) \arrow[r, "j_*"] \arrow[d, "\epsilon_*"] & H_n(X,A) \arrow[r, "\partial_*"] \arrow[d, "\epsilon_*"] & H_{n-1}(A) \arrow[r] \arrow[d, "\epsilon_*"] & \ldots \\
\ldots \arrow[r] & H_n(\{\mathrm{pt}\}) \arrow[r, equal] \arrow[d] & H_n(\{\mathrm{pt}\}) \arrow[r] \arrow[d] & 0 \arrow[r] \arrow[d] & H_{n-1}(\{\mathrm{pt}\}) \arrow[r] \arrow[d] & \ldots \\
 & 0 & 0 & 0 & 0 &
\end{tikzcd}$$

Here the bottom two nontrivial rows are the long exact sequences of the pairs (X, A) and $(\{\mathrm{pt}\}, \{\mathrm{pt}\})$, and all columns in the diagram are short exact sequences by construction. The rest is algebra:

Proposition 29.9. *Assume the following diagram of abelian groups with homomorphisms commutes, all its columns are exact sequences, and the bottom two nontrivial rows are also exact sequences:*

$$\begin{tikzcd}
 & 0 \arrow[d] & 0 \arrow[d] & 0 \arrow[d] & 0 \arrow[d] & 0 \arrow[d] & \\
\ldots & A_{n+2} \arrow[d, "\iota_{n+2}"] & A_{n+1} \arrow[d, "\iota_{n+1}"] & A_n \arrow[d, "\iota_n"] & A_{n-1} \arrow[d, "\iota_{n-1}"] & A_{n-2} \arrow[d, "\iota_{n-2}"] & \ldots \\
\ldots \arrow[r] & B_{n+2} \arrow[r, "g_{n+2}"] \arrow[d, "\epsilon_{n+2}"] & B_{n+1} \arrow[r, "g_{n+1}"] \arrow[d, "\epsilon_{n+1}"] & B_n \arrow[r, "g_n"] \arrow[d, "\epsilon_n"] & B_{n-1} \arrow[r, "g_{n-1}"] \arrow[d, "\epsilon_{n-1}"] & B_{n-2} \arrow[r] \arrow[d, "\epsilon_{n-2}"] & \ldots \\
\ldots \arrow[r] & C_{n+2} \arrow[r, "h_{n+2}"] \arrow[d] & C_{n+1} \arrow[r, "h_{n+1}"] \arrow[d] & C_n \arrow[r, "h_n"] \arrow[d] & C_{n-1} \arrow[r, "h_{n-1}"] \arrow[d] & C_{n-2} \arrow[r] \arrow[d] & \ldots \\
 & 0 & 0 & 0 & 0 & 0 &
\end{tikzcd}$$

Then the top nontrivial row can be endowed uniquely with maps $f_n : A_n \to A_{n-1}$ such that the diagram still commutes, and these make that row into an exact sequence.

Proof. If $f_n : A_n \to A_{n-1}$ can be defined so that the diagram commutes, then for $a \in A_n$ we need $f_n(a) \in \iota_{n-1}^{-1}(g_n \iota_n(a))$, and this condition will fully determine $f_n(a) \in A_{n-1}$ since ι_{n-1} is injective due to the exactness of columns. To see that the condition can be achieved, notice

$$\epsilon_{n-1} g_n \iota_n = h_n \epsilon_n \iota_n = 0,$$

thus $g_n \iota_n(a) \in \ker \epsilon_{n-1} = \operatorname{im} \iota_{n-1}$. This gives an element $x \in A_{n-1}$ such that $\iota_{n-1}(x) = g_n \iota_n(a)$, so we can set $f_n(a) = x$.

The goal is now to show that $\ldots A_{n+1} \stackrel{f_{n+1}}{\to} A_n \stackrel{f_n}{\to} A_{n-1} \to \ldots$ is an exact sequence. For each n, commutativity of the diagram gives

$$\iota_{n-2} f_{n-1} f_n = g_{n-1} g_n \iota_n = 0$$

since the middle row is exact, and the exactness of the columns implies in turn that ι_{n-2} is injective, thus $f_{n-1} f_n = 0$. To finish, we need to prove that every $a \in A_n$ satisfying $f_n(a) = 0$ also satisfies

$a = f_{n+1}(x)$ for some $x \in A_{n+1}$. Using commutativity, we have

$$0 = \iota_{n-1} f_n(a) = g_n \iota_n(a),$$

thus the exactness of the middle row gives an element $b \in B_{n+1}$ such that $g_{n+1}(b) = \iota_n(a)$. If we knew $\epsilon_{n+1}(b) = 0$, then we could at this point appeal to the exactness of the columns and write $b = \iota_{n+1}(x)$ for some $x \in A_{n+1}$, which would then satisfy $\iota_n f_{n+1}(x) = g_{n+1}\iota_{n+1}(x) = g_{n+1}(b) = \iota_n(a)$ and therefore $f_{n+1}(x) = a$ since ι_n is injective. But $\epsilon_{n+1}(b)$ might not be 0, so to finish the proof, we claim instead that b can be replaced by another element $b' \in B_{n+1}$ that satisfies both $g_{n+1}(b') = \iota_n(a)$ and $\epsilon_{n+1}(b') = 0$.

To find b', observe that by commutativity and the exactness of the columns,

$$h_{n+1}\epsilon_{n+1}(b) = \epsilon_n g_{n+1}(b) = \epsilon_n \iota_n(a) = 0,$$

thus by the exactness of the bottom row, $\epsilon_{n+1}(b) = h_{n+2}(c)$ for some $c \in C_{n+2}$. Appealing again to the exactness of the columns, ϵ_{n+2} is surjective, so we have $c = \epsilon_{n+2}(y)$ for some $y \in B_{n+2}$. Set

$$b' := b - g_{n+2}(y).$$

This satisfies $g_{n+1}(b') = g_{n+1}(b) - g_{n+1}g_{n+2}(y) = g_{n+1}(b) = \iota_n(a)$, and using commutativitiy again,

$$\epsilon_{n+1}(b') = \epsilon_{n+1}(b) - \epsilon_{n+1}g_{n+2}(y) = \epsilon_{n+1}(b) - h_{n+2}\epsilon_{n+2}(y) = \epsilon_{n+1}(b) - h_{n+2}(c) = 0.$$

□

We have proved:

THEOREM 29.10. *For any pair of spaces (X, A), there is a long exact sequence of reduced homology groups*

$$\ldots \to \widetilde{H}_{n+1}(X, A) \xrightarrow{\partial_*} \widetilde{H}_n(A) \xrightarrow{i_*} \widetilde{H}_n(X) \xrightarrow{j_*} \widetilde{H}_n(X, A) \xrightarrow{\partial_*} \widetilde{H}_{n-1}(A) \to \ldots,$$

where $i : A \hookrightarrow X$ and $j : (X, \varnothing) \hookrightarrow (X, A)$ are the obvious inclusions and $\partial_* : \widetilde{H}_n(X, A) \to \widetilde{H}_{n-1}(A)$ is the same map as the usual connecting homomorphism $H_n(X, A) \to H_{n-1}(A)$. □

Here's the upshot: if we now redo the argument behind Theorem 28.20, but replacing H_* with $\widetilde{H}_*$ at every step, then it still works, and it *also* works for $k = 0$ and $k = -1$ since $\widetilde{H}_0(C_\pm X)$ also vanishes. We conclude:

THEOREM 29.11. *For every space X and integer $k \in \mathbb{Z}$, and all choices of coefficient group, there is a natural isomorphism*

$$S_* : \widetilde{H}_k(X) \to \widetilde{H}_{k+1}(SX).$$

□

EXERCISE 29.12. Let us clarify the meaning of the word "natural" in Theorem 29.11.

(a) Show that for any continuous map $f : X \to Y$, the map $Sf : SX \to SY : [(x,t)] \mapsto [(f(x),t)]$ is well defined and continuous, and moreover, that $S(\mathrm{Id}_X) = \mathrm{Id}_{SX}$ and $S(f \circ g) = Sf \circ Sg$ whenever f and g can be composed. In other words, show that the suspension defines a functor $S : \mathsf{Top} \to \mathsf{Top}$.

(b) Denote by $\widetilde{H}^S_{n+1} : \mathsf{Top} \to \mathsf{Ab}$ the composition of the functor $S : \mathsf{Top} \to \mathsf{Top}$ in part (a) with the functor $\widetilde{H}_{n+1} : \mathsf{Top} \to \mathsf{Ab}$ which sends X to $\widetilde{H}_{n+1}(X)$. Show that there exists a natural transformation from $\widetilde{H}_n$ to $\widetilde{H}^S_{n+1}$ which associates to each space X the isomorphism $S_* : \widetilde{H}_n(X) \to \widetilde{H}_{n+1}(SX)$.

The only detail we have not yet addressed about reduced homology is the formula

$$\widetilde{H}_0(S^0;G) \cong G,$$

which was stated in Theorem 28.12 and was the starting point for the induction that led in Corollary 28.13 to $H_n(S^n;G) \cong G$ for all $n \geqslant 1$. The gap is filled by the following exercise.

EXERCISE 29.13. For any two spaces X and Y with maps $\epsilon^X : X \to \{\mathrm{pt}\}$ and $\epsilon^Y : Y \to \{\mathrm{pt}\}$, show that the natural isomorphism $H_*(X \amalg Y;G) \cong H_*(X;G) \oplus H_*(Y;G)$ identifies $\widetilde{H}_*(X \amalg Y;G)$ with $\ker(\epsilon_*^X \oplus \epsilon_*^Y) \subset H_*(X;G) \oplus H_*(Y;G)$. Then apply this in the case $X = Y = \{\mathrm{pt}\}$ to identify $\widetilde{H}_0(\{\mathrm{pt}\} \amalg \{\mathrm{pt}\};G)$ with the kernel of the map

$$\mathbb{1} \oplus \mathbb{1} : G \oplus G \to G : (g,h) \mapsto g + h,$$

which is isomorphic to G.

REMARK 29.14. It is sometimes useful to know that $\widetilde{H}_*(X)$ is also the homology of a chain complex. To see this, note that every element in $C_0(X)$ is a cycle, thus $H_0(X) = C_0(X)/\operatorname{im}\partial_1$ for the restriction $\partial_1 : C_1(X) \to C_0(X)$ of ∂. Thus for the unique map $\epsilon : X \to \{\mathrm{pt}\}$, the surjective homomorphism $\epsilon_* : H_0(X) \to H_0(\{\mathrm{pt}\}) = G$ can be composed with the quotient projection $C_0(X) \to C_0(X)/\operatorname{im}\partial_1 = H_0(X)$ to define a homomorphism $\epsilon_* : C_0(X) \to G$ such that $\epsilon_* \circ \partial_1 = 0$. This is equivalent to saying that ϵ_* defines a chain map $\epsilon_* : C_*(X) \to G_*$, where G_* denotes the chain complex with $G_0 := G$ and $G_n := 0$ for all $n \neq 0$, so that its boundary map is necessarily trivial.

In general, if C_* is any chain complex with $C_n = 0$ for all $n < 0$ and G_* is the trivial complex described above, a surjective chain map

$$\epsilon : C_* \to G_*$$

is called an **augmentation** (*Augmentationsabbildung*) of C_* over G. This is equivalent to a surjective homomorphism $\epsilon : C_0 \to G$ satisfying $\epsilon \circ \partial_1 = 0$ for the boundary map $\partial_1 : C_1 \to C_0$. One can therefore define an **augmented chain complex** $\widetilde{C}_*$ in the form

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

in other words, $\widetilde{C}_n := C_n$ for all $n \neq -1$ but $\widetilde{C}_{-1} := G$, with the new boundary map $\widetilde{C}_0 \to \widetilde{C}_{-1}$ defined as $\epsilon : C_0 \to G$. The homology of this new complex is precisely the kernel of the homomorphism $\epsilon_* : H_*(C_*) \to G_*$ induced by the chain map $\epsilon : C_* \to G_*$, thus we can sensibly call it the **reduced homology of the complex** C_*,

$$\widetilde{H}_*(C_*) := H_*(\widetilde{C}_*).$$

EXERCISE 29.15. Show that the augmentation $\epsilon_* : C_0(X;G) \to G$ described in Remark 29.14 is given by the formula

$$\epsilon_*\left(\sum_i g_i\sigma_i\right) = \sum_i g_i$$

for finite sums with $g_i \in G$ and $\sigma_i : \Delta^0 \to X$.

Homotopy invariance. The next loose end to tie up is the proof of Theorem 28.15, that two maps of pairs $f,g : (X,A) \to (Y,B)$ induce the same map

$$f_* = g_* : H_*(X,A) \to H_*(Y,B)$$

whenever f and g are homotopic (as maps of pairs). This follows from:

Proposition 29.16. *Any homotopy $h : (I \times X, I \times A) \to (Y, B)$ between two maps of pairs $f, g : (X, A) \to (Y, B)$ induces a chain homotopy*

$$h_* : C_*(X, A) \to C_{*+1}(Y, B)$$

between the two chain maps $f_*, g_* : C_*(X, A) \to C_*(Y, B)$.

In this statement we are using a popular abuse of notation and writing $h_* : C_*(X, A) \to C_{*+1}(Y, B)$ to emphasize the fact that h_* is a map *of degree* 1 between $\mathbb{Z}$-graded abelian groups, i.e. it satisfies $h_*(C_n(X, A)) \subset C_{n+1}(Y, B)$. Let's focus first on the case where $A = B = \varnothing$ and $G = \mathbb{Z}$, and see how a chain homotopy $h_* : C_*(X; \mathbb{Z}) \to C_{*+1}(Y; \mathbb{Z})$ arises from a homotopy $h : I \times X \to Y$ between the maps $f, g : X \to Y$.

It will suffice to define h_* on each generator of $C_n(X; \mathbb{Z})$ and verify that this definition satisfies the chain homotopy relation, i.e. for each $n \geqslant 0$ and each singular n-simplex $\sigma : \Delta^n \to X$, we need to specify some $(n+1)$-chain $h_*\sigma \in C_{n+1}(Y; \mathbb{Z})$ such that the relation

$$\partial(h_*\sigma) = g_*\sigma - f_*\sigma - h_*(\partial\sigma) \in C_n(Y; \mathbb{Z}).$$

is satisfied. The key is to look at the map

$$I \times \Delta^n \xrightarrow{h_\sigma} Y : (s, t) \mapsto h(s, \sigma(t)), \tag{29.1}$$

which is a homotopy between the singular n-simplices $f_*\sigma = f \circ \sigma$ and $g_*\sigma = g \circ \sigma$ in Y, and since $I \times \Delta^n$ is an $(n+1)$-dimensional domain, it would seem natural to produce from this map a singular $(n+1)$-chain in Y. Unfortunately, we cannot interpret h_σ as a singular $(n+1)$-simplex in any obvious way, as there is (with the exception of the simplest case $n = 0$) no canonical way to identify the "prism-shaped" domain $I \times \Delta^n$ with Δ^{n+1}.[48] What we can do, however, is find natural ways to produce an $(n+1)$-chain by decomposing $I \times \Delta^n$ into a union of finitely many $(n+1)$-simplices, a process known as **subdivision**, or in more general contexts, *triangulation*. We'll discuss triangulations of topological spaces in some detail next week, but for present purposes, what we need to know about triangulations of $I \times \Delta^n$ is summarized by an algebraic result stated below as Lemma 29.17. Let me give a brief sketch of the geometric picture behind this algebraic statement; the details can be outsourced to the next two lectures, where they will appear as a special case of a construction of "relative fundamental cycles" determined by triangulations.

The claim is that there exists an algorithm which subdivides $I \times \Delta^n$ for each $n \geqslant 0$ into a finite union of compact convex regions

$$\delta^{n+1} \subset I \times \Delta^n,$$

each being the convex hull of $n+2$ points in general position in $I \times \Delta^n \subset \mathbb{R} \times \mathbb{R}^{n+1} = \mathbb{R}^{n+2}$, called the **vertices** of δ^{n+1}. The words "general position" mean, in this context, that the vertices are not contained in any single n-dimensional plane, and δ^{n+1} is consequently homeomorphic to the standard simplex Δ^{n+1}; in particular, any map

$$\Delta^{n+1} \to \delta^{n+1}$$

defined by restricting a linear map $\mathbb{R}^{n+2} \to \mathbb{R} \times \mathbb{R}^{n+1}$ that sends the standard basis vectors of $\mathbb{R}^{n+2}$ bijectively to the vertices of δ^{n+1} is a homeomorphism. We shall therefore refer to the regions $\delta^{n+1} \subset I \times \Delta^n$ as the **$(n+1)$-simplices of the subdivision**, and we can similarly speak of the lower-dimensional **faces** of each $\delta^{n+1} \subset I \times \Delta^n$, which are the convex hulls in $I \times \Delta^n$ of proper subsets of the set of vertices of δ^{n+1}. In order to do what we want, our subdivision algorithm needs to have the following properties:

[48]Of course $I \times \Delta^n$ and Δ^{n+1} are both homeomorphic to the closed $(n+1)$-disk, and therefore to each other, but there is no *canonical* choice of homeomorphism, and even if we had one, it would not be useful to us unless it satisfied certain properties involving restrictions to boundary faces and so forth.

(1) For any two distinct $(n+1)$-simplices $\delta_1^{n+1}, \delta_2^{n+1} \subset I \times \Delta^n$ of the subdivision, $\delta_1^{n+1} \cap \delta_2^{n+1}$ is the convex hull of their set of common vertices, and is thus a lower-dimensional face of both simplices.
(2) Both $\{0\} \times \Delta^n$ and $\{1\} \times \Delta^n$ are boundary faces of $(n+1)$-simplices in the subdivision.
(3) The subdivision of $I \times \Delta^0$ consists of only a single 1-simplex $\delta^1 = I \times \Delta^0$.
(4) Under the canonical identification $\partial_{(k)}\Delta^n = \Delta^{n-1}$ for each $n \geqslant 1$ and $k = 0, \dots, n$, the subdivision of $I \times \Delta^{n-1} = I \times \partial_{(k)}\Delta^n$ can be obtained from the subdivision of $I \times \Delta^n$ by taking all boundary faces $\partial_{(j)}\delta^{n+1} \subset \delta^{n+1}$ of $(n+1)$-simplices $\delta^{n+1} \subset I \times \Delta^n$ such that $\partial_{(j)}\delta^{n+1} \subset I \times \partial_{(k)}\Delta^n$.

The first property in this list says in effect that our subdivision defines a triangulation of $I \times \Delta^n$, and the second says that this triangulation restricts to each of $\{0\} \times \Delta^n$ and $\{1\} \times \Delta^n$ as the trivial way of triangulating an n-simplex. The third and fourth properties give the construction an inductive quality: the beginning of the induction is to triangulate $I \times \Delta^0 \cong I \times \{\mathrm{pt}\} \cong I$ in the obvious way since $I \cong \Delta^1$, and the inductive step forces us to build the triangulation of $I \times \Delta^n$ on top of existing triangulations of $I \times \partial_{(k)}\Delta^n = I \times \Delta^{n-1}$ for each of the boundary faces of Δ^n. You should now at least be able to draw some pictures of how one might define such triangulations in practice for low values of n; there is more than one way to do it, and the precise choices will not matter so long as the lemma below is satisfied, but we will nonetheless describe a precise algorithm for this next week, as a specific example of an oriented triangulation. An additional feature of that algorithm is that it provides a recipe for turning the triangulation of $I \times \Delta^n$ into a singular $(n+1)$-chain of the form

$$\alpha_{n+1} = \sum_i \epsilon_i \sigma_i \in C_{n+1}(I \times \Delta^n; \mathbb{Z}),$$

where the sum ranges over the finite collection of $(n+1)$-simplices $\delta^{n+1} \subset I \times \Delta^n$ in the subdivision, each $\sigma_i : \Delta^{n+1} \to I \times \Delta^n$ is an explicit choice of homeomorphism of Δ^{n+1} to the corresponding δ^{n+1}, and the coefficients $\epsilon_i = \pm 1$ are suitably chosen signs. If this is done correctly, then the geometric properties listed above translate into the following algebraic statement—the point is that when computing $\partial\alpha_{n+1} \in C_n(I \times \Delta^n; \mathbb{Z})$, all boundary faces passing through the interior of $I \times \Delta^n$ cancel in pairs, leaving only those which lie entirely in

$$\partial(I \times \Delta^n) = \big(\{0\} \times \Delta^n\big) \cup \big(\{1\} \times \Delta^n\big) \cup \left(\bigcup_{k=0}^n I \times \partial_{(k)}\Delta^n \right).$$

Lemma 29.17. *One can define a sequence of singular chains $\alpha_{n+1} \in C_{n+1}(I \times \Delta^n; \mathbb{Z})$ for $n \geqslant 0$ such that $\alpha_1 = \sigma \in \mathcal{K}_1(I \times \Delta^0) \subset C_1(I \times \Delta^0; \mathbb{Z})$ is given by the usual homeomorphism* (28.1) *from Δ^1 to $I \times \Delta^0 = I \times \{\mathrm{pt}\} = I$, and for each $n \geqslant 1$,*

$$\begin{aligned} \partial\alpha_{n+1} = {} & (\{1\} \times \Delta^n \hookrightarrow I \times \Delta^n) - (\{0\} \times \Delta^n \hookrightarrow I \times \Delta^n) \\ & - \sum_{k=0}^n (-1)^k \left(I \times \partial_{(k)}\Delta^n \hookrightarrow I \times \Delta^n \right)_* \alpha_n \in C_n(I \times \Delta^n; \mathbb{Z}). \end{aligned}$$

Remark 29.18. To interpret the right hand side of the formula in Lemma 29.17 properly, one must regard $\{i\} \times \Delta^n$ for $i = 0, 1$ as canonically identified with Δ^n, so that the first two terms are singular n-simplices in $I \times \Delta^n$ (with attached coefficients ± 1). Similarly, the usual identification $\partial_{(k)}\Delta^n = \Delta^{n-1}$ makes the inclusion written after the summation into a continuous map $I \times \Delta^{n-1} \to I \times \Delta^n$, which we are using to push $\alpha_n \in C_n(I \times \Delta^{n-1}; \mathbb{Z})$ forward to an n-chain in $I \times \Delta^n$.

With Lemma 29.17 in hand, the chain homotopy $h_* : C_n(X; \mathbb{Z}) \to C_{n+1}(Y; \mathbb{Z})$ we need for the proof of Proposition 29.16 can be defined on each generator $\sigma \in \mathcal{K}_n(X)$ by using the map

$h_\sigma : I \times \Delta^n \to Y$ in (29.1) to push $\alpha_{n+1} \in C_{n+1}(I \times \Delta^n; \mathbb{Z})$ forward to Y, i.e.

$$h_*\sigma := (h_\sigma)_*\alpha_{n+1} \in C_{n+1}(Y; \mathbb{Z}).$$

Computing $\partial(h_*\sigma) \in C_n(Y; \mathbb{Z})$ is now a straightforward matter of unpacking definitions and applying the formula for $\partial\alpha_{n+1}$ stated in Lemma 29.17: the result is precisely the chain homotopy relation

$$\partial(h_*\sigma) = g_*\sigma - f_*\sigma - h_*(\partial\sigma),$$

in which each of the three terms on the right hand side arises by applying $(h_\sigma)_*$ to the corresponding term in our formula for $\partial\alpha_{n+1}$. This completes the proof of Proposition 29.16 for $A = B = \varnothing$ and $G = \mathbb{Z}$.

Extending the proof to cases with $A \ne \varnothing$ or $B \ne \varnothing$ only requires the observation that if $h(I \times A) \subset B$, then the chain homotopy h_* we constructed above sends $C_n(A; \mathbb{Z})$ into $C_{n+1}(B; \mathbb{Z})$, thus it descends to a chain homotopy on the quotient complexes $C_*(X, A; \mathbb{Z})$ and $C_*(Y, B; \mathbb{Z})$. Finally, the extension to general coefficient groups G comes for free in light of Exercises 28.5 and 28.6: the existence of the chain homotopy $h_* : C_*(X, A; \mathbb{Z}) \to C_{*+1}(Y, B; \mathbb{Z})$ gives rise to a chain homotopy

$$h_* \otimes \mathbb{1} : C_*(X, A; G) = C_*(X, A) \otimes G \to C_{*+1}(Y, B) \otimes G = C_{*+1}(Y, B; G).$$

Excision. The **excision** property (*Ausschneidungssatz*) stated in Theorem 28.17 amounts to the claim that $H_*(X, A)$ does not change if we remove from both A and X a subset with closure in the interior of A. Intuitively, this is clear since the definition of $H_*(X, A)$ is designed to ignore anything that happens completely inside of A, and the proof would be easy if we could assume that every class in $H_*(X, A)$ is representable by a relative cycle formed out of simplices $\sigma : \Delta^n \to X$ that avoid intersecting $B \subset A$. The latter is in fact true, but it will take some effort to prove it, and for the sake of further developments down the road, we will also need a somewhat stronger result than what was stated in Theorem 28.17. Concretely, we need a version that applies to chain complexes instead of homology groups.

Let us first give a suitable name to the notion of isomorphisms in the category Chain^h of chain complexes with chain homotopy classes of chain maps as morphisms.

Definition 29.19. A chain map $\Phi : A_* \to B_*$ between two chain complexes A_*, B_* of abelian groups is a **chain homotopy equivalence** (*Kettenhomotopieäquivalenz*) if there exists another chain map $\Psi : B_* \to A_*$ such that the chain maps $\Psi \circ \Phi : A_* \to A_*$ and $\Phi \circ \Psi : B_* \to B_*$ are both chain homotopic to the respective identity maps. In this situation, we call Ψ a **chain homotopy inverse** of Φ.

Here is the "chain level" version of the excision theorem:

Theorem 29.20. *If $A, B \subset X$ are subsets with $\bar{B} \subset \mathring{A}$, then the inclusion $i : (X \backslash B, A \backslash B) \hookrightarrow (X, A)$ induces a chain homotopy equivalence $i_* : C_*(X \backslash B, A \backslash B) \to C_*(X, A)$.*

As mentioned above, the proof of excision depends on the possibility of representing any given homology class in $H_*(X, A)$ by a relative cycle contained entirely in $X \backslash B$, and in fact, this is a special case of a more general phenomenon that will also be useful for other purposes later, notably when we prove the homological analogue of the Seifert-van Kampen theorem (called the Mayer-Vietoris sequence). We begin with the basic topological observation that if $\bar{B} \subset \mathring{A}$, then X is covered by the two open sets $\mathring{A}$ and $X \backslash \bar{B}$. The following more general question therefore presents itself: if we are given two subsets $\mathcal{U}, \mathcal{V} \subset X$ such that

$$X = \mathring{\mathcal{U}} \cup \mathring{\mathcal{V}},$$

to what extent can $H_*(X)$ be described purely in terms of chains that are contained in $\mathcal{U}$ or $\mathcal{V}$? To state an answer to this question, we note that the subgroups $C_*(\mathcal{U}), C_*(\mathcal{V}) \subset C_*(X)$ are both subcomplexes, i.e. the boundary operator $\partial : C_*(X) \to C_{*-1}(X)$ maps each of these subgroups to itself, and the same is therefore true of the sum of these two subgroups

$$C_*(\mathcal{U}) + C_*(\mathcal{V}) \subset C_*(X).$$

A slightly fancier way of saying the same thing is that $C_*(\mathcal{U}) + C_*(\mathcal{V})$ has a natural chain complex structure such that its inclusion into $C_*(X)$ is a chain map. The following lemma then tells us, among other things, that the smaller chain complex $C_*(\mathcal{U}) + C_*(\mathcal{V})$ can be used as a *substitute* for $C_*(X)$ when computing $H_*(X)$. The additional details stated in the lemma about the chain homotopy inverse of this inclusion will be needed when we apply it to prove Theorem 29.20.

Lemma 29.21. *For any subsets $\mathcal{U}, \mathcal{V} \subset X$ with $X = \mathring{\mathcal{U}} \cup \mathring{\mathcal{V}}$, the inclusion map*

$$j : C_*(\mathcal{U}) + C_*(\mathcal{V}) \hookrightarrow C_*(X)$$

is a chain homotopy equivalence. Moreover, it admits a chain homotopy inverse

$$\rho : C_*(X) \to C_*(\mathcal{U}) + C_*(\mathcal{V})$$

such that $\rho \circ j = \mathbb{1}$ and there is a chain homotopy $h : C_*(X) \to C_{*+1}(X)$ of $j \circ \rho$ to the identity such that h vanishes on the subgroup $C_*(\mathcal{U}) + C_*(\mathcal{V}) \subset C_*(X)$.

We focus in the following on the special case

$$G := \mathbb{Z},$$

as extending the result to general coefficient groups will be an easy algebraic exercise once the case with integer coefficients is understood. Constructing a chain homotopy inverse $\rho : C_*(X;\mathbb{Z}) \to C_*(\mathcal{U};\mathbb{Z}) + C_*(\mathcal{V};\mathbb{Z})$ requires a procedure to replace any chain in X by one that is a sum of two chains that are fully contained in $\mathcal{U}$ and $\mathcal{V}$ respectively. Even for a chain $c = \sigma \in C_n(X;\mathbb{Z})$ given by a single generator $\sigma \in \mathcal{K}_n(X)$, it is not immediately obvious how one should do this, because the map $\sigma : \Delta^n \to X$ need not have its image contained in either $\mathcal{U}$ or $\mathcal{V}$. One can easily imagine, however, that if we break up Δ^n into sufficiently smaller pieces and replace σ by a chain defined as a linear combination of its restrictions to those pieces, then since the interiors of $\mathcal{U}$ and $\mathcal{V}$ cover X, we could arrange for each of the individual pieces to be contained in one or the other, producing a chain that lives in $C_n(\mathcal{U}) + C_n(\mathcal{V})$. This leads us once again to the topic of *subdivision*: as we did with the prisms $I \times \Delta^n$ in order to prove homotopy invariance, we now need a suitable algorithm for subdividing a single simplex Δ^n into a finite collection of strictly smaller n-simplices.

There are in principle many conceivable ways of triangulating Δ^n that could work for this purpose, but the most popular is an algorithm called **barycentric subdivision** (*baryzentrische Zerlegung*), which does not require any arbitrary choices. In order to describe it, we first define a special point at the center of the standard n-simplex for each $n \geqslant 0$, called the **barycenter**

$$b_n := \left(\frac{1}{n+1}, \dots, \frac{1}{n+1}\right) \in \Delta^n \subset \mathbb{R}^{n+1}.$$

The subdivision algorithm decomposes Δ^n for each $n \geqslant 0$ into a union of finitely-many regions

$$\delta^n \subset \Delta^n,$$

each being the convex hull of a set of $n+1$ vertices that lie in general position in $\Delta^n \subset \mathbb{R}^{n+1}$, so that any choice of ordering for those vertices determines a homeomorphism $\Delta^n \to \delta^n$. The following inductive procedure determines the subdivision uniquely for every n:

(1) For $n = 0$, the subdivision of Δ^0 contains (obviously) only a single 0-simplex $\delta^0 = \Delta^0$.

(2) For each $n \geqslant 1$, the n-simplices $\delta^n \subset \Delta^n$ and $(n-1)$-simplices $\delta^{n-1} \subset \partial_{(k)}\Delta^n = \Delta^{n-1}$ in the barycentric subdivisions of Δ^n and Δ^{n-1} respectively are related by a bijective correspondence

$$\{\delta^n \subset \Delta^n\} \xleftrightarrow{1:1} \left\{\delta^{n-1} \subset \partial_{(k)}\Delta^n \;\middle|\; k = 0, \dots, n\right\},$$

in which each $(n-1)$-simplex $\delta^{n-1} \subset \partial_{(k)}\Delta^n$ on each of the $n+1$ boundary faces determines an n-simplex $\delta^n \subset \Delta^n$ formed as the convex hull of the vertices of δ^{n-1} plus the barycenter b_n.

EXERCISE 29.22. Draw pictures of the barycentric subdivisions of Δ^n for $n = 1, 2$, then look in [**Hat02**] to find a picture for $n = 3$. If you ever find a convincing picture of the case $n = 4$, let me know.

As a special case of the construction of relative fundamental cycles from triangulations described in the next two lectures, the geometric description of barycentric subdivision can be translated into the following algebraic statement, just as we previously did with our triangulation of $I \times \Delta^n$.

LEMMA 29.23. *One can define a sequence of singular chains $\beta_n \in C_n(\Delta^n;\mathbb{Z})$ for $n \geqslant 0$ such that β_0 is given by the unique singular 0-simplex in $\Delta^0 = \{\text{pt}\}$, and for each $n \geqslant 1$,*

$$\partial\beta_n = \sum_{k=0}^{n} (-1)^k \left(\partial_{(k)}\Delta^n \hookrightarrow \Delta^n\right)_* \beta_{n-1} \in C_{n-1}(\Delta^n;\mathbb{Z}).$$

Using the chains $\beta_n \in C_n(\Delta^n;\mathbb{Z})$ in Lemma 29.23, we now define for any space X the unique homomorphism

$$S : C_*(X;\mathbb{Z}) \to C_*(X;\mathbb{Z})$$

whose definition on singular n-chains $\sigma : \Delta^n \to X$ is

$$S(\sigma) := \sigma_*\beta_n \in C_n(X;\mathbb{Z}).$$

Since the notation "σ_*" may seem strange in this context, let me spell out what it means: on the right hand side of the formula, we are viewing $\sigma \in \mathcal{K}_n(X)$ not as a generator of $C_n(X;\mathbb{Z})$ but simply as a continuous map $\sigma : \Delta^n \to X$, which induces a chain map $\sigma_* : C_*(\Delta^n;\mathbb{Z}) \to C_*(X;\mathbb{Z})$ pushing the n-chain $\beta_n \in C_n(\Delta^n;\mathbb{Z})$ forward to an n-chain in X. In practice, the chain β_n can be written in the form

$$\beta_n = \sum_i \epsilon_i \left(\Delta^n \xrightarrow{\cong} \delta_i^n \hookrightarrow \Delta^n\right) \in C_n(\Delta^n;\mathbb{Z}),$$

where the sum ranges over the finite collection of smaller n-simplices $\delta_i^n \subset \Delta^n$ in the barycentric subdivision, for each of which we have chosen suitable homeomorphisms $\Delta^n \to \delta_i^n$ and signs $\epsilon_i = \pm 1$, and $S(\sigma) = \sigma_*\beta_n$ then takes the form

$$S(\sigma) = \sum_i \epsilon_i \left(\sigma|_{\delta_i^n}\right) \in C_n(X;\mathbb{Z}),$$

in which the chosen homeomorphisms $\Delta^n \to \delta_i^n$ are used in order to interpret the restricted maps $\sigma|_{\delta_i^n} : \delta_i^n \to X$ as singular n-simplices in X. The immediate consequence of the formula for $\partial\beta_n$ in Lemma 29.23 is then

$$\partial S(\sigma) = S(\partial\sigma),$$

so $S : C_*(X;\mathbb{Z}) \to C_*(X;\mathbb{Z})$ is a chain map.

The next statement tells us that, at the level of homology, acting on the chain complex $C_*(X;\mathbb{Z})$ with the chain map S changes nothing:

Proposition 29.24. *The chain map $S : C_*(X;\mathbb{Z}) \to C_*(X;\mathbb{Z})$ defined via barycentric subdivision is chain homotopic to the identity.*

By this point, you will perhaps not be surprised to learn that the chain homotopy needed for Proposition 29.24 can be constructed out of yet another subdivision algorithm. The object we need to subdivide this time is again $I \times \Delta^n$ for each $n \geqslant 0$, but the particular subdivision we used for proving homotopy invariance doesn't do the trick: what we need instead is a subdivision that matches the trivial triangulation of Δ^n at one end of the prism while reproducing the barycentric subdivision at the other. The new subdivision will thus decompose $I \times \Delta^n$ into finitely many convex $(n+1)$-simplices $\delta^{n+1} \subset I \times \Delta^n$ satisfying the same geometric properties that were stated for our previous subdivision of $I \times \Delta^n$, with only the following modification to the second property:

(2) $\{0\} \times \Delta^n$ is a boundary face of an $(n+1)$-simplex in the subdivision of $I \times \Delta^n$, and so is $\delta^n \subset \{1\} \times \Delta^n$ for every n-simplex in the barycentric subdivision of Δ^n.

The new subdivision of $I \times \Delta^0 = I \times \{\text{pt}\} = I$ is thus the same as the old one, and there is a simple inductive algorithm to produce a subdivision of $I \times \Delta^n$ with the desired properties for every $n \geqslant 1$: its $(n+1)$-simplices consist of

- The convex hull of the $n+1$ vertices of $\{0\} \times \Delta^n$ plus the extra vertex $(1, b_n)$; and
- The convex hull of the $n+1$ vertices of $\delta^n \subset I \times \partial_{(k)}\Delta^n$ plus the extra vertex $(1, b_n)$ for each n-simplex δ^n in the subdivision of $I \times \partial_{(k)}\Delta^n = I \times \Delta^{n-1}$ for each boundary face $\partial_{(k)}\Delta^n \subset \Delta^n$.

Translating this new subdivision into a singular chain via the usual trick gives:

Lemma 29.25. *One can define a sequence of singular chains $\gamma_{n+1} \in C_{n+1}(I \times \Delta^n;\mathbb{Z})$ for $n \geqslant 0$ such that $\gamma_1 = \alpha_1$ as in Lemma 29.17, while for each $n \geqslant 1$,*

$$\partial\gamma_{n+1} = (\{1\} \times \Delta^n \hookrightarrow I \times \Delta^n)_* \beta_n - (\{0\} \times \Delta^n \hookrightarrow I \times \Delta^n)$$
$$- \sum_{k=0}^{n} (-1)^k \left(I \times \partial_{(k)}\Delta^n \hookrightarrow I \times \Delta^n\right)_* \gamma_n \in C_n(I \times \Delta^n;\mathbb{Z}),$$

where $\beta_n \in C_n(\Delta^n;\mathbb{Z})$ is as in Lemma 29.23.

Proof of Proposition 29.24. Let $\text{pr} : I \times \Delta^n \to \Delta^n$ denote the natural projection map, so every singular n-simplex $\sigma : \Delta^n \to X$ gives rise to a map

$$\sigma \circ \text{pr} : I \times \Delta^n \to X.$$

We then define $h_* : C_n(X;\mathbb{Z}) \to C_{n+1}(X;\mathbb{Z})$ for each $n \geqslant 0$ as the unique homomorphism whose value on generators $\sigma \in \mathcal{K}_n(X)$ is

$$h_*\sigma := (\sigma \circ \text{pr})_*\gamma_{n+1} \in C_{n+1}(X;\mathbb{Z}).$$

The formula for $\partial\gamma_{n+1}$ in Lemma 29.25 then implies that h_* satisfies the chain homotopy relation

$$\partial(h_*\sigma) = S(\sigma) - \sigma - h_*(\partial\sigma).$$

□

With barycentric subdivision in hand, we now have the ability to subdivide the simplices we use for representing homology classes and make them smaller without changing the roles that they play in homology; in fact, the iterates $S^m : C_*(X;\mathbb{Z}) \to C_*(X;\mathbb{Z})$ for each $m \geqslant 1$ are clearly also chain homotopic to the identity, so by choosing m large enough, we can replace any chain with one whose constituent simplices have images as small as we like.

Proof of Lemma 29.21 with integer coefficients. Let me first point out how one would intuitively wish to construct the chain homotopy inverse $\rho : C_*(X;\mathbb{Z}) \to C_*(\mathcal{U};\mathbb{Z}) + C_*(\mathcal{V};\mathbb{Z})$, and why it does not quite work. As observed above, any chain $c \in C_*(X;\mathbb{Z})$ can be mapped into $C_*(\mathcal{U};\mathbb{Z}) + C_*(\mathcal{V};\mathbb{Z})$ via S^m if the integer m is sufficiently large, so S^m seems at first like a good candidate for the chain homotopy inverse ρ. The problem however is that we don't know in general how large m needs to be, and the answer is likely to depend on the chain c: in typical situations, one can always find for any fixed integer m a singular n-simplex $\sigma : \Delta^n \to X$ whose boundary is close enough to the boundary of $\mathcal{U}$ or $\mathcal{V}$ so that the m-fold subdivision $S^m(\sigma)$ includes some simplex that is not fully contained in either one. This means that regardless of how large we make m, S^m can never map *all* of $C_*(X;\mathbb{Z})$ into $C_*(\mathcal{U};\mathbb{Z}) + C_*(\mathcal{V};\mathbb{Z})$. Morally, what we really wish we could do is define

$$\text{“}\rho := \lim_{m\to\infty} S^m\text{”},$$

but since there is no obvious way to make sense of the limit on the right hand side, we will have to do something a bit less direct. Instead of trying to write down ρ explicitly, we will first write down a (somewhat naive) candidate for the chain homotopy h between $j \circ \rho$ and the identity $\mathbb{1}$, using as ingredients the chain homotopies between S^m and $\mathbb{1}$ for varying values of m. We will then be able to verify that h really is a chain homotopy between $\mathbb{1}$ and something; that so-called “something” will be defined to be ρ, whose further properties we can then verify.

Let $h_1 : C_*(X;\mathbb{Z}) \to C_{*+1}(X;\mathbb{Z})$ denote the chain homotopy provided by Proposition 29.24 for the barycentric subdivision chain map $S : C_*(X;\mathbb{Z}) \to C_*(X;\mathbb{Z})$, i.e. it satisfies $S - \mathbb{1} = \partial h_1 + h_1\partial$. We claim that for all integers $m \geqslant 0$, the map

$$h_m := h_1 \sum_{k=0}^{m-1} S^k : C_*(X;\mathbb{Z}) \to C_{*+1}(X;\mathbb{Z})$$

then satisfies

$$S^m - \mathbb{1} = \partial h_m + h_m\partial,$$

so h_m is a chain homotopy between S^m and the identity. Note that the case $m = 0$ is included here, with $S^0 = \mathbb{1}$ and $h_0 = 0$, so the claim is trivial in that case, and the definition of h_1 establishes it for $m = 1$. If we now use induction and assume that the claim holds for powers of S up to $m - 1 \geqslant 1$, then since S commutes with ∂,

$$\begin{aligned} S^m - \mathbb{1} &= (S^{m-1} - \mathbb{1})S + (S - \mathbb{1}) = (\partial h_{m-1} + h_{m-1}\partial)S + \partial h_1 + h_1\partial \\ &= \left(\partial h_1 \sum_{k=0}^{m-2} S^k + h_1 \sum_{k=0}^{m-2} S^k\partial\right) S + \partial h_1 + h_1\partial = \partial h_1 \sum_{k=1}^{m-1} S^k + h_1 \sum_{k=1}^{m-1} S^k\partial + \partial h_1 + h_1\partial \\ &= \partial h_1 \sum_{k=0}^{m-1} S^k + h_1 \sum_{k=0}^{m-1} S^k\partial = \partial h_m + h_m\partial. \end{aligned}$$

For any given $\sigma : \Delta^n \to X$, the iterated subdivision maps S^m can be assumed to satisfy

$$S^m(\sigma) \in C_*(\mathcal{U};\mathbb{Z}) + C_*(\mathcal{V};\mathbb{Z}) \tag{29.2}$$

if m is large enough, so for each each $n \geqslant 0$ and $\sigma : \Delta^n \to X$, let $m_\sigma \geqslant 0$ denote the smallest integer for which (29.2) holds with $m = m_\sigma$. We can then define a homomorphism $h : C_n(X;\mathbb{Z}) \to C_{n+1}(X;\mathbb{Z})$ for each $n \geqslant 0$ via

$$h(\sigma) := h_{m_\sigma}(\sigma).$$

Let us see whether this is a chain homotopy. We have

$$\begin{aligned} (\partial h + h\partial)(\sigma) &= \partial h_{m_\sigma}(\sigma) + h_{m_\sigma}(\partial\sigma) + (h - h_{m_\sigma})(\partial\sigma) \\ &= (S^{m_\sigma} - \mathbb{1})(\sigma) + (h - h_{m_\sigma})(\partial\sigma) = \left([S^{m_\sigma} + (h - h_{m_\sigma})\partial] - \mathbb{1}\right)(\sigma). \end{aligned}$$

Use this to define $\rho : C_*(X;\mathbb{Z}) \to C_*(X;\mathbb{Z})$ by

$$\rho(\sigma) := S^{m_\sigma}(\sigma) + (h - h_{m_\sigma})(\partial\sigma),$$

so the relation

$$\partial h + h\partial = \rho - \mathbb{1} \tag{29.3}$$

is satisfied. The latter implies that ρ is a chain map since applying ∂ from either the left or right on the left hand side of (29.3) gives $\partial h \partial$, thus on the right hand side we obtain $(\rho - \mathbb{1})\partial = \partial(\rho - \mathbb{1})$. To understand ρ better, we need to observe that each boundary face τ appearing in $\partial\sigma$ satisfies $m_\tau \leqslant m_\sigma$ since m_σ is clearly enough (but need not be the minimal number of) iterations of S to put σ (and therefore also τ) in $C_*(\mathcal{U};\mathbb{Z}) + C_*(\mathcal{V};\mathbb{Z})$. Now if $\sigma \in C_*(\mathcal{U};\mathbb{Z}) + C_*(\mathcal{V};\mathbb{Z})$, then $S^{m_\sigma}(\sigma) = \sigma$ since $m_\sigma = 0$, and the above remarks imply $h(\partial\sigma) = h_0(\partial\sigma) = 0$ as well, thus $\rho(\sigma) = \sigma$ and we conclude

$$\rho \circ j = \mathbb{1}.$$

It remains to show that for all $\sigma : \Delta^n \to X$, $\rho(\sigma)$ is a linear combination of simplices that are each contained in either $\mathcal{U}$ or $\mathcal{V}$. We have $S^{m_\sigma}(\sigma) \in C_*(\mathcal{U};\mathbb{Z}) + C_*(\mathcal{V};\mathbb{Z})$ by the definition of m_σ, so it suffices to inspect the other term $(h - h_{m_\sigma})(\partial\sigma)$. Here again we observe that $\partial\sigma$ is a linear combination of singular $(n-1)$-simplices τ for which $m_\tau \leqslant m_\sigma$, and

$$(h - h_{m_\sigma})\tau = (h_{m_\tau} - h_{m_\sigma})\tau = -h_1 \sum_{k=m_\tau}^{m_\sigma - 1} S^k(\tau) \in C_n(\mathcal{U};\mathbb{Z}) + C_n(\mathcal{V};\mathbb{Z}).$$

This last conclusion requires you to recall how h_1 was constructed in the proof of Proposition 29.24: in particular, it maps any simplex that is contained in either $\mathcal{U}$ or $\mathcal{V}$ to a linear combination of simplices that have this same property.

One last detail: the chain homotopy $h : C_*(X;\mathbb{Z}) \to C_{*+1}(X;\mathbb{Z})$ vanishes on $C_*(\mathcal{U};\mathbb{Z}) + C_*(\mathcal{V};\mathbb{Z})$ since every singular n-simplex $\sigma : \Delta^n \to X$ with image in either $\mathcal{U}$ or $\mathcal{V}$ satisfies $m_\sigma = 0$, thus $h(\sigma) = h_{m_\sigma}(\sigma) = h_0(\sigma) = 0$. ☐

Exercise 29.26. Use tensor products to deduce Lemma 29.21 for an arbitrary coefficient group G from the special case $G = \mathbb{Z}$.

We can now complete the proof of the chain-level excision theorem.

Proof of Theorem 29.20. We start with the observation that the inclusion $C_*(X \backslash B) \hookrightarrow C_*(X \backslash B) + C_*(A)$ descends to a chain map

$$C_*(X \backslash B, A \backslash B) \to \left(C_*(X \backslash B) + C_*(A)\right)/C_*(A)$$

which is an isomorphism of chain complexes, as the quotient complexes on the left and right hand side can be described as tensor products of the coefficient group G with free abelian groups having exactly the same sets of generators, namely the singular simplices in $X \backslash B$ that are not fully contained in A. With this identification understood, it suffices to prove that the chain map

$$\left(C_*(X \backslash B) + C_*(A)\right)/C_*(A) \xrightarrow{j} C_*(X)/C_*(A) = C_*(X, A) \tag{29.4}$$

induced on the quotients by the inclusion

$$C_*(X \backslash B) + C_*(A) \overset{j}{\hookrightarrow} C_*(X)$$

is a chain homotopy equivalence. Lemma 29.21 provides a chain homotopy inverse $\rho : C_*(X) \to C_*(X \backslash B) + C_*(A)$ for j, which can be assumed to satisfy $\rho \circ j = \mathbb{1}$, so that ρ restricts to the identity on $C_*(A) \subset C_*(X)$ and thus descends to the quotient, producing the desired chain homotopy inverse of (29.4). ☐

30. Simplicial complexes in singular homology (October 27, 2023)

Before further developing the theory of singular homology, I would like to pause and address a question that is important for intuition: how can we *visualize* a singular homology class? While this question is not always answerable, there is a standard answer that suffices in most situations: elements of $H_n(X)$ can often be viewed as *triangulated closed n-dimensional submanifolds* of X, where two such submanifolds represent the same homology class whenever their disjoint union (with appropriate orientations) bounds some triangulated compact $(n+1)$-dimensional submanifold with boundary. To explain in precise terms what this means, I need to digress a little bit into the subject of simplicial complexes and simplicial homology, which will also come in useful for other purposes later. Much of what I will say in this and the next lecture repeats things that I said or hinted at last semester (see especially Lecture 21), but I am now in a position to say those things more precisely.

You will recall that an **n-dimensional (topological) manifold with boundary** is a second countable Hausdorff space M in which every point has a neighborhood homeomorphic to an open subset of the n-dimensional half-space

$$\mathbb{H}^n := [0, \infty) \times \mathbb{R}^{n-1}.$$

The **boundary** $\partial M \subset M$ is defined as the set of points that are sent to $\partial \mathbb{H}^n := \{0\} \times \mathbb{R}^{n-1} \subset \mathbb{H}^n$ under these homeomorphisms, hence ∂M itself is an $(n-1)$-dimensional topological manifold (with empty boundary: $\partial(\partial M) = \varnothing$). A manifold M is called **closed** (*geschlossen*)[49] if it is compact and $\partial M = \varnothing$. All manifolds in this lecture will be compact, so we need not worry about second countability, and we will usually also omit the word "topological". (The word is often included in order to distinguish topological manifolds from *smooth* manifolds, but smooth structures will not have any role to play in our discussion.) We will usually also omit the words "with boundary" and keep in mind that all manifolds in principle have boundary, but the boundary may be empty.

Simplicial complexes and polyhedra. We usually picture a *simplicial complex* as a space decomposed into a union of simplices. Strictly speaking, a simplicial complex is a purely combinatorial object, and the topological space that we build out of it is called its *polyhedron*. Here are the precise definitions.

Definition 30.1. A **simplicial complex** (*Simplizialkomplex*) K consists of two sets V and S, called the sets of **vertices** (*Eckpunkte*) and **simplices** (*Simplizes*) respectively, where S is a subset of the set of all finite subsets of V, and $\sigma \in S$ is called an **n-simplex** of K if it has $n+1$ elements. We require the following conditions:

(1) Every vertex $v \in V$ gives rise to a 0-simplex in K, i.e. $\{v\} \in S$;
(2) If $\sigma \in S$ then every subset $\sigma' \subset \sigma$ is also an element of S.

For any n-simplex $\sigma \in S$, its subsets are called its **faces** (*Seiten* or *Facetten*), and in particular the subsets that are $(n-1)$-simplices are called **boundary faces** (*Seitenflächen*) of σ. The second condition above thus says that for every simplex in the complex, all of its faces also belong to the complex. With this condition in place, the first condition is then equivalent to the requirement that every vertex in the set V belongs to at least one simplex.

The complex K is said to be **finite** if V (and therefore also S) is finite, and its **dimension** is

$$\dim K := \sup_{\sigma \in S} \dim \sigma \in \{0, 1, 2, \ldots, \infty\},$$

where we write $\dim \sigma = n$ whenever σ is an n-simplex.

[49] Not to be confused with "*abgeschlossen*," which is what we call a closed *subset* of a topological space. These two meanings of the English word "closed" are defined in different contexts and are not equivalent.

The **polyhedron** (*Polyeder*) of a simplicial complex $K = (V, S)$ is a topological space $|K|$ defined as follows. We denote by I^V the set of all functions $V \to I := [0, 1]$, i.e. each element $t \in I^V$ is determined by a set of real numbers $t_v \in [0, 1]$ associated to the vertices $v \in V$, which we can think of as the *coordinates* of t. For each n-simplex $\sigma = \{v_0, \ldots, v_n\}$ in K, we define the set

$$|\sigma| := \left\{ t \in I^V \;\middle|\; \sum_{v \in \sigma} t_v = 1 \text{ and } t_v = 0 \text{ for all } v \notin \sigma \right\}.$$

This set is a copy of the standard n-simplex living in the finite-dimensional vector space $\mathbb{R}^\sigma \cong \mathbb{R}^{n+1}$, and we shall assign it the topology that it inherits naturally from this finite-dimensional vector space. As a set, the polyhedron $|K|$ is then defined by

$$|K| = \bigcup_{\sigma \in S} |\sigma| \subset I^V.$$

If K is finite, then $|K|$ lives inside the finite-dimensional vector space $\mathbb{R}^V$, and therefore has an obvious topology for which the topology we already defined on each of the subsets $|\sigma| \subset |K|$ matches the subspace topology. A bit more thought is required at this step if K is infinite. One possible choice would be to endow I^V with the product topology (via its obvious identification with $\prod_{v \in V} I$) and then take the subspace topology on $|K| \subset I^V$, but the product topology turns out not to be the most natural choice here. We will instead let the topology of $|K|$ be determined by that of the individual simplices: define a subset $\mathcal{U} \subset |K|$ to be open if and only if for every $\sigma \in S$, $\mathcal{U} \cap |\sigma| \subset |\sigma|$ is open for the topology on $|\sigma|$ defined above. In other words, $|K|$ carries the strongest[50] topology for which the inclusions $|\sigma| \hookrightarrow |K|$ are continuous for all σ. You should take a moment to convince yourself that this matches what was already said for the case where K is finite.

EXERCISE 30.2. Show that for any simplicial complex $K = (V, S)$ and any space X, a map $f : |K| \to X$ is continuous if and only if $f|_{|\sigma|} : |\sigma| \to X$ is continuous for every simplex $\sigma \in S$.

DEFINITION 30.3. For each integer $n \geqslant 0$, the ***n*-skeleton** (*n-Skelett* or *n-Gerüst*) of a polyhedron $X = |K|$ is the subspace $X^n \subset X$ consisting of the union of all $|\sigma| \subset X$ for k-simplices σ in K with $k \leqslant n$.

By this definition, a polyhedron is n-dimensional (i.e. corresponds to an n-dimensional simplicial complex) if and only if it is equal to its n-skeleton. The 0-skeleton of any polyhedron is just the union of all its vertices—one can show that this is always a discrete set.

While $|K|$ was defined above as a subset of a vector space whose dimension may in general be quite large (or infinite), visualizing $|K|$ in concrete examples is often easier than one might expect.

EXAMPLE 30.4. Suppose $V = \{v_0, v_1, v_2, v_3\}$ and S contains the subsets $A := \{v_0, v_1, v_2\}$ and $B := \{v_1, v_2, v_3\}$, plus all of their respective subsets. Then $|K|$ contains two copies of the triangle Δ^2, and they intersect each other along a single common edge connecting the vertices labeled v_1 and v_2. The complex is 2-dimensional, and its 1-skeleton is the union of all the edges of the triangles.

EXAMPLE 30.5. If V has $n + 1$ elements and S consists of all subsets of V except for V itself, then $|K|$ is homeomorphic to $\partial\Delta^n$, i.e. the union of all the boundary faces of Δ^n. In particular, this is homeomorphic to S^{n-1}.

[50] For some unfathomable reason, the topology on $|K|$ has traditionally been referred to in the literature as the "weak" topology, and the same strange choice of nomenclature plagues the theory of CW-complexes, which we will discuss in a few weeks. It is a question of perspective: since $|K|$ has a lot of open sets, it is fairly difficult for sequences in $|K|$ to converge, or for maps into $|K|$ to be continuous, but on the flip side, it is relatively easy for functions defined on $|K|$ to be continuous (see Exercise 30.2).

EXAMPLE 30.6. Suppose $V = \{v_0, \ldots, v_n\}$ and S is defined to consist of all the one-element subsets $\{v_i\}$ plus the 1-simplices $\{v_i, v_{i+1}\}$ for $i = 0, \ldots, n-1$ and $\{v_n, v_0\}$. Then $|K|$ is a 1-dimensional polyhedron homeomorphic to S^1.

EXAMPLE 30.7. Taking $V = \mathbb{Z}$ with S as the set of all 0-simplices $\{n\}$ plus 1-simplices of the form $\{n, n+1\}$ for $n \in \mathbb{Z}$ gives an infinite (but 1-dimensional) simplicial complex whose polyhedron is homeomorphic to $\mathbb{R}$.

EXAMPLE 30.8. If $V = \mathbb{N}$ and S is the set of all finite subsets of $\mathbb{N}$, then K is an infinite-dimensional simplicial complex. Every simplex in this complex is a face of $\{1, \ldots, n\}$ for n sufficiently large, thus you can try to picture $|K|$ as the union of an infinite nested sequence of simplices $\Delta^0 \subset \Delta^1 \subset \Delta^2 \subset \ldots$, where each Δ^k is a boundary face of Δ^{k+1}.

A **subcomplex** (*Unterkomplex* or *Teilkomplex*) of a simplicial complex $K = (V, S)$ is a simplicial complex $K' = (V', S')$ such that $V' \subset V$ and $S' \subset S$. We then call (K, K') a **simplicial pair** (*simpliziales Paar*). The polyhedron $|K'|$ can be regarded naturally as a subspace of $|K|$ via the obvious inclusion $I^{V'} \hookrightarrow I^V$ that sets all coordinates t_v for $v \in V \backslash V'$ to zero. In this way, every simplicial pair (K, K') gives rise to a pair of spaces $(|K|, |K'|)$. Note that the empty set also defines a simplicial complex (whose polyhedron is empty), thus every complex K can be identified with the simplicial pair $(K, \varnothing)$.

DEFINITION 30.9. Given two simplicial complexes $K_1 = (V_1, S_1)$ and $K_2 = (V_2, S_2)$, a **simplicial map** (*simpliziale Abbildung*) from K_1 to K_2 is a function $f : V_1 \to V_2$ such that $f(\sigma) \in S_2$ for every $\sigma \in S_1$. A map of simplicial pairs $(K_1, K_1') \to (K_2, K_2')$ is then a simplicial map $K_1 \to K_2$ that restricts to a simplicial map $K_1' \to K_2'$.

Note that a simplicial map $K_1 \to K_2$ need not be injective on any given simplex, i.e. it can send an n-simplex of K_1 onto a k-simplex of K_2 for any $k \leqslant n$. There is a natural way to turn any simplicial map into a continuous map of the polyhedra $|K_1| \to |K_2|$. Indeed, denote by $\{e_v\}_{v \in V}$ the natural basis vectors of $\mathbb{R}^V$ so that every element $t \in \mathbb{R}^V$ can be written uniquely as a formal[51] sum $\sum_{v \in V} t_v e_v$ with coordinates $t_v \in \mathbb{R}$. Then since every element $t \in |K_1|$ is of the form $\sum_{v \in V_1} t_v e_v$ where only finitely many of the coordinates are nonzero and they all add up to 1, we can define

$$f : |K_1| \to |K_2| : \sum_{v \in V_1} t_v e_v \mapsto \sum_{v \in V_1} t_v e_{f(v)} \in I^{V_2}.$$

In other words, for each simplex $\sigma \in S_1$, f maps $|\sigma|$ onto $|f(\sigma)|$ via the restriction of the obvious linear map $\mathbb{R}^\sigma \to \mathbb{R}^{f(\sigma)}$ that sends basis vectors e_v to $e_{f(v)}$ for $v \in \sigma$. If $f : (K_1, K_1') \to (K_2, K_2')$ is a map of simplicial pairs, then it induces in this way a continuous map of pairs $(|K_1|, |K_1'|) \to (|K_2|, |K_2'|)$. We have thus defined a functor

$$\mathsf{Simp}_{\mathrm{rel}} \to \mathsf{Top}_{\mathrm{rel}} : (K, K') \mapsto (|K|, |K'|),$$

where $\mathsf{Simp}_{\mathrm{rel}}$ is the category of simplicial pairs with morphisms defined to be maps of simplicial pairs. Notice that $f : |K_1| \to |K_2|$ always maps the n-skeleton of $|K_1|$ into the n-skeleton of $|K_2|$ for every $n \geqslant 0$.

Since we will mainly be concerned with compact manifolds, the following result enables us to restrict attention to finite simplicial complexes:

PROPOSITION 30.10. *A simplicial complex $K = (V, S)$ is finite if and only if its polyhedron $|K|$ is compact.*

[51] The word "formal" means in this context that we do not require the sum to converge in any sense, as it is a purely algebraic object. In practice, we are only going to consider points $t \in \mathbb{R}^V$ that have finitely many nonzero coordinates, thus the sums converge trivially.

This will follow from a more general theorem about CW-complexes that we shall prove in a few weeks, so for now, we'll settle for proving a special case which happens to cover most of the interesting examples, and is quite easy:

Proof of Proposition 30.10 for finite-dimensional complexes. If K is finite, then $|K|$ is a closed and bounded subset of the finite-dimensional vector space $\mathbb{R}^V$, and is therefore compact.

Conversely, if K is infinite but $\dim K < \infty$, there exists an infinite sequence of distinct simplices $\sigma_1, \sigma_2, \ldots \in S$ with the property that each σ_i is not a face of any other simplex in K. Now for each $i \in \mathbb{N}$, pick a point $x_i \in |\sigma_i|$ along with an open neighborhood $\mathcal{U}_i \subset |\sigma_i|$ of x_i that is contained in the interior of $|\sigma_i|$. Since σ_i is not a face of any other simplex, we have $\mathcal{U}_i \cap |\sigma| = \varnothing$ for all simplices $\sigma \neq \sigma_i$, thus $\mathcal{U}_i$ defines an open subset of $|K|$ that contains x_i but none of the other points in the sequence $x_1, x_2, \ldots$. This proves that the infinite subset $\{x_1, x_2, \ldots\} \subset |K|$ is discrete, hence $|K|$ cannot be compact. □

Simplicial homology. We now describe two versions of the so-called **simplicial chain complex** (*simplizialer Kettenkomplex*) and **simplicial homology** (*simpliziale Homologie*) that can be associated to every simplicial complex or simplicial pair. The first version is algebraically simpler than the second, while the second is easier to interpret geometrically; in practice, one can freely choose to use one or the other, because they end up (for slightly nontrivial reasons) being equivalent at the level of homology. For those who saw a definition of simplicial homology in last semester's *Topologie I* course (cf. Lecture 21): the second complex defined below is cosmetically different from the one that was defined there, but is easily seen to be isomorphic to it (see Remark 30.13). The main difference is that our previous definition required fixing an arbitrary choice of orientation for each simplex, and the definition below avoids making any such choices.

Convention. As we have already done with singular homology, for simplicial homology we shall fix an arbitrary choice of abelian coefficient group G, and omit G from the notation in most situations where this choice does not matter.

Given a simplicial complex $K = (V, S)$ let

$$\mathcal{K}_n^o(K) := \left\{ (v_0, \ldots, v_n) \in V^{\times(n+1)} \;\middle|\; \text{there exists a } \sigma \in S \text{ with } v_i \in \sigma \text{ for all } i = 0, \ldots, n \right\}$$

for each $n \geqslant 0$. The elements of $\mathcal{K}_n^o(K)$ are thus ordered $(n+1)$-tuples of vertices such that some simplex of the complex contains all of them. Note that in this definition, we are *not* assuming the $v_0, \ldots, v_n$ are all distinct, though if they are, then it means $\{v_0, \ldots, v_n\} \in S$ is an n-simplex of the complex K, and the ordered tuple $(v_0, \ldots, v_n)$ is then called an **ordered n-simplex**. The **ordered simplicial chain complex** (with coefficients in G)

$$C_*^o(K) = C_*^o(K; G) = \bigoplus_{n \in \mathbb{Z}} C_n^o(K; G) = \bigoplus_{n \in \mathbb{Z}} C_n^o(K)$$

is defined in the same way as the singular chain complex, except that instead of using singular n-simplices in a topological space as generators of the n-chain group for each $n \geqslant 0$, we use the elements of $\mathcal{K}_n^o(K)$, i.e.

$$C_n^o(K) = \bigoplus_{\sigma \in \mathcal{K}_n^o(K)} G,$$

so that n-chains can be written as finite sums $\sum_i a_i \sigma_i$ with $a_i \in G$ and $\sigma_i \in \mathcal{K}_n^o(K)$. A boundary operator $\partial : C_n^o(K) \to C_{n-1}^o(K)$ on this complex is determined for each $n \geqslant 1$ by the formula

$$\partial(v_0, \ldots, v_n) := \sum_{k=0}^{n} (-1)^k (v_0, \ldots, v_{k-1}, v_{k+1}, \ldots, v_n), \tag{30.1}$$

and ∂ is trivial for $n \leqslant 0$ due to the convention of defining $C_n^o(K) := 0$ for $n < 0$. We will refer to the resulting homology groups

$$H_*^o(K) := H_*^o(K; G) := H_*\big(C_*^o(K; G), \partial\big)$$

as the **ordered simplicial homology** of K with coefficients in G. There is similarly a relative complex

$$C_*^o(K, K') := C_*^o(K) / C_*^o(K')$$

associated to any simplicial pair (K, K'), giving rise to the ordered relative simplicial homology

$$H_*^o(K, K') := H_*\big(C_*^o(K, K'), \partial\big),$$

of which we can view $H_*^o(K)$ as a special case by identifying K with the simplicial pair $(K, \varnothing)$. In order to view ordered simplicial homology as a functor, we associate to each simplicial map $f : K_1 \to K_2$ the unique chain map

$$f_* : C_*^o(K_1) \to C_*^o(K_2)$$

whose definition on the generators $(v_0, \ldots, v_n) \in \mathcal{K}_n^o(K_1)$ is given by

$$f_*(v_0, \ldots, v_n) := (f(v_0), \ldots, f(v_n)).$$

This also descends to a chain map of relative complexes $C_*^o(K_1, K_1') \to C_*^o(K_2, K_2')$ if $f : (K_1, K_1') \to (K_2, K_2')$ is a map of simplicial pairs, so we obtain induced homomorphisms

$$f_* : H_n^o(K_1, K_1') \to H_n^o(K_2, K_2')$$

for every n, making ordered simplicial homology into a functor

$$H_*^o : \mathsf{Simp}_{\mathrm{rel}} \to \mathsf{Ab}_{\mathbb{Z}}.$$

The second version of the simplicial chain complex has a similar but smaller set of generators, because it excludes tuples $(v_0, \ldots, v_n)$ that contain repeats of the same vertex, and instead of keeping track of their orders, it keeps track of orientations. The following combinatorial exercise makes this possible:

Exercise 30.11. Show that for each $n \geqslant 1$, the boundary operator $\partial : C_n^o(K; \mathbb{Z}) \to C_{n-1}^o(K; \mathbb{Z})$ defined via (30.1) preserves the subgroup of $C_*^o(K; \mathbb{Z})$ generated by all elements of the form

$$\big(v_0, \ldots, v_n\big) \in C_n^o(K; \mathbb{Z}) \qquad \text{with } v_i = v_j \text{ for some } i \neq j \tag{30.2}$$

or of the form

$$\big(v_0, \ldots, v_n\big) - (-1)^{|\tau|}\big(v_{\tau(0)}, \ldots, v_{\tau(n)}\big) \in C_n^o(K; \mathbb{Z}) \tag{30.3}$$

for arbitrary $(v_0, \ldots, v_n) \in \mathcal{K}_n^o(K)$ and permutations $\tau \in S_{n+1}$, where $(-1)^{|\tau|} = \pm 1$ denotes the sign of the permutation.
Hint: One only really needs to check cases where $v_k = v_{k+1}$ for some k, and permutations that interchange two neighboring elements.

Definition 30.12. An **orientation** (*Orientierung*) of an n-simplex $\sigma \in S$ for $n \geqslant 1$ in a complex $K = (V, S)$ is an equivalence class of orderings of the vertices of σ, where two orderings are considered equivalent if they differ by an even permutation. The case $n = 0$ is special: an orientation of a 0-simplex is simply a choice of sign $+1$ or -1, called the **positive** or **negative orientation** respectively.

A simplex endowed with an orientation is called an **oriented simplex** (*orientiertes Simplex*), and any oriented simplex with vertices $v_0, \ldots, v_n$ can be written with the notation

$$\pm[v_0, \ldots, v_n],$$

which is understood to mean the simplex $\{v_0, \ldots, v_n\}$ with orientation determined by the ordering $v_0, \ldots, v_n$ if the sign in front is positive, and the opposite of that orientation if the sign is negative. So for example, the symbols $[v_0, v_1]$ and $-[v_1, v_0]$ represent the same oriented 1-simplex, while that simplex with the opposite orientation can be written as either $-[v_0, v_1]$ or $[v_1, v_0]$. For an oriented 0-simplex $\pm[v_0]$, there is only one possible ordering, and the orientation is thus determined entirely by the initial sign. For a 2-simplex $\{v_0, v_1, v_2\}$, the fact that cyclic permutations of three elements are always even means

$$[v_0, v_1, v_2] = [v_1, v_2, v_0] = [v_2, v_0, v_1] = -[v_1, v_0, v_2] = -[v_0, v_2, v_1] = -[v_2, v_1, v_0].$$

In pictures of 2-dimensional polyhedra, one can usefully employ arrows on 1-simplices to specify orientations by ordering the two vertices, and circular arrows in 2-simplices to indicate the cyclic orderings that determine their orientations (see Figure 15).

Thanks to Exercise 30.11, the **oriented simplicial chain complex** of $K = (V, S)$ can be defined as a quotient

$$C^{\Delta}_*(K) := C^o_*(K)/D^o_*(K),$$

where we denote by $D^o_*(K) \subset C^o_*(K)$ the subgroup generated products of arbitrary coefficients $g \in G$ with elements of the form (30.2) or (30.3); this is sometimes called the group of **degenerate chains**. For each generator $(v_0, \ldots, v_n)$ of $C^o_n(K; \mathbb{Z})$, we shall denote the equivalence class that it represents in the quotient complex by

$$[v_0, \ldots, v_n] \in C^{\Delta}_n(K; \mathbb{Z}),$$

which means

$$[v_0, \ldots, v_n] = 0 \qquad \text{if } v_i = v_j \text{ for some } i \neq j,$$

whereas if the vertices $v_0, \ldots, v_n$ are all distinct, $[v_0, \ldots, v_n]$ can be interpreted as an oriented n-simplex, and the equivalence relation in the quotient complex then reproduces our previous notational convention for oriented simplices, namely

$$[v_0, \ldots, v_i, \ldots, v_j, \ldots, v_n] = -[v_0, \ldots, v_j, \ldots, v_i, \ldots, v_n]$$

for each pair $i \neq j$ in $\{0, \ldots, n\}$. The boundary operator $\partial : C^{\Delta}_n(K) \to C^{\Delta}_{n-1}(K)$ is thus determined by the formula

$$\partial[v_0, \ldots, v_n] = \sum_{k=0}^{n} (-1)^k [v_0, \ldots, v_{k-1}, v_{k+1}, \ldots, v_n],$$

and Exercise 30.11 guarantees that this formula is independent of the order in which the vertices are written. This definition extends in the obvious way to an oriented relative simplicial chain complex $C^{\Delta}_*(K, K')$ for any simplicial pair (K, K'), and we will denote the resulting **oriented simplicial homology** by

$$H^{\Delta}_*(K, K') := H_*\big(C^{\Delta}_*(K, K')\big).$$

For a simplicial map $f : K_1 \to K_2$, it is easy to check that the chain map $f_* : C^o_*(K_1) \to C^o_*(K_2)$ descends to the quotient, and thus determines a chain map $f_* : C^{\Delta}_*(K_1) \to C^{\Delta}_*(K_2)$. In the notation of oriented n-simplices, the latter is determined by the formula

$$f_*[v_0, \ldots, v_n] = [f(v_0), \ldots, f(v_n)],$$

in which we should keep in mind that the right hand side is not always an oriented n-simplex, but will instead be 0 if the vertices $f(v_0), \ldots, f(v_n)$ are not all distinct, which happens for instance if f maps an n-simplex $\{v_0, \ldots, v_n\}$ to a simplex of strictly smaller dimension. With this understood, the chain map $f_* : C^\Delta_*(K_1) \to C^\Delta_*(K_2)$ also descends to $C^\Delta_*(K_1, K_1') \to C^\Delta_*(K_2, K_2')$ whenever $f : (K_1, K_1') \to (K_2, K_2')$ is a map of simplicial pairs, and the induced maps

$$f_* : H^\Delta_*(K_1, K_1') \to H^\Delta_*(K_2, K_2')$$

make the oriented simplicial homology into another functor

$$H^\Delta_* : \mathsf{Simp}_{\mathrm{rel}} \to \mathsf{Ab}_{\mathbb{Z}}.$$

REMARK 30.13. While it is not so obvious from the definition above, $C^\Delta_n(K;\mathbb{Z})$ for a simplicial complex K can in fact be identified with the free abelian group generated by the set of all n-simplices in the complex K, though this identification requires a choice. Indeed, if we arbitrarily fix an orientation for each n-simplex in K, then every element of $C^\Delta_n(K;\mathbb{Z})$ has a unique representation as a linear combination of n-simplices with those chosen orientations. Presenting $C^\Delta_*(K;\mathbb{Z})$ in this way then requires inserting some extra signs into the formula for $\partial : C^\Delta_n(K) \to C^\Delta_{n-1}(K)$, to account for the fact that the specific oriented $(n-1)$-simplices in our usual formula for $\partial[v_0, \dots, v_n]$ might appear with different orientations than the ones we have arbitrarily chosen. This is essentially the definition of $H^\Delta_*(K)$ that we gave in Lecture 21 last semester, and it is also the description that typically seems most convenient for actual computations of simplicial homology. The alternative formulation as a quotient complex shows why it does not actually depend on the choices of orientations.

Since $C^\Delta_*(K, K')$ is a quotient of $C^o_*(K, K')$, the quotient projection

$$C^o_*(K, K') \to C^\Delta_*(K, K') \tag{30.4}$$

is a chain map, uniquely determined by the formula

$$(v_1, \dots, v_n) \mapsto [v_1, \dots, v_n],$$

and this chain map is also natural in the technical sense of that term, meaning it defines a natural transformation from the functor $C^o_* : \mathsf{Simp}_{\mathrm{rel}} \to \mathsf{Chain}$ to the functor $C^\Delta_* : \mathsf{Simp}_{\mathrm{rel}} \to \mathsf{Chain}$. After composing both with $H_* : \mathsf{Chain} \to \mathsf{Ab}_{\mathbb{Z}}$, it therefore also determines a natural transformation from $H^o_* : \mathsf{Simp}_{\mathrm{rel}} \to \mathsf{Ab}_{\mathbb{Z}}$ to $H^\Delta_* : \mathsf{Simp}_{\mathrm{rel}} \to \mathsf{Ab}_{\mathbb{Z}}$, meaning that for every simplicial pair (K, K') and every $n \in \mathbb{Z}$, we have a natural homomorphism

$$H^o_n(K, K') \to H^\Delta_n(K, K'). \tag{30.5}$$

Here is the reason why the ordered and oriented variants of simplicial homology can, in practice, be used interchangeably:

THEOREM 30.14. *For every simplicial pair (K, K') and every choice of coefficient group, the natural chain map (30.4) is a chain homotopy equivalence. In particular, the induced maps (30.5) from ordered to oriented simplicial homology groups are all isomorphisms.*

The easiest proof I know for this theorem uses the method of acyclic models, which we will cover later in this course, and at that point the proof will become an interesting exercise that is not obvious but not especially difficult (see Exercise 45.9). I will postpone it until then, since for the near future, I have no intention of actually *using* this theorem for anything important—my purpose in stating it was simply to give you the correct intuition about what simplicial homology is. The oriented chain complex has a more obvious geometric interpretation, because its generators are in bijective correspondence with actual simplices (see Remark 30.13), which makes it extremely convenient for computations, as sketched for instance in Figure 15. By contrast, the ordered chain complex has a lot of redundant information, since each simplex gives rise to several generators corresponding to the different possible orderings of its vertices. But as we will see below, the *ordered* complex is the one that admits a straightforward relationship with the singular homology of a

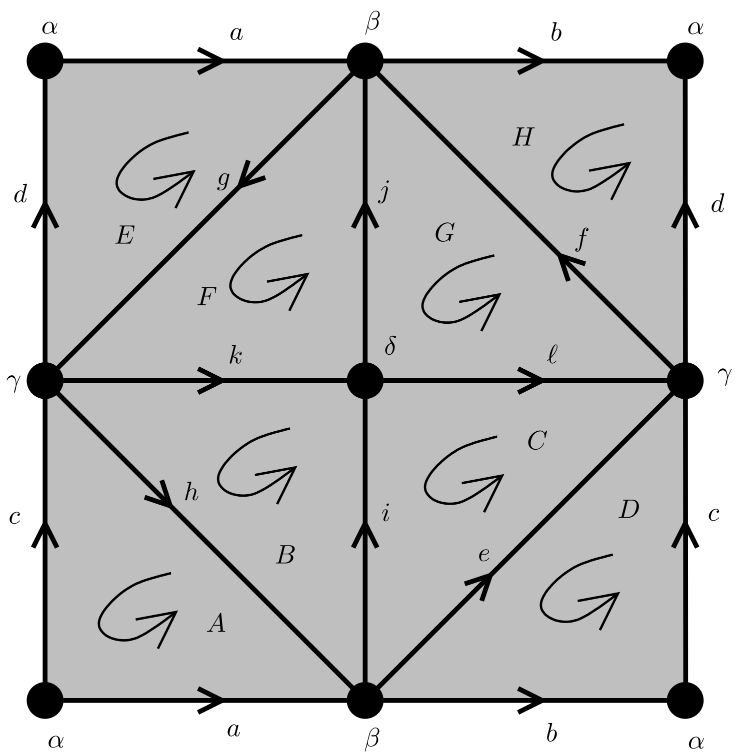

FIGURE 15. The picture shows a simplicial complex K with $|K| \cong \mathbb{T}^2$, and choices of orientations on each simplex indicated via arrows (defining cyclic orderings of three vertices in the case of each 2-simplex). With these orientations fixed, plugging in the definition of $\partial : C_n^\Delta(K;\mathbb{Z}) \to C_{n-1}^\Delta(K;\mathbb{Z})$ gives e.g. $\partial A = a - h - c$, $\partial B = h + i - k$, $\partial a = \beta - \alpha$, $\partial b = \alpha - \beta$ and so forth. The complete computation of $H_*^\Delta(K;\mathbb{Z})$ was carried out near the end of Lecture 21 last semester, with $H_2^\Delta(K;\mathbb{Z}) \cong \mathbb{Z}$ generated by the sum of the eight 2-simplices in the complex, $H_1^\Delta(K;\mathbb{Z}) \cong \mathbb{Z}^2 \cong \pi_1(\mathbb{T}^2) \cong H_1(\mathbb{T}^2;\mathbb{Z})$, and $H_0^\Delta(K;\mathbb{Z}) \cong \mathbb{Z} \cong H_0(\mathbb{T}^2;\mathbb{Z})$.

polyhedron, and while we will not immediately need to know this, it then follows via Theorem 30.14 that there is also a nice relationship between $H_*^\Delta(K, K')$ and $H_*(|K|, |K'|)$ for every simplicial pair (K, K').

There is a canonical chain map

$$C_*^o(K) \to C_*(|K|)$$

for every simplicial complex $K = (V, S)$, due to the following observation: for every ordered tuple $(v_0, \ldots, v_n)$ of vertices in K that all belong to a single simplex $\sigma \in S$, the unique linear map $\mathbb{R}^{n+1} \to \mathbb{R}^V$ that sends the standard basis vectors $e_0, \ldots, e_n \in \mathbb{R}^{n+1}$ to $v_0, \ldots, v_n$ in the same order restricts to a continuous map

$$\Delta^n \to |\sigma| \subset |K| \subset \mathbb{R}^V.$$

Interpreting this map as a singular n-simplex in $|K|$, this associates a generator of $C_*(|K|)$ to each generator of $C_*^o(K)$, and it is straightforward to check that the resulting homomorphism

$C_*^o(K) \to C_*(|K|)$ is a chain map. For a simplicial pair (K, K'), this chain map clearly also descends to a chain map $C_*^o(K, K') \to C_*(|K|, |K'|)$, and therefore induces homomorphisms

$$H_n^o(K, K') \to H_n(|K|, |K'|) \tag{30.6}$$

for every $n \in \mathbb{Z}$. This too can be regarded as a natural transformation, namely from the functor $H_*^o : \mathsf{Simp}_{\mathrm{rel}} \to \mathsf{Ab}_{\mathbb{Z}}$ to the composition of the two functors $\mathsf{Simp}_{\mathrm{rel}} \to \mathsf{Top}_{\mathrm{rel}} : (K, K') \mapsto (|K|, |K'|)$ and $H_* : \mathsf{Top}_{\mathrm{rel}} \to \mathsf{Ab}_{\mathbb{Z}}$. For our purposes in the present lecture, we only need to know how the map (30.6) is defined, but as a quick preview of things to come, we can at least state the most important theorem about it:

THEOREM 30.15. *The natural map $H_*^o(K, K') \to H_*(|K|, |K'|)$ described in (30.6) is an isomorphism for every simplicial pair (K, K') and every choice of coefficient group.*

This isomorphism will follow in a few weeks from a more general theorem about the homology of CW-complexes, which tells us that simplicial and cell complexes provide a practical combinatorial algorithm for computing the singular homology of "reasonable" spaces. Thanks to Theorem 30.14, one also obtains from this an isomorphism

$$H_*^\Delta(K, K') \cong H_*(|K|, |K'|),$$

obtained by inverting the natural isomorphism $H_*^o(K, K') \to H_*^\Delta(K, K')$ and then composing that with (30.6).

Triangulated manifolds. We would now like to discuss how simplicial complexes give rise to geometrically meaningful homology classes on manifolds. The focus for now will be on homology with $\mathbb{Z}_2$ coefficients; the next lecture will then introduce oriented triangulations in order to extend the discussion to coefficients in $\mathbb{Z}$.

DEFINITION 30.16. A **triangulation** (*Triangulierung*) of a pair of spaces (X, A) is a homeomorphism of $(|K|, |K'|)$ to (X, A) for some simplicial pair (K, K'). If M is a manifold with boundary, then a triangulation of M will be understood to mean a triangulation of the pair $(M, \partial M)$.

REMARK 30.17. Definition 30.16 is perhaps stricter than some other sensible definitions of the term "triangulation" that one could imagine. What everyone can agree upon is that a triangulation makes X homeomorphic to the polyhedron $|K|$ of some simplicial complex K in such a way that each simplex of the complex corresponds to a specific subset of X. It is debateable however whether one should really include a specific choice of homeomorphism $|K| \to X$ as part of the data defining a triangulation, e.g. one might want to regard two homeomorphisms $f_1, f_2 : |K| \to X$ as defining the same triangulation of X if $f_1(|\sigma|) = f_2(|\sigma|)$ for every simplex σ of K. For our purposes in the present lecture, it will be convenient not to do this, because the extra data of a specific homeomorphism $|K| \to X$ gives X some extra structure that is useful for writing down singular chains in X.

Given a triangulated space X, we shall often identify X with the associated polyhedron and refer accordingly to the various *skeleta* of X and its constituent simplices. If X is an n-manifold, then its local structure produces the following important observation:

PROPOSITION 30.18. *If M is a triangulated n-dimensional manifold with boundary, then the associated simplicial complex is n-dimensional, and every $(n-1)$-simplex σ in the complex is a boundary face of exactly either one or two n-simplices, where the former is the case if and only if σ belongs to the subcomplex triangulating ∂M.* □

In general, it is a subtle question whether a given manifold admits a triangulation. It is known to be true for all *smooth* manifolds, and also for topological manifolds of dimension at most three (see [Moi77]), but not in general for dimensions four and above (see [Man14]). We will not concern ourselves with such questions here, as for our purposes, it is already helpful to consider explicit examples of manifolds with triangulations, such as the picture of $\mathbb{T}^2$ in Figure 15. Our immediate motivation for doing so is to give explicit constructions of some important homology classes. The idea is to turn a triangulation $M \cong |K|$ of an n-manifold into an n-chain in the ordered simplicial complex, which gives rise via the natural chain map $C^o_*(K) \to C_*(M)$ to a singular n-chain in M which, if we do this correctly, will also be a relative n-cycle and thus represent a homology class. Turning the simplices of a triangulation into generators of $C^o_n(K)$ requires choosing a bit of extra data:

Definition 30.19. An **admissible ordering** on a simplicial complex $K = (V, S)$ assigns to each simplex $\sigma \in S$ a total order on its set of vertices such that the inclusion $\tau \hookrightarrow \sigma$ of each of its faces $\tau \subset \sigma$ is an order-preserving map.

It is easy to see that every simplicial complex admits an admissible ordering, e.g. one can simply choose a total order on the entire set of vertices V and define the total orders on every simplex $\sigma \subset V$ so that the inclusion $\sigma \hookrightarrow V$ is order preserving. In the following, we shall only be concerned with finite complexes, so you don't even need to appeal to any abstract set-theoretic machinery (e.g. the axiom of choice) before choosing a total order on V. There are also situations where establishing a rule to determine total orders on every simplex $\sigma \in S$ is more convenient than choosing a total order on V itself.

Suppose now that M is a compact n-manifold triangulated by a simplicial pair (K, K'), i.e. we are given a homeomorphism $(M, \partial M) \cong (|K|, |K'|)$. Compactness implies via Proposition 30.10 that the simplicial complex K is finite. Working with $\mathbb{Z}_2$ coefficients, any choice of admissible ordering for K determines an ordered simplicial n-chain of the form

$$c_M := \sum_\sigma \mathbf{v}_\sigma \in C^o_n(K; \mathbb{Z}_2), \tag{30.7}$$

where the sum ranges over the set of all n-simplices σ of K, and $\mathbf{v}_\sigma = (v_0, \dots, v_n)$ denotes the vertices of $\sigma = \{v_0, \dots, v_n\}$ arranged in increasing order. If $\partial M \neq \varnothing$, then the subcomplex $K' \subset K$ in this situation also defines a triangulation of the closed $(n-1)$-manifold ∂M, and the admissible order on K restricts to an admissible order on K', and thus similarly determines a simplicial $(n-1)$-chain

$$c_{\partial M} \in C^o_{n-1}(K'; \mathbb{Z}_2) \subset C_{n-1}(K; \mathbb{Z}_2).$$

If $\partial M = \varnothing$, then the recipe above defines the trivial $(n-1)$-chain, and we can therefore sensibly write

$$c_{\partial M} = 0 \in C^o_{n-1}(K; \mathbb{Z}_2) \qquad \text{if } \partial M = \varnothing.$$

Lemma 30.20. $\partial c_M = c_{\partial M}$.

Proof. By Proposition 30.18, applying ∂ to the right hand side of (30.7) produces exactly two copies of each $(n-1)$-simplex of K that is not in K', both occurring with the same ordering of the vertices, so with $\mathbb{Z}_2$ coefficients they cancel each other. What remains is a single term for each $(n-1)$-simplex in the triangulation of ∂M, which produces $c_{\partial M}$. □

Lemma 30.20 implies that c_M is a relative n-cycle for the simplicial pair (K, K'), so it represents a homology class

$$[M] := [c_M] \in H^o_n(K, K'; \mathbb{Z}_2).$$

We will refer to $[M]$ as a **fundamental class** (*Fundamentalklasse*) for the triangulated manifold M, and we will call c_M a (relative) **fundamental cycle** (*Fundamentalzykel*). If M is a closed manifold, then the subcomplex $K' \subset K$ is empty and Lemma 30.20 becomes $\partial c_M = 0$, so c_M is an absolute cycle in $C_n^o(K;\mathbb{Z}_2)$ and represents an absolute fundamental class $[M] \in H_n^o(K;\mathbb{Z}_2)$. We will use the same notation and terminology for the result of feeding c_M and $[M]$ into the natural chain maps from $C_*^o(K;\mathbb{Z}_2)$ to $C_*^\Delta(K;\mathbb{Z}_2)$ or $C_*(M;\mathbb{Z}_2)$, or the maps that these induce on relative homology: in this way, we obtain a relative fundamental cycle and fundamental class in oriented simplicial homology

$$c_M \in C_n^\Delta(K;\mathbb{Z}_2), \qquad [M] \in H_n^\Delta(K,K';\mathbb{Z}_2),$$

and also in singular homology

$$c_M \in C_n(M;\mathbb{Z}_2), \qquad [M] \in H_n(M,\partial M;\mathbb{Z}_2),$$

and the relation in Lemma 30.20 carries over immediately to both of these settings. In particular, $[M]$ is an *absolute* homology class in $H_n(M;\mathbb{Z}_2)$ whenever $\partial M = \varnothing$.

We will be able to show later in this course that in singular homology, the fundamental class $[M] \in H_n(M,\partial M;\mathbb{Z}_2)$ depends only on the compact n-manifold M, and not on the choice of triangulation—in fact, there is a canonical fundamental class in $H_n(M,\partial M;\mathbb{Z}_2)$ for every compact topological n-manifold, even those that do not admit triangulations. For the moment, we will not worry about whether $[M]$ is independent of the triangulation, but it is worth observing that in both $C_n^o(K;\mathbb{Z}_2)$ and $C_n(M;\mathbb{Z}_2)$, the triangulation on its own does not uniquely determine the fundamental cycle c_M, because the latter also depends on the chosen admissible ordering for K, which determines how each n-simplex of the triangulation gets represented as a generator of $C_n^o(K;\mathbb{Z}_2)$ and therefore also as a map $\Delta^n \to M$. The situation in $C_n^\Delta(K;\mathbb{Z}_2)$ is nicer, because the map $C_n^o(K;\mathbb{Z}_2) \to C_n^\Delta(K;\mathbb{Z}_2)$ eliminates the distinction between different orderings, and thus produces a fundamental cycle that depends *only* on the triangulation. In all three settings, the following theorem tells us at least that the fundamental *class* depends at most on the triangulation, and not on any additional choices.

THEOREM 30.21. *For any compact n-manifold M with a triangulation $(M,\partial M) \cong (|K|,|K'|)$ given by a simplicial pair (K,K'), the relative fundamental class $[M] \in H^o(K,K';\mathbb{Z}_2)$ in ordered simplicial homology is independent of choices, and it is nontrivial. Moreover, its image under the canonical map to singular homology is also a nontrivial element $[M] \in H(M,\partial M;\mathbb{Z}_2)$, and the image of the latter under the connecting homomorphism $\partial_* : H_n(M,\partial M;\mathbb{Z}_2) \to H_{n-1}(\partial M;\mathbb{Z}_2)$ for the long exact sequence of the pair $(M,\partial M)$ is*

$$\partial_*[M] = [\partial M] \in H_{n-1}(\partial M;\mathbb{Z}_2),$$

where $[\partial M]$ denotes the fundamental class determined by the triangulation of ∂M by the subcomplex $K' \subset K$.

PARTIAL PROOF. Not every detail in this statement can easily be proved without assuming Theorem 30.14 on the isomorphism $H_*^o(K,K') \cong H_*^\Delta(K,K')$, but let us first say what can be said without that result. The canonical chain map (30.4) sends the relative fundamental cycle $c_M \in C_n^o(K;\mathbb{Z}_2)$ to a relative cycle in the oriented chain complex $C_*^\Delta(K;\mathbb{Z}_2)$, given by

$$\sum_\sigma [\mathbf{v}_\sigma] \in C_n^\Delta(K;\mathbb{Z}_2),$$

where for each n-simplex σ of K, the ordered n-simplex $\mathbf{v}_\sigma = (v_0,\dots,v_n)$ becomes the oriented n-simplex $[\mathbf{v}_\sigma] := [v_0,\dots,v_n]$. This cycle cannot be a boundary in $C_*^\Delta(K,K';\mathbb{Z}_2)$ for a very simple reason: since K has no $(n+1)$-simplices, the generators $[v_0,\dots,v_{n+1}]$ of $C_{n+1}^\Delta(K)$ are all trivial; equivalently, the entirety of $C_{n+1}^o(K)$ belongs to the subcomplex $D_*^o(K)$ of degenerate chains, thus $C_{n+1}^\Delta(K) = 0$ and every nontrivial cycle in $C_n^\Delta(K,K')$ therefore represents a nontrivial homology class. It follows that $[M] \in H_n^o(K,K';\mathbb{Z}_2)$ also cannot be trivial.

In light of the explicit formula (28.6) for connecting homomorphisms in singular homology, the relation $\partial_*[M] = [\partial M]$ follows directly from the formula $\partial c_M = c_{\partial M}$ in Lemma 30.20. If we were willing to take Theorem 30.15 for now as a black box and believe that the natural map $H_n^o(K, K'; \mathbb{Z}_2) \to H_n(M, \partial M; \mathbb{Z}_2)$ is an isomorphicm, it would follow immediately that $[M] \neq 0 \in H_n(M, \partial M; \mathbb{Z}_2)$, but the latter can also be proved in easier ways. We'll defer this detail for now, but come back to it in the next lecture, where an analogous statement will need to be proved with oriented triangulations and integer coefficients.

If you are already content to accept Theorem 30.14 and believe that the natural map from $H_*^o(K, K'; \mathbb{Z}_2)$ to $H_*^\Delta(K, K'; \mathbb{Z}_2)$ is an isomorphism, then the independence of $[M] \in H_n^o(K, K'; \mathbb{Z}_2)$ on choices follows, because any two fundamental cycles $c_M \in C_n^o(K; \mathbb{Z}_2)$ defined via different choices of admissible orderings have the same image under the natural chain map $C_*^o(K; \mathbb{Z}_2) \to C_*^\Delta(K; \mathbb{Z}_2)$. □

The next exercise tells us that for compact and *connected* triangulated manifolds $(M, \partial M) \cong (|K|, |K'|)$, the image of $[M]$ in $H_n^\Delta(K, K'; \mathbb{Z}_2)$ is not just a distinguished homology class, but is in fact the *only* one that is nontrivial.

Exercise 30.22. Show that for any simplicial pair (K, K') whose polyhedron $|K|$ is a compact connected n-manifold with boundary $|K'|$, $H_n^\Delta(K, K'; \mathbb{Z}_2) \cong \mathbb{Z}_2$.

31. Oriented triangulations and fundamental cycles (October 31, 2023)

Oriented triangulations. Our first order of business in this lecture is to extend the construction of fundamental cycles on triangulated manifolds from $\mathbb{Z}_2$ to integer coefficients. This does not work for every manifold—it requires a bit of extra structure.

Definition 31.1. Suppose $n \geqslant 1$ and $\pm[v_0, \dots, v_n]$ is an oriented n-simplex in a simplicial complex. The induced **boundary orientation** on the boundary face $\{v_1, \dots, v_n\}$ is then given by the oriented $(n-1)$-simplex $\pm[v_1, \dots, v_n]$.

Note that the oriented simplex $\pm[v_0, \dots, v_n]$ can typically be written in multiple distinct ways with the vertex v_0 appearing first and the other vertices permuted, but the same permutation then applies to the oriented boundary face $\pm[v_1, \dots, v_n]$ and causes the same sign change, so that Definition 31.1 does not depend on any choices. Moreover, the definition determines an orientation on *every* boundary face of $\sigma = \{v_0, \dots, v_n\}$, because for any $k = 0, \dots, n$, one can always apply a permutation to rewrite $\pm[v_0, \dots, v_n]$ with v_k in front; in particular, $[v_0, \dots, v_n] = (-1)^k[v_k, v_0, \dots, v_{k-1}, v_{k+1}, \dots, v_n]$, so that endowing the face $\{v_0, \dots, v_{k-1}, v_{k+1}, \dots, v_n\}$ with the boundary orientation determined by $[v_0, \dots, v_n]$ produces the oriented simplex

$$(-1)^k[v_0, \dots, v_{k-1}, v_{k+1}, \dots, v_n].$$

The formula for $\partial[v_0, \dots, v_n]$ in the oriented simplicial chain complex can thus be interpreted as the sum of the $n+1$ boundary faces of $[v_0, \dots, v_n]$ endowed with their boundary orientations.

Remark 31.2. In addition to being consistent with our usual formulas for boundary operators on chain complexes, there is some geometric motivation behind Definition 31.1. In differential geometry, an oriented n-manifold M induces a natural boundary orientation on ∂M, and if M has a triangulation, the orientation of M also induces orientations of the n-simplices in its triangulation. One can check that if the polyhedron $|\sigma|$ of an oriented n-simplex σ in a complex K is viewed as an oriented n-manifold, then the geometric notion of boundary orientation on $\partial|\sigma| \cong S^{n-1}$ matches the induced orientations (according to Definition 31.1) of the boundary faces of σ, which form a triangulation of $\partial|\sigma|$.

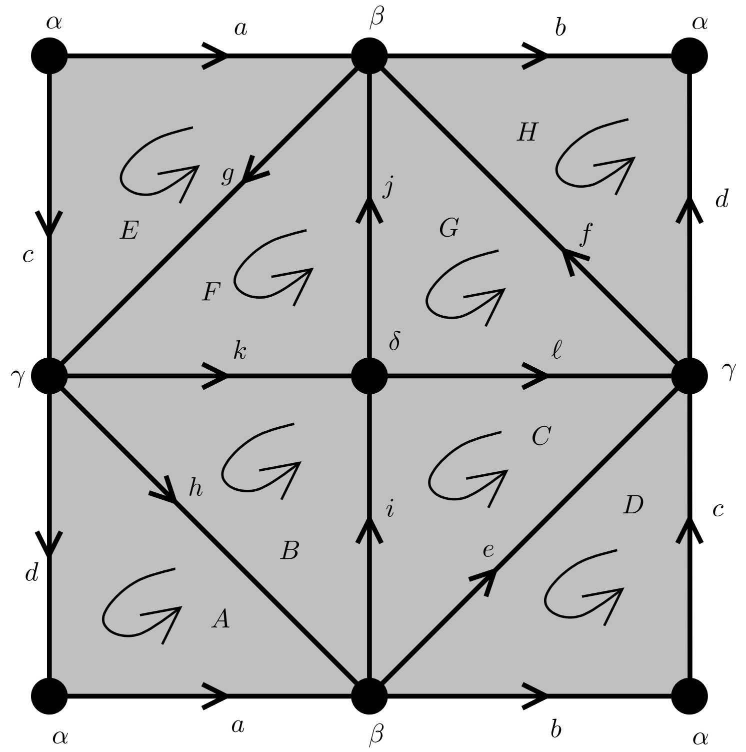

Figure 16. A triangulation of the Klein bottle that fails to be oriented.

Definition 31.3. For an n-dimensional manifold M, an **oriented triangulation** (*orientierte Triangulierung*) of M is a triangulation in which every n-simplex is endowed with an orientation such that for every $(n-1)$-simplex σ not contained in ∂M, the two boundary orientations it inherits as a boundary face of two distinct oriented n-simplices (cf. Prop. 30.18) are opposite.

I recommend now taking another look at Figure 15 to verify that the orientations of 2-simplices depicted in this picture define an oriented triangulation of $\mathbb{T}^2$. Then, contrast it with Figure 16, which shows a triangulation of the Klein bottle in which orientations of the 2-simplices have been chosen but they fail to satisfy the conditions of Definition 31.3. (The trouble is with the 1-simplices labeled c and d.) The problem with the Klein bottle is of course that it is a *non-orientable* manifold, and it turns out that only orientable manifolds can admit oriented triangulations—we sketched a proof of this for surfaces last semester in Lecture 20, and we will be able to prove it for all manifolds later in this course using homology.

Example 31.4. The triangulation of S^{n-1} described in Example 30.5 can be oriented by choosing an ordering of the vertex set V, regarding this as an oriented n-simplex σ and then endowing each of its boundary faces with the boundary orientation. The cancelation condition on $(n-2)$-simplices in this case is roughly equivalent to the fact that $\partial^2 = 0$ in the singular and simplicial chain complexes; see Lemma 31.5 below.

Integral fundamental cycles. We now consider whether a version of the fundamental cycle $c_M \in C_n^o(K; \mathbb{Z}_2)$ for a compact triangulated n-manifold $(M, \partial M) \cong (|K|, |K'|)$ with an admissible ordering of K can also be defined with integer coefficients. A natural way to generalize (30.7) to a simplicial n-chain with integer coefficients would be

$$c_M := \sum_\sigma \epsilon_\sigma \mathbf{v}_\sigma \in C_n^o(K;\mathbb{Z}), \tag{31.1}$$

where the coefficients $\epsilon_\sigma = \pm 1$ are signs that will need to be chosen carefully if we want c_M to be a relative n-cycle. The ability to do this turns out to be equivalent to the orientability of the triangulation. Let us adopt the following convention: in addition to an admissible ordering of K, suppose that an orientation has been chosen for each of the n-simplices σ of K. Independently of whether these chosen orientations satisfy the condition on boundary orientations stated in Definition 31.3, we can then associate to each n-simplex $\sigma = \{v_0, \ldots, v_n\}$ a sign $\epsilon_\sigma = \pm 1$ such that the oriented n-simplex

$$\epsilon_\sigma [v_0, \ldots, v_n]$$

matches the chosen orientation of σ whenever its vertices are arranged in increasing order. In other words, after writing the vertices $v_0, \ldots, v_n$ in accordance with the admissible ordering, ϵ_σ will be $+1$ if the oriented n-simplex $[v_0, \ldots, v_n]$ matches the chosen orientation of σ, and -1 otherwise. In the same manner, the admissible ordering together with an arbitrary choice of sign $\epsilon_\sigma = \pm 1$ for any n-simplex σ of K determines an orientation for σ.

Lemma 31.5. *Suppose $(M, \partial M) \cong (|K|, |K'|)$ is a triangulation of a compact n-manifold with a choice of admissible ordering. Under the correspondence described above between signs $\epsilon_\sigma = \pm 1$ and orientations for the n-simplices σ of K, the n-chain $c_M \in C_n^o(K;\mathbb{Z})$ defined in (31.1) satisfies*

$$\partial c_M \in C_{n-1}^o(K';\mathbb{Z}) \subset C_{n-1}^o(K;\mathbb{Z})$$

if and only if the corresponding orientations of the n-simplices satisfy the conditions of an oriented triangulation.

Proof. The point is simply that by Proposition 30.18, every $(n-1)$-simplex not belonging to K' will appear exactly twice in the computation of ∂c_M, arising from two distinct n-simplices in the sum $\sum_\sigma \epsilon_\sigma \mathbf{v}_\sigma$, with cancelation occurring if and only if those two appear with opposite signs, and the latter is equivalent to the condition on boundary faces in Definition 31.3. □

Lemma 31.5 gives both a criterion to recognize the orientability of a triangulation (such as in Example 31.4) and a recipe for constructing out of any oriented triangulation of a compact n-manifold an integral fundamental cycle $c_M \in C_n^o(K;\mathbb{Z})$.

Proposition 31.6. *Any oriented triangulation $(M, \partial M) \cong (|K|, |K'|)$ on a compact n-manifold M determines an oriented triangulation on ∂M, defined by endowing each $(n-1)$-simplex of K' with its boundary orientation as a boundary face of an oriented n-simplex of K. For any choice of admissible ordering for K and the induced admissible ordering for K', the resulting simplicial chains $c_M \in C_n^o(K;\mathbb{Z})$ and $c_{\partial M} \in C_{n-1}^o(K';\mathbb{Z}) \subset C_{n-1}^o(K';\mathbb{Z})$ satisfy $\partial c_M = c_{\partial M}$.*

Proof. That the prescription described in the statement defines a valid orientation for the triangulation of ∂M follows from Lemma 31.5 after observing that the relation $\partial c_M = c_{\partial M}$ holds by construction, and implies $\partial c_{\partial M} = \partial^2 c_M = 0$. □

The rest of the discussion is completely analogous to what was done for the unoriented case in the previous lecture: we call $c_M \in C_n^o(K;\mathbb{Z})$ a **fundamental cycle** and use it to define a **fundamental class**

$$[M] := [c_M] \in H_n^o(K, K';\mathbb{Z}),$$

which then gives rise to corresponding fundamental cycles/classes in both oriented simplicial homology

$$c_M \in C_n^\Delta(K;\mathbb{Z}), \qquad [M] \in H_n^\Delta(K, K';\mathbb{Z})$$

and singular homology

$$c_M \in C_n(M;\mathbb{Z}), \qquad [M] \in H_n(M, \partial M; \mathbb{Z}).$$

THEOREM 31.7. *For a compact n-manifold M with an oriented triangulation, Theorem 30.21 also holds with the coefficient group $\mathbb{Z}_2$ replaced throughout by $\mathbb{Z}$.*

Local homology. Let us now address a particular detail of the proof of Theorem 30.21 that was deferred in the previous lecture: how do we know that the singular homology class

$$[M] \in H_n(M, \partial M)$$

determined by a triangulation is nontrivial?

This fact can be deduced from the isomorphisms $H_*^\Delta(K, K') \cong H_*^\circ(K, K') \cong H_*(M, \partial M)$ after verifying via direct computation (cf. Exercise 30.22) that $[M]$ is nontrivial in $H_n^\Delta(K, K')$, but at least one of those isomorphisms is a deep theorem that we have not yet proved, and in fact, there is a simpler way to prove $[M] \neq 0 \in H_n(M, \partial M)$ that also contains some useful insights. The idea is to pick a point $x \in M$ lying in the interior of one of the n-simplices of the triangulation, and see what happens to $[M] \in H_n(M, \partial M)$ under the homomorphism

$$H_n(M, \partial M) \to H_n(M, M\backslash\{x\})$$

induced by the inclusion $(M, \partial M) \hookrightarrow (M, M\backslash\{x\})$. One of the reasons this idea works is that the group $H_n(M, M\backslash\{x\})$ can be computed quite explicitly: it turns out to depend not at all on the global structure of M, but only on the fact that M is an n-manifold.

Indeed, forgetting about triangulations for a moment, suppose M is any topological manifold of dimension $n \in \mathbb{N}$, and fix a point $x \in M\backslash\partial M$. We can then find a compact neighborhood $\mathcal{D} \subset M$ of x with a homeomorphism $\varphi : \mathcal{D} \to \mathbb{D}^n \subset \mathbb{R}^n$ sending x to $0 \in \mathbb{D}^n$, so φ also restricts to a homeomorphism $\mathcal{D}\backslash\{x\} \to \mathbb{D}^n\backslash\{0\}$. Using an arbitrary coefficient group for singular homology, we can consider the following string of maps:

(31.2)

$$H_n(M, M\backslash\{x\}) \xleftarrow{i_*} H_n(\mathcal{D}, \mathcal{D}\backslash\{x\}) \xrightarrow{\varphi_*} H_n(\mathbb{D}^n, \mathbb{D}^n\backslash\{0\}) \xleftarrow{j_*} H_n(\mathbb{D}^n, S^{n-1}) \xrightarrow{\partial_*} \widetilde{H}_{n-1}(S^{n-1}),$$

where $i : (\mathcal{D}, \mathcal{D}\backslash\{x\}) \hookrightarrow (M, M\backslash\{x\})$ and $j : (\mathbb{D}^n, \partial S^{n-1}) \hookrightarrow (\mathbb{D}^n, \mathbb{D}^n\backslash\{0\})$ are the obvious inclusions, and ∂_* denotes the connecting homomorphism in the reduced long exact sequence of the pair $(\mathbb{D}^n, S^{n-1})$. We claim that all of these maps are isomorphisms. For φ_* this is immediate since φ is a homeomorphism, and for i_* it follows from excision since $M\backslash\mathcal{D}$ has closure $M\backslash\mathring{\mathcal{D}}$ contained in the open subset $M\backslash\{x\}$.[52] For j_*, you hopefully have some intuition telling you that the reason involves homotopy invariance, though we have to be a bit careful: there is no deformation retraction of pairs from $(\mathbb{D}^n, \mathbb{D}^n\backslash\{0\})$ to $(\mathbb{D}^n, S^{n-1})$, even though there is a deformation retraction from $\mathbb{D}^n\backslash\{0\}$ to S^{n-1}. What is true however is that since the inclusion j defines homotopy equivalences $\mathbb{D}^n \to \mathbb{D}^n$ and $S^{n-1} \hookrightarrow \mathbb{D}^n\backslash\{0\}$, we can put the long exact sequences of $(\mathbb{D}^n, S^{n-1})$ and $(\mathbb{D}^n, \mathbb{D}^n\backslash\{0\})$ together in a commutative diagram

$$\begin{tikzcd}
\dots \arrow[r] & H_n(S^{n-1}) \arrow[r] \arrow[d, "j_*"] & H_n(\mathbb{D}^n) \arrow[r] \arrow[d, "j_*"] & H_n(\mathbb{D}^n, S^{n-1}) \arrow[r] \arrow[d, "j_*"] & H_{n-1}(S^{n-1}) \arrow[r] \arrow[d, "j_*"] & H_{n-1}(\mathbb{D}^n) \arrow[r] \arrow[d, "j_*"] & \dots \\
\dots \arrow[r] & H_n(\mathbb{D}^n\backslash\{0\}) \arrow[r] & H_n(\mathbb{D}^n) \arrow[r] & H_n(\mathbb{D}^n, \mathbb{D}^n\backslash\{0\}) \arrow[r] & H_{n-1}(\mathbb{D}^n\backslash\{0\}) \arrow[r] & H_{n-1}(\mathbb{D}^n) \arrow[r] & \dots
\end{tikzcd}$$

where both rows are exact and all of the five vertical maps except the middle one are already known to be isomorphisms. This establishes the hypotheses for a popular diagram-chasing exercise known as the *five-lemma*:

[52] Notice how we just used the assumption that manifolds are Hausdorff?

EXERCISE 31.8 (the five-lemma). Suppose the following diagram commutes and that both of its rows are exact, meaning $\operatorname{im} f = \ker g$, $\operatorname{im} g' = \ker h'$ and so forth:

$$\begin{tikzcd} A \arrow[r, "f"] \arrow[d, "\alpha"] & B \arrow[r, "g"] \arrow[d, "\beta"] & C \arrow[r, "h"] \arrow[d, "\gamma"] & D \arrow[r, "i"] \arrow[d, "\delta"] & E \arrow[d, "\varepsilon"] \\ A' \arrow[r, "f'"] & B' \arrow[r, "g'"] & C' \arrow[r, "h'"] & D' \arrow[r, "i'"] & E' \end{tikzcd}$$

Prove that if α, β, δ and ε are all isomorphisms, then so is γ.

We conclude via the five-lemma that $j_* : H_n(\mathbb{D}^n, S^{n-1}) \to H_n(\mathbb{D}^n, \mathbb{D}^n\backslash\{0\})$ is an isomorphism, and finally, the reduced exact sequence

$$0 = \widetilde{H}_n(\mathbb{D}^n) \longrightarrow H_n(\mathbb{D}^n, S^{n-1}) \xrightarrow{\partial_*} \widetilde{H}_{n-1}(S^{n-1}) \longrightarrow \widetilde{H}_{n-1}(\mathbb{D}^n) = 0$$

implies that ∂_* is also an isomorphism.

In summary, the maps described above give us a natural isomorphism

$$H_n(M, M\backslash\{x\}; G) \cong \widetilde{H}_{n-1}(S^n; G) \cong G \tag{31.3}$$

for any choice of coefficient group G and any $x \in M\backslash\partial M$. The group $H_n(M, M\backslash\{x\})$ is sometimes referred to as the **local homology** of the n-manifold M at the point x.

We will be able to show later in the semester that the map

$$H_n(M, \partial M; \mathbb{Z}_2) \to H_n(M, M\backslash\{x\}; \mathbb{Z}_2) \cong \mathbb{Z}_2$$

induced by the obvious inclusion is an isomorphism for every compact connected n-manifold M with $x \in M\backslash\partial M$, and the map $H_n(M, \partial M; \mathbb{Z}) \to H_n(M, M\backslash\{x\}; \mathbb{Z})$ is also an isomorphism if M admits an orientation. (One can deduce from the latter that the statement also remains true if one replaces $\mathbb{Z}$ with any coefficient group—there is only a restriction on the coefficients if M is not orientable.) For compact manifolds with triangulations, we will be able to prove this once we have shown that the singular homology of a polyhedron is isomorphic to the simplicial homology of its underlying simplicial complex, due in part to the following extension of Exercise 30.22:

EXERCISE 31.9. Show that for any simplicial pair (K, K') defining an oriented triangulation of a compact connected n-manifold $M \cong |K|$ with boundary $\partial M \cong |K'|$, $H_n^\Delta(K, K'; G) \cong G$ for all choices of coefficient group G.

For now, what we can prove without assuming any big theorems from the future is the following. Recall that for any abelian group G, a nontrivial element $h \in G$ is called **primitive** if it is not mg for some $g \in G$ and integer $m \geqslant 2$.

LEMMA 31.10. *For a compact n-manifold M endowed with an oriented triangulation $(M, \partial M) \cong (|K|, |K'|)$ inducing a fundamental class $[M] \in H_n(M, \partial M; \mathbb{Z})$, the map*

$$H_n(M, \partial M; \mathbb{Z}) \to H_n(M, M\backslash\{x\}; \mathbb{Z})$$

defined via the inclusion $(M, \partial M) \hookrightarrow (M, M\backslash\{x\})$ for any point $x \in M$ in the interior of an n-simplex of the triangulation sends $[M]$ to a primitive element of the local homology group $H_n(M, M\backslash\{x\}; \mathbb{Z}) \cong \mathbb{Z}$. Similarly, for an arbitrary (not necessarily oriented) triangulation and fundamental class $[M] \in H_n(M, \partial M; \mathbb{Z}_2)$, the map $H_n(M, \partial M; \mathbb{Z}_2) \to H_n(M, M\backslash\{x\}; \mathbb{Z}_2)$ sends $[M]$ to the nontrivial element of $H_n(M, M\backslash\{x\}; \mathbb{Z}_2) \cong \mathbb{Z}_2$.

PROOF. We will consider here only the oriented case; the difference in the unoriented case is that $\mathbb{Z}_2$ coefficients must be used in order for c_M to be a cycle, but the argument beyond this single detail is the same.

The first observation is that if we feed the fundamental cycle $c_M = \sum_\sigma \epsilon_\sigma \mathbf{v}_\sigma$ into the chain map $C_n(M, \partial M; \mathbb{Z}) \to C_n(M, M\backslash\{x\}; \mathbb{Z})$ induced by $(M, \partial M) \hookrightarrow (M, M\backslash\{x\})$, then every term in the summation except one can be ignored; we must keep the one corresponding to the specific n-simplex $|\sigma| \subset M$ that contains x, but the rest are n-simplices contained in $M\backslash\{x\}$. The second observation is that $\mathcal{D} := |\sigma| \cong \mathbb{D}^n$ can play the role of the compact neighborhood of x used in defining the natural isomorphisms

$$H_n(M, M\backslash\{x\}; \mathbb{Z}) \cong H_n(\mathcal{D}, \mathcal{D}\backslash\{x\}\mathcal{D}; \mathbb{Z}) \cong H_n(\Delta^n, \partial\Delta^n; \mathbb{Z}) \cong \widetilde{H}_{n-1}(\partial\Delta^n; \mathbb{Z}) \cong \mathbb{Z}$$

described above, with the standard n-simplex Δ^n and its boundary $\partial\Delta^n$ now playing the role of the disk $\mathbb{D}^n$ and sphere S^{n-1} respectively. At the chain level, $\epsilon_\sigma \mathbf{v}_\sigma = \pm \mathbf{v}_\sigma$ lands in $H_n(\Delta^n, \partial\Delta^n; \mathbb{Z})$ as a relative n-cycle of the form $\pm \mathrm{Id} \in C_n(\Delta^n; \mathbb{Z})$ defined via the singular simplex $\mathrm{Id} : \Delta^n \to \Delta^n$. Feeding this into $\partial_* : H_n(\Delta^n, \partial\Delta^n; \mathbb{Z}) \to \widetilde{H}_{n-1}(\partial\Delta^n; \mathbb{Z}) \cong \widetilde{H}_{n-1}(S^{n-1}; \mathbb{Z})$ produces a fundamental cycle corresponding to the natural oriented triangulation of $S^{n-1} \cong \partial\Delta^n$ described in Examples 30.5 and 31.4, thus the map $H_n(M, \partial M; \mathbb{Z}) \to \widetilde{H}_{n-1}(S^{n-1}; \mathbb{Z})$ sends $[M]$ to another fundamental class $[S^{n-1}] \in \widetilde{H}_{n-1}(S^{n-1}; \mathbb{Z}) \subset H_{n-1}(S^{n-1}; \mathbb{Z})$.

We complete the proof via induction on the dimension $n \in \mathbb{N}$. For $n = 1$, we can describe S^0 explicitly as the space of two points $\{1, -1\} \subset \mathbb{R}$, identify singular 0-simplices with points and observe that the class $[S^0] \in H_{n-1}(S^0; \mathbb{Z})$ is represented by a 0-cycle of the form $\pm(\{1\} - \{-1\})$, giving a generator of $\widetilde{H}_{n-1}(S^0; \mathbb{Z})$. If we then assume the result is known for all manifolds up dimension $n-1$, the n-dimensional case follows because we have a homomorphism $H_n(M, \partial M; \mathbb{Z}) \to H_{n-1}(S^{n-1}; \mathbb{Z})$ sending $[M]$ to a primitive element $[S^{n-1}] \in H_{n-1}(S^{n-1}; \mathbb{Z})$, implying that $[M]$ is also primitive. □

Triangulated bordism. If (X, A) is any pair of spaces (not necessarily manifolds), we can now picture fairly general homology classes in $H_k(X, A)$ in terms of continuous maps $f : (M, \partial M) \to (X, A)$, where M is a compact k-manifold with an oriented triangulation: any such map defines a class

$$[f] := f_*[M] \in H_k(X, A; \mathbb{Z}).$$

Moreover, we can allow a wider class of (not necessarily orientable) manifolds if we are only interested in $\mathbb{Z}_2$ coefficients: given any triangulation (oriented or not) of a compact k-manifold M, any continuous map $f : (M, \partial M) \to (X, A)$ defines a class

$$[f] := f_*[M] \in H_k(X, A; \mathbb{Z}_2).$$

In the most interesting applications, X is typically a smooth n-manifold, and we are mainly interested in *absolute* homology classes in $H_k(X)$. An important special case of the above construction is then when $M \subset X$ is a closed *submanifold* of X and we take $f : M \hookrightarrow X$ to be the inclusion: in this way, triangulated closed k-dimensional submanifolds $M \subset X$ define homology classes $[M] \in H_k(X; \mathbb{Z}_2)$, or $[M] \in H_k(X; \mathbb{Z})$ whenever the triangulation is oriented.

REMARK 31.11. We are not claiming that *every* class in $H_*(X; \mathbb{Z})$ or $H_*(X; \mathbb{Z}_2)$ in an arbitrary space X can be represented by a map of a manifold into X. It is in general a rather subtle question which homology classes can be represented in this way, but some useful results on this subject are known: if X is a compact polyhedron, then a famous theorem of Thom [Tho54] states that for any $A \in H_k(X; \mathbb{Z})$, there exists a closed oriented smooth (and therefore triangulable) k-manifold M, a map $f : M \to X$ and a number $m \in \mathbb{N}$ such that $mA = f_*[M]$. This implies in particular that maps from manifolds to X are completely sufficient to describe homology classes in X with *rational* coefficients. For $k \leqslant 6$, it is even possible to represent all classes in $H_k(X; \mathbb{Z})$ in this way, but counterexamples are known for every $k \geqslant 7$.

It is now natural to ask: under what circumstances can we say that two maps $f : M \to X$ and $g : N \to X$ defined on triangulated k-manifolds determine the same homology class? The natural answer to this question makes singular homology look a lot like the *bordism* theory that we sketched last semester in Lecture 21. The details are worked out in the following exercise.

EXERCISE 31.12. In each of the following, (X, A) is a pair of spaces, W is a compact triangulated $(k+1)$-manifold with boundary, M, M_0 and M_1 are compact triangulated k-manifolds with $\partial M_0 = \partial M_1 = \varnothing$, all maps are continuous, and all triangulations are assumed oriented (except in part (e)).

(a) Show that if $M = \partial W$ and $f = F|_M : M \to X$ for some map $F : W \to X$, then $[f] = 0 \in H_k(X; \mathbb{Z})$.
(b) Generalizing part (a), suppose $f = F|_M : M \to X$ for some map $F : W \to X$, but M is a compact subset of ∂W that is both a subcomplex and a k-dimensional submanifold with boundary, such that $F(\overline{\partial W \backslash M}) \subset A$, so in particular $f(\partial M) \subset A$. (See Figure 17.) Show that $[f] = 0 \in H_k(X, A; \mathbb{Z})$.
(c) Given maps $f_0 : M_0 \to X$ and $f_1 : M_1 \to X$, let $f : M_0 \amalg M_1 \to X$ denote the map that restricts to M_i as f_i for $i = 0, 1$. Show that $[f] = [f_0] + [f_1] \in H_k(X; \mathbb{Z})$.
(d) Assuming $\partial M = \varnothing$, show that for any map $f : M \to X$, reversing the orientations of all n-simplices in an oriented triangulation of M reverses the sign of $[f] \in H_k(X; \mathbb{Z})$.
(e) Show that parts (a), (b) and (c) remain valid with $\mathbb{Z}_2$ coefficients if all orientation hypotheses are dropped.

EXERCISE 31.13. Let $\Sigma_{1,2}$ denote the 2-torus with two holes cut out, and suppose $\alpha, \beta : S^1 \hookrightarrow \partial\Sigma_{1,2}$ are loops parametrizing its two boundary components, with α following the boundary orientation of $\partial\Sigma_{1,2}$ and β following the opposite orientation (see Figure 18). Show that if we choose an oriented triangulation of $\Sigma_{1,2}$ so that α and β inherit oriented triangulations (with the orientation on β reversed), then α and β represent the same class in $H_1(\Sigma_{1,2}; \mathbb{Z})$. (One says in this case that α and β are **homologous** (*homolog*). One can show that they are not homotopic.)

Subdivision. The proofs of homotopy invariance and excision in Lecture 29 made use of three specific relative fundamental cycles arising from oriented triangulations: one on Δ^n and two on $I \times \Delta^n$ for each $n \geqslant 0$. The following discussion is intended to spell out the details of those constructions. In the case of homotopy invariance, the required triangulation on $I \times \Delta^n$ is actually a special case of a triangulation that can be defined on $\Delta^k \times \Delta^\ell$ for any pair of integers $k, \ell \geqslant 0$; setting $k = 1$ recovers $I \times \Delta^n$ due to the usual identification of I with Δ^1. Triangulations of $\Delta^k \times \Delta^\ell$ for general $k, \ell \geqslant 0$ will be useful to have on hand when we later introduce product structures on singular homology and cohomology. They also arise in some other contexts that we will not discuss in this course, e.g. in studying the classifying space of a category.

The spaces $\Delta^n \subset \mathbb{R}^{n+1}$, $I \times \Delta^n \subset \mathbb{R} \times \mathbb{R}^{n+1} = \mathbb{R}^{n+2}$ and $\Delta^k \times \Delta^\ell \subset \mathbb{R}^{k+1} \times \mathbb{R}^{\ell+1} = \mathbb{R}^{k+\ell+2}$ are all compact topological manifolds with boundary that additionally have the following convenient feature: they are all *convex* subsets of vector spaces. For any n-manifold M that is also a convex subset of some vector space, a workable strategy for defining a triangulation on M is to specify the vertices $v_0, \ldots, v_n \in M$ of each n-simplex of the triangulation, and then require the corresponding region $\delta^n \subset M$ to be the convex hull of $v_0, \ldots, v_n$. There is of course a unique n-dimensional abstract simplicial complex K whose n-simplices have the specified vertices. Since ∂M will typically be nonempty in these situations, one must also consider the subcomplex

$$K' \subset K$$

consisting of all simplices spanned by vertices with convex hull contained in ∂M. Convexity implies that this is indeed a subcomplex, and if the conditions below can be verified so that $|K| \cong M$,

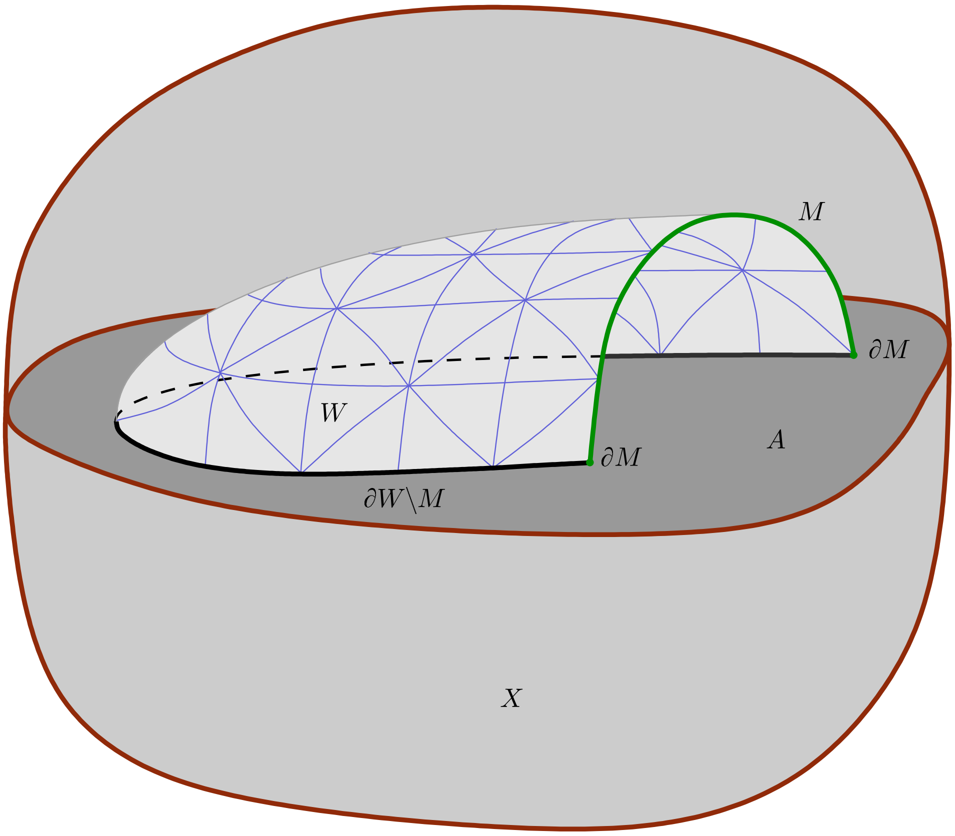

FIGURE 17. The picture shows a scenario as in Exercise 31.12 part (b), where M and W are triangulated submanifolds of X, both with nonempty boundary, and the maps f and F are defined as inclusions. The consequence is that $f : (M, \partial M) \to (X, A)$ represents the trivial relative homology class in $H_k(X, A; \mathbb{Z})$.

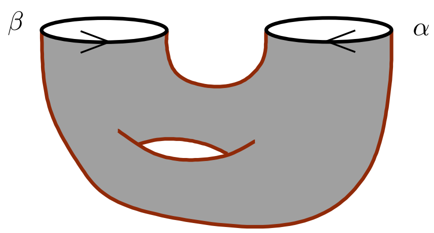

FIGURE 18. The surface in Exercise 31.13

it will follow automatically that $|K'| \cong \partial M$ and the complex K' is therefore $(n-1)$-dimensional. There is in any case a unique map

$$(|K|, |K'|) \to (M, \partial M)$$

that sends $|\sigma| \ni \sum_{v\in\sigma} t_v e_v \mapsto \sum_{v\in\sigma} t_v v \in M$ for each simplex σ of K, and the construction will be a triangulation if and only if this map is a homeomorphism. For this to hold, it is necessary and sufficient to establish two conditions:

(1) For each n-simplex $\sigma = \{v_0, \ldots, v_n\}$ of K, the points $v_0, \ldots, v_n \in M$ are in general position, so that their convex hull really is homeomorphic to Δ^n instead of a simplex of lower dimension. This ensures that the map $|K| \to M$ is a homeomorphism onto its image.
(2) Every point $p \in M$ lies in the convex hull of $v_0, \ldots, v_k \in M$ for some simplex $\{v_0, \ldots, v_k\}$ of K. This ensures that the map $|K| \to M$ is surjective.

Barycentric subdivision. The abstract simplicial complex K arising from the barycentric subdivision of Δ^n admits an easy inductive description: each n-simplex has vertices $b_n, v_1, \ldots, v_n$, where $b_n \in \Delta^n$ is the barycenter and $v_1, \ldots, v_n$ are the vertices of an $(n-1)$-simplex in the barycentric subdivision of one of the boundary faces $\partial_{(k)}\Delta^n = \Delta^{n-1}$. Note that the starting point of this inductive description is trivial, as the one-point space Δ^0 has only one possible triangulation.

Exercise 31.14. For the abstract simplicial complex K described above, whose vertices are points in Δ^n, show:

(a) The vertices of each n-simplex are in general position.
(b) Every point $p \in \Delta^n$ lies in the convex hull of the points $v_0, \ldots, v_k \in \Delta^n$ for some simplex $\{v_0, \ldots, v_k\}$ of K.
Hint: Draw a straight line from the barycenter b_n through p. What can you say about the point where this line exits Δ^n?

Exercise 31.14 establishes that barycentric subdivision defines a triangulation $(|K|, |K'|) \cong (\Delta^n, \partial\Delta^n)$ for every $n \geqslant 0$. The singular chain $\beta_n \in C_n(\Delta^n; \mathbb{Z})$ in Lemma 29.23 can now be defined as the image under the natural chain map $C_*^o(K; \mathbb{Z}) \to C_*(\Delta^n; \mathbb{Z})$ of the relative fundamental cycle

$$c_{\Delta^n} \in C_n^o(K; \mathbb{Z})$$

arising from this triangulation after fixing a few more choices, namely an orientation of the triangulation and an admissible ordering. There is surely more than one possible recipe for this, but here is one that works. Inductively, assume an admissible ordering and an orientation have already been chosen for the barycentric subdivision of Δ^{n-1}; for the case $n = 0$, there is no choice of ordering to be made, and we can fix the positive orientation on the unique 0-simplex. Now if $v_1, \ldots, v_n$ are the vertices of an $(n-1)$-simplex on one of the boundary faces $\partial_{(k)}\Delta^n = \Delta^{n-1}$ arranged in increasing order, and $\pm[v_1, \ldots, v_n]$ is its chosen orientation, define the ordering and orientation of the n-simplex $\{b_n, v_1, \ldots, v_n\}$ in Δ^n to be given by

$$b_n, v_1, \ldots, v_n \qquad \text{and} \qquad \pm(-1)^k [b_n, v_1, \ldots, v_n]$$

respectively. Following our usual prescription to define a simplicial n-chain $c_{\Delta^n} = \sum_\sigma \epsilon_\sigma \mathbf{v}_\sigma \in C_n^o(K; \mathbb{Z})$ out of this data, the next exercise implies via Lemma 31.5 that orienting the n-simplices in this way produces an oriented triangulation of Δ^n, and the formula for ∂c_{Δ^n} implies the formula for $\partial\beta_n$ in Lemma 29.23 after feeding it through the natural chain map $C_*^o(K; \mathbb{Z}) \to C_*(\Delta^n; \mathbb{Z})$.

Exercise 31.15. Prove that under the usual identification of each boundary face of Δ^n with Δ^{n-1}, the fundamental cycles of Δ^n and its boundary faces arising from the construction above are related by

$$\partial c_{\Delta^n} = \sum_{k=0}^{n} (-1)^k c_{\partial_{(k)}\Delta^n} \in C_{n-1}^o(K'; \mathbb{Z}).$$

The proof of excision also required a closely related triangulation of the prism $I \times \Delta^n$ for each $n \geqslant 0$, one that interpolates between the trivial triangulation of $\{0\} \times \Delta^n$ and the barycentric subdivision of $\{1\} \times \Delta^n$. This one also admits an inductive description: for $n = 0$, one takes the obvious triangulation of $I \times \Delta^0 \cong I$ with a single 1-simplex. Assuming that a suitable triangulation of the n-manifold $I \times \Delta^{n-1}$ for some $n \geqslant 1$ has already been constructed, the $(n+1)$-simplices of our triangulation of $I \times \Delta^n$ then come in two types:

- One whose vertices are $(0, e_0), \dots, (0, e_n)$ and $(1, b_n)$, where $e_0, \dots, e_n$ are the standard basis vectors of $\mathbb{R}^{n+1}$ (i.e. the vertices of Δ^n) and $b_n \in \Delta^n$ is the barycenter.
- For each $k = 0, \dots, n$ and each n-simplex $\{v_0, \dots, v_n\}$ in the triangulation of $I \times \partial_{(k)}\Delta^n = I \times \Delta^{n-1}$, one with vertices $v_0, \dots, v_n$ and $(1, b_n)$.

EXERCISE 31.16. Let L denote the $(n+1)$-dimensional abstract simplicial complex formed by the sets of vertices in $I \times \Delta^n$ described above, and define $L' \subset L$ to be the subcomplex of simplices whose vertices have convex hulls lying in $\partial(I \times \Delta^n)$.

(a) Carry out the analogue of Exercise 31.14 to show that the simplicial pair (L, L') defines a triangulation of $I \times \Delta^n$.
Hint: To show that every point $p \in I \times \Delta^n$ lies in one of the $(n+1)$-simplices described, draw a line from $(1, b_n)$ through p and see where it exits through $\partial(I \times \Delta^n)$.

(b) Describe an inductive algorithm to produce suitable admissible orderings and orientations for this triangulation of $I \times \Delta^n$ for each $n \geqslant 0$, and for the resulting simplicial $(n+1)$-chain $c_{I \times \Delta^n}$, derive a formula for $\partial c_{I \times \Delta^n}$ that implies Lemma 29.25.

Products of simplices. For any pair of integers $k, \ell \geqslant 0$, let $n := k + \ell$, so $\Delta^k \times \Delta^\ell \subset \mathbb{R}^{n+2}$ is both a convex set and a compact topological n-manifold with boundary

$$\partial(\Delta^k \times \Delta^\ell) = \left(\partial\Delta^k \times \Delta^\ell\right) \cup \left(\Delta^k \times \partial\Delta^\ell\right).$$

Denote the standard basis of $\mathbb{R}^{n+2}$ by

$$(e_0, 0), \dots, (e_k, 0), (0, f_0), \dots, (0, f_\ell) \in \mathbb{R}^{k+1} \times \mathbb{R}^{\ell+1} = \mathbb{R}^{n+2},$$

so we can regard $e_0, \dots, e_k$ as the vertices of Δ^k and $f_0, \dots, f_\ell$ as the vertices of Δ^ℓ. Let us describe a triangulation of $\Delta^k \times \Delta^\ell$ whose vertices are the points $(e_i, f_j) \in \Delta^k \times \Delta^\ell$ for all $i = 0, \dots, k$ and $j = 0, \dots, \ell$. Like barycentric subdivision, the triangulation can be characterized inductively with respect to the dimensions k and ℓ, but it also admits the following more direct description. Endow the set $\{0, \dots, k\} \times \{0, \dots, \ell\}$ with the total order such that $(i, j) \leqslant (i', j')$ if and only if $i \leqslant i'$ and $j \leqslant j'$, and define an n-dimensional simplicial complex K whose m-simplices for $m = 0, \dots, n$ are the sets of the form

$$\sigma = \{(e_{i_0}, f_{j_0}), \dots, (e_{i_m}, f_{j_m})\} \subset \Delta^k \times \Delta^\ell$$

for all possible strictly increasing sequences

$$(i_0, j_0) < \dots < (i_m, j_m)$$

in $\{0, \dots, k\} \times \{0, \dots, \ell\}$. We shall endow K with the canonical admissible ordering determined by the total order on $\{0, \dots, k\} \times \{0, \dots, \ell\}$. Observe that the n-simplices of K all correspond to sequences $(i_0, j_0) < \dots < (i_n, j_n)$ that begin with $(i_0, j_0) = (0, 0)$ and end with $(i_n, j_n) = (k, \ell)$, thus all of them contain the two specific vertices (e_0, f_0) and (e_k, f_ℓ). Boundary faces σ of these n-simplices come in three types, corresponding to sequences $(i_0, j_0) < \dots < (i_{n-1}, j_{n-1})$ that satisfy the following conditions:

(1) The sequence $j_0, \dots, j_{n-1}$ takes every value in $\{0, \dots, \ell\}$ but $i_0, \dots, i_{n-1}$ misses exactly one value $i \in \{0, \dots, k\}$.
(2) The sequence $i_0, \dots, i_{n-1}$ takes every value in $\{0, \dots, k\}$ but $j_0, \dots, j_{n-1}$ misses exactly one value $j \in \{0, \dots, \ell\}$.

(3) There are two consecutive terms of the form (i,j), $(i+1,j+1)$.

In the first two cases, the n vertices of σ all lie in one of the convex sets

$$\partial_{(i)}\Delta^k \times \Delta^\ell \qquad \text{or} \qquad \Delta^k \times \partial_{(j)}\Delta^\ell.$$

The union of these sets for all $i = 0,\dots,k$ and $j = 0,\dots,\ell$ is $\partial(\Delta^k \times \Delta^\ell)$, and these boundary faces thus determine an $(n-1)$-dimensional subcomplex $K' \subset K$ in which the convex hull of the vertices of each simplex is contained in $\partial(\Delta^k \times \Delta^\ell)$. It is easy to check that *all* simplices of K with convex hull contained in $\partial(\Delta^k \times \Delta^\ell)$ are of this form, because for any two points $p, q \in \partial(\Delta^k \times \Delta^\ell)$ that do not both belong to the same one of the $n+2$ convex subsets mentioned above, the line segment from p to q passes through the interior of $\Delta^k \times \Delta^\ell$. In particular, boundary faces of the third type in the list above do not belong to the subcomplex K'.

EXERCISE 31.17. For the simplicial complex K with admissible ordering described above, prove that the canonical map $(|K|,|K'|) \to (\Delta^k \times \Delta^\ell, \partial(\Delta^k \times \Delta^\ell))$ is a homeomorphism.
Hint: This is probably not the only possible approach, but here an inductive argument as in Exercises 31.14 and 31.16 is also possible. Use the fact that certain points are contained in all the n-simplices.

In order to define a suitable orientation and an integral fundamental cycle $c_{\Delta^k\times\Delta^\ell} \in C_n^o(K;\mathbb{Z})$, let $\mathbf{S}(k,\ell)$ denote the set of all strictly increasing sequences $(i_0,j_0) < \ldots < (i_n,j_n)$ of $n+1$ elements in $\{0,\dots,k\}\times\{0,\dots,\ell\}$, and write $\sigma_{\mathbf{s}}$ for the n-simplex of K determined by each $\mathbf{s} \in \mathbf{S}(k,\ell)$. Denote by $\mathbf{s}_0 \in \mathbf{S}(k,\ell)$ the specific sequence

$$(0,0) < (1,0) < \ldots < (k,0) < (k,1) < \ldots < (k,\ell),$$

and define the **parity** $|\mathbf{s}| \in \mathbb{Z}_2$ of any element $\mathbf{s} \in \mathbf{S}(k,\ell)$ to be the number of steps (modulo 2) required in order to transform $\mathbf{s}_0$ into $\mathbf{s}$ via operations that modify three consecutive terms of a sequence like so:

$$(i-1,j) < (i,j) < (i,j+1) \qquad \rightsquigarrow \qquad (i-1,j) < (i-1,j+1) < (i,j+1).$$

EXERCISE 31.18. Show that the parity $|\mathbf{s}| \in \mathbb{Z}_2$ of elements $\mathbf{s} \in \mathbf{S}(k,\ell)$ is well defined by interpreting $(-1)^{|\mathbf{s}|} \in \{1,-1\}$ as the sign of a permutation of n elements, which include k copies of the letter R (for "right") and ℓ copies of the letter U (for "up").

Using the canonical admissible ordering for K, we now define a simplicial n-chain by

$$c_{\Delta^k\times\Delta^\ell} := \sum_{\mathbf{s}\in\mathbf{S}(k,\ell)} (-1)^{|\mathbf{s}|}\mathbf{v}_{\mathbf{s}} \in C_n^o(K;\mathbb{Z}),$$

where $\mathbf{v}_{\mathbf{s}}$ denotes the $(n+1)$-tuple formed by the vertices of the n-simplex $\sigma_{\mathbf{s}}$ arranged in increasing order. The following exercise implies via Lemma 31.5 that the signs $(-1)^{|\mathbf{s}|}$ determine an orientation and thus a relative fundamental cycle for our triangulation of $\Delta^k \times \Delta^\ell$.

EXERCISE 31.19. For the simplicial n-chain $c_{\Delta^k\times\Delta^\ell}$ defined above for each $k,\ell \geqslant 0$:

(a) Prove that when $k,\ell \geqslant 1$, the following formula holds in the ordered simplicial chain complex of K under the usual identification between boundary faces and simplices of one dimension lower:

$$\partial c_{\Delta^k\times\Delta^\ell} = \sum_{i=0}^k (-1)^i c_{\partial_{(i)}\Delta^k\times\Delta^\ell} + (-1)^k \sum_{j=0}^\ell (-1)^j c_{\Delta^k\times\partial_{(j)}\Delta^\ell} \in C_{n-1}^o(K';\mathbb{Z}) \subset C_{n-1}^o(K;\mathbb{Z}).$$

(b) Deduce Lemma 29.17 from the special case $k = 1$.

REMARK 31.20. If we regard $\partial\Delta^k \times \Delta^\ell$ and $\Delta^k \times \partial\Delta^\ell$ as compact topological $(n-1)$-manifolds with matching boundary $\partial\Delta^k \times \partial\Delta^\ell$ and endow both with the obvious oriented triangulations and admissible orderings that they inherit from $\Delta^k \times \Delta^\ell$, the formula in Exercise 31.19 takes the slightly prettier form

$$\partial c_{\Delta^k \times \Delta^\ell} = c_{\partial\Delta^k \times \Delta^\ell} + (-1)^k c_{\Delta^k \times \partial\Delta^\ell}.$$

When we introduce the homological cross product later in this course, the singular homology version of this relation will take the form

$$\partial(c_{\Delta^k} \times c_{\Delta^\ell}) = \partial c_{\Delta^k} \times c_{\Delta^\ell} + (-1)^k c_{\Delta^k} \times \partial c_{\Delta^\ell},$$

which is written in terms of the obvious relative fundamental cycle $c_{\Delta^n} \in C_n(\Delta^n; \mathbb{Z})$ for the standard simplex of each dimension with its trivial triangulation, and a bilinear product operation

$$C_*(X) \otimes C_*(Y) \to C_*(X \times Y) : A \otimes B \mapsto A \times B$$

that relates the singular chain complexes of any two spaces X, Y and sends $C_k(X) \otimes C_\ell(Y)$ in general to $C_{k+\ell}(X \times Y)$. We will have plenty to say about this product later, but the detail I want to comment on right now is the sign $(-1)^k$ appearing on the right hand side of the formula. This is an instance of a general pattern known as the **Koszul sign convention**, which we will see many more examples of in this course. In a nutshell, the rule is that whenever objects carry natural gradings in either $\mathbb{Z}$ or $\mathbb{Z}_2$, exchanging the order of two objects with odd degree causes a sign change. In the present context, the "objects" to which this rule applies are not only the chains of certain degrees in Δ^k and Δ^ℓ but also the operator ∂, which we regard as having degree -1 since it maps k-chains to $(k-1)$-chains for every k. This means that no sign change is necesary when writing $\partial c_{\Delta^k} \times c_{\Delta^\ell}$, since the three objects ∂, c_{Δ^k} and c_{Δ^ℓ} appear here in the same order as on the left hand side, but writing $c_{\Delta^k} \times \partial c_{\Delta^\ell}$ exchanges the order of ∂ and c_{Δ^k}, and since ∂ has odd degree, a sign change must then result if and only if c_{Δ^k} also has odd degree, meaning k is odd. One could presumably state some general theorem in category-theoretic terms to explain why and in what contexts this particular way of dealing with signs gives the results we want, but I personally would consider writing down that theorem to be more trouble than it is worth. If you haven't seen the Koszul convention before in one of the many other contexts (e.g. the exterior algebra of differential forms on smooth manifolds) where it naturally arises, then I think that you will in any case learn through experience during the remainder of this course why it is good and useful.

32. The Eilenberg-Steenrod axioms, triples, and good pairs (November 3, 2023)

In the computation of $H_*(S^n; G)$ in Lectures 28 and 29, we never had to make any specific reference to singular simplices or any other aspects of the definition of singular homology. We did need to know that singular homology has a particular set of formal properties, e.g. functoriality, homotopy invariance, long exact sequences and excision, and we needed to understand a variant of the theory known as *reduced* homology, which was defined in terms of the unique map from any space to the one-point space. It will turn out that almost all computations of $H_*(X; G)$ we can carry out for a "reasonable" class of spaces depend on exactly this same list of properties. This realization motivated Eilenberg and Steenrod in the 1950's to codify a set of axioms for so-called "homology theories". In the early days of homology, there were in fact several competing homology theories that differed substantially in their definitions but nonetheless seemed *mostly* to be measuring the same information about topological spaces. The Eilenberg-Steenrod axioms provided an explanation for this similarity. While singular homology is now easily the most popular of the original homology theories, the others did not completely die out—I have seen some of the others used from time to time in research papers on subjects I cared about, and there are still authors who argue that other theories are preferable from certain points of view. We will sketch some examples of alternative homology theories later in this course.

DEFINITION 32.1. An **axiomatic homology theory** h_* is a covariant functor

$$\mathsf{Top}_{\mathrm{rel}} \to \mathsf{Ab}_{\mathbb{Z}} : (X, A) \mapsto h_*(X, A) = \bigoplus_{n \in \mathbb{Z}} h_n(X, A),$$

together with a natural transformation ∂_* from the functor $\mathsf{Top}_{\mathrm{rel}} \to \mathsf{Ab} : (X, A) \mapsto h_n(X, A)$ to the functor $\mathsf{Top}_{\mathrm{rel}} \to \mathsf{Ab} : (X, A) \mapsto h_{n-1}(A)$ for each $n \in \mathbb{Z}$ such that the following axioms are satisfied:

- (EXACTNESS) For all pairs (X, A) with inclusion maps $i : A \hookrightarrow X$ and $j : (X, \varnothing) \hookrightarrow (X, A)$, the sequence
$$\ldots \longrightarrow h_{n+1}(X, A) \xrightarrow{\partial_*} h_n(A) \xrightarrow{i_*} h_n(X) \xrightarrow{j_*} h_n(X, A) \xrightarrow{\partial_*} h_{n-1}(A) \longrightarrow \ldots$$
is exact.
- (HOMOTOPY) For any two homotopic maps $f, g : (X, A) \to (Y, B)$, the induced morphisms $f_*, g_* : h_*(X, A) \to h_*(Y, B)$ are identical.
- (EXCISION) For any pair (X, A) and any subset $B \subset X$ with closure in the interior of A, the inclusion $(X \backslash B, A \backslash B) \hookrightarrow (X, A)$ induces an isomorphism
$$h_*(X \backslash B, A \backslash B) \xrightarrow{\cong} h_*(X, A).$$
- (DIMENSION) For any space $\{\mathrm{pt}\}$ containing only one point, $h_n(\{\mathrm{pt}\}) = 0$ for all $n \neq 0$.
- (ADDITIVITY) For any collection of spaces $\{X_\alpha\}_{\alpha \in J}$ with inclusion maps $i^\alpha : X_\alpha \hookrightarrow \coprod_{\beta \in J} X_\beta$, the induced homomorphisms $i^\alpha_* : h_*(X_\alpha) \to h_*\left(\coprod_{\beta \in J} X_\beta\right)$ determine an isomorphism
$$\bigoplus_{\alpha \in J} i^\alpha_* : \bigoplus_{\alpha \in J} h_*(X_\alpha) \xrightarrow{\cong} h_*\left(\coprod_{\beta \in J} X_\beta\right).$$

A few comments are in order.

REMARK 32.2. The original list in [ES52] includes three other axioms before exactness, but the first two of these are equivalent to the statement that $h_* : \mathsf{Top}_{\mathrm{rel}} \to \mathsf{Ab}_{\mathbb{Z}}$ is a functor, and the third simply requires ∂_* to be a natural transformation. We could equally well have chosen to hide the homotopy axiom by calling h_* a functor $\mathsf{Top}^h_{\mathrm{rel}} \to \mathsf{Ab}_{\mathbb{Z}}$ instead of $\mathsf{Top}_{\mathrm{rel}} \to \mathsf{Ab}_{\mathbb{Z}}$.

REMARK 32.3. The additivity axiom did not appear in [ES52] but was added later by Milnor [Mil62]. One can show in fact that for *finite* disjoint unions, additivity follows as a consequence of the other axioms (see Exercise 32.6), thus Eilenberg and Steenrod did not need it because they were mainly concerned with computations for compact polyhedra—these come from finite simplicial complexes, so no infinite disjoint unions are allowed.

REMARK 32.4. I am cheating slightly by stating a stronger variant of the excision axiom than appeared in the original list by Eilenberg and Steenrod. The version in [ES52] reads as follows:

- (EXCISION′) For any pair (X, A) and any *open* subset $B \subset X$ with closure in the interior of A, the inclusion $(X \backslash B, A \backslash B) \hookrightarrow (X, A)$ induces an isomorphism
$$h_*(X \backslash B, A \backslash B) \xrightarrow{\cong} h_*(X, A).$$

This means that there might in principle exist theories that satisfy the original Eilenberg-Steenrod axioms but not ours, because $h_*(X \backslash B, A \backslash B) \to h_*(X, A)$ might fail to be an isomorphism in cases where $\bar{B} \subset \mathring{A}$ but B is not open. Note that we already have applied this axiom in cases where B is not open, e.g. in the argument of Lectures 29 and 30 to prove $\widetilde{H}_k(X) \cong \widetilde{H}_{k+1}(SX)$, we considered the inclusion $(SX\backslash\{p_-\}, C_-X\backslash\{p_-\}) \hookrightarrow (SX, C_-X)$. This is fine in singular homology because H_* satisfies the stronger axiom, but even for some hypothetical theory h_* that satisfies (EXCISION′) and not (EXCISION), one could prove that $h_*(SX\backslash\{p_-\}, C_-X\backslash\{p_-\}) \to h_*(SX, C_-X)$ is an isomorphism by relating it via homotopy invariance to the map

$$h_*(SX\backslash B_\epsilon(p_-), C_-X\backslash B_\epsilon(p_-)) \to h_*(SX, C_-X)$$

for some small open neighborhood $B_\epsilon(p_-) \subset C_-X$ of p_-, and this map definitely is an isomorphism since $B_\epsilon(p_-)$ is open. In practice, some trick of this sort will be available in every important situation where we need to apply excision, so that it will not really matter which version of the axiom we adopt. We'll opt for the stronger one in this course since several arguments would become slightly longer without it.

REMARK 32.5. It is sometimes useful to expand the definition and allow an axiomatic homology theory to be a functor $\mathscr{C} \to \mathsf{Ab}_{\mathbb{Z}}$ defined on a suitable subcategory $\mathscr{C}$ of $\mathsf{Top}_{\mathrm{rel}}$, so that we need not define $h_*(X, A)$ for *all* pairs (X, A) but only a subclass. One important example we will see later is the category of **compact pairs**, which are simply pairs of spaces (X, A) such that X is compact Hausdorff and $A \subset X$ is closed. (In reference to Remark 32.4, notice that the category of compact pairs (X, A) requires the weaker version of the excision axiom since $(X\backslash B, A\backslash B)$ will not be an object in the category unless $B \subset X$ is open.) When allowing restrictions of this type, one must take care so that all of the maps needed for expressing the axioms—e.g. the inclusions $A \hookrightarrow X$ and $(X, \varnothing) \hookrightarrow (X, A)$—are actually morphisms. In [ES52], this concern motivates the definition of the notion of an *admissible* category of pairs.

EXERCISE 32.6. Assume $h_* : \mathsf{Top}_{\mathrm{rel}} \to \mathsf{Ab}_{\mathbb{Z}}$ is a functor satisfying all of the Eilenberg-Steenrod axioms for homology theories except possibly the additivity axiom. Given two spaces X and Y, use excision and the long exact sequences of the pairs $(X \amalg Y, X)$ and $(X \amalg Y, Y)$ to prove that for the natural inclusions $i^X : X \hookrightarrow X \amalg Y$ and $i^Y : Y \hookrightarrow X \amalg Y$, the map

$$i^X_* \oplus i^Y_* : h_*(X) \oplus h_*(Y) \to h_*(X \amalg Y) : (x, y) \mapsto i^X_* x + i^Y_* y$$

is an isomorphism. Deduce that h_* does satisfy the additivity axiom for all *finite* disjoint unions.

You may notice that Definition 32.1 above makes no mention of any coefficient group. It's there, actually—it's just hidden.

DEFINITION 32.7. The **coefficient group**[53] of an axiomatic homology theory is defined to be the group $h_0(\{\mathrm{pt}\})$.

The formal properties of singular homology listed in Lecture 28 can now be summarized thus:

THEOREM 32.8. *For any abelian group G, singular homology $H_* := H_*(\cdot; G)$ is an axiomatic homology theory with coefficient group G.* □

REMARK 32.9. The axioms do *not* imply the first two properties of H_* that we discussed in Lecture 28: the relation of $H_0(X)$ and $H_1(X; \mathbb{Z})$ to $\pi_0(X)$ and $\pi_1(X)$ respectively. We will later see an example of a theory that satisfies all of the Eilenberg-Steenrod axioms, but not these two properties. One has to look at fairly strange spaces in order to see this difference, e.g. spaces that are connected but not path-connected.

---

[53]There is a slightly awkward semantic issue in this definition: strictly speaking, what we are calling "$\{\mathrm{pt}\}$" is not a unique space, but simply *any* choice of space that happens to contain only one element. It follows that the coefficient group $h_0(\{\mathrm{pt}\})$ is not a uniquely defined group, but is an isomorphism class of groups. Any two choices of one-point spaces P_0 and P_1 are related by a unique homeomorphism $P_0 \to P_1$, which induces a canonical isomorphism $h_0(P_0) \to h_0(P_1)$.

EXERCISE 32.10. Show that if the dimension axiom is dropped from Definition 32.1, then for any axiomatic homology theory h_*, the functor $\mathsf{Top} \to \mathsf{Ab}_{\mathbb{Z}}$ sending each pair (X, A) to the graded abelian group $G_* = \bigoplus_{n \in \mathbb{Z}} G_n$ with $G_n := h_{n+1}(X, A)$ is also an axiomatic homology theory. This explains why the dimension axiom is called what it is: it guarantees that there is some connection between the subscript n in the notation $h_n(X, A)$ and our intuitive notion of "dimension," i.e. $\{\mathrm{pt}\}$ is clearly a "zero-dimensional" space.

REMARK 32.11. One can define singular homology with a trivial coefficient group: it still satisfies the axioms, but it is not very interesting, because $H_*(X, A; G)$ is then trivial for all pairs. Weirdly, it is possible in general for an axiomatic homology theory to be nontrivial on some spaces (though not on "nice" spaces like polyhedra) even if its coefficient group $h_0(\{\mathrm{pt}\})$ is trivial. Look at the axioms again: you'll see that there is no obvious reason why this couldn't be allowed. My first instinct when I learned about these axioms was to try to prove as an exercise that $h_0(\{\mathrm{pt}\}) = 0$ implies $h_0(X, A) = 0$ for all (X, A), but fortunately I did not spend much time on this exercise—it turns out that someone else thought about it in 1958 and came up with counterexamples [JW58].

DEFINITION 32.12. Any axiomatic homology theory h_* has a corresponding **reduced** theory $\widetilde{h}_*$, defined by $\widetilde{h}_*(X, A) := h_*(X, A)$ whenever $A \neq \varnothing$ and

$$\widetilde{h}_*(X) := \ker \epsilon_* \subset h_*(X)$$

for the unique map $\epsilon : X \to \{\mathrm{pt}\}$.

The next result follows by exactly the same arguments as in the case of reduced singular homology, cf. Propositions 29.6, 29.7 and 29.8, and Theorem 29.10.

THEOREM 32.13. *For any axiomatic homology theory h_* with coefficient group G:*

(1) $\widetilde{h}_*(X) = 0$ for all contractible spaces X.

(2) There is a split exact sequence $0 \to \widetilde{h}_*(X) \hookrightarrow h_*(X) \xrightarrow{\epsilon_*} h_*(\{\mathrm{pt}\}) \to 0$ for all spaces X, giving rise to isomorphisms

$$h_n(X) \cong \begin{cases} \widetilde{h}_0(X) \oplus G & \text{for } n = 0, \\ \widetilde{h}_n(X) & \text{for } n \neq 0. \end{cases}$$

(3) Morphisms $h_*(X, A) \to h_*(Y, B)$ induced by maps of pairs $(X, A) \to (Y, B)$ restrict to $\widetilde{h}_*(X, A) \to \widetilde{h}_*(Y, B)$, so that $\widetilde{h}_*$ defines a functor $\mathsf{Top}^h_{\mathrm{rel}} \to \mathsf{Ab}_{\mathbb{Z}}$.

(4) All connecting homomorphisms $\partial_* : h_n(X, A) \to h_{n-1}(A)$ have image in $\widetilde{h}_{n-1}(A)$ and the resulting sequence

$$\ldots \longrightarrow \widetilde{h}_{n+1}(X, A) \xrightarrow{\partial_*} \widetilde{h}_n(A) \xrightarrow{i_*} \widetilde{h}_n(X) \xrightarrow{j_*} \widetilde{h}_n(X, A) \xrightarrow{\partial_*} \widetilde{h}_{n-1}(A) \longrightarrow \ldots$$

is exact.

□

All together this is enough information to repeat nearly verbatim the computations of Lecture 28 involving suspensions and spheres. We obtain:

THEOREM 32.14. *For any axiomatic homology theory h_* and any space X, there is a natural isomorphism $\widetilde{h}_n(X) \to \widetilde{h}_{n+1}(SX)$ for every $n \in \mathbb{Z}$.* □

THEOREM 32.15. *For every $n \in \mathbb{N}$ and every axiomatic homology theory h_* with coefficient group G,*

$$h_k(S^n) \cong \begin{cases} G & \text{for } k = 0, n, \\ 0 & \text{for all other } k. \end{cases}$$

□

EXERCISE 32.16. Here are two applications of the five-lemma (see Exercise 31.8).

(a) Given an axiomatic homology theory h_* and a map of pairs $f : (X, A) \to (Y, B)$, show that if any two of the induced maps $h_k(X) \xrightarrow{f_*} h_k(Y)$, $h_k(A) \xrightarrow{f_*} h_k(B)$ and $h_k(X, A) \xrightarrow{f_*} h_k(Y, B)$ are isomorphisms for every k, then so is the third.

(b) Given a collection of pairs of spaces $\{(X_\alpha, A_\alpha)\}_{\alpha \in J}$, consider the pair

$$\coprod_{\alpha \in J}(X_\alpha, A_\alpha) := \left(\coprod_{\alpha \in J} X_\alpha, \coprod_{\alpha \in J} A_\alpha\right)$$

with the natural inclusion maps $i^\alpha : (X_\alpha, A_\alpha) \hookrightarrow \coprod_{\beta \in J}(X_\beta, A_\beta)$. Prove that for any axiomatic homology theory h_*, the additivity axiom generalizes to pairs, producing an isomorphism

$$\bigoplus_{\alpha \in J} i^\alpha_* : \bigoplus_{\alpha \in J} h_*(X_\alpha, A_\alpha) \xrightarrow{\cong} h_*\Big(\coprod_{\alpha \in J}(X_\alpha, A_\alpha)\Big).$$

A large portion of the theorems we prove about singular homology in this course will be based only on the axioms, and will thus be valid for any axiomatic homology theory. Proofs based on the axioms are traditionally considered more elegant than those that require explicit reference to the definition of $H_*(X, A)$. On the other hand, there are a few cases in which both types of proof are possible but the one that doesn't use the axioms is much easier. The exact sequence explained below is certainly an example of this.

Suppose $B \subset A \subset X$, so (X, A), (X, B) and (A, B) are all pairs of spaces, with obvious inclusion maps of pairs

$$i : (A, B) \hookrightarrow (X, B) \quad \text{and} \quad j : (X, B) \hookrightarrow (X, A).$$

These then induce a short exact sequence of relative singular chain complexes

$$0 \longrightarrow C_*(A, B) \xrightarrow{i_*} C_*(X, B) \xrightarrow{j_*} C_*(X, A) \longrightarrow 0.$$

The special case of this with $B = \varnothing$ reproduces the usual short exact sequence for the pair (X, A). Applying Proposition 28.18 as before gives the so-called **long exact sequence of the triple** (X, A, B):

$$\ldots \to H_{n+1}(X, A) \xrightarrow{\partial_*} H_n(A, B) \xrightarrow{i_*} H_n(X, B) \xrightarrow{j_*} H_n(X, A) \xrightarrow{\partial_*} H_{n-1}(A, B) \to \ldots, \tag{32.1}$$

which directly generalizes the long exact sequence of (X, A). It is also not hard to show that the connecting homomorphism $\partial_* : H_n(X, A) \to H_{n-1}(A, B)$ satisfies a naturality property, i.e. given any map of pairs $f : (X, A) \to (X', A')$ and a subset $B' \subset A'$ such that $f(B) \subset B'$, we have commutative diagrams

$$\begin{array}{ccc} H_n(X, A) & \xrightarrow{\partial_*} & H_{n-1}(A, B) \\ \downarrow{\scriptstyle f_*} & & \downarrow{\scriptstyle f_*} \\ H_n(X', A') & \xrightarrow{\partial_*} & H_{n-1}(A', B') \end{array}$$

for all $n \in \mathbb{Z}$. Moreover, close inspection of the usual diagram chasing argument yields an explicit formula for $\partial_* : H_n(X, A) \to H_{n-1}(A, B)$: if $[c] \in H_n(X, A)$ is represented by a relative cycle $c \in C_n(X)$, then $\partial c \in C_{n-1}(A)$ is a cycle in A and therefore also a relative cycle for the pair (A, B), so that it represents a class in $H_{n-1}(A, B)$ and the formula $\partial_*[c] = [\partial c]$ thus makes sense. All of this is proved by nearly the same arguments as in the case $B = \varnothing$, so I will spare you the details.

We will use the exact sequence of the triple several times in our efforts to simplify computations of singular homology, and for this purpose the presentation above is certainly sufficient. Nonetheless, you might now be wondering: is this sequence really a property distinctive to the singular chain complex, or does it work for every axiomatic homology theory, and if so, how can one derive a connecting homomorphism $\partial_* : h_n(X, A) \to h_{n-1}(A, B)$ from the axioms? The answer to this question requires some cleverness, and I'm at a loss to conjecture how anyone might have come up with it for the first time, but here it is: given an axiomatic homology theory h_* and a triple (X, A, B) with $B \subset A \subset X$, we can consider the following "braid" diagram:

$$\begin{tikzcd}
\dots h_{n+1}(X,A) \arrow[rr, bend left, "\partial"] \arrow[dr, "\partial_1"] & & h_n(A,B) \arrow[rr, bend left, "\partial_3"] \arrow[dr, "i"] & & h_{n-1}(B) \arrow[rr, bend left, "i_2"] \arrow[dr, "i_3"] & & h_{n-1}(X) \dots \\
& h_n(A) \arrow[ur, "j_3"] \arrow[dr, "i_1"] & & h_n(X,B) \arrow[ur, "\partial_2"] \arrow[dr, "j"] & & h_{n-1}(A) \arrow[ur, "i_1"] \arrow[dr, "j_3"] & \\
\dots h_n(B) \arrow[ur, "i_3"] \arrow[rr, bend right, "i_2"'] & & h_n(X) \arrow[ur, "j_2"] \arrow[rr, bend right, "j_1"'] & & h_n(X,A) \arrow[ur, "\partial_1"] \arrow[rr, bend right, "\partial"'] & & h_{n-1}(A,B) \dots
\end{tikzcd}$$

The braid consists of four "strands," three of which you may recognize as the long exact sequences of the pairs (X, A), (X, B) and (A, B). The fourth strand is the sequence

$$\dots \longrightarrow h_{n+1}(X, A) \xrightarrow{\partial} h_n(A, B) \xrightarrow{i} h_n(X, B) \xrightarrow{j} h_n(X, A) \xrightarrow{\partial} h_{n-1}(A, B) \longrightarrow \dots, \tag{32.2}$$

which we would like to prove is exact. Here the map $\partial := j_3 \circ \partial_1$ is defined via the commutativity of the diagram, while all other maps are either induced by the obvious inclusions or are connecting homomorphisms from long exact sequences of pairs. The whole diagram commutes due to the commutativity of the obvious inclusions plus the naturality of the connecting homomorphisms.

EXERCISE 32.17. Deduce via the following steps that the sequence (32.2) appearing as the fourth strand in the braid diagram above is exact:

(a) Use the commutativity of the diagram to show that $i \circ \partial = 0$ and $\partial \circ j = 0$.
Hint: Each can be expressed as a different composition that includes two successive maps in an exact sequence.

(b) Prove that $j \circ i = 0$ by factoring it through the group $h_*(A, A)$, which is always zero. (Why?)

(c) Use a purely algebraic diagram-chasing argument to prove that the kernel of each map in the sequence (32.2) is contained in the image of the previous one.

Here is a useful application of the long exact sequence of a triple. Since $H_*(X, A)$ is defined so as to measure the topology of X while ignoring anything that happens entirely in A, it is natural to expect some relationship between this and the *absolute* homology of the space X/A defined by collapsing $A \subset X$ to a point. Here we should restrict attention to the case where $A \subset X$ is closed, since X/A may otherwise be a horrible (e.g. non-Hausdorff) space. It turns out that under a further assumption on the pair (X, A), the relative homology $H_*(X, A)$ is naturally isomorphic to $\widetilde{H}_*(X/A)$. To see this, we start by observing that there is a natural isomorphism between the reduced homology of X/A and the relative homology of the pair $(X/A, A/A)$, in which the subset $A/A \subset X/A$ is actually just a single point. Indeed:

LEMMA 32.18. *For any space X and a point $x \in X$, the inclusion of pairs $(X, \emptyset) \hookrightarrow (X, \{x\})$ induces an isomorphism*

$$\widetilde{h}_*(X) \xrightarrow{\cong} \widetilde{h}_*(X, \{x\}) = h_*(X, \{x\}),$$

where h_* is any axiomatic homology theory.

PROOF. This is immediate from the long exact sequence of $(X, \{x\})$ in reduced homology since $\widetilde{h}_*(\{x\}) = 0$. $\square$

Now, observe that the quotient projection $q : X \to X/A$ is also a map of pairs $(X, A) \to (X/A, A/A)$ and thus induces a morphism $h_*(X, A) \to h_*(X/A, A/A)$. Can we expect this map to be an isomorphism? The intuition here is that if we were allowed to remove the subset A and consider the restricted map

$$(X \backslash A, A \backslash A) \xrightarrow{q} ((X/A) \backslash (A/A), (A/A) \backslash (A/A)),$$

then it becomes a homeomorphism, and thus induces an isomorphism between two homology groups that we expect should match $h_*(X, A)$ and $h_*(X/A, A/A)$ due to excision. But we aren't quite allowed to apply excision in this way: normally, the set we're removing needs to have its closure contained in the interior of the smaller set in the pair, which is usually not true if those two sets are the same. Conclusion: we need to impose a condition on (X, A) so that A lies strictly inside of something else that will allow us to apply excision. The following bit of informal terminology is borrowed from [**Hat02**].

DEFINITION 32.19. A pair of spaces (X, A) will be called **good** if $A \subset X$ is a closed subset and is a deformation retract of some neighborhood $V \subset X$ of itself.

EXAMPLE 32.20. $(\mathbb{D}^n, \partial \mathbb{D}^n)$ is a good pair since $\partial \mathbb{D}^n = S^{n-1}$ has a neighborhood homeomorphic to $(-1, 0] \times S^{n-1}$ which deformation retracts to $\{0\} \times S^{n-1}$.

EXAMPLE 32.21. The pair (X, A) with $X = [0, 1]$ and $A = \{1, 1/2, 1/3, 1/4, \dots, 0\}$ is not good. The easiest way to prove this is probably by showing that it does not satisfy Theorem 32.23 below, due to the following exercise:

EXERCISE 32.22. Show that for the pair (X, A) in Example 32.21, $H_1(X/A; \mathbb{Z}) \not\cong H_1(X, A; \mathbb{Z})$. *Hint: $H_1(X, A; \mathbb{Z})$ is not too hard to compute from the long exact sequence of (X, A), and in particular it is an infinitely-generated but countable group. To compute $H_1(X/A; \mathbb{Z})$, you might notice that X/A is homeomorphic to the so-called* Hawaiian earring, *which we examined in Exercise 13.2 last semester as an example of an "unreasonable" space. We saw in particular that $\pi_1(X/A)$ admits a surjective homomorphism to the uncountable abelian group $\prod_{n \in \mathbb{N}} \mathbb{Z}$.*

As you might extrapolate from the two examples just mentioned, most pairs you will encounter in nature are good; it takes some creativity to come up with examples that are not.

THEOREM 32.23. *If (X, A) is a good pair, then for every axiomatic homology theory h_*, the natural quotient map $q : (X, A) \to (X/A, A/A)$ induces an isomorphism*

$$q_* : h_*(X, A) \xrightarrow{\cong} h_*(X/A, A/A),$$

implying via Lemma 32.18 that $h_*(X, A) \cong \widetilde{h}_*(X/A)$.

PROOF. Fix a neighborhood $V \subset X$ of A such that A is a deformation retract of V. Then the inclusion $(A, A) \hookrightarrow (V, A)$ is a homotopy equivalence of pairs, thus

$$h_*(V, A) \cong h_*(A, A) = 0,$$

where the latter vanishes due to the long exact sequence of (A, A). Writing down the long exact sequence of the triple (X, V, A) then gives

$$0 = h_k(V, A) \to h_k(X, A) \to h_k(X, V) \to h_{k-1}(V, A) = 0,$$

so that the map $h_*(X, A) \xrightarrow{i_*} h_*(X, V)$ induced by the inclusion $i : (X, A) \hookrightarrow (X, V)$ is an isomorphism.

One can carry out the same argument after taking the quotient of all spaces by A: the deformation retraction of V to A implies that V/A is contractible and thus $h_*(V/A, A/A) \cong \widetilde{h}_*(V/A) =$

0, so the exact sequence of $(X/A, V/A, A/A)$ then implies that the map $h_*(X/A, A/A) \xrightarrow{j_*} h_*(X/A, V/A)$ induced by the inclusion $j : (X/A, A/A) \hookrightarrow (X/A, V/A)$ is an isomorphism.

Now consider the commutative diagram

$$\begin{tikzcd}
h_*(X,A) \arrow[r, "i_*"] \arrow[d, "q_*"] & h_*(X,V) \arrow[d, "q_*"] & h_*(X\backslash A, V\backslash A) \arrow[l, "k_*"'] \arrow[d, "q_*"] \\
h_*(X/A, A/A) \arrow[r, "j_*"] & h_*(X/A, V/A) & h_*\big((X/A)\backslash(A/A), (V/A)\backslash(A/A)\big), \arrow[l, "\ell_*"]
\end{tikzcd}$$

where i_* and j_* have already been shown to be isomorphisms, and k_* and ℓ_* are also induced by the obvious inclusions. The excision axiom implies that both of the latter are isomorphisms; note that this is where it is crucial for V to be a *neighborhood* of the closed subset A, as the interior of V therefore contains the closure of A. The rightmost map labeled q_* in this diagram is also an isomorphism since it is induced by the map

$$(X\backslash A, V\backslash A) \xrightarrow{q} ((X/A)\backslash(A/A), (V/A)\backslash(A/A)),$$

which is a homeomorphism. We can now follow a path of isomorphisms from $h_*(X, A)$ all the way to the right of the diagram, then down, then back all the way to $h_*(X/A, A/A)$ at the left, proving that the leftmost map labeled q_* is also an isomorphism. $\square$

Remark 32.24. We used the strong version of the excision axiom in the above proof since the interior subsets A or A/A being removed from the pairs (X, V) or $(X/A, V/A)$ respectively are closed and typically not open. To make the argument work with the weaker version of excision mentioned in Remark 32.4, one would need to impose an extra requirement on the pair (X, A) so that instead of removing A we could remove a homotopy-equivalent open neighborhood of it that lives inside the neighborhood V. This would make the conditions defining a *good pair* slightly stricter, but they would still be satisfied by almost all pairs we will ever care about.

The following simple example will appear frequently when we compute the homology of CW-complexes.

Example 32.25. Since collapsing the boundary of the disk $\mathbb{D}^n$ produces a sphere $\mathbb{D}^n/\partial\mathbb{D}^n \cong S^n$, the theorem implies

$$h_k(\mathbb{D}^n, \partial\mathbb{D}^n) \cong \widetilde{h}_k(\mathbb{D}^n/\partial\mathbb{D}^n) \cong \widetilde{h}_k(S^n) \cong \begin{cases} G & \text{if } k = n, \\ 0 & \text{otherwise,} \end{cases}$$

where $G := h_0(\{\text{pt}\})$ is the coefficient group. (Of course it is also not hard to compute this more directly using the reduced long exact sequence of $(\mathbb{D}^n, \partial\mathbb{D}^n)$, in which the connecting homomorphism $h_k(\mathbb{D}^n, \partial\mathbb{D}^n) \to \widetilde{h}_{k-1}(S^{n-1})$ is an isomorphism.)

33. The Mayer-Vietoris sequence (November 7, 2023)

It is time to discuss the analogue in homology of the Seifert-van Kampen theorem.

Absolute Mayer-Vietoris in singular homology. The problem is as follows: we are given a space $X = A \cup B$ such that the interiors of A and B cover X, and we would like to compute $H_*(X)$ in terms of $H_*(A)$, $H_*(B)$ and $H_*(A \cap B)$. Let us first consider this problem specifically for singular homology, and after that consider the generalization to arbitrary axiomatic homology theories.

In singular homology, we already know from Lemma 29.21 that the inclusion of subgroups

$$C_*(A) + C_*(B) \hookrightarrow C_*(X)$$

is a chain homotopy equivalence, so it induces an isomorphism of homology groups

$$H_*(C_*(A) + C_*(B)) \xrightarrow{\cong} H_*(X). \tag{33.1}$$

The question is then how to relate the homology group on the left hand side of this isomorphism to the homologies of the individual chain complexes $C_*(A)$ and $C_*(B)$. By now, you will perhaps not be surprised to learn that the answer involves an exact sequence. We can define a short exact sequence of chain complexes

$$0 \longrightarrow C_*(A \cap B) \xrightarrow{(i^A_*, -i^B_*)} C_*(A) \oplus C_*(B) \xrightarrow{k^A \oplus k^B} C_*(A) + C_*(B) \longrightarrow 0,$$

where $i^A : A \cap B \hookrightarrow A$ and $i^B : A \cap B \hookrightarrow B$ are the obvious inclusions of spaces, and $k^A : C_*(A) \hookrightarrow C_*(A) + C_*(B)$ and $k^B : C_*(B) \hookrightarrow C_*(A) + C_*(B)$ are the inclusions of subgroups of $C_*(X)$. It is trivial to verify that this sequence is exact; the crucial detail for this is the sign reversal in the map $C_*(A \cap B) \to C_*(A) \oplus C_*(B)$, which sends chains $c \in C_n(A \cap B)$ to $(c, -c) \in C_n(A) \oplus C_n(B)$, where we have simplified the notation since singular chains in $A \cap B$ are naturally also singular chains in each of the spaces A and B. The short exact sequence induces a long exact sequence of homology groups by Proposition 28.18, but we need to take a moment to consider what the homology of each of these chain complexes actually is. The first is straightforward: it is simply $H_*(A \cap B)$. For the middle term, we are defining the boundary map on the chain complex $C_*(A) \oplus C_*(B)$ in the obvious way that preserves the direct sum splitting, so the homology of this complex is also a direct sum, namely $H_*(A) \oplus H_*(B)$. The last term has already been discussed: its homology is isomorphic via the obvious inclusion of chain complexes to $H_*(X)$, and the resulting map $H_*(A) \oplus H_*(B) \to H_*(X)$ is then simply $j^A_* \oplus j^B_*$, induced by the continuous inclusion maps $j^A : A \hookrightarrow X$ and $j^B : B \hookrightarrow X$. We've proved:

THEOREM 33.1. *If $A, B \subset X$ are subsets such that $X = \mathring{A} \cup \mathring{B}$ and*

$$i^A : A \cap B \hookrightarrow A, \quad i^B : A \cap B \hookrightarrow B, \quad j^A : A \hookrightarrow X, \quad j^B : B \hookrightarrow X,$$

denote the obvious inclusions, then there exists for each $n \in \mathbb{Z}$ a connecting homomorphism $\partial_* : H_n(X) \to H_{n-1}(A \cap B)$ such that the sequence

$$\begin{aligned} \ldots \longrightarrow H_{n+1}(X) \xrightarrow{\partial_*} H_n(A \cap B) &\xrightarrow{(i^A_*, -i^B_*)} H_n(A) \oplus H_n(B) \\ &\xrightarrow{j^A_* \oplus j^B_*} H_n(X) \xrightarrow{\partial_*} H_{n-1}(A \cap B) \to \ldots \end{aligned}$$

is exact. Moreover, the connecting homomorphisms are natural in the sense that for any map $f : X \to X'$ such that $X' = \mathring{A}' \cup \mathring{B}'$ for subsets $A', B' \subset X'$ with $f(A) \subset A'$ and $f(B) \subset B'$, the diagram

$$\begin{tikzcd} H_n(X) \arrow[r, "\partial_*"] \arrow[d, "f_*"] & H_{n-1}(A \cap B) \arrow[d, "f_*"] \\ H_n(X') \arrow[r, "\partial_*"] & H_{n-1}(A' \cap B') \end{tikzcd}$$

commutes. □

The exact sequence in this theorem is known as the **Mayer-Vietoris sequence** for the homology of $X = \mathring{A} \cup \mathring{B}$. Note that the naturality part of the statement follows from the functoriality of the short→long exact sequence construction, cf. Proposition 28.18 and Exercise 28.21.

EXERCISE 33.2. Prove that a formula for $\partial_* : H_n(X) \to H_{n-1}(A \cap B)$ can be written as follows. After applying barycentric subdivision to n-chains in X as in Lecture 29, every element of $H_n(X)$ can be written in the form $[a+b]$ for singular n-chains $a \in C_n(A)$ and $b \in C_n(B)$, so the condition that $a+b$ is a cycle implies $\partial a = -\partial b \in C_{n-1}(A \cap B)$. The formula for ∂_* is then

$$\partial_*[a+b] = [\partial a] = -[\partial b] \in H_{n-1}(A \cap B).$$

Axiomatic and reduced Mayer-Vietoris. I promise to look at an example or two, but first we have a theoretical issue to deal with. The construction above seems to depend heavily on the details of the singular chain complex rather than just the axioms of a homology theory—but does it really? As with the exact sequence of triples in the previous lecture, the answer is no, but a rather clever diagram-chase is required in order to prove it. I will show you the diagram and let you work out the details as an exercise.

Assume h_* is an arbitrary axiomatic homology theory and $X = \mathring{A} \cup \mathring{B}$. Notice first that the inclusion of pairs $(B, A \cap B) \hookrightarrow (X, A)$ induces an excision isomorphism

$$h_*(B, A \cap B) \xrightarrow{\cong} h_*(X, A).$$

Indeed, $(B, A \cap B)$ is obtained from (X, A) by removing the subset $A \backslash (A \cap B) = A \cap (X \backslash B)$, whose closure is in $\mathring{A}$ because there must otherwise exist a point x that is not in $\mathring{A}$ and therefore is in $\mathring{B}$, while simultaneously $x \in \overline{A \backslash (A \cap B)}$, meaning every neighborhood of x contains points that are in A and not in B, contradicting $x \in \mathring{B}$. We can now place the long exact sequences of the pairs $(B, A \cap B)$ and (X, A) into the top and bottom rows of the diagram

$$\begin{tikzcd}
\ldots \arrow[r] & h_{n+1}(B) \arrow[r] \arrow[d, "j_*^B"] & h_{n+1}(B, A \cap B) \arrow[r] \arrow[d, "\cong"] & h_n(A \cap B) \arrow[r, "i_*^B"] \arrow[d, "i_*^A"] & h_n(B) \arrow[r] \arrow[d, "j_*^B"] & h_n(B, A \cap B) \arrow[r] \arrow[d, "\cong"] & \ldots \\
\ldots \arrow[r] & h_{n+1}(X) \arrow[r] & h_{n+1}(X, A) \arrow[r] & h_n(A) \arrow[r, "j_*^A"] & h_n(X) \arrow[r] & h_n(X, A) \arrow[r] & \ldots
\end{tikzcd}$$

Here the rows are both exact, all maps in the diagram are either the connecting homomorphisms from long exact sequences of pairs or are induced by obvious inclusion maps, and as usual the diagram commutes due to the commutativity of those inclusion maps plus the naturality of connecting homomorphisms. The following diagram-chasing exercise transforms this diagram into a Mayer-Vietoris sequence for $h_*(X)$:

EXERCISE 33.3. Assume the following diagram commutes and that its top and bottom rows (both including the A_* terms) are exact.

$$\begin{tikzcd}
\cdots \arrow[r] & C_{n+1} \arrow[dr, "f"] \arrow[dd, "\ell"] & & B_n \arrow[r, "j"] \arrow[dd, "m"] & C_n \arrow[dr, "f"] \arrow[dd, "\ell"] & & B_{n-1} \arrow[r] \arrow[dd, "m"] & \cdots \\
& & A_n \arrow[ur, "h"] \arrow[dr, "i"] & & & A_{n-1} \arrow[ur, "h"] \arrow[dr, "i"] & & \\
\cdots \arrow[r] & E_{n+1} \arrow[ur, "g"] & & D_n \arrow[r, "k"] & E_n \arrow[ur, "g"] & & D_{n-1} \arrow[r] & \cdots
\end{tikzcd}$$

Prove that

$$\cdots \longrightarrow E_{n+1} \xrightarrow{h \circ g} B_n \xrightarrow{(m,-j)} D_n \oplus C_n \xrightarrow{k \oplus \ell} E_n \xrightarrow{h \circ g} B_{n-1} \longrightarrow \cdots$$

is then an exact sequence.

The exercise produces not only an axiomatic version of Theorem 33.1, but also a formula for $\partial_* : h_n(X) \to h_{n-1}(A \cap B)$. In the abstract diagram of the exercise it is the composition $h \circ g : E_n \to B_{n-1}$. In our situation, $g : E_n \to A_{n-1}$ is just the map $h_n(X) \to h_n(X, A)$ induced by the inclusion $(X, \varnothing) \hookrightarrow (X, A)$. We then need to replace $h_n(X, A)$ by $h_n(B, A \cap B)$ using the inverse of the excision map $h_n(B, A \cap B) \xrightarrow{\cong} h_n(X, A)$, so that $h : A_{n-1} \to B_{n-1}$ becomes the connecting homomorphism $h_n(B, A \cap B) \to h_{n-1}(A \cap B)$ from the exact sequence of $(B, A \cap B)$. To summarize:

THEOREM 33.4. *In the setting of Theorem 33.1, any axiomatic homology theory h_* also admits a Mayer-Vietoris exact sequence*

$$\ldots \longrightarrow h_n(A \cap B) \xrightarrow{(i_*^A, -i_*^B)} h_n(A) \oplus h_n(B) \xrightarrow{j_*^A \oplus j_*^B} h_n(X) \xrightarrow{\partial_*} h_{n-1}(A \cap B) \to \ldots,$$

with natural connecting homomorphisms $\partial_* : h_n(X) \to h_{n-1}(A \cap B)$ determined by the diagram

$$\begin{tikzcd} h_n(X) \arrow[r, "\partial_*"] \arrow[d] & h_{n-1}(A \cap B) \\ h_n(X, A) & h_n(B, A \cap B) \arrow[l, hook', "\cong"'] \arrow[u] \end{tikzcd} \tag{33.2}$$

which is formed out of the obvious inclusions $(B, A \cap B) \hookrightarrow (X, A)$ and $(X, \varnothing) \mapsto (X, A)$ together with the connecting homomorphism from the long exact sequence of the pair $(B, A \cap B)$. □

REMARK 33.5. The naturality of the connecting homomorphism $\partial_* : h_n(X) \to h_{n-1}(A \cap B)$ can be deduced directly from the diagram (33.2).

There is also an analogue for reduced homology, which follows from a similar argument to the one we used for the long exact sequence of the pair. Indeed, we can put the Mayer-Vietoris sequence for $X = \mathring{A} \cup \mathring{B}$ into a commutative diagram together with the Mayer-Vietoris sequence for $\{\text{pt}\} = \{\text{pt}\} \cup \{\text{pt}\}$ as follows:

$$\begin{tikzcd}[column sep=small] & 0 \arrow[d] & 0 \arrow[d] & 0 \arrow[d] & 0 \arrow[d] & \\ \ldots & \widetilde{h}_n(A \cap B) \arrow[d, hook] & \widetilde{h}_n(A) \oplus \widetilde{h}_n(B) \arrow[d, hook] & \widetilde{h}_n(X) \arrow[d, hook] & \widetilde{h}_{n-1}(A \cap B) \arrow[d, hook] & \ldots \\ \ldots \arrow[r] & h_n(A \cap B) \arrow[r] \arrow[d, "\epsilon_*"] & h_n(A) \oplus h_n(B) \arrow[r] \arrow[d, "\epsilon_* \oplus \epsilon_*"] & h_n(X) \arrow[r] \arrow[d, "\epsilon_*"] & h_{n-1}(A \cap B) \arrow[r] \arrow[d, "\epsilon_*"] & \ldots \\ \ldots \arrow[r] & h_n(\{\text{pt}\} \cap \{\text{pt}\}) \arrow[r] \arrow[d] & h_n(\{\text{pt}\}) \oplus h_n(\{\text{pt}\}) \arrow[r] \arrow[d] & h_n(\{\text{pt}\}) \arrow[r] \arrow[d] & h_{n-1}(\{\text{pt}\} \cap \{\text{pt}\}) \arrow[r] \arrow[d] & \ldots \\ & 0 & 0 & 0 & 0 & \end{tikzcd}$$

Since all columns of this diagram are exact and so are the bottom two nontrivial rows, Proposition 29.9 provides uniquely determined maps on the top row that preserve the commutativity of the diagram and make the top row exact:

THEOREM 33.6. *In the setting of Theorem 33.4, all maps can be restricted to the respective reduced homology groups to produce an exact sequence*

$$\ldots \longrightarrow \widetilde{h}_n(A \cap B) \xrightarrow{(i_*^A, -i_*^B)} \widetilde{h}_n(A) \oplus \widetilde{h}_n(B) \xrightarrow{j_*^A \oplus j_*^B} \widetilde{h}_n(X) \xrightarrow{\partial_*} \widetilde{h}_{n-1}(A \cap B) \to \ldots,$$

□

REMARK 33.7. If one prefers to assume only the weaker version of the excision axiom as mentioned in Remark 32.4, then some additional condition must be imposed on the subsets $A, B \subset X$ in Theorem 33.4 to make sure that $(B, A \cap B) \hookrightarrow (X, A)$ is a valid excision map, or equivalently, that $A \backslash (A \cap B) = A \cap (X \backslash B)$ is open. It suffices for instance to assume that B is closed, so in particular, everything is fine in the category of compact pairs, where only closed subsets $A, B \subset X$ are allowed.

Some sample computations. Now let's see the Mayer-Vietoris sequence in action.

EXAMPLE 33.8. Here's the *easy* way to see the isomorphism $\widetilde{h}_n(X) \cong \widetilde{h}_{n+1}(SX)$. Write $SX = C_+X \cup_X C_-X$ as usual where $C_+X = (X \times [0,1])/(X \times \{1\})$ and $C_-X = (X \times [-1,0])/(X \times \{-1\})$, pick $\epsilon > 0$ small and define two new subsets of SX by

$$A = (X \times [-\epsilon, 1])/(X \times \{1\}), \qquad B = (X \times [-1, \epsilon])/(X \times \{-1\}).$$

In other words, $A \subset SX$ is a neighborhood of C_+X that also deformation-retracts to C_+X, and B is similarly a neighborhood of C_-X, so that the interiors of A and B cover SX. Notice also that the intersection $A \cap B = X \times [-\epsilon, \epsilon]$ deformation-retracts to X. We can then write down the Mayer-Vietoris sequence

$$\ldots \to \widetilde{h}_{n+1}(A) \oplus \widetilde{h}_{n+1}(B) \to \widetilde{h}_{n+1}(SX) \to \widetilde{h}_n(A \cap B) \to \widetilde{h}_n(A) \oplus \widetilde{h}_n(B) \to \ldots$$

By homotopy invariance, $\widetilde{h}_*(A) \cong \widetilde{h}_*(C_+X) = 0$ and $\widetilde{h}_*(B) \cong \widetilde{h}_*(C_-X) = 0$, while $\widetilde{h}_*(A \cap B) \cong \widetilde{h}_*(X)$, so we obtain from this an exact sequence $0 \to h_{n+1}(SX) \to h_n(X) \to 0$ and thus an isomorphism $h_{n+1}(SX) \to h_n(X)$.

EXAMPLE 33.9. Let's compute the singular homology $H_*(\mathbb{T}^2; \mathbb{Z})$ of the 2-torus with integer coefficients. We can decompose $\mathbb{T}^2 = S^1 \times S^1$ as the union of two overlapping annuli: simply write S^1 as the union of two intervals $I_+, I_- \subset S^1$ whose interiors cover the whole circle, then set $A = I_+ \times S^1$ and $B = I_- \times S^1$. This makes both A and B homotopy equivalent to S^1, and since the intersection $I_+ \cap I_-$ necessarily contains two disjoint intervals, $A \cap B$ is homotopy equivalent to $S^1 \amalg S^1$. We thus have

$$H_k(A; \mathbb{Z}) \cong H_k(B; \mathbb{Z}) \cong H_k(S^1; \mathbb{Z}) \cong \mathbb{Z}, \quad \text{and} \quad H_k(A \cap B; \mathbb{Z}) \cong H_k(S^1; \mathbb{Z}) \oplus H_k(S^1; \mathbb{Z}) \cong \mathbb{Z} \oplus \mathbb{Z}$$

for $k = 0, 1$, where the two factors in $H_k(A \cap B; \mathbb{Z})$ correspond to the two connected components of the overlap $A \cap B$. For $k \geqslant 2$, we have $H_k(A; \mathbb{Z}) = H_k(B; \mathbb{Z}) = H_k(A \cap B; \mathbb{Z}) = 0$, so the Mayer-Vietoris sequence contains a segment of the form

$$0 = H_k(A; \mathbb{Z}) \oplus H_k(B; \mathbb{Z}) \to H_k(\mathbb{T}^2; \mathbb{Z}) \to H_{k-1}(A \cap B; \mathbb{Z}) = 0 \quad \text{for } k \geqslant 3,$$

implying $H_k(\mathbb{T}^2; \mathbb{Z}) = 0$ for all $k \geqslant 3$.

To understand $H_k(\mathbb{T}^2; \mathbb{Z})$ for $k \leqslant 2$, we need to know more about the actual maps in the Mayer-Vietoris sequence. A crucial observation here is that the inclusions $i^A : A \cap B \hookrightarrow A$ and $i^B : A \cap B \hookrightarrow B$ restrict to each of the connected components of $A \cap B$ as homotopy equivalences, thus they induce isomorphisms

$$\mathbb{Z} \xrightarrow{\cong} \mathbb{Z} \cong H_k(A; \mathbb{Z}) \text{ or } H_k(B; \mathbb{Z})$$

when restricted to each of the factors in $H_k(A \cap B; \mathbb{Z}) \cong \mathbb{Z} \oplus \mathbb{Z}$. We are free to choose the isomorphisms of $H_k(A; \mathbb{Z})$ and $H_k(B; \mathbb{Z})$ with $\mathbb{Z}$ so that each of these maps $\mathbb{Z} \to \mathbb{Z}$ is the identity, and we can therefore write the map $\Phi_k := (i^A_*, -i^B_*) : H_k(A \cap B; \mathbb{Z}) \to H_k(A; \mathbb{Z}) \oplus H_k(B; \mathbb{Z})$ for $k = 0, 1$ in the form

$$\Phi_k : \mathbb{Z} \oplus \mathbb{Z} \to \mathbb{Z} \oplus \mathbb{Z} : (m, n) \mapsto (m + n, -m - n). \tag{33.3}$$

Since $H_2(A; \mathbb{Z}) = H_2(B; \mathbb{Z}) = H_2(S^1; \mathbb{Z}) = 0$, let us use this term to begin the sequence and write it as

$$0 = H_2(A; \mathbb{Z}) \oplus H_2(B; \mathbb{Z}) \longrightarrow H_2(\mathbb{T}^2; \mathbb{Z}) \xrightarrow{\partial_2} H_1(A \cap B; \mathbb{Z}) \xrightarrow{\Phi_1} H_1(A; \mathbb{Z}) \oplus H_1(B; \mathbb{Z})$$
$$\longrightarrow H_1(\mathbb{T}^2; \mathbb{Z}) \xrightarrow{\partial_1} H_0(A \cap B; \mathbb{Z}) \xrightarrow{\Phi_0} H_0(A; \mathbb{Z}) \oplus H_0(B; \mathbb{Z}) \longrightarrow H_0(\mathbb{T}^2; \mathbb{Z}) \longrightarrow 0.$$

The first thing we can deduce from this is that the map labeled ∂_2 is injective, and its image is $\ker \Phi_1$, which by the formula in (33.3), is

$$\ker \Phi_1 \cong \{(m,n) \in \mathbb{Z} \oplus \mathbb{Z} \mid m + n = 0\} \cong \mathbb{Z}.$$

This proves $H_2(\mathbb{T}^2;\mathbb{Z}) \cong \mathbb{Z}$.

To compute $H_1(\mathbb{T}^2;\mathbb{Z})$, we can use a convenient trick for turning long exact sequences into short ones: observe first that while the map $H_1(A;\mathbb{Z}) \oplus H_1(B;\mathbb{Z}) \to H_1(\mathbb{T}^2;\mathbb{Z})$ cannot be assumed injective, it will become injective if we quotient its domain by its kernel, which is precisely the image of Φ_1. Similarly, $\partial_1 : H_1(\mathbb{T}^2;\mathbb{Z}) \to H_0(A \cap B;\mathbb{Z})$ may fail to be surjective, but it trivially becomes surjective if we replace its target space with $\operatorname{im} \partial_1$, which equals $\ker \Phi_0$. We therefore have a short exact sequence

$$0 \to \operatorname{coker} \Phi_1 \to H_1(\mathbb{T}^2;\mathbb{Z}) \to \ker \Phi_0 \to 0,$$

where the **cokernel** of Φ_1 is defined as the quotient of its target by its image, i.e.

$$\operatorname{coker} \Phi_1 := (H_1(A;\mathbb{Z}) \oplus H_1(B;\mathbb{Z})) \,/\, \operatorname{im} \Phi_1.$$

Using (33.3) again, $\operatorname{coker} \Phi_1$ is then the quotient of $\mathbb{Z} \oplus \mathbb{Z}$ by the subgroup $\{(m,-m) \mid m \in \mathbb{Z}\}$, and this quotient is isomorphic to $\mathbb{Z}$. As luck would have it, the same subgroup is also $\ker \Phi_0$, and it is also isomorphic to $\mathbb{Z}$, so our short exact sequence now looks like

$$0 \to \mathbb{Z} \to H_1(\mathbb{T}^2;\mathbb{Z}) \to \mathbb{Z} \to 0. \tag{33.4}$$

At this point there is a cheap trick available to finish the job: $\mathbb{Z}$ is a free group, thus the sequence splits by Exercise 29.3, and Exercise 29.1 then provides an isomorphism $H_1(\mathbb{T}^2;\mathbb{Z}) \cong \mathbb{Z} \oplus \mathbb{Z}$.

Finally, the sequence also gives us a surjective map $H_0(A;\mathbb{Z}) \oplus H_0(B;\mathbb{Z}) \to H_0(\mathbb{T}^2;\mathbb{Z})$ whose kernel is $\operatorname{im} \Phi_0$, so it descends to an isomorphism $\operatorname{coker} \Phi_0 \to H_0(\mathbb{T}^2;\mathbb{Z})$. Here (33.3) implies once again that $\operatorname{coker} \Phi_0 \cong \mathbb{Z}$.

The end result is:

$$H_k(\mathbb{T}^2;\mathbb{Z}) \cong \begin{cases} \mathbb{Z} & \text{for } k = 0, 2, \\ \mathbb{Z} \oplus \mathbb{Z} & \text{for } k = 1, \\ 0 & \text{for all other } k \in \mathbb{Z}. \end{cases}$$

It's worth noting that we did not use any properties of singular homology in the above computation beyond the axioms, so the result is equally valid for any axiomatic homology theory with coefficient group $\mathbb{Z}$. We did use a specific property of the coefficient group at one step: we assumed it was *free* in order to conclude that the sequence (33.4) splits. More generally, if h_* is any homology theory with a free coefficient group G, then since $h_0(S^1) \cong h_1(S^1) \cong G$, the same argument gives the result $h_0(\mathbb{T}^2) \cong h_2(\mathbb{T}^2) \cong G$ and $h_1(\mathbb{T}^2) \cong G \oplus G$. Actually, this is true for *all* coefficient groups, but we cannot see it so easily from this argument—we will see a much easier proof of this once we've learned how to compute cellular homology, and for singular homology it will also follow by combining the above computation with the universal coefficient theorem.

One drawback of the above method for computing $H_*(\mathbb{T}^2)$ is that Mayer-Vietoris does not make it very easy to say e.g. what the two generators of $H_1(\mathbb{T}^2;\mathbb{Z}) \cong \mathbb{Z} \oplus \mathbb{Z}$ look like *geometrically.* Of course, you already know the answer to this question for other reasons: its generators are the same as those of $\pi_1(\mathbb{T}^2)$, i.e. they can be represented by loops of the form $S^1 \times \{\text{const}\}$ and $\{\text{const}\} \times S^1$ in $\mathbb{T}^2$. It is not impossible, but by no means straightforward, to find both of those loops hiding in the exact sequence argument we used above.

The situation is slightly better for $H_2(\mathbb{T}^2;\mathbb{Z})$. As we proved in Lecture 31, any fundamental class $[\mathbb{T}^2] \in H_2(\mathbb{T}^2;\mathbb{Z})$ constructed via an oriented triangulation of $\mathbb{T}^2$ will be a primitive element, and since we know in the mean time that $H_2(\mathbb{T}^2;\mathbb{Z}) \cong \mathbb{Z}$, that means that $[\mathbb{T}^2]$ generates it. One can also see this more concretely as follows. Choose a pair of parallel circles $S_+, S_- \subset A \cap B$ such

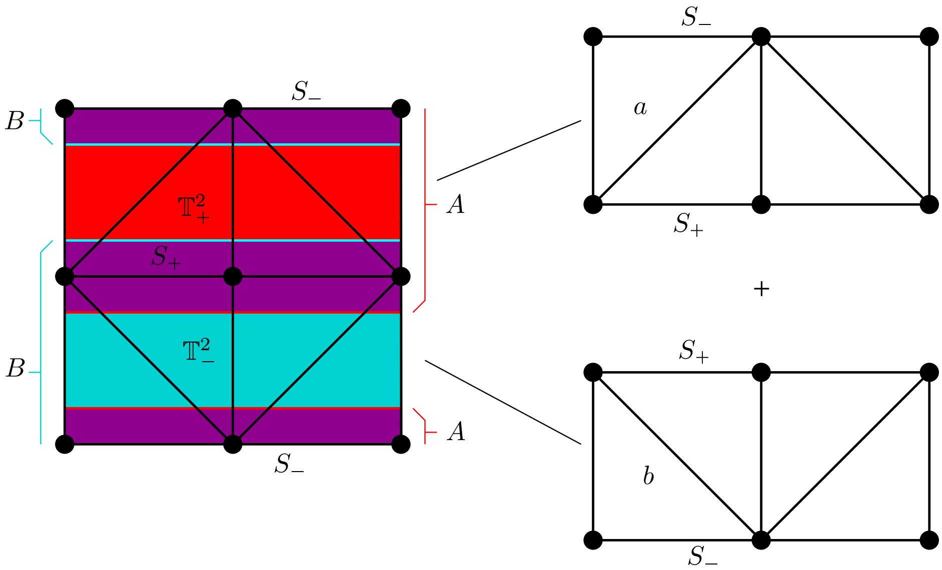

FIGURE 19. The square at the left (with opposite edges identified) represents $\mathbb{T}^2$, which is split by two parallel circles $S_+, S_- \subset \mathbb{T}^2$ into a pair of annuli $\mathbb{T}^2_+$ and $\mathbb{T}^2_-$ such that $A \subset \mathbb{T}^2$ is a neighborhood of $\mathbb{T}^2_+$ and $B \subset \mathbb{T}^2$ is a neighborhood of $\mathbb{T}^2_-$. The chosen triangulation of $\mathbb{T}^2$ respects this splitting and produces a 2-cycle of the form $a + b \in C_2(\mathbb{T}^2;\mathbb{Z})$ where $a \in C_2(A;\mathbb{Z})$ and $b \in C_2(B;\mathbb{Z})$, such that $[a + b] \in H_2(\mathbb{T}^2;\mathbb{Z}) \cong \mathbb{Z}$ is a generator.

that each connected component of $A \cap B$ admits a deformation retraction onto one of them. Since $H_1(S_\pm;\mathbb{Z}) \cong \pi_1(S_\pm)$, any loop parametrizing $S_\pm$ can then be viewed as a 1-cycle representing a generator of one of the factors in $H_1(A \cap B;\mathbb{Z}) \cong \mathbb{Z} \oplus \mathbb{Z}$; more generally, any oriented triangulation of $S_\pm$ does the same thing, since a triangulation in this case just means a string of 1-simplices that can be concatenated to form a loop parametrizing the whole circle. Let us denote the generators defined in this way by $[S_+], [S_-] \in H_1(A \cap B;\mathbb{Z})$. The Mayer-Vietoris sequence told us that the connecting homomorphism

$$\partial_* : H_2(\mathbb{T}^2;\mathbb{Z}) \to H_1(A \cap B;\mathbb{Z})$$

maps $H_2(\mathbb{T}^2;\mathbb{Z}) \cong \mathbb{Z}$ isomorphically to the kernel of $\Phi_1 : H_1(A \cap B;\mathbb{Z}) \to H_1(A;\mathbb{Z}) \oplus H_1(B;\mathbb{Z})$, which we can deduce from (33.3) is generated by $[S_+] - [S_-]$. According to Exercise 33.2, a generator of $H_2(\mathbb{T}^2;\mathbb{Z})$ can thus be represented by any singular 2-cycle of the form $a + b$ where $a \in C_2(A;\mathbb{Z})$, $b \in C_2(B;\mathbb{Z})$, and $\partial a = -\partial b$ is a 1-cycle in $A \cap B$ representing $[S_+] - [S_-]$. To find such a 2-cycle, observe that $\mathbb{T}^2$ is the union of a pair of annuli $\mathbb{T}^2_+, \mathbb{T}^2_- \subset \mathbb{T}^2$ that each have boundary $S_+ \amalg S_-$ and intersect each other only along that boundary. Choose an oriented triangulation of $\mathbb{T}^2$ in which every 2-simplex is in either $\mathbb{T}^2_+$ or $\mathbb{T}^2_-$; an example is shown in Figure 19. As explained in Lecture 31, the sum of these oriented simplices defines a singular 2-cycle after choosing an ordering of all the vertices; that cycle represents a fundamental class $[\mathbb{T}^2] \in H_2(\mathbb{T}^2;\mathbb{Z})$. That cycle clearly also has the property that (up to a sign) $\partial_*[\mathbb{T}^2] = [S_+] - [S_-]$, so if we didn't already know that $[\mathbb{T}^2]$ is a generator of $H_2(\mathbb{T}^2;\mathbb{Z})$, we could now also deduce that fact from the Mayer-Vietoris sequence.

EXERCISE 33.10. Use Mayer-Vietoris sequences to compute $H_*(X;\mathbb{Z})$ and $H_*(X;\mathbb{Z}_2)$, where X is

(a) The projective plane $\mathbb{RP}^2$.
(b) The Klein bottle.

Hint: $\mathbb{RP}^2$ is the union of a disk with a Möbius band, and the latter admits a deformation retraction to S^1. The Klein bottle, in turn, is the union of two Möbius bands, also known as $\mathbb{RP}^2 \# \mathbb{RP}^2$.

Exercise 33.11. Recall that given two connected topological n-manifolds X and Y, their **connected sum** $X \# Y$ is defined by deleting an open n-disk $\mathring{\mathbb{D}}^n$ from each of X and Y and then gluing $X \backslash \mathring{\mathbb{D}}^n$ and $Y \backslash \mathring{\mathbb{D}}^n$ together along an identification of their boundary spheres:

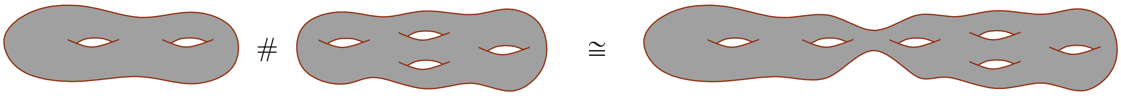

More precisely, we can choose topological embeddings $\iota_X : \mathbb{D}^n \hookrightarrow X$, $\iota_Y : \mathbb{D}^n \hookrightarrow Y$ of the closed unit n-disk $\mathbb{D}^n \subset \mathbb{R}^n$ and then define

$$X \# Y := \left(X \backslash \iota_X(\mathring{\mathbb{D}}^n) \right) \cup_{S^{n-1}} \left(Y \backslash \iota_Y(\mathring{\mathbb{D}}^n) \right),$$

where the gluing identifies the boundaries of both pieces in the obvious way with $S^{n-1} = \partial \mathbb{D}^n$. There are one or two subtle issues about the extent to which $X \# Y$ is (up to homeomorphism) independent of choices, e.g. in general this need not be true without an extra condition involving orientations, but don't worry about this for now. Last semester (see Exercise 14.12) we used the Seifert-van Kampen theorem to show that $\pi_1(X \# Y) \cong \pi_1(X) * \pi_1(Y)$ whenever $n \geqslant 3$. We can now use the Mayer-Vietoris sequence to derive a similar formula for the homology of a connected sum.

(a) Prove that for any $k = 1, \dots, n-2$ and any coefficient group, $H_k(X \# Y) \cong H_k(X) \oplus H_k(Y)$.
Hint: There are two steps, as you first need to derive a relation between $H_k(X)$ and $H_k(X \backslash \mathring{\mathbb{D}}^n)$, and then see what happens when you glue $X \backslash \mathring{\mathbb{D}}^n$ and $Y \backslash \mathring{\mathbb{D}}^n$ together.
(b) It turns out that the formula $H_{n-1}(X \# Y; \mathbb{Z}) \cong H_{n-1}(X; \mathbb{Z}) \oplus H_{n-1}(Y; \mathbb{Z})$ also holds if X and Y are both closed orientable n-manifolds with $n \geqslant 2$, and without orientability we still have $H_{n-1}(X \# Y; \mathbb{Z}_2) \cong H_{n-1}(X; \mathbb{Z}_2) \oplus H_{n-1}(Y; \mathbb{Z}_2)$. Prove this under the following additional assumption: $X \backslash \mathring{\mathbb{D}}^n$ and $Y \backslash \mathring{\mathbb{D}}^n$ both admit (possibly oriented) triangulations for which the induced triangulations of $\partial(X \backslash \mathring{\mathbb{D}}^n) = \partial(Y \backslash \mathring{\mathbb{D}}^n) = S^{n-1}$ each define generators of $H_{n-1}(S^{n-1}; \mathbb{Z}_2)$ or (in the oriented case) $H_{n-1}(S^{n-1}; \mathbb{Z})$.
(c) Find a counterexample to the formula $H_1(X \# Y; \mathbb{Z}) \cong H_1(X; \mathbb{Z}) \oplus H_1(Y; \mathbb{Z})$ where X and Y are both closed (but not necessarily orientable) 2-manifolds.

The relative version. We will not need the relative Mayer-Vietoris sequence until much later in this course, and we will only need it for singular homology in particular (not for other axiomatic homologies), so let us sketch where it comes from in that case. Assume (X, Y), (A, C) and (B, D) are pairs of spaces such that $X = \mathring{A} \cup \mathring{B}$ and $Y = \mathring{C} \cup \mathring{D}$. The subcomplex $C_*(A+B) = C_*(A) + C_*(B) \subset C_*(X)$ can be quotiented by $C_*(C+D) = C_*(C) + C_*(D) \subset C_*(Y)$ to define another chain complex

$$C_*(A+B, C+D) := C_*(A+B)/C_*(C+D) = \frac{C_*(A) + C_*(B)}{C_*(C) + C_*(D)},$$

for which the inclusion $C_*(A+B) \hookrightarrow C_*(X)$ descends to a chain map

$$C_*(A+B, C+D) \to C_*(X, Y). \tag{33.5}$$

We can now write down a commutative diagram of chain complexes and chain maps in the form
(33.6)

$$
\begin{tikzcd}
 & 0 \arrow[d] & 0 \arrow[d] & 0 \arrow[d] & \\
0 \arrow[r] & C_*(C \cap D) \arrow[r] \arrow[d] & C_*(C) \oplus C_*(D) \arrow[r] \arrow[d] & C_*(C + D) \arrow[r] \arrow[d] & 0 \\
0 \arrow[r] & C_*(A \cap B) \arrow[r] \arrow[d] & C_*(A) \oplus C_*(B) \arrow[r] \arrow[d] & C_*(A + B) \arrow[r] \arrow[d] & 0 \\
0 \arrow[r] & C_*(A \cap B, C \cap D) \arrow[r] \arrow[d] & C_*(A, C) \oplus C_*(B, D) \arrow[r] \arrow[d] & C_*(A + B, C + D) \arrow[r] \arrow[d] & 0 \\
 & 0 & 0 & 0 &
\end{tikzcd}
$$

The horizontal maps in this diagram are assumed to be the same ones that appear in the Mayer-Vietoris sequence, so the top two rows are just the short exact sequences that underlie the Mayer-Vietoris sequences for $Y = \mathring{C} \cup \mathring{D}$ and $X = \mathring{A} \cup \mathring{B}$ respectively. All three columns are exact for the usual reasons: the top vertical maps are actually just inclusions, and the bottom vertical maps are quotient projections (or direct sums of two such maps, in the case of the middle column). To understand the quotient complex at the lower right, we claim that the map induced on homology by the chain map (33.5) is an isomorphism, so that $C_*(A+B, C+D)$ can be used as a substitute for $C_*(X, Y)$ in computing $H_*(X, Y)$. We know already that the inclusion $C_*(A+B) \hookrightarrow C_*(X)$, which induces (33.5), is a chain homotopy equivalence, as is its restriction to the subgroup $C_*(C+D) \hookrightarrow C_*(Y)$. Moreover, there is a short exact sequence of chain complexes

$$0 \to C_*(C + D) \hookrightarrow C_*(A + B) \to C_*(A + B, C + D) \to 0,$$

which produces a long exact sequence of their corresponding homology groups. These form the top row of the following commutative diagram,

$$
\begin{tikzcd}
H_n(C_*(C + D)) \arrow[r] \arrow[d, "\cong"] & H_n(C_*(A + B)) \arrow[r] \arrow[d, "\cong"] & H_n(C_*(A + B, C + D)) \arrow[r] \arrow[d] & H_{n-1}(C_*(C + D)) \arrow[r] \arrow[d, "\cong"] & H_{n-1}(C_*(A + B)) \arrow[d, "\cong"] \\
H_n(Y) \arrow[r] & H_n(X) \arrow[r] & H_n(X, Y) \arrow[r] & H_{n-1}(Y) \arrow[r] & H_{n-1}(X)
\end{tikzcd}
$$

in which the bottom row is the long exact sequence of (X, Y). Since both rows are exact and four out of the five vertical maps are isomorphisms, the five-lemma (see Exercise 32.16) implies that the remaining vertical map is also an isomorphism, hence

$$H_*(C_*(A + B, C + D)) \cong H_*(X, Y).$$

Returning to the large diagram (33.6), it is obvious that the top two nontrivial rows are exact sequences, as are all three of the columns, and it is straightforward to deduce from this that the third row is also a chain complex, though it is not immediately obvious whether it is exact. If you recall Proposition 29.9 however, it should not surprise you to learn that the exactness of two of the rows implies the exactness of the third. Here is the clever way to see this: think of each row in the diagram as a chain complex, so that the vertical chain maps from top to bottom define a short exact sequence of chain complexes. (We normally draw exact sequences of chain complexes with the maps oriented horizontally and not vertically, but there's no law against doing it this way instead.) From this perspective, the short exact sequence induces a long exact sequence of homology groups, but here's the point: since the top two rows are exact, the homology groups they induce are all trivial, which means that two out of every three groups in our long exact sequence will be trivial. Exactness then forces all the other homology groups to be trivial as well, which is equivalent to the statement that the third row in the diagram is exact. Cute, no?

Putting all this together, since we can now recognize the third row in (33.6) as a short exact sequence of chain complexes, and we can identify the homology of the last term in it with $H_*(X,Y)$, the induced long exact sequence is what we will call the **relative Mayer-Vietoris sequence**

$$\begin{aligned} \ldots \to H_{n+1}(X,Y) \to H_n(A\cap B, C\cap D) \to H_n(A,C)\oplus H_n(B,D) \\ \to H_n(X,Y) \to H_{n-1}(A\cap B, C\cap D) \to \ldots \end{aligned}$$

We'll need this in the proof of Poincaré duality, but until then you are free to forget about it.

34. Mapping tori and maps between spheres (November 10, 2023)

The two topics of today's lecture are essentially independent of each other. The first is another computational tool involving an exact sequence, which permits an easy extension of our previous Mayer-Vietoris-based computation of $H_*(\mathbb{T}^2;\mathbb{Z})$ to $\mathbb{T}^n$ for all $n\in\mathbb{N}$. The second topic is the beginning of a larger discussion about the degrees of maps between closed equidimensional manifolds, which will extend into the next lecture and subsequently play an important role in the development of cellular homology.

Mapping tori. I'd like to talk about another way of computing the homology of $\mathbb{T}^2$ (and many other things), by viewing it as an example of a *mapping torus*.

Given a space X and a map $f : X \to X$, the **mapping torus** (*Abbildungstorus*) of f is defined to be the quotient space

$$X_f := (X\times I)/\sim, \qquad \text{where } (x,0)\sim(f(x),1) \text{ for all } x\in X.$$

We can regard X itself as a subspace of X_f via the inclusion map[54]

$$i : X \hookrightarrow X_f : x \mapsto [(x,1)].$$

THEOREM 34.1. *For any map $f : X\to X$ and its mapping torus X_f, every axiomatic homology theory h_* admits a long exact sequence*

$$\ldots \longrightarrow h_{k+1}(X_f) \longrightarrow h_k(X) \xrightarrow{\mathbb{1}-f_*} h_k(X) \xrightarrow{i_*} h_k(X_f) \longrightarrow h_{k-1}(X) \longrightarrow \ldots$$

Let's do an example before talking about the proof.

EXAMPLE 34.2. For each $n\in\mathbb{N}$, the n-torus $\mathbb{T}^n = S^1\times\ldots\times S^1$ is the mapping torus of the identity map $\mathrm{Id} : \mathbb{T}^{n-1}\to\mathbb{T}^{n-1}$, so the exact sequence of the mapping torus includes segments of the form

$$\ldots \longrightarrow h_k(\mathbb{T}^{n-1}) \xrightarrow{0} h_k(\mathbb{T}^{n-1}) \xrightarrow{i_*} h_k(\mathbb{T}^n) \xrightarrow{\Phi} h_{k-1}(\mathbb{T}^{n-1}) \xrightarrow{0} h_{k-1}(\mathbb{T}^{n-1}) \longrightarrow \ldots$$

The triviality of the two maps $\mathbb{1}-\mathrm{Id}_* = 0$ here means that we actually have a short exact sequence

$$0 \longrightarrow h_k(\mathbb{T}^{n-1}) \xrightarrow{i_*} h_k(\mathbb{T}^n) \xrightarrow{\Phi} h_{k-1}(\mathbb{T}^{n-1}) \longrightarrow 0. \tag{34.1}$$

Let us apply this sequence for singular homology with integer coefficients in the case $n=2$, so $\mathbb{T}^{n-1}=S^1$, and since $H_{k-1}(S^1;\mathbb{Z})$ is free for every k, the sequence splits, giving an isomorphism

$$H_k(\mathbb{T}^2;\mathbb{Z}) \cong H_k(S^1;\mathbb{Z})\oplus H_{k-1}(S^1;\mathbb{Z})$$

[54]Of course there are also natural maps $X\to X_f : x\mapsto[(x,t)]$ for every $t\in I$, and for our purposes it will not matter which one we pick since they are all obviously homotopic. The case $t=0$ is a little bit awkward however since it might not be injective—we have $[(x,0)]=[(y,0)]\in X_f$ whenever $f(x)=f(y)$.

for every k. By induction on $n \in \mathbb{N}$, we can now prove that all homology groups of the torus $\mathbb{T}^n$ for every n are free, so the sequence (34.1) again splits and gives

$$H_k(\mathbb{T}^n;\mathbb{Z}) \cong H_k(\mathbb{T}^{n-1};\mathbb{Z}) \oplus H_{k-1}(\mathbb{T}^{n-1};\mathbb{Z}).$$

This means that each $H_k(\mathbb{T}^n;\mathbb{Z})$ is isomorphic to $\mathbb{Z}^r$ for some integer $r \geqslant 0$, the **rank** (*Rang*) of the group, and these ranks satisfy $\operatorname{rank} H_k(\mathbb{T}^n;\mathbb{Z}) = \operatorname{rank} H_k(\mathbb{T}^{n-1};\mathbb{Z}) + \operatorname{rank} H_{k-1}(\mathbb{T}^{n-1};\mathbb{Z})$, so they are precisely the numbers in *Pascal's triangle*, i.e. the familiar binomial coefficients:

$$\operatorname{rank} H_k(\mathbb{T}^n;\mathbb{Z}) = \binom{n}{k} \quad \text{for } 0 \leqslant k \leqslant n, \qquad H_k(\mathbb{T}^n;\mathbb{Z}) = 0 \text{ for } k > n.$$

Exercise 34.3. The mapping torus of $f : S^1 \to S^1 : e^{i\theta} \mapsto e^{-i\theta}$ is homeomorphic to the Klein bottle K^2. Use Theorem 34.1 to compute $H_*(K^2;\mathbb{Z})$ and $H_*(K^2;\mathbb{Z}_2)$.

To prove the theorem, we shall first state a more general result that implies it. Given two spaces X, Y and maps $f, g : X \to Y$, define the space

$$Z := ((X \times I) \amalg Y)/\sim \qquad \text{where } (x,0) \sim f(x) \text{ and } (x,1) \sim g(x) \text{ for all } x \in X.$$

This space comes with a natural inclusion

$$i : Y \hookrightarrow Z,$$

and the special case with $X = Y$ and $g = \mathrm{Id}$ reproduces the mapping torus X_f of $f : X \to X$. Theorem 34.1 follows immediately from the next statement:

Theorem 34.4. *Given $f, g : X \to Y$ and the space Z described above, there exists a long exact sequence*

$$\ldots \longrightarrow h_{k+1}(Z) \longrightarrow h_k(X) \xrightarrow{g_* - f_*} h_k(Y) \xrightarrow{i_*} h_k(Z) \longrightarrow h_{k-1}(X) \longrightarrow \ldots$$

for every axiomatic homology theory h_*.

Remark 34.5. It is not too hard to see intuitively why the composition $i_* \circ (g_* - f_*)$ in this sequence is trivial. Imagine for instance a homology class of the form $a = j_*[M] \in H_n(X;\mathbb{Z})$ defined via a closed n-manifold M with an oriented triangulation and a map $j : M \to X$. This gives rise to a map $\tilde{\jmath} : M \times I \to X \times I : (x,t) \mapsto (j(x),t)$, so that any choice of oriented triangulation on $M \times I$ turns this into a singular $(n+1)$-chain $c \in C_{n+1}(X \times I;\mathbb{Z})$. Composing $\tilde{\jmath}$ with the quotient projection sending $X \times I$ to Z then produces a chain $c' \in C_{n+1}(Z;\mathbb{Z})$ with $\partial c' = \pm(g_* a - f_* a)$, thus proving that $g_* a - f_* a \in H_n(Y;\mathbb{Z})$ becomes trivial after acting on it with the map $i_* : H_n(Y;\mathbb{Z}) \to H_n(Z;\mathbb{Z})$.

Proof of Theorem 34.1. Consider the map of pairs $q : (X \times I, X \times \partial I) \to (Z, Y)$ defined as the composition of the two maps

$$(X \times I, X \times \partial I) \hookrightarrow ((X \times I) \amalg Y, X \times \partial I) \to (Z, Y),$$

where the first is the inclusion and the second is the quotient projection. Using the naturality of connecting homomorphisms in long exact sequences of pairs, this gives rise to a commuting diagram
(34.2)

$$\begin{tikzcd}
\ldots \arrow[r] & h_{k+1}(X \times I) \arrow[r, "\beta_*"] \arrow[d, "q_*"] & h_{k+1}(X \times I, X \times \partial I) \arrow[r, "\partial_*"] \arrow[d, "q_*"] & h_k(X \times \partial I) \arrow[r, "\alpha_*"] \arrow[d, "q_*"] & h_k(X \times I) \arrow[r, "\beta_*"] \arrow[d, "q_*"] & \ldots \\
\ldots \arrow[r] & h_{k+1}(Z) \arrow[r] & h_{k+1}(Z, Y) \arrow[r, "\partial_*^-"] & h_k(Y) \arrow[r, "i_*"] & h_k(Z) \arrow[r] & \ldots
\end{tikzcd}$$

where the two rows are the exact sequences of the pairs $(X \times I, X \times \partial I)$ and (Z, Y), and the maps $\alpha : X \times \partial I \hookrightarrow X \times I$ and $\beta : (X \times I, \varnothing) \hookrightarrow (X \times I, X \times \partial I)$ are the obvious inclusions. Since $X \times \partial I = X \times \{0,1\} \cong X \amalg X$, the additivity axiom gives an isomorphism

$$j^0_* \oplus j^1_* : h_k(X) \oplus h_k(X) \xrightarrow{\cong} h_k(X \times \partial I), \tag{34.3}$$

defined in terms of the inclusions $j^i : X \hookrightarrow X \times \{0,1\} : x \mapsto (x,i)$ for $i = 0,1$. Composing this with the inclusion $\alpha : X \times \partial I \hookrightarrow X \times I$, we notice that each of the maps $\alpha \circ j^i : X \hookrightarrow X \times I$ for $i = 0,1$ is a homotopy equivalence, and they are also homotopic to each other, so by the homotopy axiom, the two maps $(\alpha \circ j^i)_* : h_k(X) \to h_k(X \times I)$ for $i = 0,1$ are both the same isomorphism. It follows that

$$\alpha_* \circ (j^0_* \oplus j^1_*) = (\alpha_* \circ j^0_*) \oplus (\alpha_* \circ j^1_*) : h_k(X) \oplus h_k(X) \to h_k(X \times I)$$

is surjective, its kernel being the group of all pairs $(c, -c)$ for $c \in h_k(X)$. In particular, α_* itself is surjective, and we have an isomorphism

$$\Psi : h_k(X) \xrightarrow{\cong} \ker \alpha_* : c \mapsto (j^0_* \oplus j^1_*)(-c, c) = j^1_* c - j^0_* c. \tag{34.4}$$

Exactness of the top row now implies $\beta_* = 0$, and the connecting homomorphism $\partial_* : H_k(X \times I, X \times \partial I) \to H_{k-1}(X \times \partial I)$ is thus injective. This makes ∂_* an isomorphism onto its image, which is $\ker \alpha_*$.

Now observe that for the map $q : X \times \partial I \to Y$, the compositions $q \circ j^i : X \to Y$ for $i = 0,1$ are the maps f and g respectively, thus we have

$$q_* \circ \Psi : h_k(X) \to h_k(Y) : c \mapsto g_* c - f_* c = (g_* - f_*)c. \tag{34.5}$$

On the other hand, the map

$$(X \times I)/(X \times \partial I) \xrightarrow{q} Z/Y$$

determined by $q : (X \times I, X \times \partial I) \to (Z, Y)$ is a homeomorphism and thus induces an isomorphism

$$q_* : \widetilde{h}_*\big((X \times I)/(X \times \partial I)\big) \xrightarrow{\cong} \widetilde{h}_*(Z/Y),$$

and since both pairs are *good* in the sense of Definition 32.19, Theorem 32.23 implies that the map $q_* : h_*(X \times I, X \times \partial I) \to h_*(Z, Y)$ is also an isomorphism. We can put all of this information together to produce a commatative diagram

$$\begin{tikzcd}
h_{k+1}(Z,Y) \arrow[r, "q_*^{-1}"] \arrow[drr, "\partial_*^-"'] & h_{k+1}(X \times I, X \times \partial I) \arrow[r, "\partial_*"] & \ker \alpha_* \arrow[r, "\Psi^{-1}"] \arrow[d, "q_*"] & h_k(X) \arrow[dl, "g_* - f_*"] \\
 & & h_k(Y) &
\end{tikzcd}$$

in which all the horizontal maps on the top row are isomorphisms. The composition of these maps therefore gives an isomorphism $h_{k+1}(Z, Y) \to h_k(X)$ that we can use to replace $h_{k+1}(Z, Y)$ by $h_k(X)$ in the bottom row of (34.2); the original connecting homomorphism $\partial_*^- : h_{k+1}(Z, Y) \to h_k(Y)$ then gets replaced by the map $g_* - f_* : h_k(X) \to h_k(Y)$, producing an exact sequence as in the statement of the theorem. □

EXERCISE 34.6. The goal of this exercise is to gain a more concrete picture of the connecting homomorphism $\Phi : H_1(X_f; \mathbb{Z}) \to H_0(X; \mathbb{Z})$ that appears in the long exact sequence of the mapping torus of a *homeomorphism* $f : X \to X$,

$$\ldots \longrightarrow H_{k+1}(X_f; \mathbb{Z}) \xrightarrow{\Phi} H_k(X; \mathbb{Z}) \xrightarrow{\mathbb{1}_* - f_*} H_k(X; \mathbb{Z}) \xrightarrow{i_*} H_k(X_f; \mathbb{Z}) \xrightarrow{\Phi} H_{k-1}(X; \mathbb{Z}) \longrightarrow \ldots$$

in singular homology with integer coefficients. It will be useful to observe first that if $f : X \to X$ is a homeomorphism, then its mapping torus admits an alternative description as the quotient

$$X_f = (X \times \mathbb{R}) \Big/ (x,t) \sim (f(x), t+1),$$

where the equivalence is defined for every $t \in \mathbb{R}$. Take a moment to convince yourself that this quotient is homeomorphic to the slightly different definition of X_f given above. The new perspective has the advantage that one can view $\widetilde{X} := X \times \mathbb{R}$ as a covering space for X_f, with the quotient projection defining a covering map $\widetilde{X} \to X_f$ of infinite degree. Writing $S^1 := \mathbb{R}/\mathbb{Z}$, we also see a natural continuous surjective map $\pi : X_f \to S^1 : [(x,t)] \mapsto [t]$, whose **fibers** $\pi^{-1}(t)$ are homeomorphic to X for all $t \in S^1$. We shall denote by $i : X \hookrightarrow X_f$ the inclusion of the fiber $\pi^{-1}([0])$.

Assume X is path-connected, so there is a natural isomorphism $H_0(X;\mathbb{Z}) = \mathbb{Z}$, and notice that X_f is then also path-connected. Since $H_1(X_f;\mathbb{Z})$ is isomorphic to the abelianization of $\pi_1(X_f, x)$ for any choice of base point $x \in X_f$, we can identify X with $\pi^{-1}([0]) \subset X_f$, fix a base point $x \in X \subset X_f$ and represent any class in $H_1(X_f;\mathbb{Z})$ by a loop $\gamma : [0,1] \to X_f$ with $\gamma(0) = \gamma(1) = x$. Now let $\tilde{\gamma} : [0,1] \to \widetilde{X}$ denote the unique lift of γ to the cover $\widetilde{X} = X \times \mathbb{R}$ such that $\tilde{\gamma}(0) = (x, 0)$. Since γ is a loop, it follows that $\tilde{\gamma}(1) = (f^m(x), m)$ for some $m \in \mathbb{Z}$.

(a) Prove that under the natural identification of $H_0(X;\mathbb{Z})$ with $\mathbb{Z}$, the connecting homomorphism $\Phi : H_1(X_f;\mathbb{Z}) \to \mathbb{Z}$ can be chosen[55] such that

$$\Phi([\gamma]) = m,$$

so in particular, $[\gamma] \in \ker \Phi$ if and only if the lift of γ to the cover $\widetilde{X}$ is a loop.

(b) Prove directly from the characterization in part (a) that $\Phi : H_1(X_f;\mathbb{Z}) \to H_0(X;\mathbb{Z})$ is surjective.
Remark: Of course this can also be deduced less directly from the exact sequence.

One last comment about mapping tori: the usefulness of the exact sequence in Theorem 34.1 depends heavily on how easy it is to compute the homomorphism $f_* : h_*(X) \to h_*(X)$. This is not always easy, but sometimes it is, particularly in cases where $h_*(X)$ is relatively simple.

Degrees of maps between spheres. For the second topic today, let's talk about the homomorphism $f_* : H_n(S^n) \to H_n(S^n)$ induced by a map $f : S^n \to S^n$. For $n \geqslant 1$, the following definition makes sense due to the fact that $H_n(S^n;\mathbb{Z}) \cong \mathbb{Z}$ and $H_n(S^n;\mathbb{Z}_2) \cong \mathbb{Z}_2$ and homomorphisms $\mathbb{Z} \to \mathbb{Z}$ or $\mathbb{Z}_2 \to \mathbb{Z}_2$ admit very simple characterizations. In order to accommodate the case $n = 0$ under the same umbrella, we can use reduced homology since $\widetilde{H}_0(S^0;\mathbb{Z}) \cong \mathbb{Z}$ and $\widetilde{H}_0(S^0;\mathbb{Z}_2) \cong \mathbb{Z}_2$.

Definition 34.7. The **mapping degree** (*Abbildungsgrad*)

$$\deg(f) \in \mathbb{Z}$$

of a map $f : S^n \to S^n$ is the unique integer k such that the homomorphism $f_* : \widetilde{H}_n(S^n;\mathbb{Z}) \to \widetilde{H}_n(S^n;\mathbb{Z})$ is given by $c \mapsto kc$. Analogously, the **mod-2 mapping degree**

$$\deg_2(f) \in \mathbb{Z}_2$$

is the unique $k \in \mathbb{Z}_2$ such that $f_* : \widetilde{H}_n(S^n;\mathbb{Z}_2) \to \widetilde{H}_n(S^n;\mathbb{Z}_2)$ is the map $c \mapsto kc$.

[55] There is a bit of freedom allowed in the definition of Φ, e.g. we could replace it with $-\Phi$ and the sequence would still be exact since $\ker \Phi$ and $\operatorname{im} \Phi$ would not change.

We will not do much with $\deg_2(f)$ in the present lecture, but most of the basic properties of $\deg(f)$ as in the following proposition have obvious analogues for the mod-2 degree. The latter becomes more important when one also wants to consider maps $f : M \to M$ on a closed connected manifold M that need not be orientable, and we will touch upon this in the next lecture.

PROPOSITION 34.8. *The integer-valued degree for maps $S^n \to S^n$ has the following properties.*

(1) If $f, g : S^n \to S^n$ are homotopic, then $\deg(f) = \deg(g)$.
(2) For any $f, g : S^n \to S^n$, $\deg(f \circ g) = \deg(f) \cdot \deg(g)$.
(3) The identity map $S^n \to S^n$ has $\deg(\mathrm{Id}) = 1$.
(4) If f is constant, then $\deg(f) = 0$.
*(5) The degree of any map $f : S^1 \to S^1$ matches its **winding number** (Windungszahl), i.e. it is the unique $k \in \mathbb{Z}$ such that any continuous function $\theta : [0, 1] \to \mathbb{R}$ with $f(e^{2\pi it}) = e^{2\pi i\theta(t)}$ satisfies $\theta(1) - \theta(0) = k$.*

PROOF. The first three properties are immediate from the homotopy invariance of $\widetilde{H}_*(\cdot; \mathbb{Z})$ and the fact that it is a functor. For the fourth, observe that any constant map $f : S^n \to S^n$ can be factored as $i \circ \epsilon$ for the unique map $\epsilon : S^n \to \{\mathrm{pt}\}$ and a suitable inclusion $i : \{\mathrm{pt}\} \hookrightarrow S^n$, thus $f_* : \widetilde{H}_n(S^n; \mathbb{Z}) \to \widetilde{H}_n(S^n; \mathbb{Z})$ factors through $\widetilde{H}_n(\{\mathrm{pt}\}; \mathbb{Z}) = 0$. Finally, the fifth property follows from standard facts about $\pi_1(S^1)$ and the natural isomorphism $\widetilde{H}_1(S^1; \mathbb{Z}) = H_1(S^1; \mathbb{Z}) \cong \pi_1(S^1)$. □

EXERCISE 34.9. There are only four possible maps $f : S^0 \to S^0$. What are their degrees?

Recall from Exercise 29.12 that the suspension $SX = C_+X \cup_X C_-X$ of a space X can be regarded as a functor $\mathsf{Top} \to \mathsf{Top}$ sending objects X to SX, where maps $f : X \to Y$ are transformed to maps

$$Sf : SX \to SY : [(x, t)] \mapsto [(f(x), t)].$$

In particular, any map $f : S^n \to S^n$ gives rise to a map $Sf : S^{n+1} \to S^{n+1}$ using the identification $SS^n \cong S^{n+1}$.

PROPOSITION 34.10. *For any $f : S^n \to S^n$, $\deg(f) = \deg(Sf)$.*

PROOF. Recall from Example 33.8 that the isomorphism $\widetilde{H}_{n+1}(SX; \mathbb{Z}) \to \widetilde{H}_n(X; \mathbb{Z})$ can always be constructed as the connecting homomorphism in a Mayer-Vietoris exact sequence for SX. Given a map $f : S^n \to S^n$, the naturality of this connecting homomorphism produces a commuting diagram

$$\begin{tikzcd}
\widetilde{H}_{n+1}(S^{n+1}; \mathbb{Z}) \arrow[r, "\partial_*"] \arrow[d, "(Sf)_*"] & \widetilde{H}_n(S^n; \mathbb{Z}) \arrow[d, "f_*"] \\
\widetilde{H}_{n+1}(S^{n+1}; \mathbb{Z}) \arrow[r, "\partial_*"] & \widetilde{H}_n(S^n; \mathbb{Z})
\end{tikzcd}$$

where the two maps labeled ∂_* are the same isomorphism. Now if $(Sf)_*c = kc$ for some nontrivial $c \in \widetilde{H}_{n+1}(S^{n+1}; \mathbb{Z})$, it follows that $\partial_*(Sf)_*c = k\partial_*c = f_*\partial_*c$, where $\partial_*c \in \widetilde{H}_n(S^n; \mathbb{Z})$ is also nontrivial, hence $\deg(Sf) = k = \deg(f)$. □

PROPOSITION 34.11. *If $f : S^n \to S^n$ is the restriction to the unit sphere $S^n \subset \mathbb{R}^{n+1}$ of an orthogonal linear transformation $\mathbf{A} \in \mathrm{O}(n+1)$, then $\deg(f) = \det(\mathbf{A}) = \pm 1$.*

PROOF. Recall that $\mathrm{O}(n+1)$ has exactly two path-components, which can be labeled according to whether their elements have determinant $+1$ or -1. A given $\mathbf{A} \in \mathrm{O}(n+1)$ thus admits a continuous path in $\mathrm{O}(n+1)$ to the identity matrix $\mathbb{1}$ if and only if $\det(\mathbf{A}) = 1$, whereas if $\det(\mathbf{A}) = -1$, then it admits a path to the reflection matrix

$$R_{n+1} = \begin{pmatrix} -1 & & & \\ & 1 & & \\ & & \ddots & \\ & & & 1 \end{pmatrix}$$

It follows that $f : S^n \to S^n$ is homotopic to the identity and thus has degree 1 if $\det(\mathbf{A}) = 1$, and otherwise f is homotopic to a reflection map. What remains to be shown is that reflection maps always have degree -1. For $n = 0$ or $n = 1$, this is easy to check by direct calculation, e.g. a reflection on S^1 produces a map $S^1 \to S^1$ with winding number -1, so the claim follows from Proposition 34.8. Now if we assume the claim is true for reflections $f : S^n \to S^n$, it suffices to observe that $Sf : SS^n \to SS^n$ is also a reflection under a suitable identification $SS^n \cong S^{n+1}$, so by induction, the result follows from Proposition 34.10. □

The basic properties of $\deg(f)$ established thus far already have some quite nontrivial consequences about the topology of spheres. Here are two such results.

THEOREM 34.12. *Every map $f : S^n \to S^n$ with $\deg(f) \neq (-1)^{n+1}$ has a fixed point.*

PROOF. It is easy to think of a specific map $f : S^n \to S^n$ that has no fixed point: the antipodal map $x \mapsto -x$ has this property, and its degree according to Proposition 34.11 is $(-1)^{n+1}$. The theorem will follow from the claim that, in fact, *every* map $f : S^n \to S^n$ with no fixed point is homotopic to the antipodal map, and therefore also has degree $(-1)^{n+1}$.

Indeed, if $f : S^n \to S^n$ has no fixed point, then $f(x)$ and $-x$ are never antipodal points for any $x \in S^n$, thus the line in $\mathbb{R}^{n+1}$ connecting them does not pass through the origin. We can parametrize this line by

$$g_t(x) = (1-t)f(x) - tx \quad \text{for } t \in [0,1],$$

thus defining a continuous 1-parameter family of maps $g_t : S^n \to \mathbb{R}^{n+1}\backslash\{0\}$ with $g_0 = f$ and $g_1(x) = -x$. Since $g_t(x)$ is never zero, we can then define a homotopy $h : I \times S^n \to S^n$ in S^n from f to the antipodal map g_1 by

$$h(t,x) = \frac{g_t(x)}{|g_t(x)|}.$$

□

THEOREM 34.13 (the "hairy sphere" theorem). *If $n \in \mathbb{N}$ is even, then there does not exist any continuous nowhere zero vector field on S^n, i.e. there is no map $V : S^n \to \mathbb{R}^{n+1}\backslash\{0\}$ such that $V(x)$ is orthogonal to x for all $x \in S^n \subset \mathbb{R}^{n+1}$.*

PROOF. If such a map V exists, then for each $x \in S^n$, we can define $P_x \subset \mathbb{R}^{n+1}$ as the 2-dimensional plane spanned by x and $V(x)$, so that $P_x \cap S^n$ is a circle in S^n. The idea is then to follow a path along this circle from x through $V(x)/|V(x)|$ ending at $-x$. Concretely, such a path is given by the formula

$$t \mapsto f_t(x) := (\cos \pi t)x + (\sin \pi t)\frac{V(x)}{|V(x)|} \in S^n \quad \text{for } t \in [0,1],$$

which defines a homotopy from $f_0 = \mathrm{Id}$ to the antipodal map $f_1(x) = -x$. The degree of the latter was observed in the previous theorem to be $(-1)^{n+1}$, so we conclude $1 = \deg(f_0) = \deg(f_1) = (-1)^{n+1}$, implying n must be odd. □

35. Local and global mapping degree (November 14, 2023)

We would now like to generalize the mapping degree beyond spheres, while also giving it a more concrete geometric interpretation. The degree of a map $f : X \to Y$ in general is meant to be an answer to the following question: *for each $y \in Y$, how many points are there in $f^{-1}(y)$?* For arbitrary spaces, the answer of course depends on our choice of the point $y \in Y$, e.g. any bounded function $f : \mathbb{R} \to \mathbb{R}$ has the property that $f^{-1}(y)$ is empty for some points $y \in \mathbb{R}$ and not for others. It is perhaps surprising that if we are somewhat more restrictive about the class of spaces we consider, and we interpret the question "how many?" in the right way, then the answer no longer depends on y, and in fact, it depends on f only up to homotopy. We are already familiar with one situation where at least the first statement is true: if $f : X \to Y$ is a finite covering map and Y is connected, then every fiber $f^{-1}(y) \subset X$ contains the same finite number of points, called the *degree* of the cover (see Theorem 15.8). We will eventually be able to show that a reasonable generalization of this statement is true whenever X and Y are both closed, connected and oriented topological manifolds of the same dimension. The present lecture will prove this, modulo a couple of black boxes involving the computation of $H_n(M)$ when M is an arbitrary closed n-manifold; the general version of that computation will be dealt later in the course, when we construct fundamental classes without triangulations.

Local orientations and degree. Suppose M is a topological manifold of dimension $n \in \mathbb{N}$. We saw in Lecture 31 that the so-called *local homology* groups $H_n(M, M\backslash\{x\}; G)$ are isomorphic for every point $x \in M\backslash\partial M$ to the coefficient group G. This depends on the fact that every interior point x in an n-manifold M admits a so-called **Euclidean neighborhood**; we shall use this term in the following to refer to any (open or compact) neighborhood $\mathcal{U} \subset M$ of x together with a choice of homeomorphism

$$\varphi : \mathcal{U} \xrightarrow{\cong} \mathbb{R}^n \text{ or } \mathbb{D}^n, \qquad \text{such that} \qquad \varphi(x) = 0,$$

i.e. a coordinate chart identifying x with the origin. The choice of whether to use $\mathbb{R}^n$ or $\mathbb{D}^n$ as a local model for M near x is a matter of taste, and we shall generally use whichever seems more convenient in any given situation. While the local homology groups of Lecture 31 were defined specifically within singular homology, the arguments there were based only on its formal properties, and are thus equally valid for any axiomatic homology theory h_*. Indeed, if h_* has coefficient group $h_0(\{\text{pt}\}) = G$, then a Euclidean neighborhood $\mathcal{U} \subset M$ of x with coordinate chart $\varphi : \mathcal{U} \to \mathbb{D}^n$ determines a string of isomorphisms

(35.1)

$$h_n(M, M\backslash\{x\}) \longleftarrow h_n(\mathcal{U}, \mathcal{U}\backslash\{x\}) \xrightarrow{\varphi_*} h_n(\mathbb{D}^n, \mathbb{D}^n\backslash\{0\}) \longleftarrow h_n(\mathbb{D}^n, S^{n-1}) \xrightarrow{\partial_*} \widetilde{h}_{n-1}(S^{n-1}) \cong G,$$

in which every arrow without a label is induced by a continuous inclusion map. The same trick can be done with an open Euclidean neighborhood and chart $\varphi : \mathcal{U} \to \mathbb{R}^n$ after replacing the middle term of (35.1) with $h_n(\mathbb{R}^n, \mathbb{R}^n\backslash\{0\})$.

Specializing again to singular homology with integer coefficients, the local homology can be used to define a notion of *orientations* for topological manifolds, without needing to mention triangulations.

Definition 35.1. A **local orientation** of an n-manifold M at an interior point $x \in M\backslash\partial M$ is a choice of generator $[M]_x$ for the group $H_n(M, M\backslash\{x\}; \mathbb{Z}) \cong \mathbb{Z}$.

Note that in light of the excision isomorphism

$$H_n(\mathcal{U}, \mathcal{U}\backslash\{x\}; \mathbb{Z}) \xrightarrow{\cong} H_n(M, M\backslash\{x\}; \mathbb{Z})$$

defined for any Euclidean neighborhood $\mathcal{U} \subset M$ of x, a local orientation can equivalently be regarded as a generator of $H_n(\mathcal{U}, \mathcal{U}\backslash\{x\}; \mathbb{Z}) \cong \mathbb{Z}$.

Example 35.2. If M is a surface without boundary and $x \in M$ is a point, then a specific relative 2-cycle generating $H_2(M, M\backslash\{x\}; \mathbb{Z})$ can be defined via a single singular 2-simplex $\sigma : \Delta^2 \to M$ that embeds the triangle Δ^2 onto a neighborhood of x. Indeed, $\sigma \in C_2(M)$ is clearly a relative cycle in $(M, M\backslash\{x\})$ since σ maps $\partial\Delta^2$ to $M\backslash\{x\}$, and to see that it generates $H_2(M, M\backslash\{x\}; \mathbb{Z})$, one can follow the string of isomorphisms (35.1): they map $[\sigma]$ to the homology class of a 1-cycle in $S^1 \cong \partial\Delta^2$ consisting of the three edges of the triangle, a loop that clearly generates $\pi_1(\partial\Delta^2) \cong H_1(\partial\Delta; \mathbb{Z})$. In this picture, we can think of a local orientation at x as a choice (up to homotopy) of a small embedded loop in M about x: since there are two directions that such a loop can wind around x, there are two choices of local orientation.

Definition 35.3. Suppose M and N are manifolds of dimension $n \in \mathbb{N}$, $f : M \to N$ is a map, and $x \in M\backslash\partial M$ and $y = f(x) \in N\backslash\partial N$ are interior points such that x is an isolated point in the set $f^{-1}(y)$, i.e. there exists an open neighborhood $\mathcal{U} \subset M$ of x such that $f^{-1}(y) \cap \mathcal{U} = \{x\}$. Assume without loss of generality that $\mathcal{U} \subset M\backslash\partial M$ is a Euclidean neighborhood. Given local orientations $[M]_x \in H_n(\mathcal{U}, \mathcal{U}\backslash\{x\}; \mathbb{Z})$ and $[N]_y \in H_n(N, N\backslash\{y\}; \mathbb{Z})$, the **local degree**

$$\deg(f; x) \in \mathbb{Z}$$

of f and x is then defined as the unique integer $k \in \mathbb{Z}$ such that the map $H_n(\mathcal{U}, \mathcal{U}\backslash\{x\}; \mathbb{Z}) \to H_n(N, N\backslash\{y\}; \mathbb{Z})$ induced by $f : (\mathcal{U}, \mathcal{U}\backslash\{x\}) \to (N, N\backslash\{y\})$ sends $[M]_x$ to $k[N]_y$.

Under the same assumptions, the **mod 2 local degree**

$$\deg_2(f; x) \in \mathbb{Z}_2$$

is similarly defined to be the unique $k \in \mathbb{Z}_2$ such that $f_* : H_n(\mathcal{U}, \mathcal{U}\backslash\{x\}; \mathbb{Z}_2) \to H_n(N, N\backslash\{y\}; \mathbb{Z}_2)$ sends $[M]_x$ to $k[N]_y$, where $[M]_x$ and $[N]_y$ are now taken to be the unique nontrivial elements of $H_n(\mathcal{U}, \mathcal{U}\backslash\{x\}; \mathbb{Z}_2) \cong \mathbb{Z}_2$ and $H_n(N, N\backslash\{y\}; \mathbb{Z}_2) \cong \mathbb{Z}_2$ respectively.

Notice that there are no choices involved in the definition of $\deg_2(f; x)$, whereas $\deg(f; x)$ will change sign whenever we change the choice of one of the local orientations.

As explained in (35.1), any choice of Euclidean neighborhood $\mathcal{U} \subset M$ of x gives rise to an isomorphism of $H_n(\mathcal{U}, \mathcal{U}\backslash\{x\}; \mathbb{Z})$ with $\widetilde{H}_{n-1}(S^{n-1}; \mathbb{Z}) \cong \mathbb{Z}$. We can use this isomorphism to transform the definition above into a condition about maps between spheres:

Proposition 35.4. *In the setting of Definition 35.3, fix a generator $[S^{n-1}] \in \widetilde{H}_{n-1}(S^{n-1}; \mathbb{Z})$ and Euclidean neighborhoods $\mathcal{U} \subset M$ of x and $\mathcal{V} \subset N$ of y such that the resulting isomorphisms of $H_n(\mathcal{U}, \mathcal{U}\backslash\{x\}; \mathbb{Z})$ and $H_n(\mathcal{V}, \mathcal{V}\backslash\{y\}; \mathbb{Z})$ to $\widetilde{H}_{n-1}(S^{n-1}; \mathbb{Z})$ send $[M]_x$ and $[N]_y$ to $[S^{n-1}]$. Then if $\hat{f}$ denotes the map f written in the chosen local coordinates as a map between neighborhoods of 0 in $\mathbb{R}^n$, we have*

$$\deg(f; x) = \deg\left(\left.\frac{\hat{f}}{|\hat{f}|}\right|_{\partial\mathbb{D}^n_\epsilon} : \partial\mathbb{D}^n_\epsilon \to S^{n-1} \right)$$

for all $\epsilon > 0$ sufficiently small, where $\mathbb{D}^n_\epsilon$ denotes the closed ϵ-disk and its boundary is identified in the obvious way with S^{n-1}, so that the right hand side is the degree of a map $S^{n-1} \to S^{n-1}$. Similarly, $\deg_2(f; x)$ is related in the same say to the mod 2 degree of the same map $\partial\mathbb{D}^n_\epsilon \to S^{n-1}$. □

Corollary 35.5. *Suppose $\{f_t : M \to N\}_{t\in[0,1]}$ is a continuous family of maps between two manifolds of dimension $n \in \mathbb{N}$, with interior points $x \in M\backslash\partial M$ and $y \in N\backslash\partial N$ such that x is an isolated point of $f_t^{-1}(y)$ for every t. Then for any fixed choice of local orientations at x and y, $\deg(f_0; x) = \deg(f_1; x)$, and similarly, $\deg_2(f_0; x) = \deg_2(f_1; x)$.*

PROOF. We can interpret both local degrees via Proposition 35.4 as degrees of maps $S^{n-1} \to S^{n-1}$, and the assumption about the family f_t implies that these two maps between spheres are homotopic. $\square$

EXAMPLE 35.6. Continuing the discussion of Example 35.2, suppose $f : M \to N$ is a map between surfaces such that x is an isolated point in $f^{-1}(y)$, and suppose we have fixed local orientations $[M]_x \in H_2(M, M\backslash\{x\}; \mathbb{Z})$ and $[N]_y \in H_2(N, N\backslash\{y\}; \mathbb{Z})$. Choose small Euclidean neighborhoods $\mathcal{U} \subset M$ of x and $\mathcal{V} \subset N$ of y such that $f(\mathcal{U}) \subset \mathcal{V}$ and $f^{-1}(y) \cap \mathcal{U} = \{x\}$. Then $[M]_x$ determines a homotopy class of embedded loops $\alpha : S^1 \hookrightarrow \mathcal{U}\backslash\{x\}$ winding once around x, so that $f \circ \alpha : S^1 \to \mathcal{V}\backslash\{y\}$ is also uniquely determined up to homotopy. The winding number of $f \circ \gamma$ is then the local degree $\deg(f; x)$; its definition requires a local orientation at y in order to decide which winding numbers are positive and which are negative, i.e. those that wind in the same direction as the loops $S^1 \hookrightarrow \mathcal{V}\backslash\{y\}$ determined by $[N]_y$ are considered positive.

Let us discuss more concretely how local degrees of maps from $\mathbb{R}^n$ to itself can be computed. There is a natural way to choose local orientations $[\mathbb{R}^n]_x \in H_n(\mathbb{R}^n, \mathbb{R}^n\backslash\{x\}; \mathbb{Z})$ at every point $x \in \mathbb{R}^n$: if $\mathbb{D}^n_x \subset \mathbb{R}^n$ denotes the closed unit disk about x and we identify its boundary in the obvious way with S^{n-1}, then we obtain as in (35.1) a string of natural isomorphisms

$$H_n(\mathbb{R}^n, \mathbb{R}^n\backslash\{x\}; \mathbb{Z}) \cong H_n(\mathbb{D}^n_x, \partial\mathbb{D}^n_x; \mathbb{Z}) \cong \widetilde{H}_{n-1}(S^{n-1}; \mathbb{Z}),$$

so that any choice of generator $[S^{n-1}] \in \widetilde{H}_{n-1}(S^{n-1}; \mathbb{Z})$ determines local orientations $[\mathbb{R}^n]_x \in H_n(\mathbb{R}^n, \mathbb{R}^n\backslash\{x\}; \mathbb{Z})$ for all $x \in \mathbb{R}^n$ simultaneously. With this choice in place, any continuous map $f : \mathcal{U} \to \mathbb{R}^n$ defined on an open subset $\mathcal{U} \subset \mathbb{R}^n$ has a well-defined local degree at any point $x \in \mathcal{U}$ that is isolated in $f^{-1}(f(x))$, and we notice that $\deg(f; x)$ does not depend on our arbitrary choice of generator $[S^{n-1}]$ since reversing this would reverse *both* of the local orientations $[\mathbb{R}^n]_x$ and $[\mathbb{R}^n]_{f(x)}$. We can now prove:

PROPOSITION 35.7. *Suppose local orientations $[\mathbb{R}_n]_x$ for points $x \in \mathbb{R}^n$ are fixed according to the prescription above, $\mathcal{U} \subset \mathbb{R}^n$ is an open subset and $f : \mathcal{U} \to \mathbb{R}^n$ is a map that is differentiable at a point $x \in \mathcal{U}$ such that its derivative $df(x) : \mathbb{R}^n \to \mathbb{R}^n$ is an isomorphism. Then x is an isolated point of $f^{-1}(f(x))$, and $\deg(f; x) = \pm 1$, with sign matching the sign of $\det df(x)$.*

PROOF. We can write $f : \mathcal{U} \to \mathbb{R}^n$ near x as

$$f(x + h) = y + df(x)h + |h|\eta(h)$$

for sufficiently small $h \in \mathbb{R}^n$, where $y := f(x)$ and $\eta(h)$ is an $\mathbb{R}^n$-valued function satisfying $\lim_{h\to 0} \eta(h) = 0$. If $df(x) : \mathbb{R}^n \to \mathbb{R}^n$ is invertible, then there exists a constant $c > 0$ such that $|df(x)h| \geqslant c|h|$ for all $h \in \mathbb{R}^n$, so

$$|f(x + h) - y| = |df(x)h + |h|\eta(h)| \geqslant |df(x)h| - |h||\eta(h) \geqslant (c - |\eta(h)|)\,|h|,$$

and the right hand side is positive for all $|h|$ sufficiently small since $\eta(h) \to 0$. This proves that x is isolated in $f^{-1}(y)$. Now modify f near x by

$$f_t(x + h) = y + df(x)h + \rho_t(h)|h|\eta(h),$$

where $\rho_t(h) \in [0, 1]$ is a family of cutoff functions that equal 1 away from $h = 0$ such that $\rho_0 \equiv 1$ and ρ_1 vanishes on a smaller neighborhood of $h = 0$. This changes f by a homotopy through maps in which x remains an isolated point of $f_t^{-1}(y)$, so in light of Corollary 35.5, we can now assume without loss of generality that the remainder term vanishes completely, i.e. $f(x + h) = y + df(x)$. Now observe that if we modify f by a further homotopy of the form

$$f_t(x + h) = y + A_t h,$$

where $A_t : \mathbb{R}^n \to \mathbb{R}^n$ is a family of invertible linear transformations, then the local degree still will not change due to Corollary 35.5, thus we are free to assume without loss of generality that $df(x)$ is an *orthogonal* transformation. The corresponding map $S^{n-1} \to S^{n-1}$ is then of the type considered in Proposition 34.11, so its degree is the determinant of the orthogonal transformation, which is $+1$ if the original derivative $df(x)$ had positive determinant and -1 otherwise. □

Many applications of the local degree are based on the following exercise, as it provides a criterion for existence of solutions to equations of the form $f(x) = y$ that are *stable* under small perturbations of f:

Exercise 35.8. Prove that if $\mathcal{U} \subset \mathbb{R}^n$ is open and $f : \mathcal{U} \to \mathbb{R}^n$ is a continuous map with $f(x) = y$ and either $\deg(f; x)$ or $\deg_2(f; x)$ is nonzero for some $x \in \mathcal{U}$, then for any neighborhood $\mathcal{U}_x \subset \mathcal{U}$ of x, there exists an $\epsilon > 0$ such that every continuous map $\hat{f} : \mathcal{U} \to \mathbb{R}^n$ satisfying $|\hat{f} - f| < \epsilon$ maps some point in $\mathcal{U}_x$ to y.
Hint: Consider the restriction of $\hat{f}$ to the boundary of a small ball about x, and normalize it so that it maps to the sphere surrounding a small ball about y. What can you say about the degree of this map between spheres if $\hat{f}$ maps the ball about x to $\mathbb{R}^n \backslash \{y\}$?

Exercise 35.9. Find an example of a smooth map $f : \mathbb{R}^2 \to \mathbb{R}^2$ that has an isolated zero at the origin with $\deg(f; 0) = 0$ and admits arbitrarily small perturbations that are nowhere zero.

Mapping degree for closed manifolds. We will now extend the global mapping degree previously defined for maps $f : S^n \to S^n$ to maps $f : M \to N$ between more general closed connected n-manifolds. Most of what follows can also be extended to maps $f : (M, \partial M) \to (N, \partial N)$ between compact n-manifolds with boundary, so long as $f(\partial M) \subset \partial N$, but we will leave this extension as an exercise for the reader.

Our definition of degree for maps $f : M \to N$ will necessitate imposing a condition on the manifolds that we consider. It will later turn out that this condition is satisfied for *all* closed and connected manifolds—possibly with a condition on orientations,[56] depending which coefficients we want to use—though it will be a while before we are in a position to fully prove this.

Definition 35.10. Given an axiomatic homology theory h_*, a topological manifold M of dimension $n \in \mathbb{N}$ will be called **h_*-admissible**[57] if M is closed and the obvious inclusion $i^x : (M, \varnothing) \hookrightarrow (M, M\backslash\{x\})$ induces an isomorphism

$$i^x_* : h_n(M) \xrightarrow{\cong} h_n(M, M\backslash\{x\})$$

for every point $x \in M$. For the case $h_* = H_*(\cdot; G)$, we shall abbreviate the terminology and say that M is **G-admissible**.

Clearly an h_*-admissible n-manifold must have $h_n(M)$ isomorphic to the coefficient group, so there are in general some nontrivial computations of homology to be done before we can prove that any given manifold is admissible. To start with, it will be useful to note that we already know how to do this for spheres of arbitrary dimension:

Exercise 35.11. Prove that S^n is h_*-admissible for every $n \in \mathbb{N}$ and every axiomatic homology theory h_*.
Hint: For any $x \in S^n$, one can pick an open Euclidean neighborhood $\mathcal{U} \subset S^n$ of x and use

[56] We will define later what it means in general for a topological n-manifold to be orientable; you may already be able to intuit part of the definition from the notion of "local" orientations introduced in Definition 35.1. A notion of orientations for topological manifolds of dimension $n \leqslant 2$ was discussed last semester in Lecture 20, and Example 35.2 hints at the relationship between that notion and local orientations in the case $n = 2$.

[57] This is not a universally standard term, but it is convenient for our purposes at the moment.

$h_n(S^n, S^n \backslash \mathcal{U})$ as a substitute for the local homology group $h_n(S^n, S^n \backslash \{x\})$ (why?). What does the space $S^n \backslash \mathcal{U}$ look like, and what does the long exact sequence of the pair $(S^n, S^n \backslash \mathcal{U})$ tell you?

Just so it's clear how widely applicable the mapping degree is, let us state a result whose proof in full generality will have to wait until after the general construction of fundamental classes later in this course. Certain special cases of it, however, are already within reach, and a lot more will be so in the near future, once we've discussed cellular homology.

PROPOSITION 35.12. *Every closed and connected topological manifold M is $\mathbb{Z}_2$-admissible, and if M is also orientable, then it is $\mathbb{Z}$-admissible.*

PARTIAL PROOF. The cases for which we already know how to prove $\mathbb{Z}$-admissibility include the spheres S^n and the tori $\mathbb{T}^n$, and we can also prove $\mathbb{Z}_2$-admissibility for these, plus certain non-orientable examples such as the projective plane and the Klein bottle. The point is that all of these examples of closed connected n-manifolds M have the following two features in common:

(1) M admits a triangulation;
(2) $H_n(M; \mathbb{Z}_2) \cong \mathbb{Z}_2$, and in cases where the triangulation is orientable, $H_n(M; \mathbb{Z}) \cong \mathbb{Z}$.

The latter fact about S^n was our first serious computation of singular homology, and for the other manifolds mentioned above, it can be proved using Mayer-Vietoris sequences (see Lecture 33, especially Exercise 33.10), or in certain cases also the exact sequence of a mapping torus (Lecture 34). These two conditions in tandem give us the following. In Lectures 30 and 31, we saw that any triangulation of a closed n-manifold M gives rise to a fundamental class $[M] \in H_n(M; \mathbb{Z}_2)$, and for any given point $x \in M$, we are free to assume after perhaps a small perturbation of the triangulation that x lies in the interior of one of its n-simplices. Lemma 31.10 then implies that the homomorphism $i^x_* : H_n(M; \mathbb{Z}_2) \to H_n(M, M \backslash \{x\}; \mathbb{Z}_2)$ sends $[M]$ to the unique nontrivial element $[M]_x \in H_n(M, M \backslash \{x\}; \mathbb{Z}_2)$, so if we also know $H_n(M; \mathbb{Z}_2) \cong \mathbb{Z}_2$, it follows that the map $H_n(M; \mathbb{Z}_2) \to H_n(M, M \backslash \{x\}; \mathbb{Z}_2)$ is an isomorphism. If additionally the triangulation is oriented, then we also have an integral fundamental class $[M] \in H_n(M; \mathbb{Z})$ and the stronger statement that $i^x_* : H_n(M; \mathbb{Z}) \to H_n(M, M \backslash \{x\}; \mathbb{Z}) \cong \mathbb{Z}$ sends $[M]$ to a *primitive* element, which in this case means a local orientation $[M]_x$ at x. The condition $H_n(M; \mathbb{Z}) \cong \mathbb{Z}$ then similarly implies that $[M]$ generates $H_n(M; \mathbb{Z})$ and the map $H_n(M; \mathbb{Z}) \to H_n(M, M \backslash \{x\}; \mathbb{Z})$ is an isomorphism.

By the end of next week, we will also be able to deduce condition (2) above from condition (1), using the deep theorem that $h_*(|K|) \cong H^\Delta_*(K; G)$ for every simplicial complex K and every axiomatic homology theory h_* with coefficient group G. For the computation of $H^\Delta_n(K; G)$ when K is the simplicial complex that triangulates a closed connected n-manifold (with or without orientation), see Exercises 30.22 and 31.9. We will then be able to say without any black boxes that every closed connected manifold admitting a triangulation is $\mathbb{Z}_2$-admissible (and $\mathbb{Z}$-admissible if the triangulation is orientable), which in particular includes all (closed and connected) *smooth* manifolds. For closed topological manifolds without triangulations, the result will follow from the general construction of fundamental classes. □

DEFINITION 35.13. Assume M and N are $\mathbb{Z}$-admissible manifolds of dimension $n \in \mathbb{N}$, and choose generators $[M] \in H_n(M; \mathbb{Z}) \cong \mathbb{Z}$ and $[N] \in H_n(N; \mathbb{Z}) \cong \mathbb{Z}$. We then define the **degree** (*Grad*) of any map $f : M \to N$ to be the unique integer $\deg(f) = k \in \mathbb{Z}$ such that

$$f_*[M] = k[N].$$

If M and N are $\mathbb{Z}_2$-admissible (but not necessarily $\mathbb{Z}$-admissible), one can similarly define the **mod 2 degree** of f as the unique $k \in \mathbb{Z}_2$ such that $f_*[M] = k[N]$ where $[M] \in H_n(M; \mathbb{Z}_2) \cong \mathbb{Z}_2$ and $[N] \in H_n(N; \mathbb{Z}_2) \cong \mathbb{Z}_2$ are the unique nontrivial elements.

Note that the sign of $\deg(f)$ depends in general on the choices of generators $[M]$ and $[N]$, but if $M = N$, then it is natural to choose $[M] = [N]$, and $\deg(f)$ is then independent of choices since reversing the signs of $[M]$ and $[N]$ simultaneously changes nothing in the relation $f_*[M] = k[N]$. In this way, our new definition recovers the old one for maps $S^n \to S^n$. The mod 2 degree is in any case defined with no need for choices, since the generators $[M]$ and $[N]$ are unique in homology with $\mathbb{Z}_2$-coefficients. It is again easy to check that the obvious analogues of items (1)–(4) in Proposition 34.8 are satisfied for this new definition.

We can now state the main result relating global and local degrees.

THEOREM 35.14. *Suppose M and N are $\mathbb{Z}$-admissible manifolds of dimension $n \in \mathbb{N}$, fix generators $[M] \in H_n(M;\mathbb{Z})$ and $[N] \in H_n(N;\mathbb{Z})$ and use these to determine local orientations $[M]_x := i_*^x[M]$ and $[N]_y := i_*^y[N]$ at all points $x \in M$ and $y \in N$. Then for any map $f : M \to N$ and any point $y \in N$ such that $f^{-1}(y)$ is a finite set,*

$$\deg(f) = \sum_{x \in f^{-1}(y)} \deg(f;x). \tag{35.2}$$

Similarly, if M and N are $\mathbb{Z}_2$-admissible and $f : M \to N$ is any map with a point $y \in N$ such that $f^{-1}(y)$ is finite, we have

$$\deg_2(f) = \sum_{x \in f^{-1}(y)} \deg_2(f;x).$$

We sometimes refer to the expression on the right hand side of (35.2) as the *algebraic count* of points in $f^{-1}(y)$. One can check that if $f : M \to N$ happens to be a covering map, then for suitable choices of the generators $[M]$ and $[N]$, the local degrees $\deg(f;x)$ are all 1 and the algebraic count is thus the actual count of points. In more general situations, the points must be counted with signs and “weights” determined by the local degree, but the advantage is that the result does not depend on the point $y \in N$, and it only depends on f up to homotopy.

REMARK 35.15. For a given map $f : M \to N$, there is no guarantee in general that $f^{-1}(y)$ will be a finite set for any choice of $y \in N$; if it isn't, then the statement of Theorem 35.14 becomes vacuous. Since M and N have the same dimension, however, one can reasonably hope to encounter situations in which f is a local homeomorphism near each $x \in f^{-1}(y)$, implying that $f^{-1}(y) \subset M$ is a discrete subset, and since M is compact, the finiteness of $f^{-1}(y)$ follows. If M and N have smooth structures, then standard results in differential topology imply that this is in fact the “generic” situation, i.e. every continuous map $f : M \to N$ can be perturbed (without changing its homotopy class) to one that is smooth, and by Sard's theorem, almost every point $y \in N$ will then be a regular value of the smoothened version of f, making f a local diffeomorphism near every point of $f^{-1}(y)$. In this case the local degree $\deg(f;x)$ at every point $x \in f^{-1}(y)$ can be deduced from the derivative $df(x)$ via Proposition 35.7, and will always be ± 1.

Theorem 35.14 has a wide range of applications, but it also establishes an important theoretical connection between algebraic and *differential* topology. In the setting of closed differentiable manifolds and smooth maps $f : M \to N$, there is a natural way to define $\deg(f)$ using transversality results for smooth maps, e.g. one can use the perturbation trick mentioned in Remark 35.15 above to restrict attention to cases in which f is a local diffeomorphism with $\deg(f;x) = \pm 1$ for each $x \in f^{-1}(y)$. One then defines $\deg(f)$ essentially as the right hand side of (35.2) and interprets it as “counting $f^{-1}(y)$ with signs”; the interesting part is then to prove that the result does not depend on y or on f beyond its homotopy class. Without knowing Theorem 35.14 or anything else about homology, the latter can also be proven as a consequence of transversality results—the main point is that if f_0 and f_1 are homotopic, then a generic choice of smooth homotopy $\{f_t : M \to N\}_{t \in [0,1]}$

between them gives rise to a compact oriented 1-manifold

$$Q := \{(t,x) \in [0,1] \times M \mid f_t(x) = y\}$$

for which the difference $\#f_0^{-1}(y) - \#f_1^{-1}(y)$ between the two signed counts of preimages of y is interpreted as a signed count of the points in the oriented 0-manifold ∂Q. The classification of 1-manifolds implies that every component of a compact oriented 1-manifold with nonempty boundary has exactly one boundary point that counts positively and one that counts negatively, hence the total count is always zero. This perspective on the degree is explained beautifully in the classic book by Milnor [**Mil97**].[58] It is by no means easy however to see from the differentiable viewpoint what the mapping degree has to do with the homology of manifolds, i.e. why the right hand side of (35.2) matches the left hand side. The proof of that requires the formal properties of homology theories.

Proof of Theorem 35.14. For later convenience, we shall carry out most of the proof in the framework of an arbitrary axiomatic homology theory h_*, assuming M and N to be h_*-admissible. Write

$$f^{-1}(y) = \{x_1, \dots, x_\ell\},$$

fix a Euclidean neighborhood $\mathcal{V} \subset N$ of y, along with Euclidean neighborhoods $\mathcal{U}_k \subset M$ of the individual points x_k for $k = 1, \dots, \ell$ such that

$$f(\mathcal{U}_k) \subset \mathcal{V} \quad \text{and} \quad \mathcal{U}_k \cap \mathcal{U}_j = \varnothing \text{ for } j \neq k.$$

These assumptions guarantee that $f(\mathcal{U}_k \backslash \{x_k\}) \subset \mathcal{V}\backslash\{y\}$, hence f also defines a map of pairs $(\mathcal{U}_k, \mathcal{U}_k\backslash\{x_k\}) \to (\mathcal{V}, \mathcal{V}\backslash\{y\})$ for every $k = 1, \dots, \ell$. Now consider the diagram

$$\begin{tikzcd}
 & h_n(\mathcal{U}_k, \mathcal{U}_k\backslash\{x_k\}) \arrow[r, "f_*"] \arrow[dl, "\alpha^k_*"'] \arrow[d, "\gamma^k_*"] & h_n(\mathcal{V}, \mathcal{V}\backslash\{y\}) \arrow[d, "\beta"] \\
h_n(M, M\backslash\{x_k\}) & h_n(M, M\backslash f^{-1}(y)) \arrow[l, "p^k_*"] \arrow[r, "f_*"] & h_n(N, N\backslash\{y\}) \\
 & h_n(M) \arrow[r, "f_*"] \arrow[ul, "i^{x_k}_*"] \arrow[u, "j_*"] & h_n(N) \arrow[u, "i^y_*"']
\end{tikzcd} \tag{35.3}$$

where the maps α^k, p^k, γ^k, j and β are all inclusions. By the admissibility assumption, $i^{x_k}_*$ and i^y_* are isomorphisms, and α^k_* and β_* are also isomorphisms by excision. To understand the maps p^k_* for $k = 1, \dots, \ell$, observe that these can all be combined to define a product map

$$p := (p^1_*, \dots, p^\ell_*) : h_n(M, M\backslash f^{-1}(y)) \to \bigoplus_{k=1}^{\ell} h_n(M, M\backslash\{x_k\}),$$

which fits into the following diagram:

$$\begin{tikzcd}
h_n(M, M\backslash f^{-1}(y)) \arrow[r, "p"] & \bigoplus_{k=1}^{\ell} h_n(M, M\backslash\{x_k\}) \\
h_n\Big(\coprod_{k=1}^{\ell}(\mathcal{U}_k, \mathcal{U}_k\backslash\{x_k\})\Big) \arrow[u, "\cong"] & \bigoplus_{k=1}^{\ell} h_n(\mathcal{U}_k, \mathcal{U}_k\backslash\{x_k\}) \arrow[l, "\cong"] \arrow[u, "\cong"']
\end{tikzcd}$$

Here the maps are all induced by obvious inclusions, the two vertical maps are isomorphisms by excision, and the bottom horizontal map is an isomorphism due to a combination of the additivity axiom with the five-lemma (see Exercise 32.16), thus p is also an isomorphism. If we use this to

[58]The differentiable approach to the mapping degree was also sketched in Exercise 19.14 last semester.

replace $h_n(M, M\backslash f^{-1}(y))$ in (35.3) by $\bigoplus_{k=1}^{\ell} h_n(M, M\backslash\{x_k\})$, then the map p_*^k becomes simply the projection of $\bigoplus_{k=1}^{\ell} h_n(M, M\backslash\{x_k\})$ to the factor $h_n(M, M\backslash\{x_k\})$. With this replacement understood, we have

$$j_* = (i_*^{x_1}, \ldots, i_*^{x_\ell}) : h_n(M) \to \bigoplus_{k=1}^{\ell} h_n(M, M\backslash\{x\}),$$

and the commutativity of the bottom right square in (35.3) then gives the formula

$$f_* j_* = \sum_{k=1}^{\ell} f_* i_*^{x_k} = i_*^y f_* : h_n(M) \to h_n(N, N\backslash\{y\}). \tag{35.4}$$

If h_* is $H_*(\cdot;\mathbb{Z})$ and we apply this formula to the chosen generator $[M] \in H_n(M;\mathbb{Z})$ with $i_*^{x_k}[M] = [M]_{x_k}$, the result is

$$\sum_{k=1}^{\ell} f_*[M]_{x_k} = \sum_{k=1}^{\ell} \deg(f;x_k)[N]_y = i_*^y f_*[M] = \deg(f) i_*^y[N] = \deg(f)[N]_y,$$

from which the formula for the integer-valued degree follows. The formula for the mod 2 degree follows in the same way using $h_* = H_*(\cdot;\mathbb{Z}_2)$. □

It is easy to see that a non-surjective map $f : S^n \to S^n$ must have degree 0, because its image then lies in the contractible space $S^n\backslash\{\mathrm{pt}\} \cong \mathbb{R}^n$, making f homotopic to a constant. The same argument does not work for maps $f : M \to N$ when N is a more general closed n-manifold, but Theorem 35.14 nonetheless gives us an easy proof of the same statement:

Corollary 35.16. *If M and N are $\mathbb{Z}$-admissible n-manifolds with $n \geqslant 1$ and $f : M \to N$ is not surjective, then $\deg(f) = 0$. Similarly, if both manifolds are $\mathbb{Z}_2$-admissible and f is not surjective, then $\deg_2(f) = 0$.*

Proof. Apply Theorem 35.14 to identify $\deg(f)$ or $\deg_2(f)$ with a suitable count of points in $f^{-1}(y)$ where $y \notin f(M)$. □

In the axiomatic setting, the proof of Theorem 35.14 also gives rise to the following result, which will be of some theoretical importance when we develop cellular homology. Recall from Exercise 35.11 that we already know S^n to be h_*-admissible for every h_* and every $n \in \mathbb{N}$.

Theorem 35.17. *For any map $f : S^n \to S^n$ with $n \in \mathbb{N}$ and any axiomatic homology theory h_*, the induced homomorphism $f_* : h_n(S^n) \to h_n(S^n)$ takes the form $c \mapsto \deg(f)c$.*

Proof. One can verify explicitly that the corresponding statement about reduced homology holds for all maps $f : S^0 \to S^0$; this is easy to check because there exist only four distinct maps from S^0 to itself, and the reduced homology of S^0 can be derived directly from the additivity and dimension axioms (see Exercise 29.13). We now argue by induction on the dimension, assuming for a given n that homomorphisms $f_* : \widetilde{h}_{n-1}(S^{n-1}) \to \widetilde{h}_{n-1}(S^{n-1})$ are always given by multiplication with the integer-valued degree of maps $f : S^{n-1} \to S^{n-1}$. Using perturbation results from differential topology as mentioned in Remark 35.15, we can assume after a small perturbation of any given map $f : S^n \to S^n$ within its homotopy class that $f^{-1}(y)$ is a finite set for some $y \in S^n$. Now write $f^{-1}(y) = \{x_1, \ldots, x_\ell\}$ and, given $c \in h_n(S^n)$, use (35.4) to write

$$i_*^y f_* c = \sum_{k=1}^{\ell} f_* i_*^{x_k} c \in h_n(S^n, S^n\backslash\{y\}),$$

where the individual terms on the right hand side involve the homomorphisms

$$f_* : h_n(S^n, S^n\backslash\{x_k\}) \to h_n(S^n, S^n\backslash\{y\}).$$

Using excision and connecting homomorphisms as in Proposition 35.4, one can identify both the domain and target of this map with $\widetilde{h}_{n-1}(S^{n-1})$ so that f_* is equivalent to the homomorphism $\widetilde{h}_{n-1}(S^{n-1}) \to \widetilde{h}_{n-1}(S^{n-1})$ induced by a map $S^{n-1} \to S^{n-1}$, whose degree is precisely $\deg(f; x_k)$. The inductive hypothesis thus expresses the homomorphism as multiplication by $\deg(f; x_k)$, giving a commutative diagram

$$\begin{tikzcd} h_n(S^n, S^n \backslash \{x_k\}) \arrow[r, "f_*"] \arrow[d, "\cong"] & h_n(S^n, S^n \backslash \{y\}) \arrow[d, "\cong"] \\ h_{n-1}(S^{n-1}) \arrow[r, "\cdot \deg(f; x_k)"] & h_{n-1}(S^{n-1}) \end{tikzcd}$$

Adding up these contributions for every $x_k \in f^{-1}(y)$ produces multiplication by $\deg(f)$ according to Theorem 35.14. □

One consequence of this result is that the definition of $\deg(f)$ for maps $f : S^n \to S^n$ does not actually depend on the choice to use *singular* homology in particular—we could have replaced $H_*(\cdot; \mathbb{Z})$ with any other axiomtaic homology theory with coefficient group $\mathbb{Z}$ and would thus obtain an equivalent definition.

One can use a similar inductive argument to prove a straightforward relationship between $\deg(f)$ and $\deg_2(f)$; we will later also be to give a purely algebraic proof of the following result, using the universal coefficient theorem.

Corollary 35.18. *If M and N are both $\mathbb{Z}$-admissible and $\mathbb{Z}_2$-admissible, then for every map $f : M \to N$, $\deg_2(f)$ is the image of $\deg(f)$ under the natural projection $\mathbb{Z} \to \mathbb{Z}_2$.* □

36. CW-complexes (November 17, 2023)

Let's clear up one thing straightaway: the "CW" in "CW-complex" does not stand for my name.

If you must know, the "C" stands for "closure-finite," and the "W" for "weak topology". Both of these terms refer to slightly subtle issues involving the definition and properties of the topology on a CW-complex. We'll get to that.

But first, I should tell you what they are. The informal answer is that CW-complexes are spaces that we can construct by gluing disks (of various dimensions) to things along their boundaries. It turns out that almost all spaces of importance in geometric settings can be constructed in this way, so understanding the algebraic topology of CW-complexes opens the way toward an enormous range of applications. The motivation to focus on CW-complexes rather than more general spaces is practical: in essence, CW-complexes are the class of topological spaces for which the subject of algebraic topology is doable.

Definition 36.1 (CW-complexes, part 1 of 2). A **CW-complex** (*CW-Komplex*) or **cell complex** (*Zellkomplex*) is a topological space X that is presented as the union of a nested sequence of subspaces

$$X^0 \subset X^1 \subset X^2 \subset \ldots \subset \bigcup_{n \geqslant 0} X^n = X$$

constructed by the following inductive procedure:

- X^0 is a space with the discrete topology;
- For each $n \in \mathbb{N}$, there is a set $\mathcal{K}^n$ and a collection of maps $\{\varphi_\alpha : S^{n-1} \to X^{n-1}\}_{\alpha \in \mathcal{K}^n}$ such that X^n is the result of attaching n-disks $\mathbb{D}^n$ along their boundaries to X^{n-1} via the maps φ_α for every $\alpha \in \mathcal{K}^n$, i.e.

$$X^n = X^{n-1} \cup_{\varphi^n} \coprod_{\alpha \in \mathcal{K}^n} \mathbb{D}^n, \quad \text{where} \quad \varphi^n := \coprod_{\alpha \in \mathcal{K}^n} \varphi_\alpha : \coprod_{\alpha \in \mathcal{K}^n} \partial \mathbb{D}^n \to X^{n-1}. \tag{36.1}$$

We call X^n the **n-skeleton** (*n-Skelett* or *n-Gerüst*) of X. We call the individual points of X^0 the 0-**cells** (0-*Zellen*) of the complex, and it will be convenient to also denote $\mathcal{K}^0 := X^0$. For each $n \in \mathbb{N}$ and $\alpha \in \mathcal{K}^n$, the interior of the copy of $\mathbb{D}^n$ associated to α in the disjoint union defines an open subset

$$e^n_\alpha \subset X^n,$$

which is called an **n-cell** (*n-Zelle*) of the complex, and the associated map $\varphi_\alpha : S^{n-1} \to X^{n-1}$ is called its **attaching map** (*Anklebeabbildung*). The map

$$\Phi^n_\alpha : \mathbb{D}^n \to X$$

that satisfies $\Phi^n_\alpha|_{\partial\mathbb{D}^n} = \varphi_\alpha$ and restricts to the interior of the disk as the inclusion $e^n_\alpha \hookrightarrow X^n$ is called the **characteristic map** (*charakteristische Abbildung*) of the cell e^n_α. The complex is called **n-dimensional** if n is the largest number for which it contains an n-cell, i.e. $\mathcal{K}^m = \varnothing$ for all $m > n$ but $\mathcal{K}^n \neq \varnothing$.

Let us recall quickly what the notation in (36.1) means: we are defining X^n as a quotient of a disjoint union,

$$X^n = X^{n-1} \amalg \left(\coprod_{\alpha \in \mathcal{K}^n} \mathbb{D}^n \right) \Big/ \sim,$$

where $x \sim \varphi^n(x)$ for every $x \in \coprod_{\alpha\in\mathcal{K}^n} \partial\mathbb{D}^n$. The topology of X^n is implicit in this definition: if we know the topology of X^{n-1}, then the topology of X^n is determined via the quotient topology and the disjoint union topology, so in this way one can start from the discrete space X^0 and deduce the topology of every individual skeleton X^n one by one. Now, I'm not sure if you noticed this, but nothing we've said so far specifies the topology of X itself, at least not in the most general cases—it may well happen that $X = X^n$ for some $n \geqslant 0$ because the complex is finite-dimensional, so then the topology of X^n defines the topology of X, but more needs to be said if the complex is infinite dimensional.

Definition 36.2 (CW-complexes, part 2 of 2). The topology of a CW-complex $X = X^0 \cup X^1 \cup X^2 \cup \ldots$ is defined by the condition that a subset $\mathcal{U} \subset X$ is open if and only if $\mathcal{U} \cap X^n$ is an open subset of X^n for every $n \geqslant 0$.

Exercise 36.3. Show that a subset $\mathcal{U} \subset X$ in a CW-complex is open if and only if for every $n \geqslant 0$ and every n-cell e^n_α, $\Phi^{-1}_\alpha(\mathcal{U})$ is an open subset of $\mathbb{D}^n$. In other words, the topology of a CW-complex is the strongest possible topology for which all characteristic maps are continuous.

Exercise 36.4. Show that for any CW-complex X and any space Y, a map $f : X \to Y$ is continuous if and only if its restriction to the n-skeleton of X is continuous for every $n \geqslant 0$, or equivalently, if $f \circ \Phi_\alpha : \mathbb{D}^n \to Y$ is continuous for every $n \geqslant 0$ and $\alpha \in \mathcal{K}^n$.

Remark 36.5. You may by now have noticed an awkward problem with our terminology: the "W" in "CW" supposedly stands for "weak topology," yet the topology described in Definition 36.2 is not weak at all, but is the *strongest* with a given property. This discrepancy is apparently the fault of J.H.C. Whitehead,[59] whose influence on the subject was so substantial that many authors still refer to the topology of CW-complexes as "the weak topology" in the literature. Exercise 36.4 at least provides an argument for this term, as a CW-complex X is "weak" in the sense that it is fairly easy for functions defined on X to be continuous.

Definition 36.6. A **cell decomposition** (*Zellenzerlegung*) of a space X is a choice of homeomorphism from X to a CW-complex.

[59]This terminological awkwardness has motivated various popular jokes among mathematicians, one of which is that Whitehead's initials actually stand for "Jesus, he's confused!"

Example 36.7. Since the standard n-simplex Δ^n is homeomorphic to $\mathbb{D}^n$, the polyhedron $X = |K|$ of any simplicial complex K is also a CW-complex, whose n-cells are the interiors of the n-simplices, and the n-skeleton is thus the union of all the k-simplices for $k \leqslant n$. The attaching map $\partial\Delta^n \cong S^{n-1} \to X^{n-1}$ of each n-cell $e^n_\alpha \subset X$ is then a homeomorphism to the polyhedron of the $(n-1)$-dimensional subcomplex determined by the boundary faces of the corresponding n-simplex, and it follows that all of the characteristic maps in this case are inclusions. A cell decomposition of this type is equivalent to a triangulation.

Example 36.8. Recall that $S^n \cong \mathbb{D}^n/\partial\mathbb{D}^n$ for $n \geqslant 1$. This picture of the sphere defines a cell decomposition of S^n with one 0-cell and one n-cell: the 0-cell is the point $e^0 \in \mathbb{D}^n/\partial\mathbb{D}^n$ represented by any point in $\partial\mathbb{D}^n$, and the characteristic map of the n-cell e^n is the quotient map $\Phi : \mathbb{D}^n \to \mathbb{D}^n/\partial\mathbb{D}^n$. This identifies S^n with an n-dimensional CW-complex whose k-skeleton for each $k < n$ is a single point.

Note that in Example 36.8, the attaching map of the n-cell is very far from being injective, thus its characteristic map is also not injective at $\partial\mathbb{D}^n$, in contrast to Example 36.7. On the other hand, the restriction of a characteristic map to the interior of the disk is *always* injective, and is a homeomorphism onto its (open) image.

Example 36.9. There is another favorite cell decomposition of S^n in which the k-skeleton for each $k = 0, \dots, n$ is homeomorphic to S^k. The idea is to start with two points $X^0 := S^0$, and then inductively define X^k for each $k = 1, \dots, n$ by regarding $X^{k-1} = S^{k-1}$ as an equator and gluing two cells to it to form the "northern" and "southern" hemispheres of S^k:

$$S^k = S^{k-1} \cup_{\varphi^k} (\mathbb{D}^k_+ \amalg \mathbb{D}^k_-).$$

In this case there are exactly two k-cells for each $k = 0, \dots, n$, all attaching maps $S^{k-1} \to X^{k-1}$ are homeomorphisms, and all characteristic maps are inclusions.

Example 36.10. It is natural to define the decomposition $S^n = \mathbb{D}^n_+ \cup_{S^{n-1}} \mathbb{D}^n_-$ used in the previous example such that the antipodal map $S^n \to S^n$ sends $\mathbb{D}^n_\pm$ to $\mathbb{D}^n_\mp$ and restricts to the equator S^{n-1} as the antipodal map, which we can then assume satisfies the same condition with respect to the decomposition $S^{n-1} = \mathbb{D}^{n-1}_+ \cup_{S^{n-2}} \mathbb{D}^{n-1}_-$ and so forth. In this way, Example 36.9 also gives rise to a cell decomposition of $\mathbb{RP}^n = S^n/\mathbb{Z}_2$ with exactly one k-cell for each $k = 0, \dots, n$. The k-skeleton of $\mathbb{RP}^n$ is then a submanifold of the form

$$X^k = \left\{[(x_0, \dots, x_n)] \in \mathbb{RP}^n = S^n/\mathbb{Z}_2 \;\middle|\; x_{k+1} = \dots = x_n = 0\right\} \cong \mathbb{RP}^k.$$

In contrast to Example 36.9, the characteristic maps $\mathbb{D}^k \to \mathbb{RP}^n$ for this cell decomposition are not injective: indeed, the k-cells in Example 36.9 are attached to the $(k-1)$-skeleton S^{k-1} via a homeomorphism $S^{k-1} \to S^{k-1}$, but in $\mathbb{RP}^n$ this must be understood as a map to $X^{k-1} = \mathbb{RP}^{k-1} = S^k/\mathbb{Z}_2$, thus the homeomorphism $S^{k-1} \to S^{k-1}$ from Example 36.9 gets composed with the quotient projection $S^{k-1} \to \mathbb{RP}^{k-1}$ and becomes a covering map of degree 2.

Example 36.11. This will be harder to picture, but one can adjust Example 36.9 by following the same procedure of attaching two k-cells along homeomorphisms $S^{k-1} \to X^{k-1}$ for *every* $k \in \mathbb{N}$, without stopping when $k = n$. The result is an infinite-dimensional CW-complex called S^∞. The best way to picture it is probably as a subset of the infinite-dimensional vector space $\mathbb{R}^\infty := \bigoplus_{k=1}^\infty \mathbb{R}$, consisting of all sequences of real numbers $(x_1, x_2, x_3, \dots)$ that have only finitely many nonzero terms. Here we can identify $\mathbb{R}^n$ for each $n \geqslant 1$ with the subspace $\{(x_1, \dots, x_n, 0, 0, \dots) \in \mathbb{R}^\infty\}$, so that $S^k \subset \mathbb{R}^{k+1}$ becomes a subset of $\mathbb{R}^\infty$ that also happens to be contained in S^{k+1}, and S^∞ is the union of the nested sequence of spaces

$$S^0 \subset S^1 \subset S^2 \subset S^3 \subset \dots \subset \bigcup_{k \geqslant 0} S^k = S^\infty.$$

More concretely, S^∞ is just the subset of $\mathbb{R}^\infty$ defined by the condition $\sum_{i=1}^\infty x_i^2 = 1$, where there is no question about convergence since only finitely many terms can be nonzero. As the next exercise shows, there is something a bit subtle about the topology of S^∞.

EXERCISE 36.12. Show that if $x_k \in S^\infty$ is a convergent sequence, then there exists $n \in \mathbb{N}$ such that $x_k \in S^n$ for every k.
Hint: Given $x \in S^n \subset S^\infty$ and a sequence $x_k \in S^\infty$ such that $x_k \notin S^k$ for all k, construct a neighborhood $\mathcal{U} \subset S^\infty$ of x such that $x_k \notin \mathcal{U}$ for all k.

REMARK 36.13. The exercise reveals that S^∞ is in some sense quite different from any "infinite-dimensional sphere" that one would be likely to study in functional analysis. For instance, if S is the set of unit vectors in the infinite-dimensional Hilbert space

$$\ell^2 := \left\{ \mathbf{x} = (x_1, x_2, \dots) \in \prod_{i=1}^\infty \mathbb{R} \;\middle|\; \sum_{i=1}^\infty x_i^2 < \infty \right\}$$

with inner product $\langle \mathbf{x}, \mathbf{y} \rangle := \sum_i x_i y_i$, then there is no reason for the terms in a convergent sequence in S to belong to any particular finite-dimensional subspace. One can show however that S and S^∞ are nonetheless homotopy equivalent—in fact, both are contractible! (A proof of this for S^∞ can be found in [**Hat02**, p. 88].)

REMARK 36.14. Combining Examples 36.10 and 36.11 in the obvious way produces another infinite-dimensional CW-complex called $\mathbb{RP}^\infty$, which has exactly one k-cell for every $k \geqslant 0$. This space is of great theoretical importance, as it arises e.g. as the so-called *classifying space* of the group $\mathbb{Z}_2$, meaning that classification questions for certain classes of vector bundles over reasonable spaces X can be reduced to computations of the set of homotopy classes of maps $X \to \mathbb{RP}^\infty$. The theory of characteristic classes is founded in large part on understanding the homotopy types of certain infinite-dimensional CW-complexes such as this one; see e.g. [**MS74**].

EXAMPLE 36.15. Recall that the closed oriented surface Σ_g of genus $g \geqslant 0$ can be presented as a polygon with $4g$ sides, with certain pairs of sides identified as dictated by the word $a_1, b_1, a_1^{-1}, b_1^{-1}, a_2, b_2, a_2^{-1}, b_2^{-1}, \dots, a_g, b_g, a_g^{-1}, b_g^{-1}$ (see Definition 14.6 in last semester's Lecture 14). This defines a CW-complex in which there is one 0-cell (the vertices of the polygon are all identified with the same point), $2g$ one-cells which can be labeled $a_1, b_1, \dots, a_g, b_g$ and are attached along the unique map $S^0 \to X^0$, and a single 2-cell attached via a map $S^1 \to X^1$ that defines the concatenation of loops indicated by the above word.

DEFINITION 36.16. A **subcomplex** of a CW-complex X is a subset $A \subset X$ that is also a CW-complex with n-skeleton $A^n = A \cap X^n$ for all $n \geqslant 0$, such that every cell in A is also a cell in X with the same characteristic map.

Our goal in this lecture is to get as quickly as possible to the definition of cellular homology so that we can compute some examples. For this definition to make sense in full generality, we need an observation about the point-set topology of CW-complexes that is vacuous in the case of finite complexes but nontrivial for infinite complexes. We will postpone its proof until the next lecture, and since most of the interesting examples we consider will be finite complexes anyway, you are safe in ignoring it most of the time.

PROPOSITION 36.17. *For any CW-complex X, any compact subspace $K \subset X$ is contained in a finite subcomplex of X, i.e. in a subcomplex with only finitely many cells.*

The following consequence is the reason for the term "closure-finite":

COROLLARY 36.18. *The closure of each cell in a CW-complex intersects only finitely many other cells.* □

Remark 36.19. Notice that Proposition 36.17 immediately implies the result of Exercise 36.12. It is worth trying to do the exercise independently of this, in order to develop some intuition as to why Proposition 36.17 is true.

We can now define the **cellular chain complex** (*zellulärer Kettenkomplex*) associated to a CW-complex X. As usual, we shall fix an arbitrary choice of abelian group G to use for coefficients, and omit G from the notation whenever the choice of coefficients is unimportant. For $n \in \mathbb{Z}$, define $C_n^{\mathrm{CW}}(X)$ to be the trivial group if $n < 0$, and otherwise

$$C_n^{\mathrm{CW}}(X) = C_n^{\mathrm{CW}}(X;G) := \bigoplus_{\alpha \in \mathcal{K}^n} G,$$

so e.g. $C_n^{\mathrm{CW}}(X;\mathbb{Z})$ is the free abelian group generated by the set of n-cells e_α^n in our given cell decomposition of X, and for general coefficient groups G, $C_n^{\mathrm{CW}}(X;G)$ can be identified with $C_n^{\mathrm{CW}}(X;\mathbb{Z}) \otimes G$. We shall regard the cells e_α^n as generators of $C_n^{\mathrm{CW}}(X)$, thus writing elements of $C_n^{\mathrm{CW}}(X)$ as finite sums

$$\sum_i m_i e_{\alpha_i}^n \in C_n^{\mathrm{CW}}(X)$$

for coefficients $m_i \in G$ and indices $\alpha_i \in \mathcal{K}^n$. The direct sum of all these groups produces a $\mathbb{Z}$-graded abelian group

$$C_*^{\mathrm{CW}}(X) = C_*^{\mathrm{CW}}(X;G) := \bigoplus_{n \in \mathbb{Z}} C_n^{\mathrm{CW}}(X),$$

which we shall now turn into a chain complex by defining a suitable boundary operator $\partial : C_*^{\mathrm{CW}}(X) \to C_{*-1}^{\mathrm{CW}}(X)$. There is a geometric motivation for the definition: for each generator e_α^n of $C_n^{\mathrm{CW}}(X)$, we want ∂e_α^n to be a linear combination of $(n-1)$-cells determined by the attaching map φ_α, which tells us how the closure of e_α^n is glued to the $(n-1)$-skeleton of X. For this purpose, associate to each $\alpha \in \mathcal{K}^n$ with $n \geqslant 1$ the map $p_\alpha : X^n \to S^n$ determined by the following diagram:

(36.2)
$$\begin{tikzcd} X^n \arrow[r, "p_\alpha"] \arrow[d, "\mathrm{pr}"] & \mathbb{D}^n/\partial\mathbb{D}^n \arrow[r, equal] \arrow[dl, "\Phi_\alpha"] & S^n \\ X^n/(X^n \backslash e_\alpha^n) & & \end{tikzcd}$$

Here pr denotes the quotient projection, and the fact that φ_α maps $\partial\mathbb{D}^n$ into $X^{n-1} \subset X^n \backslash e_\alpha^n$ implies that the characteristic map $\Phi_\alpha : \mathbb{D}^n \to X^n$ descends to a map of the quotients $\mathbb{D}^n/\partial\mathbb{D}^n \to X^n/(X^n \backslash e_\alpha^n)$. The key point is that the latter is a homeomorphism, thus we can invert it to define $p_\alpha = \Phi_\alpha^{-1} \circ \mathrm{pr}$ as a map from $X^n \to S^n$ after identifying S^n with $\mathbb{D}^n/\partial\mathbb{D}^n$. This doesn't quite make sense if $n = 0$ since we cannot write "$\mathbb{D}^0/\partial\mathbb{D}^0 = S^0$," nonetheless there is a natural map of $X^0/(X^0 \backslash e_\alpha^0)$ to $S^0 = \{1, -1\}$ sending the cell e_α^0 to 1 and the equivalence class represented by every other 0-cell to -1. This map is a bijection except in the special case $X^0 \cong \{\mathrm{pt}\}$, i.e. when there is only one 0-cell e_α^0 and $X^0 \backslash e_\alpha^0 = \varnothing$, but in this case the map $X^0/(X^0 \backslash e_\alpha^0) \to S^0$ is well defined nonetheless, so we will adopt the convention of using it to define

(36.3)
$$\begin{tikzcd} X^0 \arrow[r, "p_\alpha"] \arrow[d, "\mathrm{pr}"] & S^0 \\ X^0/(X^0 \backslash e_\alpha^0) \arrow[ur] & \end{tikzcd}$$

as the analogue of (36.2) for $n = 0$.

DEFINITION 36.20. Given an n-cell e^n_α and an $(n-1)$-cell e^{n-1}_β in a CW-complex X, we define the **incidence number**

$$[e^{n-1}_\beta : e^n_\alpha] \in \mathbb{Z}$$

as the degree of the map

$$S^{n-1} \xrightarrow{p_\beta \circ \varphi_\alpha} S^{n-1}$$

defined by composing the attaching map $\varphi_\alpha : S^{n-1} \to X^{n-1}$ for e^n_α with the map $p_\beta : X^{n-1} \to S^{n-1}$ defined by replacing e^n_α with e^{n-1}_β in the diagram (36.2) or (36.3). We may sometimes abbreviate the incidence number by $[\beta : \alpha]$.

Observe that whenever $\overline{e^n_\alpha} \cap e^{n-1}_\beta = \varnothing$, it follows that the image of $\varphi_\alpha : S^{n-1} \to X^{n-1}$ is disjoint from e^{n-1}_β and is thus mapped to a constant by $p_\beta : X^{n-1} \to S^{n-1}$, hence $p_\beta \circ \varphi_\alpha$ is a constant map and $[e^{n-1}_\beta : e^n_\alpha] = 0$. In light of Corollary 36.18, this implies that the sum in the following definition makes sense, because it can only have finitely-many nonzero terms.

DEFINITION 36.21. For each $n \in \mathbb{N}$, the boundary map on the cellular chain complex $C^{\mathrm{CW}}_*(X)$ is determined via the formula

$$\partial e^n_\alpha := \sum_{\beta \in \mathcal{K}^{n-1}} [e^{n-1}_\beta : e^n_\alpha] e^{n-1}_\beta$$

for each $\alpha \in \mathcal{K}^n$.

Implicit in this definition is that $\partial : C^{\mathrm{CW}}_n(X) \to C^{\mathrm{CW}}_{n-1}(X)$ is the trivial map for every $n \leqslant 0$, as it must be since its target is then the trivial group. We shall now state two important theorems whose proofs will take up most of the next two lectures: the first states simply that $(C^{\mathrm{CW}}_*(X), \partial)$ is a chain complex.

THEOREM 36.22. *The map $\partial : C^{\mathrm{CW}}_*(X) \to C^{\mathrm{CW}}_*(X)$ satisfies $\partial^2 = 0$.*

The **cellular homology** (*zelluläre Homologie*) of the CW-complex X with coefficient group G can now be defined as

$$H^{\mathrm{CW}}_*(X) = H^{\mathrm{CW}}_*(X; G) := H_*\big(C^{\mathrm{CW}}_*(X; G), \partial\big).$$

The notation $C^{\mathrm{CW}}_*(X)$ and $H^{\mathrm{CW}}_*(X)$ is in some sense slightly non-ideal, as it hides the fact that the definitions of these objects depend on more than just a space X and coefficient group G, but also on a cell decomposition of X. The next theorem reveals why this is not a big deal.

THEOREM 36.23. *For any CW-complex X and any axiomatic homology theory h_* with coefficient group G, there is an isomorphism $H^{\mathrm{CW}}_*(X; G) \cong h_*(X)$.*

We will improve this statement slightly in the next lecture by talking about CW-pairs and morphisms of CW-pairs, so that the isomorphism between $H^{\mathrm{CW}}_*(\cdot; G)$ and h_* can be regarded as a natural transformation. Theorem 36.23 has several remarkable consequences that can be recognized immediately: one is that $H^{\mathrm{CW}}_*(X)$ depends (up to isomorphism) only on the topology of X and not on its cell decomposition, and another is that all axiomatic homology theories with any given coefficient group are isomorphic if we restrict them to spaces that are nice enough to have cell decompositions. In light of Example 36.7 and the following exercise, this also tells us why the *simplicial* homology of a polyhedron depends only on its topology, and not on its simplicial decomposition.

EXERCISE 36.24. Convince yourself that if K is a simplicial complex and its polyhedron $X = |K|$ is viewed as a CW-complex as in Example 36.7, then its cellular chain complex $C^{\mathrm{CW}}_*(X)$ is isomorphic to its oriented simplicial chain complex $C^{\Delta}_*(K)$.

Before trying to explain why all this is true, let's look at a couple of examples that will make Theorem 36.23 look more plausible.

EXAMPLE 36.25. We saw in Example 36.8 that S^n for each $n \in \mathbb{N}$ has a cell decomposition with one 0-cell e^0 and one n-cell e^n, so $X^0 = X^1 = \ldots = X^{n-1} \cong \{\text{pt}\}$ and $X^n = S^n$. These two cells are thus the only generators of $C_*^{\text{CW}}(S^n)$, giving

$$C_k^{\text{CW}}(S^n; G) = \begin{cases} G & \text{if } k = 0, n, \\ 0 & \text{otherwise.} \end{cases}$$

We claim that on this chain complex, $\partial = 0$, hence $H_*^{\text{CW}}(S^n) = C_*^{\text{CW}}(S^n)$, which matches our previous computation of $h_*(S^n)$ for any axiomatic homology theory. If $n \geqslant 2$, then the claim holds trivially because for every $k \in \mathbb{Z}$, either the domain or the target of the map $\partial : C_k^{\text{CW}}(S^n) \to C_{k-1}^{\text{CW}}(S^n)$ is trivial. When $n = 1$ there is still something to check: $\partial : C_1^{\text{CW}}(S^n) \to C_0^{\text{CW}}(S^n)$ might theoretically be nontrivial since its domain and target are both G. The map will be trivial for every choice of coefficient group if and only if

$$\partial e^1 = [e^0 : e^1] e^0$$

is trivial, i.e. if the incidence number $[e^0 : e^1]$ is 0. This is the degree of a map $p \circ \varphi : S^0 \to S^0$, where $\varphi : S^0 \to X^0 \cong \{\text{pt}\}$ is the attaching map for e^1 and $p : X^0 \to S^0$ sends e^0 to $1 \in S^0$. Since both of these maps are constant, $[e^0 : e^1] = \deg(p \circ \varphi) = 0$.

EXAMPLE 36.26. We consider S^2 with the alternative cell decomposition described in Example 36.9, which has two k-cells $e_\pm^k$ for each $k = 0, 1, 2$, hence $S^2 = e_+^0 \cup e_-^0 \cup e_+^1 \cup e_-^1 \cup e_+^2 \cup e_-^2$, and the k-skeleton is $X^k = S^k \subset S^2$ for $k = 0, 1, 2$. We now have $C_k^{\text{CW}}(S^n) = 0$ for $k < 0$ or $k > 2$, while $C_k^{\text{CW}}(S^n) = G \oplus G$ for each $k = 0, 1, 2$, with the two factors of the coefficient group G corresponding to the two generators $e_+^k, e_-^k \in C_k^{\text{CW}}(S^n; \mathbb{Z})$. Denote the attaching map for $e_\pm^k$ by $\varphi_\pm^k : S^{k-1} \to X^{k-1}$, and denote the projection map as defined in (36.2) by $p_\pm^k : X^k \to S^k$, so $\partial : C_k^{\text{CW}}(S^n) \to C_{k-1}^{\text{CW}}(S^n)$ is now determined by

$$\begin{aligned} \partial e_+^k &= \deg(p_+^{k-1} \circ \varphi_+^k) e_+^{k-1} + \deg(p_-^{k-1} \circ \varphi_+^k) e_-^{k-1}, \\ \partial e_-^k &= \deg(p_+^{k-1} \circ \varphi_-^k) e_+^{k-1} + \deg(p_-^{k-1} \circ \varphi_-^k) e_-^{k-1}. \end{aligned} \tag{36.4}$$

To compute these degrees, we will need a slightly more concrete description of the maps involved. Let us regard S^2 as the unit sphere in the xyz-plane, with its 1-skeleton formed by the unit circle in the xy-plane, and the 0-skeleton consisting of the two points $(\pm 1, 0, 0)$. It is then natural to parametrize the characteristic maps $\Phi_\pm^1 : \mathbb{D}^1 \to S^2$ of the two 1-cells $e_\pm^1$ via the x coordinate, giving

$$\Phi_\pm^1 : \mathbb{D}^1 \to S^2 : x \mapsto (x, \pm\sqrt{1 - x^2}, 0),$$

so the attaching maps $\varphi_\pm^1 : S^0 \to S^0$ are the restrictions of these to $\partial \mathbb{D}^1$ and are thus both the identity map $S^0 \to S^0$. Each of the maps $p_\pm^0 : X^0 \to S^0$ is likewise a bijection in this example, sending its "favorite" 0-cell $e_\pm^0$ to $1 \in S^0$ and the other one to $-1 \in S^0$, so in fact, p_+^0 is the identity map $S^0 \to S^0$ and p_-^0 is the bijection sending ± 1 to ∓ 1. The latter has degree -1, so we can now fill in the coefficients for $k = 1$ in (36.4) and write

$$\partial e_+^1 = \partial e_-^1 = e_+^0 - e_-^0.$$

For the 2-cells $e_\pm^2$, the most obvious parametrization is defined by inverting the projection $(x, y, z) \mapsto (x, y)$, so we can define the characteristic maps by

$$\Phi_\pm^2 : \mathbb{D}^2 \to S^2 : (x, y) \mapsto (x, y, \pm\sqrt{1 - x^2 - y^2}),$$

and the attaching maps $\varphi^2_\pm : S^1 \to X^1$ thus become once again the identity map $S^1 \to S^1$. To understand the maps $p^1_\pm : X^1 \to S^1$, let us first agree that the identification of $\mathbb{D}^1/\partial\mathbb{D}^1$ with S^1 should be defined via a path $\gamma : \mathbb{D}^1 \to S^1$ that sends $\pm 1 \mapsto 1$ and traverses a loop $\gamma(t) \in S^1$ with winding number $+1$ as t goes from -1 to 1. Now, $p^1_+ : S^1 \to \mathbb{D}^1/\partial\mathbb{D}^1$ sends the top half of the circle $S^1 = X^1$ to $\mathbb{D}^1$ via the inverse of our chosen characteristic map Φ^1_+ and sends the bottom half of the circle to a constant: the resulting winding number is $\deg(p^1_+ \circ \varphi^2_\pm) = -1$. Meanwhile, $p^2_- : S^1 \to \mathbb{D}^1/\partial\mathbb{D}^1$ sends the top half of the circle to a constant but maps the bottom half to $\mathbb{D}^1$ as the inverse of Φ^1_-, producing $\deg(p^1_- \circ \varphi^2_\pm) = 1$. We thus have

$$\partial e^2_+ = \partial e^2_- = -e^1_+ + e^1_-.$$

With these formulas in place, we can compute the homology of $C^{\mathrm{CW}}_*(S^2)$ explicitly: acting with ∂ on an arbitrary 2-chain $ge^2_+ + he^2_-$ for $g,h \in G$ gives

$$\partial(ge^2_+ + he^2_-) = -(g+h)e^1_+ + (g+h)e^1_- = (g+h)(-e^1_+ + e^1_-),$$

which vanishes if and only if $g = -h$, so in terms of the obvious identification of $C^{\mathrm{CW}}_2(S^2)$ with $G \oplus G$, the group of 2-cycles takes the form

$$\ker \partial_2 = \{(g,-g) \in G \oplus G \mid g \in G\} \subset C^{\mathrm{CW}}_2(S^2),$$

which is isomorphic to G. Since $C^{\mathrm{CW}}_3(S^2) = 0$, we conclude $H^{\mathrm{CW}}_2(S^2) \cong G$. To find the 1-cycles, we similarly compute

$$\partial(ge^1_+ + he^1_-) = (g+h)e^0_+ - (g+h)e^0_- = (g+h)(e^0_+ - e^0_-),$$

and this again vanishes if and only if $g = -h$, so the 1-cycles consist of all elements of the form $g(e^1_+ - e^1_-)$. But these are also boundaries since $\partial(-ge^2_+) = g(e^1_+ - e^1_-)$, thus $H^{\mathrm{CW}}_1(S^2) = 0$. Finally, all 0-chains $ge^0_+ + he^0_-$ are cycles since $C^{\mathrm{CW}}_{-1}(S^2) = 0$, but under the obvious isomorphism $C^{\mathrm{CW}}_0(S^2) = G \oplus G$ we have

$$\operatorname{im} \partial_1 = \{(g,-g) \in G \oplus G \mid g \in G\} \subset C^{\mathrm{CW}}_0(S^2),$$

so $H^{\mathrm{CW}}_0(S^2)$ is isomorphic to the quotient of $G \oplus G$ by this subgroup, which is again G. The end result therefore matches the $n = 2$ case of Example 36.25.

It is not too hard to extend Example 36.26 to a computation of $H^{\mathrm{CW}}_*(S^n)$ for every $n \in \mathbb{N}$ in terms of the cell decomposition $S^n = e^0_+ \cup e^0_- \cup \ldots \cup e^n_+ \cup e^n_-$. Getting all the signs right is a bit of a pain, but all coefficients will again work out to ± 1 in such a way that all nontrivial k-cycles are also boundaries for $k = 1, \ldots, n-1$, but the groups $\ker \partial_n$ and $C^{\mathrm{CW}}_0(S^n)/\operatorname{im}\partial_1$ are again both G. The fact that getting all the signs right is a bit tricky is an argument for doing the computation via the simpler cell decomposition $S^n = e^0 \cup e^n$ instead, as in Example 36.25, so we will invest considerable effort over the next couple of lectures into proving that this is allowed, because the isomorphism class of $H^{\mathrm{CW}}_*(X)$ depends in general only on the topology of X and not on its cell decomposition.

Exercise 36.27. Figure 20 shows two spaces that you may recall from *Topologie I* are both homeomorphic to the Klein bottle. Each also defines a cell complex $X = X^0 \cup X^1 \cup X^2$ consisting of one 0-cell, two 1-cells (labeled a and b) and one 2-cell.

(a) Compute $H^{\mathrm{CW}}_*(X;\mathbb{Z})$, $H^{\mathrm{CW}}_*(X;\mathbb{Z}_2)$ and $H^{\mathrm{CW}}_*(X;\mathbb{Q})$ for both complexes. (You'll know you've done something wrong if the answers you get from the two complexes are not isomorphic!)

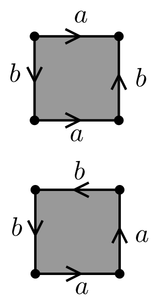

FIGURE 20. The two cell decompositions of the Klein bottle considered in Exercise 36.27.

(b) Recall that the **rank** (*Rang*) of a finitely generated abelian group G is the unique integer $k \geqslant 0$ such that $G \cong \mathbb{Z}^k \oplus T$ for some finite group T. Verify for both cell decompositions of the Klein bottle above that

$$\sum_k (-1)^k \operatorname{rank} H_k^{\mathrm{CW}}(X;\mathbb{Z}) = \sum_k (-1)^k \dim_{\mathbb{Z}_2} H_k^{\mathrm{CW}}(X;\mathbb{Z}_2) = \sum_k (-1)^k \dim_{\mathbb{Q}} H_k^{\mathrm{CW}}(X;\mathbb{Q}) = 0.$$

(Congratulations, you've just computed the Euler characteristic of the Klein bottle! A comprehensive discussion of this invariant is coming up in Lecture 40.)

37. Invariance of cellular homology, part 1 (November 21, 2023)

Having defined the cellular homology $H_*^{\mathrm{CW}}(X;G)$ for a CW-complex X in the previous lecture, we would now like to begin working toward the proof that it is isomorphic to $h_*(X)$ for any axiomatic homology theory h_* with coefficient group G. The proper statement of that result is Theorem 37.9 below. But first, there was a more basic result about CW-complexes that we left unproved in the previous lecture, without which the cellular chain complex is not generally well defined.

PROPOSITION 37.1. *For any CW-complex X, any compact subspace $K \subset X$ is contained in a finite subcomplex of X, i.e. in a subcomplex with only finitely many cells.*

PROOF. *Step 1:* Suppose $A \subset X$ is a subset with the property that for every pair of distinct elements $x, y \in A$, x and y belong to different cells of the complex. We claim then that $A \cap X^n$ is a closed subset of X^n for every integer $n \geqslant 0$. The proof is by induction on n; for $n = 0$ it is trivially true since X^0 carries the discrete topology, so all of its subsets are closed. Now if we assume $A \cap X^{n-1} \subset X^{n-1}$ is closed, it follows that for every n-cell e_α^n with attaching map $\varphi_\alpha : S^{n-1} \to X^{n-1}$ and characteristic map $\Phi_\alpha : \mathbb{D}^n \to X$, $\varphi_\alpha^{-1}(A)$ is a closed subset of S^{n-1}. Since at most one element of A can lie in e_α^n, the set $\Phi_\alpha^{-1}(A) \subset \mathbb{D}^n$ is then either $\varphi_\alpha^{-1}(A)$ or the union of this with a single point in the interior of the disk, so in either case it is closed. Viewing X^n itself as a CW-complex in the obvious way and remembering that closed sets are complements of open sets, Exercise 36.3 now implies that $A \cap X^n \subset X^n$ is closed. By induction, this is true for every $n \geqslant 0$, and it follows via the definition of the topology of X that A is a closed subset of X.

Step 2: Given a compact subset $K \subset X$, we claim that K can intersect at most finitely many distinct cells of X. Otherwise there exists an infinite subset $A \subset K$ in which every element belongs to a different cell. Step 1 implies that $A \subset X$ is closed, and moreover, so is every subset of A, which means that the induced subspace topology on A is the discrete topology. Since K is compact, this makes $A \subset K$ a compact discrete space, contradicting the assumption that A is infinite.

Step 3: We claim that for every $n \geqslant 0$, every compact subset $K \subset X^n$ is contained in a finite subcomplex of X^n. For $n = 0$ this is obvious since the compact subsets of X^0 are finite. By induction, if the claim is known for compact subsets of X^{n-1}, then it holds in particular for the image of the attaching map $\varphi_\alpha : S^{n-1} \to X^{n-1}$ of any n-cell e^n_α, providing a finite subcomplex $A \subset X^{n-1}$ whose union with e^n_α is a finite subcomplex of X^n containing e^n_α. In light of step 2, this proves the claim for all compact subsets of X^n, as finite unions of finite subcomplexes are also finite subcomplexes.

To conclude, step 3 implies that for every cell e^n_α of the complex, the compact subset $\overline{e^n_\alpha} = \Phi_\alpha(\mathbb{D}^n) \subset X$ is contained in a finite subcomplex, and combining this with the claim in step 2 proves the result. □

Let's briefly recall how the cellular chain complex $C^{\mathrm{CW}}_*(X)$ is defined. Each chain group $C^{\mathrm{CW}}_n(X)$ is freely generated by the set of n-cells e^n_α in X, and $\partial : C^{\mathrm{CW}}_n(X) \to C^{\mathrm{CW}}_{n-1}(X)$ is determined by the formula

$$\partial e^n_\alpha = \sum_{e^{n-1}_\beta} [e^{n-1}_\beta : e^n_\alpha] e^{n-1}_\beta,$$

where the sum is over all the $(n-1)$-cells e^{n-1}_β in X, and Proposition 37.1 implies that only finitely many terms are nonzero. The most important detail here is the incidence number $[e^{n-1}_\beta : e^n_\alpha] \in \mathbb{Z}$, which is the degree of the composition of two maps

$$S^{n-1} \xrightarrow{\varphi_\alpha} X^{n-1} \xrightarrow{p_\beta} S^{n-1},$$

where φ_α is the attaching map for e^n_α, and p_β is defined by collapsing everything outside of e^{n-1}_β to a point and using the characteristic map $\Phi_\beta : \mathbb{D}^{n-1} \to X^{n-1}$ to identify the resulting quotient with $\mathbb{D}^{n-1}/\partial\mathbb{D}^{n-1} = S^{n-1}$.

As mentioned in the previous lecture, this description of p_β doesn't quite work when $n = 1$, so let us work out a more useful formula for $\partial : C^{\mathrm{CW}}_1(X) \to C^{\mathrm{CW}}_0(X)$. If $X^0 \cong \{\mathrm{pt}\}$, then $p_\beta \circ \varphi_\alpha : S^0 \to S^0$ always factors through a one-point space and is therefore a constant map, implying $[e^0_\beta : e^1_\alpha] = 0$ for all $\beta \in \mathcal{K}^0$ and $\alpha \in \mathcal{K}^1$, so $\partial = 0$. If there is more than one 0-cell, then $p_\beta : X^0 \to S^0$ is the map that sends e^0_β to $1 \in S^1$ and every other 0-cell to $-1 \in S^1$, so composing it with the attaching map $\varphi_\alpha : \partial\mathbb{D}^1 \to X^0$ produces the following possibilities:

- If $\varphi_\alpha(1) = e^0_\beta$ and $\varphi_\alpha(-1) \neq e^0_\beta$, then $p_\beta \circ \varphi_\alpha : S^0 \to S^0$ is the identity map and thus $[e^0_\beta : e^1_\alpha] = 1$.
- If $\varphi_\alpha(1) \neq e^0_\beta$ but $\varphi_\alpha(-1) = e^0_\beta$, then $p_\beta \circ \varphi_\alpha(\pm 1) = \mp 1$ and thus $[e^0_\beta : e^1_\alpha] = -1$.
- In all other cases, $p_\beta \circ \varphi_\alpha$ is constant and thus $[e^0_\beta : e^1_\alpha] = 0$.

Since each point of X^0 is a 0-cell, we can identify it with a generator of $C^{\mathrm{CW}}_0(X)$ and thus deduce from the remarks above the following:

PROPOSITION 37.2. *The map $\partial : C^{\mathrm{CW}}_1(X) \to C^{\mathrm{CW}}_0(X)$ is determined by the formula*

$$\partial e^1_\alpha = \varphi_\alpha(1) - \varphi_\alpha(-1).$$

□

Let's do another easy example.

EXAMPLE 37.3. We saw in Example 36.15 that the closed oriented surface Σ_g of genus $g \geqslant 0$ has a cell decomposition with one 0-cell e^0, $2g$ cells of dimension one which we can label

$$e^1_{a_1}, e^1_{b_1}, \ldots, e^1_{a_g}, e^1_{b_g},$$

and a single 2-cell e^2, which is the interior of the usual polygon with $4g$ sides. In particular, the 0-skeleton X^0 is a single point, and the 1-skeleton X^1 is a wedge of $2g$ circles labeled $a_1, b_1, \dots, a_g, b_g$ that all intersect only at X^0. Since there is only one 0-cell, all of the 1-cells are cycles in $C_1^{\mathrm{CW}}(\Sigma_g)$:

$$\partial e^1_{a_j} = \partial e^1_{b_j} = 0 \quad \text{for} \quad j = 1, \dots, g.$$

The attaching map $\varphi : S^1 \to X^1$ of the 2-cell is a loop that traverses a_1, then b_1, then a_1 again backwards and b_1 again backwards, then moves on to a_2, b_2 and so forth, ending with b_g backwards. Composing this with the projection $p_{a_1} : X^1 \to S^1$ that collapses $X^1 \backslash e^1_{a_1}$ to a point, we obtain a concatenation of the loop a_1 with a constant path and then a_1^{-1} followed by another constant path, resulting in a map $S^1 \to S^1$ with degree 0. The same happens with all the other projections p_{a_j}, p_{b_j}, so that all of the incidence numbers in the computation of ∂e^2 vanish and we obtain

$$\partial e^2 = 0.$$

This proves that $\partial = 0$ for the entire cellular chain complex with arbitrary coefficients, hence

$$H_k^{\mathrm{CW}}(\Sigma_g; G) = C_k^{\mathrm{CW}}(\Sigma_g; G) \cong \begin{cases} G & \text{for } k = 0, 2, \\ G^{2g} & \text{for } k = 1, \\ 0 & \text{for } k < 0 \text{ and } k > 2. \end{cases}$$

There is also a relative version of cellular homology. A **CW-pair** (*CW-Paar*) is a pair of CW-complexes (X, A) such that A is a subcomplex of X. In this case $C_*^{\mathrm{CW}}(A)$ is a subcomplex of $C_*^{\mathrm{CW}}(X)$, i.e. it is a subgroup preserved by the boundary map, giving rise to a quotient chain complex

$$C_*^{\mathrm{CW}}(X, A) = C_*^{\mathrm{CW}}(X, A; G) := C_*^{\mathrm{CW}}(X) / C_*^{\mathrm{CW}}(A).$$

The homology of this complex is the **relative cellular homology**

$$H_*^{\mathrm{CW}}(X, A) = H_*^{\mathrm{CW}}(X, A; G) := H_*\big(C_*^{\mathrm{CW}}(X, A; G)\big).$$

By this point you should not be surprised to learn that one can define a category $\mathsf{CW}_{\mathrm{rel}}$ whose objects are CW-pairs, such that $C_*^{\mathrm{CW}} : \mathsf{CW}_{\mathrm{rel}} \to \mathsf{Chain}$ and $H_*^{\mathrm{CW}} : \mathsf{CW}_{\mathrm{rel}} \to \mathsf{Ab}_{\mathbb{Z}}$ become functors. But I still need to tell you what the morphisms in $\mathsf{CW}_{\mathrm{rel}}$ are.

Definition 37.4. A continuous map $f : X \to Y$ between CW-complexes is called a **cellular map** (*zelluläre Abbildung*) if $f(X^n) \subset Y^n$ for every $n \geqslant 0$. More generally, if (X, A) and (Y, B) are CW-pairs, a **map of CW-pairs** is a cellular map $f : X \to Y$ such that $f(A) \subset B$. (Its restriction $f|_A : A \to B$ is then automatically a cellular map.)

Example 37.5. If X and Y are polyhedra (and therefore also CW-complexes as explained in Example 36.7), then any simplicial map $f : X \to Y$ is also a cellular map.

Unlike simplicial maps, a cellular map $f : X \to Y$ need not generally map cells of X to cells of Y. Instead, the image of an individual cell $e^n_\alpha \subset X$ may cover many n-cells $e^n_\beta \subset Y$, and it may cover some of them multiple times, which can be measured by an incidence number analogous to the one appearing in the definition of ∂. On the 0-skeleton, the situation is straightforward: since $f(X^0) \subset Y^0$ and each 0-cell e^0_α is just a single point, $f(e^0_\alpha)$ is always a specific 0-cell in Y, so that for each 0-cell e^0_β of Y we can define

$$[e^0_\beta : e^0_\alpha] := \begin{cases} 1 & \text{if } f(e^n_\alpha) = e^0_\beta, \\ 0 & \text{otherwise.} \end{cases}$$

For $n \geqslant 1$, the key observation is that since $f(X^n) \subset Y^n$ and $f(X^{n-1}) \subset Y^{n-1}$, f descends to a map of quotients, $X^n/X^{n-1} \to Y^n/Y^{n-1}$ and we can therefore consider the composition

$$S^n = \mathbb{D}^n/\partial\mathbb{D}^n \xrightarrow{\Phi_\alpha} X^n/X^{n-1} \xrightarrow{f} Y^n/Y^{n-1} \xrightarrow{\mathrm{pr}} Y^n/(Y^n \backslash e^n_\beta) \xrightarrow{\Phi_\beta^{-1}} \mathbb{D}^n/\partial\mathbb{D}^n = S^n, \tag{37.1}$$

where the map labeled pr is the natural quotient projection, and the map Φ_β on quotients is invertible for the same reason as before. We shall denote the degree of this map by

$$[e^n_\beta : e^n_\alpha] \in \mathbb{Z}.$$

This incidence number vanishes whenever $e^n_\beta \cap f\left(\overline{e^n_\alpha}\right) = \varnothing$ since the map in (37.1) is in this case constant, so Proposition 37.1 implies that for each individual $e^n_\alpha \subset X$, there are at most finitely many $e^n_\beta \subset Y$ with $[e^n_\beta : e^n_\alpha] \neq 0$. This allows us to define a homomorphism

$$f_* : C^{\mathrm{CW}}_*(X) \to C^{\mathrm{CW}}_*(Y)$$

acting on the generators $e^n_\alpha \in C^{\mathrm{CW}}_n(X)$ as

$$f_* e^n_\alpha = \sum_{e^n_\beta} [e^n_\beta : e^n_\alpha] e^n_\beta, \tag{37.2}$$

where the sum ranges over all n-cells $e^n_\beta \subset Y$ and has only finitely many nonzero terms.

Exercise 37.6. Show that if X and Y are the same CW-complex and $f : X \to Y$ is the identity map, the incidence number $[e^n_\beta : e^n_\alpha]$ is 1 for $\alpha = \beta$ and 0 otherwise, so in particular, $f_* : C^{\mathrm{CW}}_*(X) \to C^{\mathrm{CW}}_*(Y)$ is the identity homomorphism.

Exercise 37.7. Show that if $f : X \to Y$ and $g : Y \to Z$ are cellular maps then $(g \circ f)_* = g_* \circ f_* : C^{\mathrm{CW}}_*(X) \to C^{\mathrm{CW}}_*(Z)$.

This discussion of induced maps extends in an obvious way to the relative case: if $f : (X, A) \to (Y, B)$ is a map of CW-pairs, then f_* maps $C^{\mathrm{CW}}_*(A)$ into $C^{\mathrm{CW}}_*(B)$ and thus descends to a homomorphism

$$f_* : C^{\mathrm{CW}}_*(X, A) \to C^{\mathrm{CW}}_*(Y, B).$$

The proof of the next theorem will arise naturally from the proof of the much bigger theorem that follows it.

Theorem 37.8. *For any map of CW-pairs $f : (X, A) \to (Y, B)$, $f_* : C^{\mathrm{CW}}_*(X, A) \to C^{\mathrm{CW}}_*(Y, B)$ is a chain map and thus induces homomorphisms $f_* : H^{\mathrm{CW}}_n(X, A) \to H^{\mathrm{CW}}_n(Y, B)$ for every n. In particular, the cellular chain complex and cellular homology with coefficients in any given abelian group G define functors*

$$C^{\mathrm{CW}}_* = C^{\mathrm{CW}}_*(\cdot; G) : \mathsf{CW}_{\mathrm{rel}} \to \mathsf{Chain}, \qquad \textit{and} \qquad H^{\mathrm{CW}}_* = H^{\mathrm{CW}}_*(\cdot; G) : \mathsf{CW}_{\mathrm{rel}} \to \mathsf{Ab}_{\mathbb{Z}},$$

where $\mathsf{CW}_{\mathrm{rel}}$ denotes the category of CW-pairs, with morphisms defined as maps of CW-pairs.

We can now state the complete version of the theorem about cellular homology and axiomatic homology theories.

Theorem 37.9. *Suppose h_* is an axiomatic homology theory with coefficient group G. Then one can associate to any CW-pair (X, A) isomorphisms*

$$\Psi_{(X,A)} : H^{\mathrm{CW}}_n(X, A; G) \xrightarrow{\cong} h_n(X, A)$$

for every n, which are natural in the sense that for any map of CW-pairs $f : (X, A) \to (Y, B)$, the following diagram commutes:

$$\begin{array}{ccc} H_n^{\mathrm{CW}}(X, A; G) & \xrightarrow{\Psi_{(X,A)}} & h_n(X, A) \\ \Big\downarrow f_* & & \Big\downarrow f_* \\ H_n^{\mathrm{CW}}(Y, B; G) & \xrightarrow{\Psi_{(Y,B)}} & h_n(Y, B) \end{array}$$

In the language of category theory, this theorem says the following. There is a functor $\mathsf{CW}_{\mathrm{rel}} \to \mathsf{Top}_{\mathrm{rel}}$ that sends each CW-pair to the underlying pair of spaces and each map of CW-pairs to the underlying continuous map, and composing h_* with this functor produces a functor $\mathsf{CW}_{\mathrm{rel}} \to \mathsf{Ab}_{\mathbb{Z}}$. The theorem defines a natural transformation from $H_*^{\mathrm{CW}}(\cdot; G)$ to the latter functor, associating to every CW-pair (X, A) the isomorphism $\Psi_{(X;A)}$.

Let us begin setting up the proof of the theorem. We shall focus here on the case of absolute homology, i.e. pairs (X, A) with $A = \varnothing$, leaving the relative case as a (worthwhile!) exercise. The key idea is to establish a relationship between $h_*(X)$ and the homology of a chain complex built out of the long exact sequences of the pairs (X^n, X^{n-1}) and (X^{n+1}, X^n), as it will turn out that the latter chain complex can be identified naturally with $C_*^{\mathrm{CW}}(X)$.

LEMMA 37.10. *For all $n \in \mathbb{N}$, (X^n, X^{n-1}) is a good pair in the sense of Definition 32.19, i.e. X^{n-1} is a deformation retract of some neighborhood $V \subset X^n$ of X^{n-1}.*

PROOF. Since $X^n = X^{n-1} \cup_{\varphi^n} \coprod_{\alpha \in \mathcal{K}^n} \mathbb{D}^n$, it suffices to set $V := X^{n-1} \cup_{\varphi^n} \coprod_{\alpha \in \mathcal{K}^n} (\mathbb{D}^n \backslash \{0\})$. □

By Theorem 32.23, we now have a natural isomorphism $h_*(X^n, X^{n-1}) \cong \widetilde{h}_*(X^n/X^{n-1})$ for each $n \geqslant 1$. Observe next that the disjoint union of the characteristic maps of n-cells defines a map of pairs

$$\Phi^n := \coprod_{\alpha \in \mathcal{K}^n} \Phi_\alpha : \coprod_{\alpha \in \mathcal{K}^n} (\mathbb{D}^n, \partial \mathbb{D}^n) \to (X^n, X^{n-1}).$$

We claim that this map descends to a *homeomorphism* between the quotients

$$\Phi^n : \coprod_{\alpha \in \mathcal{K}^n} \mathbb{D}^n \Big/ \coprod_{\alpha \in \mathcal{K}^n} \partial \mathbb{D}^n \xrightarrow{\cong} X^n / X^{n-1}.$$

Indeed, under the usual identification $\mathbb{D}^n / \partial \mathbb{D}^n = S^n$ that regards the collapsed boundary of $\mathbb{D}^n$ as a base point in S^n, the quotient on the left hand side here becomes the wedge sum $\bigvee_{\alpha \in \mathcal{K}^n} S^n$, with all copies of S^n attached at this base point. By inspection, the right hand side is exactly the same thing: $X^n \backslash X^{n-1}$ is the union of all the n-cells, which Φ^n identifies with copies of $\mathring{\mathbb{D}}^n$, and the quotient collapses the boundaries of all these disks to a point. With this understood, it follows that the map Φ_*^n at the bottom of the following diagram is an isomorphism, and so therefore is the map at the top:

$$\begin{array}{ccc} h_*\Big(\coprod_{\alpha \in \mathcal{K}^n} (\mathbb{D}^n, \partial \mathbb{D}^n) \Big) & \xrightarrow{\Phi_*^n} & h_*(X^n, X^{n-1}) \\ \Big\downarrow \cong & & \Big\downarrow \cong \\ \widetilde{h}_*\Big(\coprod_{\alpha \in \mathcal{K}^n} \mathbb{D}^n \Big/ \coprod_{\alpha \in \mathcal{K}^n} \partial \mathbb{D}^n \Big) & \xrightarrow{\Phi_*^n} & \widetilde{h}_*(X^n / X^{n-1}). \end{array}$$

Applying the additivity axiom (in conjunction with the five-lemma as in Exercise 32.16) to identify $h_*\Big(\coprod_{\alpha \in \mathcal{K}^n} (\mathbb{D}^n, \partial \mathbb{D}^n) \Big)$ with $\bigoplus_{\alpha \in \mathcal{K}^n} h_*(\mathbb{D}^n, \partial \mathbb{D}^n)$, this proves:

Lemma 37.11. *The characteristic maps* $\Phi_\alpha : (\mathbb{D}^n, \partial\mathbb{D}^n) \to (X^n, X^{n-1})$ *determine isomorphisms*

$$\bigoplus_{\alpha\in\mathcal{K}^n} (\Phi_\alpha)_* : \bigoplus_{\alpha\in\mathcal{K}^n} h_*(\mathbb{D}^n, \partial\mathbb{D}^n) \xrightarrow{\cong} h_*(X^n, X^{n-1})$$

for each $n \in \mathbb{N}$. □

The long exact sequence of $(\mathbb{D}^n, \partial\mathbb{D}^n)$ in reduced homology implies that the connecting homomorphisms

$$h_k(\mathbb{D}^n, \partial\mathbb{D}^n) \xrightarrow{\partial_*} \widetilde{h}_{k-1}(S^{n-1}) \cong \begin{cases} G & \text{if } k = n, \\ 0 & \text{if } k \ne n \end{cases}$$

are isomorphisms for all k and n, thus we've proved

$$h_k(X^n, X^{n-1}) \cong \begin{cases} C_n^{\mathrm{CW}}(X; G) = \bigoplus_{\alpha\in\mathcal{K}^n} G & \text{if } k = n, \\ 0 & \text{if } k \ne n. \end{cases} \tag{37.3}$$

We've been assuming $n \geqslant 1$ so far, but it is not hard to incorporate $n = 0$ into this discussion: if we set

$$X^{-1} := \varnothing,$$

then $h_k(X^0, X^{-1}) = h_k(X^0)$ is simply the homology of a discrete space, i.e. the disjoint union of one-point spaces

$$X^0 = \coprod_{\alpha\in\mathcal{K}^0} \{\mathrm{pt}\},$$

so that (37.3) is also correct in this case due to the dimension and additivity axioms. The group $h_n(X^n, X^{n-1})$ can therefore serve as a stand-in for $C_n^{\mathrm{CW}}(X; G)$ in our proof of Theorem 37.9. This proof will be the main topic of the next lecture.

38. Invariance of cellular homology, part 2 (November 24, 2023)

Let's quickly rephrase what we've done so far toward the proof of Theorem 37.9. We are assuming h_* is an axiomatic homology theory with coefficient group G. The latter means (via the additivity axiom) that $h_0(S^0)$ has a canonical isomorphism with $G \oplus G$, where the first factor corresponds to the point $1 \in S^0$ and the second to the other point $-1 \in S^0$. Applying Exercise 29.13 to h_* instead of singular homology gives

$$\widetilde{h}_0(S^0) = \{(g, -g) \in G \oplus G \mid g \in G\},$$

which we can identify with G via the injection $G \hookrightarrow G \oplus G : g \mapsto (g, -g)$, and $\widetilde{h}_k(S^0) = 0$ for all $k \ne 0$ due to the additivity and dimension axioms. Applying the suspension isomorphisms $S_* : \widetilde{h}_{k-1}(S^{n-1}) \xrightarrow{\cong} \widetilde{h}_k(S^n)$ repeatedly, we can then identify G with $\widetilde{h}_{n-1}(S^{n-1})$ for each $n \in \mathbb{N}$, and the connecting homomorphism in the reduced long exact sequence of $(\mathbb{D}^n, S^{n-1})$ then identifies G in turn with $h_n(\mathbb{D}^n, \partial\mathbb{D}^n)$, while simultaneously proving $h_k(\mathbb{D}^n, \partial\mathbb{D}^n) = 0$ for all $k \ne n$. Now if X is a CW-complex, the argument at the end of the previous lecture showed that

$$\bigoplus_{\alpha\in\mathcal{K}^n} (\Phi_\alpha)_* : \bigoplus_{\alpha\in\mathcal{K}^n} h_k(\mathbb{D}^n, \partial\mathbb{D}^n) \to h_k(X^n, X^{n-1}) \tag{38.1}$$

is an isomorphism for every $n \geqslant 1$ and $k \in \mathbb{Z}$, which proves $h_k(X^n, X^{n-1}) = 0$ for $k \ne n$ and identifies $h_n(X^n, X^{n-1})$ via our isomorphism $h_n(\mathbb{D}^n, \partial\mathbb{D}^n) = G$ with the cellular n-chain group $C_n^{\mathrm{CW}}(X) = C_n^{\mathrm{CW}}(X; G)$. The plan going forward is to use the Eilenberg-Steenrod axioms to construct a boundary map on $\bigoplus_{n\in\mathbb{Z}} h_n(X^n, X^{n-1})$ and prove that the homology of the resulting chain complex is isomorphic to $h_*(X)$. The last step will then be to show that our boundary map on $\bigoplus_{n\in\mathbb{Z}} h_n(X^n, X^{n-1})$ matches the cellular boundary map $\partial : C_*^{\mathrm{CW}}(X) \to C_{*-1}^{\mathrm{CW}}(X)$ under our identification.

Let us first derive some more consequences from the vanishing of $h_k(X^n, X^{n-1})$ for $k \neq n$. Observe that whenever either $k > n$ or $k < n-1$, the long exact sequence of (X^n, X^{n-1}) contains a segment of the form

$$0 = h_{k+1}(X^n, X^{n-1}) \to h_k(X^{n-1}) \to h_k(X^n) \to h_k(X^n, X^{n-1}) = 0, \tag{38.2}$$

implying that the inclusion $X^{n-1} \hookrightarrow X^n$ induces an isomorphism $h_k(X^{n-1}) \xrightarrow{\cong} h_k(X^n)$. This has two immediate consequences. For $k > n$, we can apply these isomorphisms repeatedly to decrease n to 0:

$$h_k(X^n) \cong h_k(X^{n-1}) \cong \ldots \cong h_k(X^0) \cong \bigoplus_{\alpha\in\mathcal{K}^0} h_k(\{\mathrm{pt}\}) = 0,$$

where at the last step we have applied the additivity and dimension axioms, using the fact that X^0 is a discrete space. This already proves a quite nontrivial fact that we did not yet know, though you may have expected it: for any homology theory, the homology groups of an n-dimensional CW-complex vanish in dimensions greater than n.

Lemma 38.1. *For every $k > n$, $h_k(X^n) = 0$.* □

Similarly, starting with $k < n$ and applying (38.2) repeatedly to increase n gives:

Lemma 38.2. *For every $k < n$, the inclusions $X^n \hookrightarrow X^{n+1} \hookrightarrow X^{n+2} \hookrightarrow \ldots$ induce isomorphisms $h_k(X^n) \cong h_k(X^{n+1}) \cong h_k(X^{n+2}) \cong \ldots$.* □

We can now proceed to the heart of the proof of Theorem 37.9. We define for each $n \geqslant 1$ a map

$$\beta_n : h_n(X^n, X^{n-1}) \to h_{n-1}(X^{n-1}, X^{n-2})$$

by combining the long exact sequences of the pairs (X^n, X^{n-1}) and (X^{n+1}, X^n) in the following diagram:

(38.3)

$$\begin{tikzcd}
0 \arrow[d, equal] & & & h_n(X^{n-1}) \arrow[d] \arrow[r, equal] & 0 & \\
h_{n+1}(X^n) \arrow[r] & h_{n+1}(X^{n+1}) \arrow[r, "j_{n+1}"] & h_{n+1}(X^{n+1}, X^n) \arrow[r, "\partial_{n+1}"] \arrow[dr, "\beta_{n+1}"'] & h_n(X^n) \arrow[r, "i_n"] \arrow[d, "j_n"] & h_n(X^{n+1}) \arrow[r] & h_n(X^{n+1}, X^n) \arrow[d, equal] \\
& & & h_n(X^n, X^{n-1}) \arrow[d, "\partial_n"] & & 0 \\
& & & h_{n-1}(X^{n-1}) & &
\end{tikzcd}$$

In other words, we define $\beta_{n+1} := j_n \circ \partial_{n+1}$ for each $n \geqslant 0$, and of course $\beta_0 := 0$. (We can use the convention $X^{-1} := \varnothing$ so that the diagram also makes sense in the case $n = 0$.) The relation $\beta_0 \circ \beta_1$ is then trivially true, while for every $n \geqslant 1$, we have

$$\beta_n \circ \beta_{n+1} = j_{n-1} \circ \partial_n \circ j_n \circ \partial_{n+1} = 0$$

since $\partial_n \circ j_n = 0$, thus we can now regard the sequence

$$\ldots \to h_n(X^n, X^{n-1}) \xrightarrow{\beta_n} h_{n-1}(X^{n-1}, X^{n-2}) \to \ldots\ldots \to h_1(X^1, X^0) \xrightarrow{\beta_1} h_0(X^0) \xrightarrow{\beta_0} 0 \to \ldots \tag{38.4}$$

as a chain complex whose individual chain groups are canonically isomorphic to the chain groups in $C_*^{\mathrm{CW}}(X)$. The exactness of the horizontal and vertical sequences in the diagram now give us the following observations: first, i_n is surjective, and thus descends to an isomorphism

$$h_n(X^n)/\ker i_n \xrightarrow[\cong]{i_n} h_n(X^{n+1}). \tag{38.5}$$

Second, j_{n-1} is injective, thus

$$\ker \beta_n = \ker(j_{n-1} \circ \partial_n) = \ker \partial_n = \operatorname{im} j_n,$$

and since j_n is also injective, it maps $h_n(X^n)$ isomorphically to $\ker \beta_n$. Moreover, it maps the subgroup $\ker i_n = \operatorname{im} \partial_{n+1}$ isomorphically to $\operatorname{im} \beta_{n+1}$, implying that j_n descends to an isomorphism

$$h_n(X^n)/\ker i_n \xrightarrow[\cong]{j_n} \ker \beta_n/\operatorname{im} \beta_{n+1}. \tag{38.6}$$

The latter is of course the nth homology group of the chain complex (38.4). Let us at this point make a simplifying assumption and suppose the CW-complex X is finite-dimensional: then there exists $N \in \mathbb{N}$ such that $X = X^N$. For any given integer $n \geqslant 0$ we can then take $N \geqslant n+1$ without loss of generality, and use Lemma 38.2 to conclude via (38.5) and (38.6) that

$$\ker \beta_n/\operatorname{im} \beta_{n+1} \cong h_n(X^{n+1}) \cong h_n(X^{n+2}) \cong \ldots \cong h_n(X^N) = h_n(X).$$

We will discuss in the next lecture how to lift the assumption $\dim X < \infty$, but if you are willing to accept this assumption for now, then the proof that $h_*(X) \cong H_*^{\mathrm{CW}}(X)$ will be complete as soon as we can show that the boundary maps β_n in (38.4) are the same as our usual cellular boundary maps. In other words, we need to prove that the diagram

$$\begin{tikzcd}
C_n^{\mathrm{CW}}(X) \arrow[r, "\cong"] \arrow[d, "\partial"] & h_n(X^n, X^{n-1}) \arrow[d, "\beta_n"] \\
C_{n-1}^{\mathrm{CW}}(X) \arrow[r, "\cong"] & h_{n-1}(X^{n-1}, X^{n-2})
\end{tikzcd}$$

commutes for every n, where the horizontal maps are the canonical isomorphisms that we discussed at the beginning of this lecture. The theorem that $\partial^2 = 0$ on $C_*^{\mathrm{CW}}(X)$ will also follow from this, since we already know $\beta_{n-1} \circ \beta_n = 0$.

Here is a useful observation: the characteristic maps $\Phi_\alpha : (\mathbb{D}^n, \partial\mathbb{D}^n) \to (X^n, X^{n-1})$ also induce maps of quotients $\mathbb{D}^n/\partial\mathbb{D}^n \to X^n/X^{n-1}$ such that the direct sum of the induced map on reduced homology

$$\bigoplus_{\alpha \in \mathcal{K}^n} (\Phi_\alpha)_* : \bigoplus_{\alpha \in \mathcal{K}^n} \widetilde{h}_n(\mathbb{D}^n/\partial\mathbb{D}^n) \to \widetilde{h}_n(X^n/X^{n-1}) \tag{38.7}$$

is an isomorphism. Indeed, under the natural isomorphisms between relative homology for good pairs and reduced homology of quotients, this is equivalent to the fact that (38.1) is an isomorphism. The advantage of rewriting this map in terms of quotients is, however, that we can explicitly write down its inverse. We recall the projections $p_\alpha : X^n \to X^n/(X^n \backslash e_\alpha^n) = \mathbb{D}^n/\partial\mathbb{D}^n$ that appear in the definition of the cellular boundary map, and notice that p_α sends X^{n-1} to the base point in $\mathbb{D}^n/\partial\mathbb{D}^n$ represented by points in the boundary, hence it descends to a map

$$p_\alpha : X^n/X^{n-1} \to \mathbb{D}^n/\partial\mathbb{D}^n.$$

Lemma 38.3. *The inverse of the map* (38.7) *is*

$$\prod_{\alpha \in \mathcal{K}^n} (p_\alpha)_* : \widetilde{h}_n(X^n/X^{n-1}) \to \bigoplus_{\alpha \in \mathcal{K}^n} \widetilde{h}_n(\mathbb{D}^n/\partial\mathbb{D}^n).$$

PROOF. Since we already know that (38.7) is an isomorphism, it will suffice to prove that $\prod_\beta (p_\beta)_* \circ \bigoplus_\alpha (\Phi_\alpha)_*$ is the identity map on $\bigoplus_\alpha \widetilde{h}_n(\mathbb{D}^n/\partial\mathbb{D}^n)$. This follows from the fact that $p_\alpha \circ \Phi_\alpha : \mathbb{D}^n/\partial\mathbb{D}^n \to \mathbb{D}^n/\partial\mathbb{D}^n$ is the identity map and thus induces the identity on $\widetilde{h}_n(\mathbb{D}^n/\partial\mathbb{D}^n)$, while for $\beta \neq \alpha$, $p_\beta \circ \Phi_\alpha$ is a constant map and thus factors through a one-point space, so the map it induces on $\widetilde{h}_n(\mathbb{D}^n/\partial\mathbb{D}^n)$ is trivial. $\square$

Here's a diagram to help us understand what β_n has to do with the cellular boundary map:

$$
\begin{tikzcd}
 & C_n^{\mathrm{CW}}(X) \arrow[d, equal] & C_{n-1}^{\mathrm{CW}}(X) \arrow[d, equal] \\
\bigoplus_{\alpha\in\mathcal{K}^n} h_n(\mathbb{D}^n, \partial\mathbb{D}^n) \arrow[r, "\bigoplus_\alpha \partial_*", "\cong"'] \arrow[d, "\cong"', "\bigoplus_\alpha (\Phi_\alpha)_*"] & \bigoplus_{\alpha\in\mathcal{K}^n} \widetilde{h}_{n-1}(\partial\mathbb{D}^n) \arrow[r, "\partial^{\mathrm{CW}}"] \arrow[d, "\bigoplus_\alpha (\varphi_\alpha)_*"] & \bigoplus_{\beta\in\mathcal{K}^{n-1}} \widetilde{h}_{n-1}(\mathbb{D}^{n-1}/\partial\mathbb{D}^{n-1}) \\
h_n(X^n, X^{n-1}) \arrow[r, "\partial_n"] \arrow[dr, "\beta_n"] & \widetilde{h}_{n-1}(X^{n-1}) \arrow[r, "q_*"] \arrow[d, "j_{n-1}"] & \widetilde{h}_{n-1}(X^{n-1}/X^{n-2}) \arrow[u, "\prod_\beta (p_\beta)_*", "\cong"'] \arrow[d, "\cong"] \\
 & h_{n-1}(X^{n-1}, X^{n-2}) \arrow[r, "\cong"] & h_{n-1}(X^{n-1}/X^{n-2}, X^{n-2}/X^{n-2})
\end{tikzcd}
\tag{38.8}
$$

The following details deserve clarification:

- The map labeled q_* is induced by the quotient projection $q : X^{n-1} \to X^{n-1}/X^{n-2}$.
- Regarding the same quotient projection as a map of pairs produces the horizontal map at the bottom, which we proved in Theorem 32.23 is an isomorphism. Composing the latter with (the inverse of) the lower right vertical isomorphism from the reduced long exact sequence of $(X^{n-1}/X^{n-2}, X^{n-2}/X^{n-2})$ produces the usual natural isomorphism $h_{n-1}(X^{n-1}, X^{n-2}) \xrightarrow{\cong} \widetilde{h}_{n-1}(X^{n-1}/X^{n-2})$.
- We have replaced $h_{n-1}(X^{n-1})$ with $\widetilde{h}_{n-1}(X^{n-1})$ for the middle term in the composition $\beta_n = j_{n-1} \circ \partial_n$, which is fine because the connecting homomorphism in the long exact sequence of a pair always has its image in redued homology anyway.
- The diagram is intended to serve as a **definition** of the map $\partial^{\mathrm{CW}} : C_n^{\mathrm{CW}}(X) \to C_{n-1}^{\mathrm{CW}}(X)$, i.e. it is what $\beta_n : h_n(X^n, X^{n-1}) \to h_{n-1}(X^{n-1}, X^{n-2})$ turns into after using canonical isomorphisms to replace its domain and target with cellular chain groups.

The point here is really just to replace the target group $h_{n-1}(X^{n-1}, X^{n-2})$ of β_n with $\widetilde{h}_{n-1}(X^{n-1}/X^{n-2})$ so that we can then Lemma 38.3 to identify the latter with $C_{n-1}^{\mathrm{CW}}(X)$ via an explicit formula. The resulting formula for ∂^{CW} is

$$\prod_{\beta\in\mathcal{K}^{n-1}} (p_\beta)_* \circ \bigoplus_{\alpha\in\mathcal{K}^n} (\varphi_\alpha)_* : \bigoplus_{\alpha\in\mathcal{K}^n} \widetilde{h}_{n-1}(S^{n-1}) \to \bigoplus_{\beta\in\mathcal{K}^{n-1}} \widetilde{h}_{n-1}(S^{n-1}).$$

This is determined by the collection of endomorphisms of $\widetilde{h}_{n-1}(S^{n-1})$ induced by $p_\beta \circ \varphi_\alpha$ for all $\alpha \in \mathcal{K}^n$ and $\beta \in \mathcal{K}^{n-1}$, and by Theorem 35.17, each of these maps is just multiplication by the degree of $p_\beta \circ \varphi_\alpha$, also known as the incidence number $[e_\beta^{n-1} : e_\alpha^n]$. This proves that ∂^{CW} is indeed simply the cellular boundary map ∂, and in particular, the latter satisfies $\partial^2 = 0$.

To complete the proof of Theorem 37.9 in the absolute case, we still need to understand how a cellular map $f : X \to Y$ between two CW-complexes interacts with the two isomorphisms $H_*^{\mathrm{CW}}(X) \cong h_*(X)$ and $H_*^{\mathrm{CW}}(Y) \cong h_*(Y)$. Being a cellular map implies that f defines a map of pairs $(X^n, X^{n-1}) \to (Y^n, Y^{n-1})$ for every n and thus induces homomorphisms from every term in the diagram (38.3) to the corresponding term in a similar diagram for Y. Something like this:

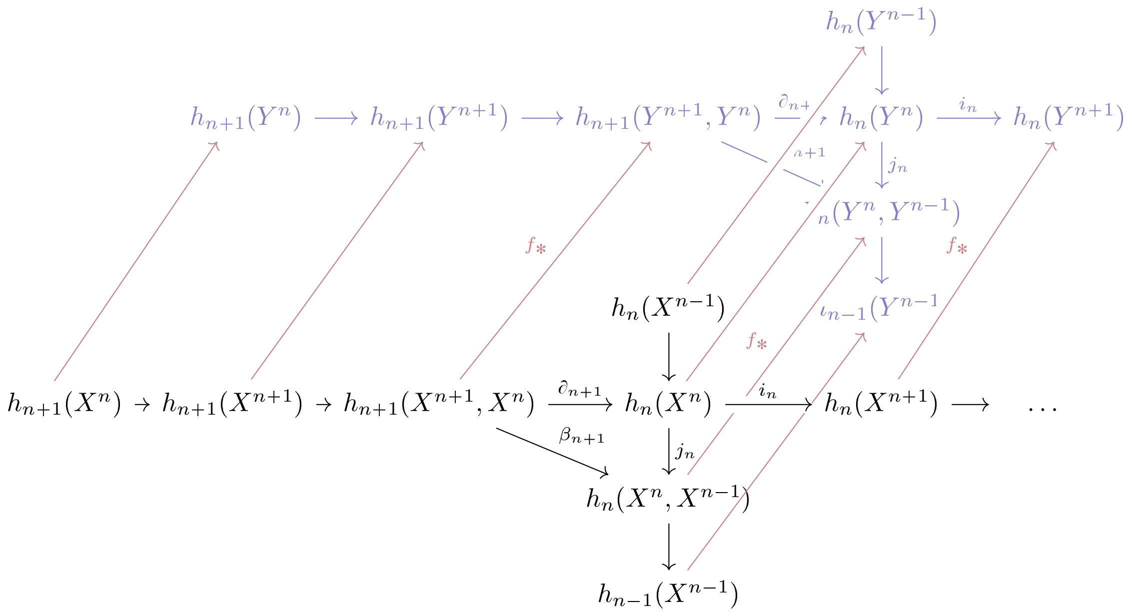

All of the red arrows in this three-dimensional diagram are maps induced by f, and the diagram commutes due to the naturality of long exact sequences. In particular, we now have

$$\begin{tikzcd} h_{n+1}(X^{n+1}, X^n) \arrow[r, "\beta_{n+1}"] \arrow[d, "f_*"] & h_n(X^n, X^{n-1}) \arrow[d, "f_*"] \\ h_{n+1}(Y^{n+1}, Y^n) \arrow[r, "\beta_{n+1}"] & h_n(Y^n, Y^{n-1}), \end{tikzcd}$$

so that f_* defines a chain map from the chain complex (38.4) to the corresponding chain complex for Y, and therefore determines a chain map $H_*^{\mathrm{CW}}(X) \to H_*^{\mathrm{CW}}(Y)$. To relate this to the map $f_* : h_*(X) \to h_*(Y)$, recall that the isomorphism $H_n^{\mathrm{CW}}(X) = \ker \beta_n / \operatorname{im} \beta_{n+1} \cong h_n(X)$ is defined in terms of the maps i_n and j_n in the diagram, along with the map induced by the inclusion $X^{n+1} \hookrightarrow X$, and all of these commute with f_*, thus we also obtain

$$\begin{tikzcd} H_n^{\mathrm{CW}}(X) \arrow[r, "\cong"] \arrow[d, "f_*"] & h_n(X) \arrow[d, "f_*"] \\ H_n^{\mathrm{CW}}(Y) \arrow[r, "\cong"] & h_n(Y). \end{tikzcd}$$

To finish, we just need to check that under the canonical identification of $h_n(X^n, X^{n-1})$ and $h_n(Y^n, Y^{n-1})$ with $C_n^{\mathrm{CW}}(X)$ and $C_n^{\mathrm{CW}}(Y)$ respectively, the map $f_* : h_n(X^n, X^{n-1}) \to h_n(Y^n, Y^{n-1})$ matches the formula we gave in (37.2) for maps $C_n^{\mathrm{CW}}(X) \to C_n^{\mathrm{CW}}(Y)$ induced by cellular maps. This will prove simultaneously the theorem that the homomorphism in (37.2) is a chain map. Here is the analogue of the diagram (38.8) for the situation at hand:

$$
\begin{tikzcd}
C_n^{\mathrm{CW}}(X) \arrow[d, equal] & C_n^{\mathrm{CW}}(Y) \arrow[d, equal] \\
\bigoplus_{e_\alpha^n \subset X} \widetilde{h}_n(\mathbb{D}^n/\partial\mathbb{D}^n) \arrow[r, "f_*^{\mathrm{CW}}"] \arrow[d, "\cong"', "\bigoplus_\alpha (\Phi_\alpha)_*"] & \bigoplus_{e_\beta^n \subset Y} \widetilde{h}_n(\mathbb{D}^n/\partial\mathbb{D}^n) \\
\widetilde{h}_n(X^n/X^{n-1}) \arrow[r, "f_*"] & \widetilde{h}_n(Y^n/Y^{n-1}) \arrow[u, "\prod_\beta (p_\beta)_*"', "\cong"]
\end{tikzcd}
\tag{38.9}
$$

The direct sums here are over the set of all n-cells e_α^n in X or e_β^n in Y, and the diagram is to be understood as a definition of the map $f_*^{\mathrm{CW}} : C_n^{\mathrm{CW}}(X) \to C_n^{\mathrm{CW}}(Y)$, which is equivalent to $f_* : h_n(X^n, X^{n-1}) \to h_n(Y^n, Y^{n-1})$ under the canonical isomorphisms. It produces the formula

$$f_*^{\mathrm{CW}} = \prod_{e_n^\beta \subset Y} (p_\beta)_* \circ f_* \circ \bigoplus_{e_\alpha^n \subset X} (\Phi_\alpha)_* : \bigoplus_{e_\alpha^n \subset X} \widetilde{h}_n(S^n) \to \bigoplus_{e_\beta^n \subset Y} \widetilde{h}_n(S^n),$$

and this map is determined by the set of all its "matrix elements"

$$(p_\beta)_* \circ f_* \circ (\Phi_\alpha)_* = (p_\beta \circ f \circ \Phi_\alpha)_* : \widetilde{h}_n(S^n) \to \widetilde{h}_n(S^n)$$

for each individual $e_\alpha^n \subset X$ and $e_\beta^n \subset Y$. Applying Theorem 35.17 again, this map is multiplication by $\deg(p_\beta \circ f \circ \Phi_\alpha) = [e_\beta^n : e_\alpha^n]$, thus f_*^{CW} does indeed match the formula given in (37.2) for $f_* : C_n^{\mathrm{CW}}(X) \to C_n^{\mathrm{CW}}(Y)$.

The proof of Theorem 37.9 is now complete except for three details, the first two of which will be left as exercises.

EXERCISE 38.4. Some portions of the discussion above do not make sense for $n = 0$, especially when $\mathbb{D}^n/\partial\mathbb{D}^n$ is mentioned. Adapt the discussion as needed for that particular case.

EXERCISE 38.5. Extend the entire discussion to the case of a CW-pair (X, A) with $A \neq \varnothing$. *Hint: Start by showing that $C_n^{\mathrm{CW}}(X, A; G)$ is canonically isomorphic to $h_n(X^n \cup A, X^{n-1} \cup A)$, and instead of the long exact sequence of the pair (X^n, X^{n-1}), consider the long exact sequence of the triple $(X^n \cup A, X^{n-1} \cup A, A)$.*
Comment: This exercise is a bit lengthy, but it is not fundamentally difficult—every step is simply a minor generalization of something that we discussed in this lecture. Working through it is one of the best ways to achieve a deeper understanding of the isomorphism $H_*^{\mathrm{CW}}(X) \cong h_*(X)$.

The third unresolved issue is the simplifying assumption $\dim X < \infty$ that we imposed in order to argue that $h_n(X^{n+1}) \cong h_n(X)$. We will discuss in the next lecture how to lift this assumption in the special case where h_* is singular homology.

EXERCISE 38.6. The **complex projective n-space** $\mathbb{CP}^n$ is a compact $2n$-manifold defined as the set of all complex lines through the origin in $\mathbb{C}^{n+1}$, or equivalently,

$$\mathbb{CP}^n = (\mathbb{C}^{n+1}\backslash\{0\})/\sim$$

where two points $z, z' \in \mathbb{C}^{n+1}\backslash\{0\}$ are equivalent if and only if $z' = \lambda z$ for some $\lambda \in \mathbb{C}$. It is conventional to write elements of $\mathbb{CP}^n$ in so-called *homogeneous coordinates*, meaning the equivalence class represented by $(z_0, \ldots, z_n) \in \mathbb{C}^{n+1}$ is written as $[z_0 : \ldots : z_n]$. Notice that $\mathbb{CP}^n$ can be partitioned into two disjoint subsets

$$\mathbb{C}^n \cong \{[1 : z_1 : \ldots : z_n] \in \mathbb{CP}^n\} \quad \text{and} \quad \mathbb{CP}^{n-1} \cong \{[0 : z_1 : \ldots : z_n] \in \mathbb{CP}^n\}.$$

(a) Show that the partition $\mathbb{CP}^n = \mathbb{C}^n \cup \mathbb{CP}^{n-1}$ gives rise to a cell decomposition of $\mathbb{CP}^n$ with one $2k$-cell for every $k = 0, \ldots, n$.

(b) Compute $H_*(\mathbb{CP}^n)$ and $H^*(\mathbb{CP}^n)$ for an arbitrary coefficient group.
Hint: This is easy.

39. Direct limits and infinite-dimensional cell complexes (November 28, 2023)

If X is an infinite-dimensional CW-complex, then the arguments of the previous lecture do not suffice to prove $H_*^{\mathrm{CW}}(X;G) \cong h_*(X)$ for every axiomatic homology theory h_* with coefficient group G. What they do prove is that for every integer $n \geqslant 0$, there are isomorphisms

$$H_n^{\mathrm{CW}}(X;G) \cong h_n(X^{n+1}) \cong h_n(X^{n+2}) \cong h_n(X^{n+3}) \cong \dots,$$

where the maps $h_n(X^{n+k}) \to h_n(X^{n+k+1})$ are induced by the inclusions $X^{n+k} \hookrightarrow X^{n+k+1}$, and moreover, these isomorphisms are natural in the sense that for any cellular map $f : X \to Y$, the induced homomorphism $f_* : H_n^{\mathrm{CW}}(X;G) \to H_n^{\mathrm{CW}}(Y;G)$ fits into a commutative diagram

$$\begin{tikzcd}
H_n^{\mathrm{CW}}(X;G) \arrow[r, "\cong"] \arrow[d, "f_*"] & h_n(X^{n+1}) \arrow[r, "\cong"] \arrow[d, "f_*"] & h_n(X^{n+2}) \arrow[r, "\cong"] \arrow[d, "f_*"] & h_n(X^{n+3}) \arrow[r, "\cong"] \arrow[d, "f_*"] & \dots \\
H_n^{\mathrm{CW}}(Y;G) \arrow[r, "\cong"] & h_n(Y^{n+1}) \arrow[r, "\cong"] & h_n(Y^{n+2}) \arrow[r, "\cong"] & h_n(Y^{n+3}) \arrow[r, "\cong"] & \dots
\end{tikzcd}$$

To get from here to a computation of $h_n(X)$, the idea is to interpret X as a "limit" of the sequence of spaces $X^0, X^1, X^2, \dots$, so that if the functor h_* can be shown to be "continuous" with respect to such limits, we would conclude

$$h_n(X) = h_n\big(\lim_{k\to\infty} X^k\big) = \lim_{k\to\infty} h_n(X^k),$$

and the value of this limit seems intuitively clear since all the groups in the sequence

$$h_n(X^{n+1}), h_n(X^{n+2}), h_n(X^{n+3}), \dots$$

are isomorphic to $H_n^{\mathrm{CW}}(X;G)$. To make all this precise, we need to explain in what sense a topological space X can be a "limit" of a sequence of spaces $\{X^n\}_{n=0}^\infty$, and similarly for a sequence of abelian groups such as $\{h_n(X^k)\}_{k=0}^\infty$.

Suppose I is a set with a pre-order $\prec$, i.e. $\prec$ is reflexive ($\alpha \prec \alpha$) and transitive ($\alpha \prec \beta$ and $\beta \prec \gamma$ implies $\alpha \prec \gamma$), but the relations $\alpha \prec \beta$ and $\beta \prec \alpha$ need not imply $\alpha = \beta$, so $\prec$ need not be a partial order. Recall that $(I, \prec)$ is called a **directed set** (*gerichtete Menge*) if for every pair $\alpha, \beta \in I$, there exists $\gamma \in I$ with $\gamma \succ \alpha$ and $\gamma \succ \beta$. The most common directed set in our examples will be $(\mathbb{N}, \leqslant)$, or sometimes $(\mathbb{N}_0, \leqslant)$ where $\mathbb{N}_0 := \{0\} \cup \mathbb{N}$. Some more interesting examples will arise when we discuss Poincaré duality and Čech (co-)homology later in this semester; see also Example 39.17 below.

In the following, we use the notation $X \xrightarrow{f} Y$ to indicate that f is a morphism from X to Y, where X and Y may be objects in an arbitrary category. In this way we can use commutative diagrams to encode relations between compositions of morphisms in any category—one should keep in mind however that the literal meaning of such a diagram may vary radically depending on the category we are working with.

Definition 39.1. Given a category $\mathscr{C}$, a **direct system** *(induktives System)* $\{X_\alpha, \varphi_{\beta\alpha}\}$ in $\mathscr{C}$ over $(I, \prec)$ associates to each $\alpha \in I$ an object X_α of $\mathscr{C}$, along with morphisms

$$\varphi_{\beta\alpha} \in \mathrm{Mor}(X_\alpha, X_\beta) \quad \text{for each} \quad \alpha \prec \beta$$

such that

$$\varphi_{\alpha\alpha} = \mathrm{Id}_{X_\alpha}$$

and the diagram

$$\begin{tikzcd} X_\alpha \arrow[r, "\varphi_{\beta\alpha}"] \arrow[rr, bend right, "\varphi_{\gamma\alpha}"'] & X_\beta \arrow[r, "\varphi_{\gamma\beta}"] & X_\gamma \end{tikzcd}$$

commutes for every triple $\alpha, \beta, \gamma \in I$ with $\alpha \prec \beta \prec \gamma$.

REMARK 39.2. Exercise 27.16 shows that a pre-order $\prec$ on a set I can be encoded by calling I the collection of objects in a category $\mathscr{I}$, such that for each pair $x, y \in I$, the set of morphisms $\mathrm{Mor}(x, y)$ contains exactly one element whenever $x \prec y$ and is otherwise empty. A direct system in $\mathscr{C}$ over $(I, \prec)$ is then nothing other than a (covariant) functor $\mathscr{I} \to \mathscr{C}$.

EXAMPLE 39.3. For any CW-complex X, its collection of skeleta $\{X^n\}_{n=0}^\infty$ forms a direct system in Top over $(\mathbb{N}_0, \leqslant)$, with the maps φ_{mn} for each $m \geqslant n$ defined as the inclusions $X^n \hookrightarrow X^m$. Similarly, the skeleta of a CW-pair define a direct system in $\mathsf{Top}_{\mathrm{rel}}$.

EXAMPLE 39.4. For any axiomatic homology theory h_*, the homology groups of the skeleta of a CW-complex from a direct system in $\mathsf{Ab}_{\mathbb{Z}}$ over $(\mathbb{N}_0, \leqslant)$: it consists of the graded abelian groups $\{h_*(X^n)\}_{n=0}^\infty$ and for each $m \geqslant n$ the map $h_*(X^n) \to h_*(X^m)$ induced by the inclusion $X^n \hookrightarrow X^m$. For each individual $k \in \mathbb{Z}$, we can also extract from this a direct system in Ab over $(\mathbb{N}_0, \leqslant)$, formed by the sequence of abelian groups $\{h_k(X^n)\}_{n=0}^\infty$.

The last example illustrates the following general observation, which is immediate from the definitions:

PROPOSITION 39.5. *If $\{X_\alpha, \varphi_{\beta\alpha}\}$ is a direct system in $\mathscr{A}$ over $(I, \prec)$, and $\mathcal{F} : \mathscr{A} \to \mathscr{B}$ is a covariant functor, then $\{\mathcal{F}(X_\alpha), \mathcal{F}(\varphi_{\beta\alpha})\}$ forms a direct system in $\mathscr{B}$ over $(I, \prec)$.* □

The notion of "convergence" for a direct system will necessarily look somewhat different from what we've seen before for sequences or nets: in most categories, there is no obvious topology or metric with which to measure how closely the objects X_α approach some limiting object X_∞ as $\alpha \in I$ becomes large. What we do have in every category is the notion of morphisms and the composition function $(f, g) \mapsto f \circ g$, so this is the structure that we will use. The idea is to measure the convergence of a direct system $\{X_\alpha, \varphi_{\beta\alpha}\}$ in terms of the morphisms from each X_α to other fixed objects in the category.

DEFINITION 39.6. For a direct system $\{X_\alpha, \varphi_{\beta\alpha}\}$ in $\mathscr{C}$ over $(I, \prec)$, a **target** $\{Y, f_\alpha\}$ of the system consists of an object Y of $\mathscr{C}$ together with associated morphisms $f_\alpha \in \mathrm{Mor}(X_\alpha, Y)$ for each $\alpha \in I$ such that the diagram

$$\begin{tikzcd} X_\alpha \arrow[rr, "\varphi_{\beta\alpha}"] \arrow[dr, "f_\alpha"'] & & X_\beta \arrow[dl, "f_\beta"] \\ & Y & \end{tikzcd}$$

commutes for every pair $\alpha, \beta \in I$ with $\alpha \prec \beta$.

DEFINITION 39.7. A target $\{X_\infty, \varphi_\alpha\}$ of the direct system $\{X_\alpha, \varphi_{\beta\alpha}\}$ is called a **direct limit**[60] *(induktiver Limes)* of the system and written as

$$X_\infty = \varinjlim \{X_\alpha\}$$

[60] Direct limits are also sometimes called **inductive limits**, and they are a special case of a more general notion in category theory called **colimits** (*Kolimes*); cf. Exercise 39.24.

if it satisfies the following "universal" property: for all targets $\{Y, f_\alpha\}$ of $\{X_\alpha, \varphi_{\beta\alpha}\}$, there exists a unique morphism $f_\infty \in \mathrm{Mor}(X_\infty, Y)$ such that the diagram

$$\begin{tikzcd} X_\alpha \arrow[r, "\varphi_\alpha"] \arrow[dr, "f_\alpha"'] & X_\infty \arrow[d, "f_\infty"] \\ & Y \end{tikzcd}$$

commutes for every $\alpha \in I$.

The essential meaning of a direct limit can be encoded in the diagram

$$\begin{tikzcd} X_\alpha \arrow[r, "\varphi_{\beta\alpha}"] \arrow[drrrr] & X_\beta \arrow[r, "\varphi_{\gamma\beta}"] \arrow[drrr] & X_\gamma \arrow[r] \arrow[drr] & \ldots \arrow[r] & \varinjlim\{X_\alpha\} \arrow[d, dashed] \\ & & & & Y \end{tikzcd}$$

where we assume $\alpha \prec \beta \prec \gamma \prec \ldots \in I$. The key feature of the object $\varinjlim\{X_\alpha\}$ is that whenever an object Y and morphisms $X_\alpha \to Y$ in a commuting diagram of this type are given, the "limit" morphism from $\varinjlim\{X_\alpha\}$ to Y indicated by the dashed arrow must also exist and be unique.

Note that from these definitions, there is generally no guarantee that a direct limit exists, and if it exists then it is generally not unique. Indeed:

EXERCISE 39.8. If $\{X, f_\alpha\}$ is a direct limit of $\{X_\alpha, \varphi_{\beta\alpha}\}$ and Y is another object such that there exists an isomorphism $\psi \in \mathrm{Mor}(X, Y)$, show that $\{Y, \psi \circ f_\alpha\}$ is also a direct limit of $\{X_\alpha, \varphi_{\beta\alpha}\}$. *Remark: The invertibility of ψ is needed only for showing that $\{Y, \psi \circ f_\alpha\}$ satisfies the universal property; it is already a target without this.*

The non-uniqueness exhibited by the exercise above is however the worst thing that can happen: if $\{X, f_\alpha\}$ and $\{Y, g_\alpha\}$ are any two direct limits of the same system $\{X_\alpha, \varphi_{\beta\alpha}\}$, then the universal property provides unique morphisms $g_\infty \in \mathrm{Mor}(X, Y)$ and $f_\infty \in \mathrm{Mor}(Y, X)$ satisfying $g_\infty \circ f_\alpha = g_\alpha$ and $f_\infty \circ g_\alpha = f_\alpha$ for every $\alpha \in I$. It follows that $f_\infty \circ g_\infty$ is the unique morphism from X to X satisfying $(f_\infty \circ g_\infty) \circ f_\alpha = f_\alpha$ for every $\alpha \in I$, which implies $f_\infty \circ g_\infty = \mathrm{Id}_X$. A similar argument shows $g_\infty \circ f_\infty = \mathrm{Id}_Y$, thus X and Y are isomorphic, and there is a distinguished isomorphism relating them. For this reason, we shall typically feel free to refer to "the" (rather than "a") direct limit of any system for which a limit exists.

The next exercise computes direct limits in a situation that is of concrete interest for the homology of a CW-complex X: recall from the previous lecture that for each $k \in \mathbb{Z}$, the sequence of homology groups $h_k(X^0) \to h_k(X^1) \to \ldots \to h_k(X^n) \to h_k(X^{n-1}) \to \ldots$ stabilizes as $n \to \infty$, i.e. the maps induced by the inclusions $X^n \hookrightarrow X^{n+1}$ all become isomorphisms as soon as n is sufficiently large. The intuition here is the same as in the elementary observation that for any sequence that is "eventually constant," its limit is what you think it should be.

EXERCISE 39.9. Suppose $\{X_\alpha, \varphi_{\beta\alpha}\}$ is a direct system in $\mathscr{C}$ over $(I, \prec)$ with the property that for some $\alpha_0 \in I$, $\varphi_{\gamma\beta} \in \mathrm{Mor}(X_\beta, X_\gamma)$ is an isomorphism for every $\beta, \gamma \in I$ with $\beta \succ \alpha_0$ and $\gamma \succ \alpha_0$. For each $\alpha \in I$, choose $\gamma \in I$ such that $\gamma \succ \alpha$ and $\gamma \succ \alpha_0$, and define

$$\varphi_\alpha := \varphi_{\gamma\alpha_0}^{-1} \circ \varphi_{\gamma\alpha} \in \mathrm{Mor}(X_\alpha, X_{\alpha_0}).$$

(a) Prove that the morphism φ_α does not depend on the choice of the element $\gamma \in I$.
(b) Prove that $\{X_{\alpha_0}, \varphi_\alpha\}$ is a target of the system.
(c) Prove that $\{X_{\alpha_0}, \varphi_\alpha\}$ also satisfies the universal property in Definition 39.7, hence $X_{\alpha_0} = \varinjlim\{X_\alpha\}$.

For the categories that we are most interested in, we will see presently that direct limits always exist and can be described in more concrete terms.

EXERCISE 39.10. Suppose $\{X_\alpha, \varphi_{\beta\alpha}\}$ is a direct system in $\mathscr{C}$ over $(I, \prec)$, where $\mathscr{C}$ is any category in which objects are sets (possibly with extra structure) and morphisms are maps between them. For any $\alpha, \beta \in I$, $x \in X_\alpha$ and $y \in X_\beta$, define the relation $x \sim y$ to mean

$$x \sim y \quad \Leftrightarrow \quad \varphi_{\gamma\alpha}(x) = \varphi_{\gamma\beta}(y) \text{ for some } \gamma \in I \text{ with } \gamma \succ \alpha \text{ and } \gamma \succ \beta.$$

Prove that $\sim$ is an equivalence relation on the set-theoretic disjoint union $\coprod_{\alpha \in I} X_\alpha$.[61]

PROPOSITION 39.11. *If $\{X_\alpha, \varphi_{\beta\alpha}\}$ is a direct system in* Top *over $(I, \prec)$, then its direct limit is the space*

$$\varinjlim\{X_\alpha\} = \coprod_{\alpha \in I} X_\alpha \Big/ \sim,$$

where the equivalence relation is defined as in Exercise 39.10, and the associated morphisms $\varphi_\alpha : X_\alpha \to \varinjlim\{X_\alpha\}$ are the compositions of the inclusions $X_\alpha \hookrightarrow \coprod_{\beta \in I} X_\beta$ with the quotient projection. Moreover, the topology on $\varinjlim\{X_\alpha\}$ is the strongest topology for which the maps $\varphi_\alpha : X_\alpha \to \varinjlim\{X_\alpha\}$ are continuous for all $\alpha \in I$.

PROOF. Abbreviate $X_\infty = \coprod_\alpha X_\alpha / \sim$. The topology of X_∞ is determined from that of the individual spaces X_α via the quotient and disjoint union topologies: concretely, this means that a set $\mathcal{U} \subset X_\infty$ is open if and only if its preimage $q^{-1}(\mathcal{U}) \subset \coprod_\beta X_\beta$ via the quotient projection $q : \coprod_\beta X_\beta \to X_\infty$ is open, and the latter is true if and only if $q^{-1}(\mathcal{U}) \cap X_\alpha$ is open in X_α for every $\alpha \in I$. Since $q^{-1}(\mathcal{U}) \cap X_\alpha = \varphi_\alpha^{-1}(\mathcal{U})$, this means that $\mathcal{U} \subset X_\infty$ is open if and only if every $\varphi_\alpha^{-1}(\mathcal{U}) \subset X_\alpha$ is open, thus characterizing the topology of X_∞ as the strongest for which every map $\varphi_\alpha : X_\alpha \to X_\infty$ is continuous. An easy corollary of this observation is that for any other space Y, a map $f : X_\infty \to Y$ is continuous if and only if the maps $f \circ \varphi_\alpha : X_\alpha \to Y$ are continuous for all $\alpha \in I$ (cf. Exercise 36.4).

It is clear that $\{X_\infty, \varphi_\alpha\}$ is a target since for any $\alpha, \beta \in I$ with $\alpha \prec \beta$, the relation

$$\varphi_\beta \circ \varphi_{\beta\alpha}(x) = \varphi_\alpha(x) \quad \text{for all} \quad x \in X_\alpha$$

follows from the fact that $x \sim \varphi_{\beta\alpha}(x)$. Now assuming $\{Y, f_\alpha\}$ is another target, we need to show that there is a unique continuous map $f_\infty : X_\infty \to Y$ satisfying the condition $f_\infty \circ \varphi_\alpha = f_\alpha$ for every $\alpha \in I$. To write down $f_\infty(x)$ for an arbitrary element $x \in X_\infty$, observe that since the quotient projection $q : \coprod_\beta X_\beta \to X_\infty$ is surjective, we have $x = q(x_\alpha) = \varphi_\alpha(x_\alpha)$ for some $\alpha \in I$ and $x_\alpha \in X_\alpha \subset \coprod_\beta X_\beta$, so in order to achieve $f_\infty \circ \varphi_\alpha = f_\alpha$, we are forced to define

$$f_\infty(x) := f_\alpha(x_\alpha).$$

We claim that $f_\infty(x)$ is then independent of the choice of element $x_\alpha \in q^{-1}(x)$. Indeed, suppose $\beta \in I$ and $x_\beta \in X_\beta \subset \coprod_\gamma X_\gamma$ such that $\varphi_\beta(x_\beta) = q(x_\beta) = x$. The equivalence $x_\alpha \sim x_\beta$ then means that for some $\gamma \in I$ satisfying $\gamma \succ \alpha$ and $\gamma \succ \beta$,

$$\varphi_{\gamma\alpha}(x_\alpha) = \varphi_{\gamma\beta}(x_\beta) =: x_\gamma \in X_\gamma,$$

and thus $f_\gamma(x_\gamma) = f_\alpha(x_\alpha) = f_\beta(x_\beta)$. This proves that a map $f_\infty : X_\infty \to Y$ with the desired properties is well defined and uniquely determined, though a remark is still required on why f_∞ is

[61]The **set-theoretic disjoint union** of a collection of sets $\{X_\alpha\}_{\alpha \in I}$ can be defined in general as the set $\{(\alpha, x) \mid \alpha \in I,\ x \in X_\alpha\}$, i.e. it is a union of all the sets X_α, but defined such that even if some pair of the sets X_α and X_β for $\alpha \neq \beta$ have elements in common, they are each identified with *disjoint* subsets of $\coprod_\gamma X_\gamma$. The disjoint union of topological spaces is defined in the same way, but with the extra structure of a topology, which for the purposes of Exercise 39.10 is not needed.

continuous: this follows from the previous paragraph since $f_\infty \circ \varphi_\alpha = f_\alpha : X_\alpha \to Y$ is continuous for every $\alpha \in I$. $\square$

REMARK 39.12. Proposition 39.11 extends in an obvious way to give a concrete description of any direct limit in the category $\mathsf{Top}_{\mathrm{rel}}$ of pairs of spaces.

Consider the specific direct system of topological spaces $\{X^n\}_{n=0}^\infty$ from Example 39.3, consisting of the skeleta of a CW-complex X with maps $X^m \hookrightarrow X^n$ for $n \geqslant m$ defined by inclusion. Considering the quotient $X^\infty := \coprod_{n=0}^\infty X^n \big/ \sim$ as in Proposition 39.11 along with the natural maps $\varphi_n : X^n \to X^\infty$, the disjoint union of the inclusion maps $i_n : X^n \hookrightarrow X$ descends to the quotient as a bijection

$$\coprod_{n=0}^\infty i_n : \coprod_{n=0}^\infty X^n \Big/ \sim \xrightarrow{\cong} X$$

which identifies φ_n with the inclusion i_n for each n. Since the topology of both X^∞ and X is the strongest for which the maps φ_n or i_n respectively are all continuous, this bijection is a homeomorphism, and we've proved:

COROLLARY 39.13. *For the direct system of Example 39.3 formed by the skeleta of a CW-complex X,*

$$\varinjlim \{X^n\} = X,$$

with the associated morphisms $X^n \to \varinjlim\{X^n\}$ defined as the inclusions $X^n \hookrightarrow X$. $\square$

We next consider the analogue of Proposition 39.11 in the category Ab of abelian groups.

PROPOSITION 39.14. *If $\{G_\alpha, \varphi_{\beta\alpha}\}$ is a direct system in Ab over $(I, \prec)$, then its direct limit is the group*

$$\varinjlim\{G_\alpha\} = \bigoplus_{\alpha \in I} G_\alpha \Big/ H,$$

were $H \subset \bigoplus_\alpha G_\alpha$ is the subgroup generated by all elements of the form $g - \varphi_{\beta\alpha}(g)$ for $g \in G_\alpha$ and $\beta \succ \alpha$, and the associated homomorphisms $\varphi_\alpha : G_\alpha \to \varinjlim\{G_\alpha\}$ are the compositions of the natural inclusions $G_\alpha \hookrightarrow \bigoplus_\beta G_\beta$ with the quotient projection.

PROOF. Abbreviating $G_\infty = \bigoplus_\alpha G_\alpha \big/ H$, it is easy to see that $\{G_\infty, \varphi_\alpha\}$ is a target. Given another target $\{A, \psi_\alpha\}$, the condition $\psi_\beta \circ \varphi_{\beta\alpha} = \psi_\alpha$ for each $\beta \succ \alpha$ implies that the homomorphism

$$\bigoplus_{\alpha \in I} \psi_\alpha : \bigoplus_{\alpha \in I} G_\alpha \to A$$

vanishes on the subgroup H and thus descends to a homomorphism $\psi_\infty : G_\infty \to A$ that satisfies $\psi_\infty \circ \varphi_\alpha = \psi_\alpha$ for all α. $\square$

EXERCISE 39.15. Prove the obvious analogues of Proposition 39.14 for direct systems in the categories $\mathsf{Ab}_\mathbb{Z}$ of $\mathbb{Z}$-graded abelian groups and Chain of chain complexes.

The following consequence of Proposition 39.14 makes proving things about direct limits of abelian groups (or the other algebraic categories mentioned in the exercise above) considerably easier.

COROLLARY 39.16. *The following statements hold for any direct system $\{G_\alpha, \varphi_{\beta\alpha}\}$ in Ab, $\mathsf{Ab}_\mathbb{Z}$ or Chain over a directed set $(I, \prec)$:*

(i) For every $x \in \varinjlim\{G_\alpha\}$, there exists $\beta \in I$ and $x_\beta \in G_\beta$ such that $x = \varphi_\beta(x_\beta)$.

(ii) For every $\beta \in I$ and $x_\beta \in G_\beta$ satisfying $\varphi_\beta(x_\beta) = 0 \in \varinjlim\{G_\alpha\}$, there exists $\gamma \succ \beta$ such that $\varphi_{\gamma\beta}(x_\beta) = 0 \in G_\gamma$.

PROOF. Writing $\varinjlim\{G_\alpha\} = \bigoplus_\alpha G_\alpha \big/ H$, any given element $x \in \varinjlim\{G_\alpha\}$ is an equivalence class represented by an element

$$\sum_{\alpha \in I_0} g_\alpha \in \bigoplus_{\alpha \in I} G_\alpha$$

for some finite subset $I_0 \subset I$. Since $(I, \prec)$ is a directed set, we can then find an element $\beta \in I$ satisfying $\beta \succ \alpha$ for every $\alpha \in I_0$, so

$$\sum_{\alpha \in I_0} g_\alpha - \sum_{\alpha \in I_0} \varphi_{\beta\alpha}(g_\alpha) \in H,$$

implying that $x_\beta := \sum_{\alpha \in I_0} \varphi_{\beta\alpha}(g_\alpha) \in G_\beta$ satisfies $\varphi_\beta(x_\beta) = x$.

For the second statement, we observe that $\varphi_\beta(x_\beta) = 0$ holds if and only if $x_\beta \in G_\beta \subset \bigoplus_\alpha G_\alpha$ belongs to the subgroup H, meaning

$$x_\beta = \sum_{i=1}^{N} (g_i - \varphi_{\beta_i \alpha_i}(g_i)) \tag{39.1}$$

for some finite collection of elements $\beta_i \succ \alpha_i \in I$ and $g_i \in G_{\alpha_i}$, $i = 1, \ldots, N$. Choose a finite subset $I_0 \subset I$ that contains all the α_i, β_i for $i = 1, \ldots, N$, along with an element $\gamma \in I$ such that $\gamma \succ \alpha$ for all $\alpha \in I_0$. Applying the homomorphism $\bigoplus_{\alpha \in I_0} \varphi_{\gamma\alpha}$ to both sides of (39.1) then produces $\varphi_{\gamma\beta}(x_\beta) \in G_\gamma$ on the left hand side and kills the right hand side since for each i,

$$\left(\bigoplus_{\alpha \in I_0} \varphi_{\gamma\alpha} \right) (g_i - \varphi_{\beta_i \alpha_i}(g_i)) = \varphi_{\gamma\alpha_i}(g_i) - \varphi_{\gamma\beta_i} \circ \varphi_{\beta_i\alpha_i}(g_i) = \varphi_{\gamma\alpha_i}(g_i) - \varphi_{\gamma\alpha_i}(g_i) = 0.$$

We have thus proved $\varphi_{\gamma\beta}(x_\beta) = 0$. □

We have seen above that any CW-complex X can be identified with the direct limit of its skeleta. Combining Exercise 39.9 with the computations of the previous lecture proves moreover that for any axiomatic homology theory h_* and any $k \in \mathbb{Z}$, the direct system of abelian groups $\{h_k(X^n)\}_{n=0}^\infty$ stabilizes as $n \to \infty$ and thus has direct limit $h_k(X^n)$ for any n sufficiently large, which matches $H_k^{\mathrm{CW}}(X^n) = H_k^{\mathrm{CW}}(X)$. This gives an isomorphism of $\mathbb{Z}$-graded abelian groups

$$H_*^{\mathrm{CW}}(X) \cong \varinjlim\{h_*(X^n)\}.$$

The isomorphism $H_*^{\mathrm{CW}}(X) \cong h_*(X)$ will therefore follow if we can prove that the functor h_* behaves "continuously" under this direct limit, i.e. the question becomes

$$\varinjlim\{h_*(X^n)\} \cong h_*\big(\varinjlim\{X^n\}\big)?$$

It is time to insert a word of caution: the next exercise shows that singular homology does *not* always behave as nicely as one would hope with respect to direct limits.

EXERCISE 39.17. Define $\{X_\alpha\}_{\alpha \in I}$ to be the collection of all countable subspaces of S^1, with a partial order assigned to the index set such that

$$\alpha \prec \beta \qquad \Leftrightarrow \qquad X_\alpha \subset X_\beta.$$

In this case we can define $\varphi_{\beta\alpha} : X_\alpha \hookrightarrow X_\beta$ to be the inclusion map and regard $\{X_\alpha, \varphi_{\beta\alpha}\}$ as a direct system of topological spaces. Prove that $\varinjlim\{X_\alpha\}$ is homeomorphic to S^1, but $\varinjlim\{H_*(X_\alpha; \mathbb{Z})\}$ is not isomorphic to $H_*(S^1; \mathbb{Z})$.

Hint 1: Describing $\varinjlim\{X_\alpha\}$ as in Proposition 39.11, it is not hard to find a natural bijection between $\varinjlim\{X_\alpha\}$ and $\bigcup_{\alpha \in I} X_\alpha = S^1$, but you need to check that the topology of this direct limit matches the standard topology of S^1.

Hint 2: What can you say about $H_1(X_\alpha; \mathbb{Z})$ for each α?

To see nonetheless why it might *sometimes* be true that $\varinjlim\{h_*(X_\alpha)\} \cong h_*\big(\varinjlim\{X_\alpha\}\big)$, let us observe first that there is always a natural morphism between these two objects. Indeed, suppose more generally that $\{X_\alpha, \varphi_{\beta\alpha}\}$ is a direct system in some category $\mathscr{A}$ over $(I, \prec)$, and $\mathcal{F} : \mathscr{A} \to \mathscr{B}$ is a covariant functor, thus producing a direct system $\{\mathcal{F}(X_\alpha), \mathcal{F}(\varphi_{\beta\alpha})\}$ in $\mathscr{B}$. If $\varinjlim\{X_\alpha\}$ exists, then the natural morphisms

$$X_\alpha \xrightarrow{\varphi_\alpha} \varinjlim\{X_\alpha\}$$

for every $\alpha \in I$ induce morphisms

$$\mathcal{F}(X_\alpha) \xrightarrow{\Phi_\alpha := \mathcal{F}(\varphi_\alpha)} \mathcal{F}(\varinjlim\{X_\alpha\})$$

which satisfy

$$\Phi_\beta \circ \mathcal{F}(\varphi_{\beta\alpha}) = \mathcal{F}(\varphi_\beta) \circ \mathcal{F}(\varphi_{\beta\alpha}) = \mathcal{F}(\varphi_\beta \circ \varphi_{\beta\alpha}) = \mathcal{F}(\varphi_\alpha) = \Phi_\alpha$$

for all $\beta \succ \alpha$ and thus make $\{\mathcal{F}(\varinjlim\{X_\alpha\}), \Phi_\alpha\}$ a target of the system $\{\mathcal{F}(X_\alpha), \mathcal{F}(\varphi_{\beta\alpha})\}$. If we assume that $\varinjlim\{\mathcal{F}(X_\alpha)\}$ also exists, then it now follows via the universal property of the direct limit that there is a limiting morphism

$$\varinjlim\{\mathcal{F}(X_\alpha)\} \xrightarrow{\Phi_\infty} \mathcal{F}(\varinjlim\{X_\alpha\}). \tag{39.2}$$

We would like to identify some situations in which Φ_∞ is guaranteed to be an isomorphism. In particular, we shall prove that this is true when $\{X_\alpha, \varphi_{\beta\alpha}\}$ is the direct system in Top formed by the skeleta of a CW-complex and $\mathcal{F}$ is the singular homology functor.

Recall that $H_* : \mathsf{Top} \to \mathsf{Ab}_{\mathbb{Z}}$ is in fact the composition of two functors: the first is $C_* : \mathsf{Top} \to \mathsf{Chain}$, which sends each space X to its singular chain complex with coefficients in a given group G, and the second is $H_* : \mathsf{Chain} \to \mathsf{Ab}_{\mathbb{Z}}$, sending a chain complex to its homology. The second of these two functors turns out to be extremely well behaved with respect to direct limits.

Proposition 39.18. *Suppose $(I, \prec)$ is a directed set, with a chain complex C_*^α associated to each $\alpha \in I$ and a chain map $\varphi_{\beta\alpha} : C_*^\alpha \to C_*^\beta$ associated to each pair $\alpha \prec \beta \in I$ such that $\{C_*^\alpha, \varphi_{\beta\alpha}\}$ is a direct system in* Chain *over $(I, \prec)$. Then choosing $\mathcal{F}$ to be the functor $H_* : \mathsf{Chain} \to \mathsf{Ab}_{\mathbb{Z}}$, the map*

$$\Phi_\infty : \varinjlim\{H_*(C_*^\alpha)\} \to H_*\big(\varinjlim\{C_*^\alpha\}\big)$$

defined as in (39.2) is an isomorphism of $\mathbb{Z}$-graded abelian groups.

Proof. We prove first that Φ_∞ is surjective. Given a homology class $[c] \in H_*\big(\varinjlim\{C_*^\alpha\}\big)$ represented by a cycle $c \in \varinjlim\{C_*^\alpha\}$, Corollary 39.16 implies $c = \varphi_\beta(c_\beta)$ for some $\beta \in I$ and $c_\beta \in C_*^\beta$, where $\varphi_\beta : C_*^\beta \to \varinjlim\{C_*^\alpha\}$ denotes the natural morphism associated to the direct limit. Since $\partial c = 0$ and φ_β is a chain map, we have $\varphi_\beta(\partial c_\beta) = 0$, so by Corollary 39.16, we can find some $\gamma \succ \beta$ and replace c_β with $c_\gamma := \varphi_{\gamma\beta}(c_\beta) \in C_*^\gamma$ such that $\varphi_\gamma(c_\gamma) = \varphi_\gamma \circ \varphi_{\gamma\beta}(c_\beta) = \varphi_\beta(c_\beta) = c$ but also $\partial c_\gamma = 0$, and c_γ thus represents a homology class $[c_\gamma] \in H_*(C_*^\gamma)$. Now let

$$\Psi_\gamma : H_*(C_*^\gamma) \to \varinjlim\{H_*(C_*^\alpha)\}$$

denote the natural morphism associated to the direct limit of the system $\{H_*(C_*^\alpha), \Phi_{\gamma\alpha}\}$, where $\Phi_{\gamma\alpha} := (\varphi_{\gamma\alpha})_* : H_*(C_*^\alpha) \to H_*(C_*^\gamma)$ for $\gamma \succ \alpha$. Writing $\Phi_\gamma := (\varphi_\gamma)_* : H_*(C_*^\gamma) \to H_*\big(\varinjlim\{C_*^\alpha\}\big)$, the diagram

$$\begin{tikzcd} H_*(C_*^\gamma) \arrow[r, "\Psi_\gamma"] \arrow[dr, "\Phi_\gamma"'] & \varinjlim\{H_*(C_*^\alpha)\} \arrow[d, "\Phi_\infty"] \\ & H_*\big(\varinjlim\{C_*^\alpha\}\big) \end{tikzcd}$$

commutes by the definition of Φ_∞, thus $\Phi_\infty(\Psi_\gamma[c_\gamma]) = \Phi_\gamma[c_\gamma] = [\varphi_\gamma(c_\gamma)] = [c]$, proving that Φ_∞ is surjective.

The proof of injectivity uses all the same ideas, so we shall leave it as an exercise. $\square$

An essential role in the proof above was played by Corollary 39.16, which is a tool for replacing statements about direct limits with corresponding statements about individual objects in the direct system. We saw in Exercise 39.17 that the singular homology functor $H_* : \mathsf{Top} \to \mathsf{Ab}_{\mathbb{Z}}$ is not always continuous with respect to direct limits; since this is the composition of two functors $H_* : \mathsf{Chain} \to \mathsf{Ab}_{\mathbb{Z}}$ and $C_* : \mathsf{Top} \to \mathsf{Chain}$, Proposition 39.18 implies that something must go wrong in general with the continuity of $C_* : \mathsf{Top} \to \mathsf{Chain}$. The following result therefore contains an extra hypothesis that is not satisfied in pathological examples such as Exercise 39.17, but certainly is satisfied (due to Proposition 37.1) by the direct system formed by the skeleta of any CW-complex. The key point is that since every singular n-simplex $\sigma : \Delta^n \to \varinjlim\{X_\alpha\}$ has image contained in a compact set, this extra hypothesis will allow us to write it as $\varphi_\beta \circ \sigma_\beta$ for some $\beta \in I$ and a singular n-simplex $\sigma_\beta : \Delta^n \to X_\beta$, producing an analogue of Corollary 39.16 for the situation at hand.

Proposition 39.19. *Suppose $\{X_\alpha, \varphi_{\beta\alpha}\}$ is a direct system of topological spaces over $(I, \prec)$ satisfying the following conditions:*

(1) *For every $\alpha \in I$, X_α is a subspace of $X := \varinjlim\{X_\alpha\}$ and the maps $\varphi_{\beta\alpha} : X_\alpha \to X_\beta$ and $\varphi_\alpha : X_\alpha \to X$ are the natural inclusions;*
(2) *Every compact subset $K \subset X$ is contained in X_α for some $\alpha \in I$.*

Then choosing $\mathcal{F}$ to be the singular chain complex functor $C_* : \mathsf{Top} \to \mathsf{Chain}$ with an arbitrary coefficient group G, the chain map

$$\Phi_\infty : \varinjlim\{C_*(X_\alpha)\} \to C_*\big(\varinjlim\{X_\alpha\}\big)$$

defined as in (39.2) is an isomorphism of chain complexes.

Proof. For surjectivity, given $c = \sum_i g_i\sigma_i \in C_n\big(\varinjlim\{X_\alpha\}\big)$, the finitely many singular n-simplices $\sigma_i : \Delta^n \to \varinjlim\{X_\alpha\}$ can each be written as $\sigma_i = \varphi_{\alpha_i} \circ \sigma_i'$ for some $\alpha_i \in I$ and $\sigma_i' : \Delta^n \to X_{\alpha_i}$ since Δ^n is compact. We can then find $\beta \in I$ with $\beta \succ \alpha_i$ for all i and define $\sigma_i'' := \varphi_{\beta\alpha_i} \circ \sigma_i' : \Delta^n \to X_\beta$, so

$$\sigma_i = \varphi_{\alpha_i} \circ \sigma_i' = \varphi_\beta \circ \varphi_{\beta\alpha_i} \circ \sigma_i' = \varphi_\beta \circ \sigma_i'',$$

producing an element $c_\beta := \sum_i g_i\sigma_i'' \in C_n(X_\beta)$ such that $(\varphi_\beta)_* c_\beta = c$. Writing $\Psi_\beta : C_*(X_\beta) \to \varinjlim\{C_*(X_\alpha)\}$ for the natural map associated to the direct limit, the diagram

$$\begin{tikzcd} C_*(X_\beta) \arrow[r, "\Psi_\beta"] \arrow[dr, "(\varphi_\beta)_*"'] & \varinjlim\{C_*(X_\alpha)\} \arrow[d, "\Phi_\infty"] \\ & C_*\big(\varinjlim\{X_\alpha\}\big) \end{tikzcd}$$

commutes by the definition of Φ_∞, and thus gives $\Phi_\infty(\Psi_\beta(c_\beta)) = c$.

Injectivity is again proved by similar arguments, which we shall leave as an exercise. $\square$

Applying Propositions 39.18 and 39.19 together, we've proved:

Theorem 39.20. *Under the same hypotheses as in Proposition 39.19, there is a natural isomorphism of $\mathbb{Z}$-graded abelian groups*

$$\varinjlim\{H_*(X_\alpha)\} \xrightarrow{\cong} H_*\big(\varinjlim\{X_\alpha\}\big)$$

for every choice of coefficient group. $\square$

COROLLARY 39.21. *For any CW-complex X and abelian group G, there is a natural isomorphism $H_*^{\mathrm{CW}}(X;G) \cong H_*(X;G)$.* □

It is a straightforward matter to extend this entire discussion to the case of a CW-pair (X,A) and prove $H_*^{\mathrm{CW}}(X,A;G) \cong H_*(X,A;G)$, and also to prove that these isomorphisms commute with the maps induced by any map of CW-pairs $f:(X,A) \to (Y,B)$. We shall leave the further details of these extensions as exercises, and thus regard the proof of Theorem 37.9 (at least for singular homology) as complete.

Corollary 39.21 is also true if singular homology H_* is replaced by an arbitrary axiomatic homology theory h_*, but proving this would take us into somewhat more abstract territory than we have time for right now, so we will settle for the special case of singular homology. The original treatment of the homology axioms by Eilenberg and Steenrod [ES52] dealt mainly with finite complexes, for which our proof of the isomorphism $H_*^{\mathrm{CW}}(X) \cong h_*(X)$ was already completed in the previous lecture. The extension of this result to infinite-dimensional complexes was accomplished originally by Milnor in [Mil62], who introduced the additivity axiom for this purpose.[62] Milnor's proof via the "mapping telescope" construction is reproduced in [Hat02, pp. 138–139].

EXERCISE 39.22. Each of the following spaces can be defined as a direct limit in terms of the natural inclusions $\mathbb{F}^m \hookrightarrow \mathbb{F}^n$ for $n \geqslant m$, where $\mathbb{F}$ is $\mathbb{R}$ or $\mathbb{C}$, and we identify $\mathbb{F}^m$ with the subspace $\mathbb{F}^m \oplus \{0\} \subset \mathbb{F}^n$. In particular, $\mathbb{R}^{m+1} \hookrightarrow \mathbb{R}^{n+1}$ gives rise to inclusions $S^m \hookrightarrow S^n$ and $\mathbb{RP}^m \hookrightarrow \mathbb{RP}^n$, and the complex version gives $\mathbb{CP}^m \hookrightarrow \mathbb{CP}^n$. Use cell decompositions to compute the homology with integer coefficients for each space:

(a) $S^\infty = \varinjlim \{S^n\}_{n\in\mathbb{N}}$
(b) $\mathbb{RP}^\infty = \varinjlim \{\mathbb{RP}^n\}_{n\in\mathbb{N}}$
(c) $\mathbb{CP}^\infty = \varinjlim \{\mathbb{CP}^n\}_{n\in\mathbb{N}}$

EXERCISE 39.23. Suppose $\{X_\alpha, \varphi_{\beta\alpha}\}$ is a direct system of topological spaces such that each X_α is a subspace of some fixed topological space X, $\beta > \alpha$ if and only if $X_\alpha \subset X_\beta$, and the maps $\varphi_{\beta\alpha} : X_\alpha \to X_\beta$ in this case are the natural inclusions. Let us use Proposition 39.11 to identify $\varinjlim\{X_\alpha\}$ with $\coprod_\alpha X_\alpha \big/ \sim$, in terms of the equivalence relation

$$X_\alpha \ni x \sim y \in X_\beta \qquad \Leftrightarrow \qquad \varphi_{\gamma\alpha}(x) = \varphi_{\gamma\beta}(y) \text{ for some } \gamma \in I \text{ with } \gamma > \alpha,\ \gamma > \beta.$$

The disjoint union of the inclusions $X_\alpha \hookrightarrow \bigcup_{\beta\in I} X_\beta$ then descends to the quotient as a bijection

$$\varinjlim\{X_\alpha\} \to \bigcup_{\alpha\in I} X_\alpha,$$

and we have seen examples where it is a homeomorphism: this is true in particular for the direct system consisting of the skeleta of a CW-complex. The following example shows however that it need not be a homeomorphism in general: let $I = (0,1)$ and consider the family of sets $X_t = \{0\} \cup (t,1] \subset \mathbb{R}$ for $t \in I$, ordered by inclusion. The union of these sets is $[0,1]$, but show that the topological space $\varinjlim\{X_t\}$ is not connected.

EXERCISE 39.24. Direct limits are a special case of a more general notion in category theory called *colimits*. In order to express the definition, recall (cf. Remark 39.2) that every pre-ordered set $(I,\prec)$ can be encoded as a category $\mathscr{I}$ in which relations $\alpha \prec \beta$ are viewed as morphisms $\alpha \to \beta$, and from this perspective, a direct system $\{X_\alpha, \varphi_{\beta\alpha}\}$ over $(I,\prec)$ in the category $\mathscr{C}$ is the same thing as a (covariant) functor $\mathcal{F} : \mathscr{I} \to \mathscr{C}$. To define targets and direct limits in this language, one can identify each object Y of $\mathscr{C}$ with the "constant" functor $\mathcal{Y} : \mathscr{I} \to \mathscr{C}$ that sends

[62] Recall that for finite disjoint unions, the additivity axiom follows from the other axioms, and one does not need any infinite disjoint unions to compute the homology of a finite CW-complex.

every $\alpha \in I$ to Y and every morphism $\alpha \to \beta$ of $\mathscr{I}$ to the identity morphism $Y \to Y$. Note that if $\mathcal{X}, \mathcal{Y} : \mathscr{I} \to \mathscr{C}$ are two such contant functors associated to objects X, Y respectively in $\mathscr{C}$, then a natural transformation from $\mathcal{X}$ to $\mathcal{Y}$ must associate to every $\alpha \in I$ the same morphism $X \to Y$, and conversely, every morphism $X \to Y$ determines a natural transformation $\mathcal{X} \to \mathcal{Y}$. A target $\{Y, f_\alpha\}$ of the system $\{X_\alpha, \varphi_{\beta\alpha}\}$ is now the same thing as a natural transformation $T_Y : \mathcal{F} \to \mathcal{Y}$, assigning to each object α of $\mathscr{I}$ the morphism $T_Y(\alpha) := f_\alpha : X_\alpha \to Y$, and in this language, a target $T_X : \mathcal{F} \to \mathcal{X}$ is called universal (and is thus a direct limit of the system) if for every target $T_Y : \mathcal{F} \to \mathcal{Y}$, there is a unique natural transformation $\Phi : \mathcal{X} \to \mathcal{Y}$ such that $T_Y = \Phi \circ T_X$.

Having expressed the definition of a direct limit in this form, the whole discussion still makes sense if one replaces the category $\mathscr{I}$ associated with the directed set $(I, \prec)$ by an arbitrary[63] category $\mathscr{A}$. For reasons that may become clearer when you look at the examples below, a functor $\mathcal{F} : \mathscr{A} \to \mathscr{C}$ is then often referred to as a *diagram* in $\mathscr{C}$ over $\mathscr{A}$. A target is again simply a natural transformation $T_X : \mathcal{F} \to \mathcal{X}$ to the constant functor $\mathcal{X} : \mathscr{A} \to \mathscr{C}$ determined by some object X of $\mathscr{C}$, and it is called a **colimit** of the diagram if it satisfies the universal property described above. In this case one often writes

$$X = \operatorname{colim} \mathcal{F},$$

though it is important to keep in mind that the colimit consists of not just the object X but also the morphisms $T_X(\alpha) : \mathcal{F}(\alpha) \to X$ associated to each object α of $\mathscr{A}$. As with direct limits, colimits are not guaranteed to exist, and they are also not generally unique, but the universal property guarantees that they are unique up to canonical isomorphisms whenever they exist.

(a) If $\mathscr{A}$ is a category whose objects form a set J and whose morphisms consist of only the identity morphism on each object, then a diagram $\mathscr{A} \to \mathscr{C}$ is simply a collection $\{X_\alpha\}_{\alpha \in J}$ of objects in $\mathscr{C}$, and a colimit of such a diagram is called a **coproduct** of the collection. Flesh out the details of the following statement: coproducts in the categories Top and $\mathsf{Top}_{\mathrm{rel}}$ are disjoint unions, and coproducts in Ab, $\mathsf{Ab}_{\mathbb{Z}}$ and Chain are direct sums.

(b) Give a concrete description of coproducts in the categories Top_* (pointed spaces) and Grp (not necessarily abelian groups).
Hint: Both answers are constructions that were introduced in last semester's Topologie I course.

(c) If $\mathscr{A}$ contains only two objects α, β and its morphisms consist only of the identity morphisms on α, β plus exactly two morphisms $\alpha \to \beta$, then a diagram $\mathcal{F} : \mathscr{A} \to \mathscr{C}$ can be described as a pair of objects in $\mathscr{C}$ with a pair of morphisms

$$\begin{tikzcd} X \arrow[r, bend left, "f"] \arrow[r, bend right, "g"'] & Y, \end{tikzcd}$$

and a colimit of such a diagram is called a **coequalizer**. Give an explicit description of coequalizers in the categories Top and Ab.
Hint: Use quotients.

(d) Prove: If $\mathscr{C}$ is a category in which coproducts and coequalizers always exist, then every direct system in $\mathscr{C}$ has a direct limit.
Hint: Special cases of this yield the explicit descriptions of direct limits in Top *and* Ab *that appear in Propositions 39.11 and 39.14.*

[63] I say "arbitrary," but in practice, $\mathscr{A}$ is almost always taken to be a *small* category, meaning that its objects form an honest set, rather than a proper class. In many important special cases, $\mathscr{A}$ contains only finitely many objects, and there are already interesting examples (as in Exercise 39.24(c)) in which it has only two.

40. The Euler characteristic (December 1, 2023)

We would now like to discuss a few applications of the isomorphism

$$H_*^{\mathrm{CW}}(X, A; G) \cong H_*(X, A; G).$$

One of the advantages of cellular homology is that for compact spaces, cell decompositions are always finite, in which case the cellular chain complex itself is finitely generated, and so therefore is its homology. This proves:

Corollary 40.1. *If (X, A) is a compact CW-pair, then its singular homology $H_*(X, A; \mathbb{Z})$ is finitely generated.* □

Recall that if G is any abelian group and $\mathbb{K}$ is a field, regarded as an abelian group with respect to its addition operation, then the tensor product $G \otimes \mathbb{K}$ inherits the structure of a vector space over $\mathbb{K}$: indeed, to define this we just need to say what scalar multiplication $\mathbb{K} \times (G \otimes \mathbb{K}) \to G \otimes \mathbb{K}$ means, and the obvious definition determined by the formula

$$\lambda(g \otimes k) := g \otimes (\lambda k) \quad \text{for} \quad \lambda, k \in \mathbb{K},\ g \in G$$

satisfies the required axioms. Moreover, if $\Phi : G \to H$ is a homomorphism between two abelian groups, then $\Phi \otimes \mathbb{1} : G \otimes \mathbb{K} \to H \otimes \mathbb{K}$ becomes a $\mathbb{K}$-linear map with respect to the natural vector space structures on its domain and target. This means that the correspondence

$$\mathsf{Ab} \to \mathsf{Vec}_{\mathbb{K}} : G \mapsto G \otimes \mathbb{K}$$

can be understood as a functor, and in particular, if we choose the field $\mathbb{K}$ itself as a coefficient group in the singular chain complex $C_*(X, A; \mathbb{K}) = C_*(X, A; \mathbb{Z}) \otimes \mathbb{K}$, the latter becomes a chain complex of vector spaces over $\mathbb{K}$ with boundary maps that are $\mathbb{K}$-linear, so that $H_*(X, A; \mathbb{K})$ is also a vector space over $\mathbb{K}$. The same applies to the cellular chain complex with coefficients in $\mathbb{K}$, but now the compactness of (X, A) makes $C_*^{\mathrm{CW}}(X, A; \mathbb{K})$ into a *finite-dimensional* vector space, and Corollary 40.1 becomes a statement of linear algebra:

Corollary 40.2. *If (X, A) is a compact CW-pair, then for any field $\mathbb{K}$, $H_*(X, A; \mathbb{K})$ is a finite-dimensional vector space over $\mathbb{K}$.* □

Note that each of these corollaries is actually two statements in one: they say on the one hand that $H_n(X, A; \mathbb{Z})$ is finitely generated or $H_n(X, A; \mathbb{K})$ is finite dimensional for every $n \in \mathbb{Z}$, but also that both are *trivial* for all but finitely many values of n. It is similarly obvious that $C_k^{\mathrm{CW}}(X; G)$ and therefore also $H_k^{\mathrm{CW}}(X; G)$ must vanish for any CW-pair that has no k-cells:

Corollary 40.3. *If (X, A) is an n-dimensional CW-pair, then $H_k(X, A; G) = 0$ for all $k > n$ and every coefficient group G.* □

Remark 40.4. As I'm sure I've mentioned a few times by now, it is not too hard to prove that every smooth n-manifold is triangulable and is therefore also an n-dimensional CW-complex, so Corollary 40.3 applies to every smooth n-manifold. It also applies to every n-dimensional *topological* manifold, though this is less easy to see—there exist manifolds that do not admit cell decompositions, but it is also known that every n-dimensional manifold is *homotopy equivalent* to a CW-complex of dimension n or less. Since singular homology depends only on homotopy type, Corollary 40.3 still applies.

For a closed n-manifold, we will see another proof that $H_k(M) = 0$ for all $k > n$ when we talk about Poincaré duality later in the course, and that proof requires no knowledge of cell decompositions. It's worth mentioning that homology is in this sense very different from the higher homotopy groups: there are plenty of n-dimensional manifolds M that have $\pi_k(M) \neq 0$ for some $k > n$, e.g. the simplest example is $\pi_3(S^2) \cong \mathbb{Z}$. This is one of the details that makes homology generally easier than homotopy theory.

REMARK 40.5. By results of Palais [Pal66] proved in 1966, it is also known that every smooth (but not necessarily finite-dimensional) Fréchet manifold is homotopy equivalent to a (not necessarily finite-dimensional) CW-complex. Fréchet manifolds are spaces that can be covered by charts identifying them locally with Fréchet spaces, a class of complete metrizable topological vector space that includes all Banach spaces, plus popular non-Banach examples like the space of C^∞-functions on a compact smooth manifold. For example, if M and N are two smooth finite-dimensional manifolds and M is compact, then $C^\infty(M, N)$ is naturally a Fréchet manifold. Since many results of algebraic topology hold only for CW-complexes, Palais's theorem makes the techniques of the subject applicable in many of the functional-analytic settings that are used to study nonlinear PDEs.

Associating a sequence of abelian groups to every topological space is a nice thing to do, but sometimes one would prefer something simpler, e.g. a *number*. There are several numerical invariants that we can now associate to spaces in terms of their homology. Recall that according to the classification of finitely-generated abelian groups, every such group G is isomorphic to

$$G \cong \mathbb{Z}^n \oplus T,$$

for a unique integer $n \geqslant 0$ and a unique finite group T. Concretely, T is the **torsion subgroup** of G, meaning the group of all elements $g \in G$ that satisfy $mg = 0$ for some $m \in \mathbb{N}$. The quotient G/T is then a finitely-generated abelian group with trivial torsion, and thus turns out to be a free abelian group; the smallest number of elements required to generate this group is the same integer $n \geqslant 0$ that appears in the isomorphism above, and is called the **rank** (*Rang*) of G,

$$\operatorname{rank} G := n \geqslant 0.$$

If like many people you prefer linear algebra to group theory, then you might prefer the following way of repackaging the definition of $\operatorname{rank} G$. Suppose $\mathbb{K}$ is a field of characteristic 0. Then $G \otimes \mathbb{K} \cong (\mathbb{Z}^n \otimes \mathbb{K}) \oplus (T \otimes \mathbb{K})$, but for every $g \in T$, we have $mg = 0$ for some $m \in \mathbb{N}$ and thus for $q \in \mathbb{K}$,

$$g \otimes q = g \otimes \left(\underbrace{\frac{q}{m} + \ldots + \frac{q}{m}}_{m} \right) = m \left(g \otimes \frac{q}{m} \right) = mg \otimes \frac{q}{m} = 0,$$

implying $T \otimes \mathbb{K} = 0$.[64] Since $\mathbb{Z} \otimes \mathbb{K} \cong \mathbb{K}$, this gives $G \otimes \mathbb{K} \cong \mathbb{K}^n$, so in terms of the natural $\mathbb{K}$-vector space structure on $G \otimes \mathbb{K}$, we have

$$\operatorname{rank} G = \dim_{\mathbb{K}}(G \otimes \mathbb{K}).$$

It will be convenient to use a related algebraic fact about homology with field coefficients that we are not yet in a position to justify, but it will follow after a couple more lectures when we prove the *universal coefficient theorem*. If C_* is any chain complex of free abelian groups and $\mathbb{K}$ is any field of characteristic zero, it turns out that there is an isomorphism

$$H_n(C_* \otimes \mathbb{K}) \cong H_n(C_*) \otimes \mathbb{K}, \tag{40.1}$$

for every $n \in \mathbb{Z}$, thus giving

$$H_n(X; \mathbb{K}) \cong H_n(X; \mathbb{Z}) \otimes \mathbb{K}$$

when applied to the singular chain complex of a space X. This statement looks relatively innocent at first, but I should point out that it becomes false in general if one tries to relax the assumption on characteristic zero, or to replace $\mathbb{K}$ with an arbitrary abelian coefficient group. Indeed, the Klein bottle (cf. Exercise 36.27) has $H_2(K^2; \mathbb{Z}) = 0$ but $H_2(K^2; \mathbb{Z}_2) \cong \mathbb{Z}_2 \not\cong H_2(K^2; \mathbb{Z}) \otimes \mathbb{Z}_2$. In

[64] Did you notice how we used the assumption here that $\mathbb{K}$ has characteristic 0? It implies in particular that $\mathbb{K}$ contains the rational numbers $\mathbb{Q}$. You cannot do the same trick e.g. with $\mathbb{K} = \mathbb{Z}_2$.

the universal coefficient theorem, the general relationship between $H_*(C_* \otimes G)$ and $H_*(C_*) \otimes G$ takes the form of a short exact sequence instead of a straightforward isomorphism, but the third term conveniently vanishes under certain circumstances, one of them being when G is a field of characteristic zero! We will assume (40.1) in any case for the rest of this lecture, as it makes it possible to erase one or two slightly subtle algebraic issues by converting them into linear algebra.

Definition 40.6. For any space X and integer $k \geqslant 0$, the **kth Betti number** of X is the nonnegative (or possibly infinite) integer

$$b_k(X) := \operatorname{rank} H_k(X;\mathbb{Z}).$$

In light of (40.1), $b_k(X)$ could equivalently be defined as

$$b_k(X) = \dim_{\mathbb{K}} H_k(X;\mathbb{K})$$

if $\mathbb{K}$ is any field of characteristic zero. The most popular choice for this purpose is $\mathbb{Q}$, though $\mathbb{R}$ and $\mathbb{C}$ work just as well.

Remark 40.7. We will see when we study the singular cohomology groups $H^k(X;G)$ that $b_k(X)$ can equally well be defined as the rank of $H^k(X;\mathbb{Z})$ or the dimension of $H^k(X;\mathbb{K})$ for any field $\mathbb{K}$ of characteristic zero, because another version of the universal coefficient theorem implies that these numbers are all the same. In differential geometry, you may also see a definition of $b_k(M)$ for smooth manifolds M as the dimension of the de Rham cohomology $H^k_{\mathrm{dR}}(M;\mathbb{R})$, which is a vector space over $\mathbb{R}$. This matches our definition above due to de Rham's theorem, which provides an isomorphism between $H^k_{\mathrm{dR}}(M;\mathbb{R})$ and the singular cohomology group $H^k(M;\mathbb{R})$ with coefficients in $\mathbb{R}$.

Definition 40.8. For any space X with finitely-generated singular homology, the **Euler characteristic** (*Eulercharakteristik*) of X is the integer[65]

$$\chi(X) = \sum_{k=0}^{\infty} (-1)^k b_k(X) \in \mathbb{Z}.$$

The usefulness of $\chi(X)$ as an invariant derives from a simple phenomenon in homological algebra that has remarkable consequences for topology:

Proposition 40.9. *If C_* is a finitely-generated chain complex of free abelian groups, then*

$$\sum_{n\in\mathbb{Z}} (-1)^n \operatorname{rank} H_n(C_*) = \sum_{n\in\mathbb{Z}} (-1)^n \operatorname{rank} C_n = \sum_{n\in\mathbb{Z}} (-1)^n \dim_{\mathbb{K}} H_n(C_* \otimes \mathbb{K})$$

for any field $\mathbb{K}$.

Proof. Since each C_n is a free abelian group for every n, we have

$$\operatorname{rank} C_n = \dim_{\mathbb{K}}(C_n \otimes \mathbb{K})$$

for any field $\mathbb{K}$. (Here there is no need for $\mathbb{K}$ to have characteristic zero because C_n by assumption has no torsion.) With this in mind, let us suppose instead for the moment that C_* is a finite-dimensional chain complex of $\mathbb{K}$-vector spaces with $\mathbb{K}$-linear boundary maps $\partial_n : C_n \to C_{n-1}$, and abbreviate

$$Z_n := \ker \partial_n, \qquad B_n := \operatorname{im} \partial_{n+1}, \qquad H_n := Z_n/B_n.$$

The dimensions of these vector spaces over $\mathbb{K}$ are then related by

$$\dim H_n = \dim Z_n - \dim B_n$$

[65] Let's be clear about this notational detail: χ is the Greek latter "chi," not a variety of the letter "X" in a strange font. The χ of course stands for "characteristic".

for every $n \in \mathbb{Z}$. Since $\partial_n : C_n \to C_{n-1}$ descends to an isomorphism $C_n/Z_n \to B_{n-1}$, we also have
$$\dim C_n = \dim Z_n + \dim B_{n-1}.$$
Combining these two relations gives
$$\begin{aligned}\sum_{n\in\mathbb{Z}}(-1)^n \dim H_n &= \ldots - (\dim Z_{-1} - \dim B_{-1}) + (\dim Z_0 - \dim B_0) - (\dim Z_1 - \dim B_1) + \ldots \\ &= \ldots - \dim Z_{-1} + (\dim B_{-1} + \dim Z_0) - (\dim B_0 + \dim Z_1) + \dim B_1 + \ldots \\ &= \ldots - \dim C_{-1} + \dim C_0 - \dim C_1 + \ldots = \sum_{n\in\mathbb{Z}}(-1)^n \dim C_n.\end{aligned}$$
Moving back to the original hypothesis where each C_n is a free abelian group, this computation implies
$$\sum_{n\in\mathbb{Z}}(-1)^n \operatorname{rank} C_n = \sum_{n\in\mathbb{Z}}(-1)^n \dim_{\mathbb{K}}(C_n \otimes \mathbb{K}) = \sum_{n\in\mathbb{Z}}(-1)^n \dim_{\mathbb{K}} H_n(C_* \otimes \mathbb{K})$$
for every choice of coefficient field $\mathbb{K}$. Now choosing $\mathbb{K} = \mathbb{Q}$ and applying (40.1) identifies the last expression with $\sum_{n\in\mathbb{Z}}(-1)^n \dim_{\mathbb{Q}} \left(H_n(C_*)\otimes\mathbb{Q}\right) = \sum_{n\in\mathbb{Z}}(-1)^n \operatorname{rank} H_n(C_*)$ and thus completes the proof. □

Notice what we gain if this result is combined with the isomorphism $H_*^{\mathrm{CW}}(X) \cong H_*(X)$ for X a finite cell complex: $H_*(X;\mathbb{Z})$ is then identified with the homology of a finitely-generated chain complex of free abelian groups $C_*^{\mathrm{CW}}(X;\mathbb{Z})$, and this knowledge alone is sufficient to compute the Euler characteristic $\chi(X)$ *without needing to compute the homology!* All we actually need to know for the computation is the rank of each chain group $C_n^{\mathrm{CW}}(X;\mathbb{Z})$, which is the same as the number of n-cells in the complex. But since $\chi(X)$ is a topological invariant, we also learn from this that the alternating sum of these counts of cells does not depend on the choice of cell decomposition:

COROLLARY 40.10. *For any compact space X that admits a cell decomposition, every such decomposition satisfies*
$$\sum_{n=0}^{\infty}(-1)^n|\mathcal{K}^n| = \chi(X),$$
where $|\mathcal{K}^n| \geqslant 0$ denotes the number of n-cells in the decomposition. □

Computing Euler characteristics of cell complexes is now quite easy.

EXAMPLE 40.11. For $n \geqslant 0$, we have $\chi(S^n) = 2$ when n is even and $\chi(S^n) = 0$ when n is odd. One can see this by writing S^n as the union of one 0-cell with one n-cell, or almost as easily, by writing S^n as the union of two k-cells for every $k = 0, \ldots, n$.

EXAMPLE 40.12. For the closed surface Σ_g of genus $g \geqslant 0$, we computed $H_*(\Sigma_g)$ in Example 37.3: the nontrivial homology groups were $H_0(\Sigma_g) \cong \mathbb{Z}$, $H_1(\Sigma_g) \cong \mathbb{Z}^{2g}$ and $H_2(\Sigma_g) \cong \mathbb{Z}$, thus
$$\chi(\Sigma_g) = 1 - 2g + 1 = 2 - 2g.$$
But one can also compute $\chi(\Sigma_g)$ without computing $H_*(\Sigma_g)$, just by observing that Σ_g has a cell decomposition with one 0-cell, one 2-cell and $2g$ cells of dimension 1; this is the same cell decomposition we used in Example 37.3, but there is no longer any need to compute the boundary map.

Here is an application of a more combinatorial nature. Recall that a **graph** (*Graph*) consists of a set V whose elements are called **vertices** (*Ecken* or *Punkte*), and a set E whose elements are called **edges** (*Kanten*), each of which is associated to a particular pair of vertices. Graphs are typically depicted by drawing a point for each vertex and drawing a curve for each edge such that its end points are the two vertices associated to that edge, and in this way every graph Γ naturally gives rise to a 1-dimensional CW-complex $|\Gamma|$ whose 0-cells are the vertices and 1-cells are the edges. The space $|\Gamma|$ is compact if and only if the graph Γ is finite, meaning both V and E are finite, and we say that Γ is **connected** if $|\Gamma|$ is a connected space. A finite connected graph is called a **tree** (*Baum*) if it contains no **cycles**, meaning there does not exist any finite sequence of distinct vertices $v_0, \dots, v_N \in V$ together with a finite sequence of distinct edges $e_0, \dots, e_N$ such that the end points of e_j are v_j and v_{j+1} for $j = 0, \dots, N-1$ but the end points of e_N are v_N and v_0. Now, since $|\Gamma|$ is a 1-dimensional CW-complex, we have $H_k(|\Gamma|) = 0$ for all k except 0 and 1. If Γ is connected, then $|\Gamma|$ is also path-connected and therefore $H_0(|\Gamma|) \cong \mathbb{Z}$. Since there are no 2-cells, $H_1(|\Gamma|)$ is isomorphic to the subgroup of 1-cycles in $C_1^{\mathrm{CW}}(|\Gamma|)$, but it is not hard to prove that if Γ is a tree, then there are also no nontrivial 1-cycles in the chain complex, so $H_1(|\Gamma|) = 0$. This proves $\chi(|\Gamma|) = 1$, and combining it with Corollary 40.10, we then have:

THEOREM 40.13. *For any finite graph Γ with v vertices and e edges, if Γ is a tree, then $v - e = 1$.* □

Let's conclude this discussion with an application to covering spaces.

EXERCISE 40.14. Suppose X is a compact cell complex and $\pi : Y \to X$ is a covering map of finite degree $d \in \mathbb{N}$. Show that Y admits a cell decomposition such that $Y^n = \pi^{-1}(X^n)$ for every n, and every individual n-cell $e_\alpha^n \subset X$ corresponds to exactly d cells in Y whose characteristic maps $\mathbb{D}^n \to Y$ are lifts of the characteristic map $\mathbb{D}^n \to X$ for e_α^n.
Hint: The key point here is that characteristic maps $\mathbb{D}^n \to X$ will always lift to the cover since $\mathbb{D}^n$ is simply connected. It's probably easiest if you argue by induction on n.

The exercise implies:

THEOREM 40.15. *If X is a finite cell complex and $\pi : Y \to X$ is a covering map of finite degree $d \in \mathbb{N}$, then $\chi(Y) = d\chi(X)$.* □

As an easy application, the fact that $\chi(X)$ is always an integer allows us to deduce that there are not very many ways for an even-dimensional sphere to be the universal cover of something else:

COROLLARY 40.16. *If $\pi : S^n \to X$ is a d-fold covering map, n is even and X is a CW-complex, then d is either 1 or 2.* □

EXAMPLE 40.17. Clearly both options in the above corollary are possible: $d = 1$ is always possible since the identity map is a covering map, and $d = 2$ occurs for the natural quotient projection $S^n \to \mathbb{RP}^n$.

41. The Lefschetz fixed point theorem (December 5, 2023)

As another application of cellular homology, in this lecture I'd like to address the following general question:

QUESTION 41.1. *What topological conditions on a map $X \xrightarrow{f} X$ are sufficient to guarantee that f has a fixed point?*

We saw one example last semester: by the Brouwer fixed point theorem, no conditions at all are needed for f if X is a disk. We also saw in Lecture 35 that for $X = S^n$, every map f that does not have degree $(-1)^{n+1}$ must have a fixed point—this is a homotopy-invariant condition, but of course it is important to include the exception in this statement, as e.g. the antipodal map does not have any fixed points.

Our goal for today is a much more general homotopy-invariant criterion for the existence of fixed points. The rough idea is as follows: if $f : X \to X$ has no fixed points but X is a nice enough space to admit cell decompositions, then we would like to find a special cell decomposition of X such that after adjusting f by a homotopy, f becomes a cellular map sending each cell to *different* cells. In other words, the induced chain map $f_* : C_*^{\mathrm{CW}}(X) \to C_*^{\mathrm{CW}}(X)$ given by

$$f_* e_\alpha^n = \sum_{e_\beta^n \subset X} [e_\beta^n : e_\alpha^n] e_\beta^n$$

will then have the property that its *diagonal* terms all vanish:

$$[e_\alpha^n : e_\alpha^n] = 0 \qquad \text{for all } n\text{-cells } e_\alpha^n \subset X.$$

In this situation, the chain map f_* is represented by a matrix that has zeroes along the diagonal, so its trace vanishes. At this point it is useful to introduce an algebraic result that has much in common with the previous lecture's Propostion 40.9:

THEOREM 41.2 (Hopf trace formula). *Suppose C_* is a finite-dimensional chain complex of vector spaces over a field $\mathbb{K}$, and $f : C_* \to C_*$ is a $\mathbb{K}$-linear chain map. Then*

$$\sum_{n \in \mathbb{Z}} (-1)^n \operatorname{tr}\left(C_n \xrightarrow{f} C_n\right) = \sum_{n \in \mathbb{Z}} (-1)^n \operatorname{tr}\left(H_n(C_*) \xrightarrow{f_*} H_n(C_*)\right).$$

PROOF. For the boundary maps $\partial_n : C_n \to C_{n-1}$ for each $n \in \mathbb{Z}$, abbreviate

$$Z_n := \ker \partial_n \subset C_n, \qquad B_n := \operatorname{im} \partial_{n+1} \subset C_n, \qquad H_n := Z_n / B_n.$$

Denote $f_{C_n} := f|_{C_n} : C_n \to C_n$, and note that since f is a chain map, it restricts to these subspaces and the quotient as linear maps

$$Z_n \xrightarrow{f_{Z_n}} Z_n, \qquad B_n \xrightarrow{f_{B_n}} B_n, \qquad H_n \xrightarrow{f_{H_n}} H_n,$$

such that the following diagram commutes

$$\begin{tikzcd}
0 \arrow[r] & Z_n \arrow[r, hook] \arrow[d, "f_{Z_n}"] & C_n \arrow[r, "\partial_n"] \arrow[d, "f_{C_n}"] & B_{n-1} \arrow[r] \arrow[d, "f_{B_{n-1}}"] & 0 \\
0 \arrow[r] & Z_n \arrow[r, hook] & C_n \arrow[r, "\partial_n"] & B_{n-1} \arrow[r] & 0
\end{tikzcd}$$

and its rows are exact. Here it is convenient to make use of the assumption that all these objects are *vector spaces*, not just abelian groups—it guarantees in particular that a short exact sequence always splits, i.e. we can choose a subspace of C_n complementary to Z_n and use the map ∂_n to identify that subspace with B_{n-1}, giving a (non-canonical) isomorphism

$$C_n \cong Z_n \oplus B_{n-1}.$$

Identifying C_n in this way with $Z_n \oplus B_{n-1}$, the map $f_{C_n} : C_n \to C_n$ becomes a matrix of the form

$$f_{C_n} = \begin{pmatrix} f_{Z_n} & g \\ 0 & f_{B_{n-1}} \end{pmatrix}$$

for some linear map $g : B_{n-1} \to Z_n$. Here the lower-left term vanishes because f_{C_n} preserves the subspace Z_n, and the other off-diagonal term might not vanish because f_{C_n} need not preserve the complementary subspace, yet if we restrict f_{C_n} to this subspace and project away the term in Z_n, what remains is the map $B_{n-1} \to B_{n-1}$ induced by the same chain map f, i.e. it is the lower-right term $f_{B_{n-1}}$. This formula proves

$$\operatorname{tr}(f_{C_n}) = \operatorname{tr}(f_{Z_n}) + \operatorname{tr}(f_{B_{n-1}}). \tag{41.1}$$

Now apply the same argument to the diagram

$$\begin{array}{ccccccccc} 0 & \longrightarrow & B_n & \hookrightarrow & Z_n & \longrightarrow & H_n & \longrightarrow & 0 \\ & & \big\downarrow{\scriptstyle f_{B_n}} & & \big\downarrow{\scriptstyle f_{Z_n}} & & \big\downarrow{\scriptstyle f_{H_n}} & & \\ 0 & \longrightarrow & B_n & \hookrightarrow & Z_n & \longrightarrow & H_n & \longrightarrow & 0 \end{array}$$

where the maps $Z_n \to H_n$ are the natural quotient projections and the rows are therefore exact. We obtain

$$\operatorname{tr}(f_{Z_n}) = \operatorname{tr}(f_{B_n}) + \operatorname{tr}(f_{H_n}),$$

and combining this with (41.1) gives

$$\sum_{n\in\mathbb{Z}}(-1)^n\left[\operatorname{tr}(f_{C_n}) - \operatorname{tr}(f_{B_{n-1}})\right] = \sum_{n\in\mathbb{Z}}(-1)^n \operatorname{tr}(f_{Z_n}) = \sum_{n\in\mathbb{Z}}(-1)^n\left[\operatorname{tr}(f_{B_n}) + \operatorname{tr}(f_{H_n})\right],$$

which implies the desired result after dropping the extraneous terms $\operatorname{tr}(f_{B_n})$ from both sides. □

Definition 41.3. For any space X and a field $\mathbb{K}$ such that $H_*(X;\mathbb{K})$ is finite dimensional, the **Lefschetz number** (*Lefschetz-Zahl*) of a map $f : X \to X$ is defined by

$$L_{\mathbb{K}}(f) := \sum_{n\in\mathbb{Z}}(-1)^n \operatorname{tr}\left(H_n(X;\mathbb{K}) \xrightarrow{f_*} H_n(X;\mathbb{K})\right) \in \mathbb{K}.$$

In the case $\mathbb{K} = \mathbb{Q}$, we denote this more simply by

$$L(f) := L_{\mathbb{Q}}(f).$$

Notice that by the homotopy axiom for homology, $L_{\mathbb{K}}(f)$ depends on f only up to homotopy.

Remark 41.4. We will not need to know this for our discussion, but it's interesting to note that while the definition above makes $L(f)$ a rational number, it is secretly always an *integer*. If X is a finite CW-complex and f a cellular map, then this follows easily from the Hopf trace formula, as $L_{\mathbb{Q}}(f)$ is then the same as the alternating sum of the traces of maps $f_* : C_n^{\mathrm{CW}}(X)\otimes\mathbb{Q} \to C_n^{\mathrm{CW}}(X)\otimes\mathbb{Q}$, each of which can be written in the canonical basis as a matrix with integer entries. Without these assumptions, it follows more generally from the universal coefficient theorem, which will give us a natural isomorphism $H_*(X;\mathbb{Q}) \cong H_*(X)\otimes\mathbb{Q}$, so that the maps $f_* : H_*(X;\mathbb{Q}) \to H_*(X;\mathbb{Q})$ can also be presented as matrices with integer entries. More precisely, every endomorphism $H_n(X) \to H_n(X)$ preserves the torsion subgroup $T_n \subset H_n(X)$ and thus descends to an endomorphism of the *free part* of $H_n(X)$,

$$H_n(X)/T_n \xrightarrow{f_*} H_n(X)/T_n,$$

which is a free abelian group. Thus f_* can again be presented as an integer matrix with respect to any basis of this free group, and the alternating sum of the traces of these matrices is the integer $L(f)$.

Exercise 41.5. Show that if X has finitely-generated homology and $f : X \to X$ is homotopic to the identity map, $L(f) = \chi(X)$.

Here is the main result of this lecture.

Theorem 41.6 (Lefschetz-Hopf). *If X is a compact polyhedron and $\mathbb{K}$ is a field, then every map $f : X \to X$ satisfying $L_{\mathbb{K}}(f) \neq 0$ has a fixed point.*

Before discussing the proof, we give one application and a few remarks. The application is an extension of the famous "hairy sphere" theorem (recall Theorem 34.13), and its proof requires some knowledge of the *flow* of a smooth vector field from differential geometry.

COROLLARY 41.7. *For any closed smooth manifold M with $\chi(M) \neq 0$, there is no continuous vector field on M that is nowhere zero.*

PROOF. If such a vector field exists, then we can approximate it with a smooth vector field X that is also nowhere zero. The flow of X for some small but nonzero time $t > 0$ is then a diffeomorphism $\varphi_X^t : M \to M$ with no fixed points, but is clearly also homotopic to the identity, thus by Exercise 41.5, $L(\varphi_X^t) = \chi(M) = 0$. □

REMARK 41.8. Another easy corollary of the theorem is that it also holds for spaces somewhat more general than compact polyhedra: it holds in particular whenever X is a compact **Euclidean neighborhood retract**, meaning X admits a topological embedding $X \hookrightarrow \mathbb{R}^N$ for some $N \in \mathbb{N}$ such that some neighborhood $\mathcal{U} \subset \mathbb{R}^N$ of X admits a retraction to X. It is not so hard to prove (see [**Hat02**, Corollary A.9]) that all compact topological manifolds have this property, even those which do not admit triangulations. In this situation, even if X does not have a triangulation, we can triangulate $\mathbb{R}^N$ finely enough so that all simplices touching $X \subset \mathbb{R}^N$ are contained in the neighborhood $\mathcal{U}$, and the retraction $r : \mathcal{U} \to X$ then makes X a retract of a compact polyhedron $K \subset \mathcal{U}$ containing X. Now if $f : X \to X$ has $L_{\mathbb{K}}(f) \neq 0$, one can consider the map

$$i \circ f \circ r : K \to K$$

where $i : X \hookrightarrow K$ is the inclusion, and use Exercise 41.9 below to prove $L_{\mathbb{K}}(i \circ f \circ r) = L_{\mathbb{K}}(f)$, so that Theorem 41.6 guarantees a fixed point for $i \circ f \circ r$. But $i \circ f \circ r(x) = x$ implies $x \in X$ and $f(x) = x$.

EXERCISE 41.9. Suppose $A \subset X$ is a subspace with inclusion $i : A \hookrightarrow X$ and a retraction $r : X \to A$, and X has finite-dimensional homology with coefficients in some field $\mathbb{K}$. Show that $H_*(A; \mathbb{K})$ is also finite dimensional, and for any map $f : A \to A$, the induced maps $f_* : H_n(A; \mathbb{K}) \to H_n(A; \mathbb{K})$ and $(i \circ f \circ r)_* : H_n(X; \mathbb{K}) \to H_n(X; \mathbb{K})$ for every $n \in \mathbb{Z}$ satisfy

$$\mathrm{tr}(f_*) = \mathrm{tr}((i \circ f \circ r)_*).$$

Hint: Write $(i \circ f \circ r)_* = i_* f_* r_*$ as the composition of two homomorphisms $f_* r_* : H_n(X; \mathbb{K}) \to H_n(A; \mathbb{K})$ and $i_* : H_n(A; \mathbb{K}) \to H_n(X; \mathbb{K})$, and recall the formula $\mathrm{tr}(\mathbf{AB}) = \mathrm{tr}(\mathbf{BA})$.

REMARK 41.10. Lefschetz's original version of the fixed point theorem applied only to manifolds and was thus more restrictive, but it has the following nice feature that Theorem 41.6 lacks. For a map $f : M \to M$ on an n-manifold with at most finitely many fixed points, the Lefschetz number $L(f)$ gives not only a sufficient condition but also an *algebraic count* of the fixed points, in the same sense that the degree of a map $f : M \to N$ counts the points in $f^{-1}(q)$ for any $q \in N$. The proof of this version is best expressed in terms of Poincaré duality and homological intersection theory; see e.g. [**Bre93**, §VI.12]. As a consequence, one can then extend Corollary 41.7 to the statement that on a closed oriented manifold M, for any vector field that has at most finitely many zeroes, the algebraic count of these zeroes is $\chi(M)$; this is known as the *Poincaré-Hopf theorem.*

REMARK 41.11. It is easy to see that the compactness of X in Theorem 41.6 is essential: for instance, $\mathbb{R}$ has finitely-generated homology and $f : \mathbb{R} \to \mathbb{R} : x \mapsto x + 1$ is homotopic to the identity, hence $L(f) = \chi(\mathbb{R}) = 1$, even though f has no fixed points.

REMARK 41.12. Figure 21 shows a compact space X that violates the Lefschetz fixed point theorem because it is not a polyhedron. Indeed, X has three path-components, two (the outer and inner circle) that are homeomorphic to S^1 and one (the spiral in between) homeomorphic to $\mathbb{R}$, thus

$$H_*(X) \cong H_*(S^1) \oplus H_*(S^1) \oplus H_*(\mathbb{R}),$$

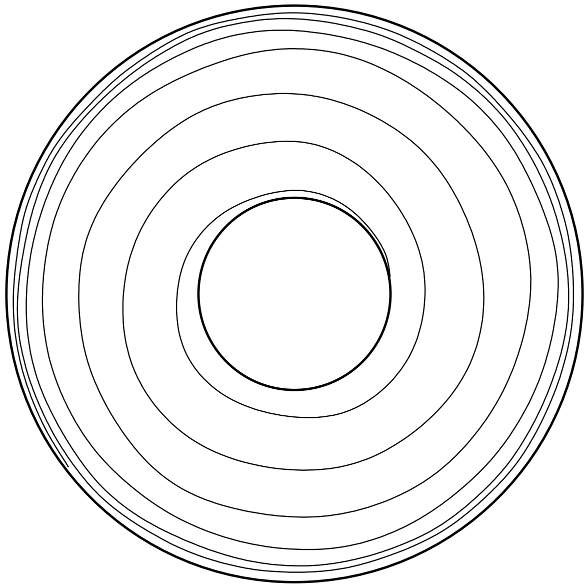

FIGURE 21. A compact space X with $\chi(X) = 1$ admitting maps homotopic to the identity that have no fixed point.

implying $\chi(X) = \chi(S^1)+\chi(S^1)+\chi(\mathbb{R}) = 0+0+1 = 1$. But it is easy to visualize a map $f : X \to X$ that is homotopic to the identity and has no fixed points, e.g. define f by a small rotation, with radii adjusted appropriately so that it preserves the spiral. (You may notice that X is also an example of a space that is connected but not path-connected—that is a property that polyhedra never have.)

To prove Theorem 41.6, we need to make precise the idea sketched at the beginning of this lecture: a map $f : X \to X$ with no fixed points can be modified to a *cellular* map whose induced chain map has no diagonal terms. Since we are working with polyhedra, it is natural to consider not just cellular but also *simplicial* maps. You may want to take a moment to review the definitions given in Lecture 30 for the combinatorial notion of a simplicial map and the induced continuous map on polyhedra, which maps each n-simplex linearly to a k-simplex for some $k \leqslant n$. In this context, the following question seems natural, though its immediate answer may disappoint you:

QUESTION 41.13. *Given two polyhedra X and Y, is every continuous map $f : X \to Y$ homotopic to a simplicial map?*

The answer is **no**: for example, if Y happens to have more vertices than X, then a simplicial map $f : X \to Y$ can never be surjective. We could for instance take X and Y to be the sphere S^n with two distinct triangulations such that Y has more vertices, in which case only maps $S^n \to S^n$ with degree zero have any hope of being homotopic to simplicial maps. The message of this answer is that we asked the wrong question. Suppose we eliminate this counterexample as follows: instead of fixing given triangulations of X and Y, we first make the triangulation of X finer by *subdividing* it until it has at least as many vertices as Y. If this extra step is allowed, then it is no longer obvious that any given map $f : X \to Y$ cannot be homotopic to a simplicial map. I remind you that when I say "subdivision," I typically mean *barycentric* subdivision, as shown in Figure 22.

Here's the main technical result we need. I will give only a sketch of the proof, but the main idea is not so hard to understand. (For a more detailed proof, see [**Hat02**, §2.C].)

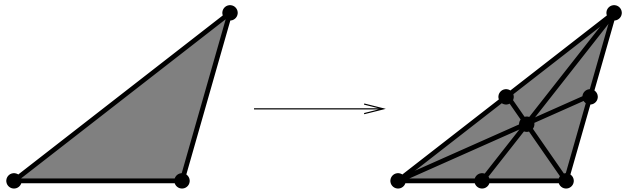

Figure 22. Barycentric subdivision of a 2-simplex.

Theorem 41.14 (simplicial approximation). *If X and Y are compact polyhedra and $f : X \to Y$ is any continuous map, then after modifying the triangulation of X by finitely many barycentric subdivisions, f is homotopic to a simplicial map $g : X \to Y$ such that for every $x \in X$, $g(x)$ is contained in the smallest simplex of Y containing $f(x)$.*

Sketch of the proof. For each vertex $v \in X$, define the so-called **open star** of v as the open neighborhood

$$\operatorname{st} v \subset X$$

of v formed by the union of the interiors of all simplices in X that have v as a vertex. Figure 23 shows the open stars of two neighboring vertices in a 2-dimensional polyhedron; notice that their intersection contains the interior of the 1-simplex bounded by these two vertices (cf. Exercise 41.15 below). The collection of all open stars of vertices defines an open covering of any polyhedron. Now given $f : X \to Y$ continuous, after subdividing the triangulation of X enough times, we can assume that for every vertex $v \in X$ there exists a vertex $w_v \in Y$ such that (see Figure 23 again)

$$\operatorname{st} v \subset f^{-1}(\operatorname{st} w_v).$$

Having associated to each $v \in X$ some $w_v \in Y$ with this property, there is a unique simplicial map $g : X \to Y$ that satisfies $g(v) = w_v$: indeed, for every simplex $\{v_0, \dots, v_n\}$ of X, the exercise below implies that the set $\{w_{v_0}, \dots, w_{v_n}\}$ is also a simplex of Y. One can now check that g is indeed an "approximation" of f in the sense that $g(x)$ is contained in the smallest simplex of Y containing $f(x)$ for every $x \in X$. In light of this, a homotopy $h : I \times X \to Y$ from f to g can be defined by choosing $h(\cdot, x) : I \to Y$ for every $x \in X$ to be the linear path from $f(x)$ to $g(x)$ in the smallest simplex containing $f(x)$. □

Exercise 41.15. Given vertices $v_0, \dots, v_k$ in a polyhedron X, show that $\bigcap_{i=0}^k \operatorname{st} v_i \neq \varnothing$ if and only if X contains a simplex whose vertices are $v_0, \dots, v_k$.

We can now prove the Lefschetz-Hopf theorem.

Proof of Theorem 41.6. Assume X is a compact polyhedron, $\mathbb{K}$ is a field and $f : X \to X$ has no fixed points. Compact polyhedra are metrizable, so we can choose a metric $d(\cdot, \cdot)$ on X and observe that since X is compact, there exists a number $\epsilon > 0$ such that

$$d(x, f(x)) \geqslant \epsilon > 0 \quad \text{for all } x \in X.$$

After repeated subdivisions, we can assume without loss of generality that every simplex in the triangulation of X has diameter less than $\epsilon/2$. Now let X' denote the same space but with its triangulation further subdivided so that the simplicial approximation theorem applies, giving a simplicial map

$$g : X' \to X$$

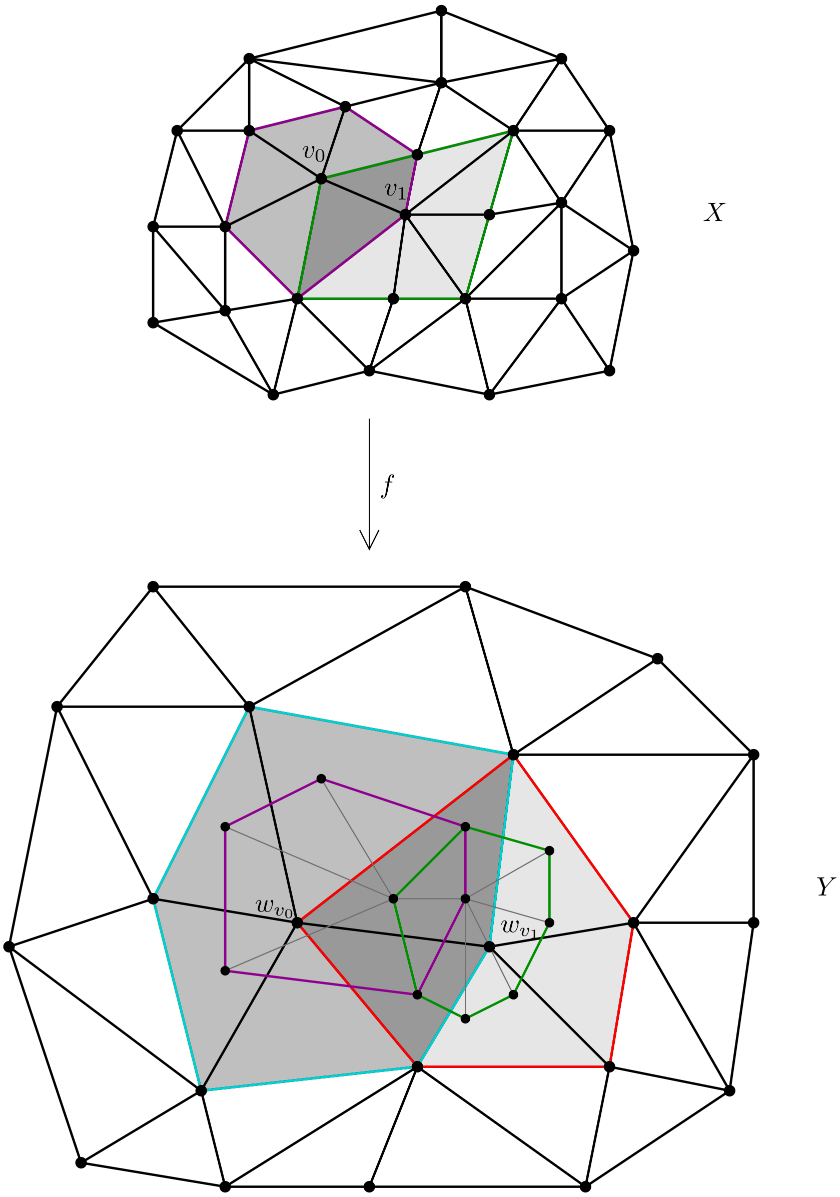

FIGURE 23. A map $f : X \to Y$ between two polyhedra, with vertices $v_0, v_1 \in X$ and $w_{v_0}, w_{v_1} \in Y$ chosen such that f maps the open star of v_i into the open star of w_{v_i} for $i = 0, 1$. The prescription in the proof of Theorem 41.14 will then produce a simplicial map $g : X \to Y$ sending $v_i \mapsto w_{v_i}$ for $i = 0, 1$, so the 1-simplex in X bounded by v_0 and v_1 is sent to the 1-simplex in Y bounded by w_{v_0} and w_{v_1}.

that is homotopic to f as a continuous map. Since the n-skeleton of X is contained in the n-skeleton of X' for every $n \geqslant 0$, one can also regard g as a *cellular* (though not simplicial) map

$$g : X' \to X'.$$

Now, every simplex in either X' or X has diameter less than $\epsilon/2$, and since $g(x)$ and $f(x)$ always lie in a common simplex of X, it follows that $d(g(x), f(x)) < \epsilon/2$ for every $x \in X$. Therefore,

$$d(x, g(x)) \geqslant d(x, f(x)) - d(f(x), g(x)) > \epsilon - \frac{\epsilon}{2} = \frac{\epsilon}{2},$$

implying that x and $g(x)$ never belong to the same simplex of X'. It follows that the diagonal incidence numbers $[e_\alpha^n : e_\alpha^n]$ vanish for every n-cell $e_\alpha^n \subset X'$ defined as the interior of an n-simplex in our subdivided triangulation, implying that the induced chain map

$$C_*^{\mathrm{CW}}(X'; \mathbb{K}) \xrightarrow{g_*} C_*^{\mathrm{CW}}(X'; \mathbb{K})$$

has only zeroes along the diagonal, and its trace in every dimension is therefore 0. By the Hopf trace formula, it follows that $L_{\mathbb{K}}(g) = L_{\mathbb{K}}(f) = 0$. □

42. The universal coefficient theorem (December 12, 2023)

Our goal in this lecture is to clarify precisely how the groups $H_*(X, A; G_1)$ and $H_*(X, A; G_2)$ are related to each other for different choices of abelian coefficient group G_1 and G_2. This is an essentially algebraic question, and the methods we need in order to answer it come from homological algebra. Recall that for any abelian group G, $H_*(\cdot; G) : \mathsf{Top}_{\mathrm{rel}} \to \mathsf{Ab}_{\mathbb{Z}}$ can be regarded as the composition of three functors:

$$\mathsf{Top}_{\mathrm{rel}} \xrightarrow{C_*(\cdot;\mathbb{Z})} \mathsf{Chain} \xrightarrow{\otimes G} \mathsf{Chain} \xrightarrow{H_*} \mathsf{Ab}_{\mathbb{Z}}, \tag{42.1}$$

where the first sends each pair of spaces (X, A) to its singular chain complex $C_*(X, A; \mathbb{Z})$ with integer coefficients, the second replaces the latter with $C_*(X, A; G) = C_*(X, A; \mathbb{Z}) \otimes G$, and the third computes the homology of the chain complex. If we were only concerned with *defining* topological invariants and not with *computing* them, then we could stop with $C_*(\cdot; \mathbb{Z}) : \mathsf{Top}_{\mathrm{rel}} \to \mathsf{Chain}$, as the singular chain complex $C_*(X, A; \mathbb{Z})$ is in itself a topological invariant, and it contains in principle all of the information that we could ever want to extract from any version of the singular homology of (X, A). The problem is that, as an invariant in itself, $C_*(X, A; \mathbb{Z})$ is horribly unwieldy and impractical: the group is absurdly large, and if you want to prove (X, A) and (Y, B) are not homeomorphic, you typically will not be able to do it by proving directly that $C_*(X, A; \mathbb{Z})$ and $C_*(Y, B; \mathbb{Z})$ are non-isomorphic chain complexes. This is where replacing the complexes with their homology groups is useful: strictly speaking, we lose a lot of information when we do this, but it's worth it if the information that remains afterwards is manageable.

The situation is slightly different for the cellular homology functor $H_*^{\mathrm{CW}}(\cdot; G) : \mathsf{CW}_{\mathrm{rel}} \to \mathsf{Ab}_{\mathbb{Z}}$, which is similarly a composition of three functors

$$\mathsf{CW}_{\mathrm{rel}} \xrightarrow{C_*^{\mathrm{CW}}(\cdot;\mathbb{Z})} \mathsf{Chain} \xrightarrow{\otimes G} \mathsf{Chain} \xrightarrow{H_*} \mathsf{Ab}_{\mathbb{Z}}.$$

For a CW-pair (X, A), the chain complex $C_*^{\mathrm{CW}}(X, A; \mathbb{Z})$ is typically much more manageable, but it depends on the cell decomposition and is thus not a topological invariant. Passing to the homology $H_*^{\mathrm{CW}}(X, A)$ is thus necessary in order to obtain something that depends only on the topology of (X, A). One could say the same thing about simplicial homology, which of course is just a special case of cellular homology.

For this lecture, we're going to focus on the purely algebraic aspects that are common to both singular and cellular homology, namely the two functors

$$\mathsf{Chain} \xrightarrow{\otimes G} \mathsf{Chain} \xrightarrow{H_*} \mathsf{Ab}_{\mathbb{Z}}. \tag{42.2}$$

Each of them destroys some information in general, e.g. the homology of a chain complex can easily be trivial even when the complex itself is not, and applying $\otimes\mathbb{Q}$ to an abelian group with nontrivial torsion elements will always kill them. The case $G = \mathbb{Z}$ is special, because $\otimes\mathbb{Z}$ is actually the identity functor, so there seems to be less potential for losing information if we stick with integer coefficients. Of course, we have also seen cases (e.g. the Klein bottle) where $H_n(X;\mathbb{Z})$ is trivial while $H_n(X;\mathbb{Z}_2)$ is not. The main result of this lecture will show however that $H_n(X;\mathbb{Z}_2)$ is nevertheless determined up to isomorphism by the two groups $H_n(X;\mathbb{Z})$ and $H_{n-1}(X;\mathbb{Z})$. Results like this should not be interpreted to mean that homology with different coefficient groups is *unnecesary*—we've seen for instance that when M is an n-manifold with a non-orientable triangulation, it is in some ways more natural to consider $H_n(M;\mathbb{Z}_2)$ than $H_n(M;\mathbb{Z})$. But as a computational device, it is also often useful to know that the homology with integer coefficients determines everything else. Moreover, the correspondence can also be made to go in the other direction, allowing information about $H_*(X;\mathbb{Z})$ to be deduced from a collection of groups $H_*(X;\mathbb{K})$ for various choices of field coefficients $\mathbb{K}$.

Categories of R-modules. Moving forward, it will be useful to consider a slight generalization of our usual algebraic setting for chain complexes of abelian groups. Assume R is a commutative ring with unit, and let

$$\mathsf{Mod}^R$$

denote the category of R-modules, in which the morphisms $\Phi : A \to B$ between two objects A, B are R-module homomorphisms, meaning that they are group homomorphisms that additionally satisfy

$$\Phi(ra) = r\Phi(a) \qquad \text{for all } r \in R,\ a \in A.$$

Recall that an R-module A is called **free** if it admits a **basis**, meaning a subset $\mathcal{B} \subset A$ such that every element $x \in A$ can be expressed in the form

$$x = \sum_{b \in \mathcal{B}} x_e e$$

for coefficients $x_e \in R$ that are uniquely determined by x and vanish except for finitely-many $e \in \mathcal{B}$. A basis $\mathcal{B} \subset A$ is thus equivalent to an R-module isomorphism between A and a direct sum of copies of R,

$$A \cong \bigoplus_{e \in \mathcal{B}} R =: F_R^{\mathrm{mod}}(\mathcal{B}),$$

and the R-module $F_R^{\mathrm{mod}}(\mathcal{B})$ defined via this direct sum (which makes sense in principle for any set $\mathcal{B}$) is also called the **free R-module on the set $\mathcal{B}$**. Taking $R = \mathbb{Z}$ as a special case, modules over $\mathbb{Z}$ are simply abelian groups, and a free $\mathbb{Z}$-module is the same thing as a free abelian group, hence $\mathsf{Mod}^{\mathbb{Z}} = \mathsf{Ab}$. On the other hand, taking R to be a field $\mathbb{K}$ makes R-modules into vector spaces over $\mathbb{K}$, so that $\mathsf{Mod}^{\mathbb{K}} = \mathsf{Vec}_{\mathbb{K}}$, and the fact that vector spaces always admit bases means that all $\mathbb{K}$-modules are free.[66]

For general choices of commutative ring R, the categories $\mathsf{Ab}_{\mathbb{Z}}$ and Chain have obvious generalizations $\mathsf{Mod}_{\mathbb{Z}}^R$ and Chain^R to categories of $\mathbb{Z}$-graded R-modules or chain complexes of R-modules respectively, where in the latter case, boundary operators and chain maps are all required to be

[66] One should not overstate how trivial this fact about vector spaces is: in the infinite-dimensional case, it depends on Zorn's lemma, and thus the axiom of choice. Note also that when we talk algebraically about a basis of an infinite-dimensional vector space, we do not mean the same thing that is meant when discussing bases of Hilbert spaces in functional analysis: in a separable Hilbert space $\mathcal{H}$, the term "basis" is typically taken to mean a countable subset $\mathcal{B} \subset \mathcal{H}$ such that every $x \in \mathcal{H}$ can be written uniquely as a *convergent* sum $\sum_{\mathcal{B} \in B} x_e e$, where it is possible for infinitely many of the coefficients x_e to be nonzero if they decay fast enough. In algebra, our vector spaces have no topologies and thus no notion of convergence, so for $\mathcal{B}$ to be a basis means that every x can be written uniquely as a sum $\sum_{e \in \mathcal{B}} x_e e$ that converges for a trivial reason, namely that at most finitely-many of its terms are nonzero.

R-module homomorphisms. The homology of a chain complex of R-modules is thus a $\mathbb{Z}$-graded R-module, generalizing the usual algebraic homology functor $H_* : \mathsf{Chain} \to \mathsf{Ab}_{\mathbb{Z}}$ to

$$H_* : \mathsf{Chain}^R \to \mathsf{Mod}^R_{\mathbb{Z}}.$$

Given two R-modules A and B, we shall often denote the set of R-module homomorphisms $A \to B$ by

$$\mathrm{Hom}_R(A, B) := \{R\text{-module homomorphisms } A \to B\},$$

which has a natural R-module structure of its own, and observe that if $R = \mathbb{Z}$, then $\mathrm{Hom}_{\mathbb{Z}}(A, B) = \mathrm{Hom}(A, B)$ is just the usual abelian group of group homomorphisms from A to B. The **tensor product** of A and B can be defined as a quotient of the free R-module on the set $A \times B$: we write

$$A \otimes_R B := F^{\mathrm{mod}}_R(A \times B)/N,$$

where $N \subset F^{\mathrm{mod}}_R(A \times B)$ is the smallest submodule containing all elements of the form $(a + a', b) - (a, b) - (a', b)$, $(a, b + b') - (a, b) - (a, b')$, $(ra, b) - r(a, b)$ and $(a, rb) - r(a, b)$ for $a, a' \in A$, $b, b' \in B$ and $r \in R$. We denote the equivalence class represented by $(a, b) \in F^{\mathrm{mod}}_R(a \times b)$ in the quotient by

$$a \otimes b := [(a, b)] \in A \otimes_R B,$$

and it is straightforward to check that if A and B are both free with bases $\mathcal{A} \subset A$ and $\mathcal{B} \subset B$, then the set of all products $a \otimes b$ for $(a, b) \in \mathcal{A} \times \mathcal{B}$ defines a basis of $A \otimes_R B$. It should be clear that if $R = \mathbb{Z}$, then this definition reproduces the usual tensor product of abelian groups $A \otimes B$, though for other choices of R, the R-module $A \otimes_R B$ is typically a different set from the abelian group $A \otimes B$, because tensor products of R-modules satisfy the additional relation

$$a \otimes (rb) = r(a \otimes b) = (ra) \otimes b \qquad \text{for all } a \in A,\ b \in B,\ r \in R, \tag{42.3}$$

which is redundant in the special case $R = \mathbb{Z}$. If G is an R-module and C_* is a $\mathbb{Z}$-graded R-module or a chain complex of R-modules, then $C_* \otimes_R G$ also admits the structure of a $\mathbb{Z}$-graded R-module or chain complex of R-modules respectively in an obvious way. In this way we obtain a generalization of the composition of functors in (42.2) to

$$\mathsf{Chain}^R \xrightarrow{\otimes_R G} \mathsf{Chain}^R \xrightarrow{H_*} \mathsf{Mod}^R_{\mathbb{Z}}.$$

CONVENTION. For the rest of this and the next lecture, we fix a commutative ring R and regard all chain complexes, homologies and tensor products as living in the category Mod^R of R-modules or the related categories $\mathsf{Mod}^R_{\mathbb{Z}}$ or Chain^R as appropriate. With this understood, we shall generally drop R from the notation whenever the choice of ring is unimportant, thus writing e.g. the module of homomorphisms and the tensor product of two R-modules with the same notation that is normally reserved for abelian groups,

$$\mathrm{Hom}(A, B) := \mathrm{Hom}_R(A, B), \qquad A \otimes B := A \otimes_R B.$$

For our purposes, the specific commutative rings R that will be important to consider are $\mathbb{Z}$ and arbitrary fields $\mathbb{K}$. These examples all have a particular property that we will need to make use of at a few crucial points, namely: they are principal ideal domains. If you've forgotten what a principal ideal domain is, then there is no pressing need to look up the definition right now, so long as you are willing to accept the following basic fact about them:

PROPOSITION 42.1 (see [**Lan02**, §III.7]). *If R is a principal ideal domain, then every submodule of a free R-module is also free.* $\square$

The universal coefficient theorem. Assume C_* is a chain complex of R-modules with boundary operator $\partial : C_* \to C_{*-1}$, and G is another R-module. The first important observation in the background of the universal coefficient theorem is that for every $n \in \mathbb{Z}$, there is a natural homomorphism

$$H_n(C_*) \otimes G \xrightarrow{h} H_n(C_* \otimes G) : [c] \otimes g \mapsto [c \otimes g]. \tag{42.4}$$

Indeed, this map is well defined since $\partial c = 0$ implies $(\partial \otimes \mathbb{1})(c \otimes g) = 0$ and $c = \partial a$ implies $c \otimes g = (\partial \otimes \mathbb{1})(a \otimes g)$. In the main cases of interest to us, e.g. when $C_* = C_*(X, A; \mathbb{Z})$ is the singular chain complex of a pair of spaces (X, A) with integer coefficients and G is a chosen coefficient group, this map becomes

$$H_n(X, A; \mathbb{Z}) \otimes G \xrightarrow{h} H_n(X, A; G),$$

and thus gives a relation between the homologies with coeffients in $\mathbb{Z}$ and G respectively. Alternatively, one could place the cellular chain complex of a CW-pair (X, A) in the role of C_* and write down a similar homomorphism relating the cellular homologies with coefficients in $\mathbb{Z}$ and G. In both of these examples, the chain groups C_n are all free, so it will do no harm to impose this assumption in our algebraic setting—more generally, we will usually assume that C_* is a chain complex of free R-modules and that R is a principal ideal domain.

A first naive hope would be for the map (42.4) to be an isomorphism, and we will see that this is indeed true in many important cases, but not always. It will turn out that it is always injective, but not always surjective: in general, $h : H_n(C_*) \otimes G \to H_n(C_* \otimes G)$ forms the first two nontrivial terms in a natural short exact sequence, whose third term will vanish if and only if h is an isomorphism.

Let us start by deriving a short exact sequence that has $H_n(C_* \otimes G)$ as the middle term. For the given boundary maps $\partial_n : C_n \to C_{n-1}$, abbreviate

$$Z_n := \ker \partial_n \subset C_n, \qquad B_n := \operatorname{im} \partial_{n+1} \subset Z_n \subset C_n,$$

and note that since we are assuming C_n is free and R is a principal ideal domain, Proposition 42.1 implies that Z_n and B_n are likewise free. Since $\partial_n : C_n \to C_{n-1}$ restricts to B_n and Z_n as the zero map, we can collect the obvious short exact sequences

$$0 \to Z_n \hookrightarrow C_n \xrightarrow{\partial} B_{n-1} \to 0, \qquad n \in \mathbb{Z}$$

together to form a short exact sequence of chain complexes

$$0 \to Z_* \hookrightarrow C_* \xrightarrow{\partial} B_{*-1} \to 0$$

after defining $Z_* := \bigoplus_{n \in \mathbb{Z}} Z_n$ and $B_{*-1} := \bigoplus_{n \in \mathbb{Z}} B_{n-1}$ as chain complexes with trivial boundary maps. In other words, we have a commuting diagram

$$\begin{tikzcd}
 & \vdots \arrow[d] & \vdots \arrow[d] & \vdots \arrow[d] & \\
0 \arrow[r] & Z_n \arrow[r, hook] \arrow[d, "0"] & C_n \arrow[r, "\partial_n"] \arrow[d, "\partial_n"] & B_{n-1} \arrow[r] \arrow[d, "0"] & 0 \\
0 \arrow[r] & Z_{n-1} \arrow[r, hook] \arrow[d, "0"] & C_{n-1} \arrow[r, "\partial_{n-1}"] \arrow[d, "\partial_{n-1}"] & B_{n-2} \arrow[r] \arrow[d, "0"] & 0 \\
 & \vdots & \vdots & \vdots &
\end{tikzcd}$$

in which each row is a short exact sequence. Since the modules B_n are all free, each row of this diagram is in fact a *split* exact sequence by Exercise 29.1,[67] and is thus isomorphic to an exact sequence of the form

$$0 \to Z_n \hookrightarrow Z_n \oplus B_{n-1} \to B_{n-1} \to 0,$$

with $Z_n \hookrightarrow Z_n \oplus B_{n-1}$ and $Z_n \oplus B_{n-1} \to B_{n-1}$ being the obvious inclusion and projection respectively. Applying the functor $\otimes G$ to such a sequence produces another split exact sequence

$$0 \to Z_n \otimes G \hookrightarrow (Z_n \oplus B_{n-1}) \otimes G = (Z_n \otimes G) \oplus (B_{n-1} \otimes G) \to B_{n-1} \otimes G \to 0,$$

so it follows that we can also apply this functor to the entirety of the diagram above and obtain another short exact sequence of chain complexes

$$0 \longrightarrow Z_* \otimes G \longrightarrow C_* \otimes G \longrightarrow B_{*-1} \otimes G \longrightarrow 0.$$

By the usual diagram-chasing result,[68] this gives rise to a long exact sequence of the homology groups of those complexes:

$$\ldots \longrightarrow B_n \otimes G \xrightarrow{\Phi} Z_n \otimes G \longrightarrow H_n(C_* \otimes G) \longrightarrow B_{n-1} \otimes G \xrightarrow{\Phi} Z_{n-1} \otimes G \longrightarrow \ldots,$$

where Φ denotes the connecting homomorphisms in this long exact sequence. If you look closely at the diagram chase required for constructing Φ, you'll find that it can be described explicitly: the map $\Phi : B_n \otimes G \to Z_n \otimes G$ for each $n \in \mathbb{Z}$ is just $i_n \otimes \mathbb{1}$, where

$$i_n : B_n \hookrightarrow Z_n$$

denotes the obvious inclusion. Now we use the standard trick for turning a long exact sequence into a short exact sequence centered on a certain term, in this case $H_n(C_* \otimes G)$: the map from this to $B_{n-1} \otimes G$ is surjective onto the kernel of Φ, while the map from $Z_n \otimes G$ preceding this descends to an injection on the quotient of $Z_n \otimes G$ by $\operatorname{im} \Phi$, giving a short exact sequence

$$0 \longrightarrow \operatorname{coker}(i_n \otimes \mathbb{1}) \longrightarrow H_n(C_* \otimes G) \longrightarrow \ker(i_{n-1} \otimes \mathbb{1}) \longrightarrow 0. \tag{42.5}$$

In order to proceed further, we need a more concrete understanding of the kernel and cokernel of the map $i_n \otimes \mathbb{1} : B_n \otimes G \to Z_n \otimes G$. To this end, we consider a second short exact sequence

$$0 \longrightarrow B_n \xhookrightarrow{i_n} Z_n \xrightarrow{\mathrm{pr}} H_n(C_*) \longrightarrow 0,$$

where $\mathrm{pr} : Z_n \to H_n(C_*) = Z_n/B_n$ denotes the quotient projection. We can act on this sequence with the functor $\otimes G$ to produce a sequence of R-module homomorphisms

$$0 \longrightarrow B_n \otimes G \xrightarrow{i_n \otimes \mathbb{1}} Z_n \otimes G \xrightarrow{\mathrm{pr} \otimes \mathbb{1}} H_n(C_*) \otimes G \longrightarrow 0, \tag{42.6}$$

but there is a difference between this and the previous time we applied $\otimes G$ to a short exact sequence: since $H_n(C_*)$ might not be free, we cannot assume that the sequence $0 \to B_n \to Z_n \to H_n(C_*) \to 0$ splits, and this makes it much less obvious whether we can also expect the sequence in (42.6) to be exact. If it is, then two useful consequences follow: first, $\mathrm{pr} \otimes \mathbb{1}$ is surjective and has kernel equal to the image of $i_n \otimes \mathbb{1}$, in which case letting $\mathrm{pr} \otimes \mathbb{1}$ descend to the quotient by its kernel gives an isomorphism

$$\operatorname{coker}(i_n \otimes \mathbb{1}) = (Z_n \otimes G)\big/\operatorname{im}(i_n \otimes \mathbb{1}) \xrightarrow{\cong} H_n(C_*) \otimes G. \tag{42.7}$$

[67]Strictly speaking, Exercise 29.1 was a result about short exact sequences of abelian groups, rather than modules over an arbitrary commutative ring R, but generalizing it from Ab to Mod^R requires no meaningful modifications to the proof.

[68]The result we are referring to here is Proposition 28.18, another result about abelian groups that generalizes in a completely straightforward manner to modules over arbitrary commutative rings, with no meaningful change in the proof.

If we use this isomorphism to replace the first term in (42.5) with $H_n(C_*)\otimes G$, one can check that the map $H_n(C_*)\otimes G \to H_n(C_*\otimes G)$ in that sequence then becomes (42.4), which is therefore injective, and our exact sequence now takes the form

$$0 \longrightarrow H_n(C_*)\otimes G \xrightarrow{h} H_n(C_*\otimes G) \longrightarrow \ker(i_{n-1}\otimes \mathbb{1}) \longrightarrow 0, \tag{42.8}$$

Strictly speaking, this conclusion depends on the exactness of (42.6) only at its second and third nontrivial terms, and we will see in Proposition 42.15 below that, in fact, this always holds.

Now, if (42.6) is also exact at its first term, or rather the analogous statement after replacing n with $n-1$, then $i_{n-1}\otimes\mathbb{1}$ is an injective map, and the third term in (42.8) therefore vanishes, implying that $h : H_n(C_*)\otimes G \to H_n(C_*\otimes G)$ is an isomorphism. This, we will see, is sometimes true but not always: in general, (42.6) may fail to be an exact sequence, because feeding an injective map $i_n : B_n \hookrightarrow Z_n$ into the functor $\otimes G$ can produce a non-injective map $i_n\otimes\mathbb{1} : B_n\otimes G \to Z_n\otimes G$. It turns out however that up to natural isomorphisms, the kernel of $i_n\otimes\mathbb{1}$ depends only on $H_n(C_*)$ and G, because it can be identified with a standard construction in homological algebra, the so-called "Tor" functor,

$$\operatorname{Tor}(H_n(C_*), G) \cong \ker(i_n\otimes\mathbb{1}) \subset B_n\otimes G.$$

Putting all of this together will give the following as our main result.

THEOREM 42.2 (universal coefficient theorem). *For any commutative ring R with unit, there exists a functor*

$$\operatorname{Tor} : \mathsf{Mod}^R \times \mathsf{Mod}^R \to \mathsf{Mod}^R,$$

covariant in both variables, such that the following holds whenever R is a principal ideal domain (e.g. in the case $R = \mathbb{Z}$). For any chain complex C_* of free R-modules, any fixed R-module G and any $n \in \mathbb{Z}$, there exists a split exact sequence

$$0 \longrightarrow H_n(C_*)\otimes G \xrightarrow{h} H_n(C_*\otimes G) \longrightarrow \operatorname{Tor}(H_{n-1}(C_*), G) \longrightarrow 0,$$

where h is the map in (42.4). Moreover, the sequence (but not its splitting) is natural in the sense that for any chain map of $\Phi : A_* \to B_*$ between two chain complexes of free R-modules, the diagram

$$\begin{tikzcd}
0 \arrow[r] & H_n(A_*)\otimes G \arrow[r, "h"] \arrow[d, "\Phi_*\otimes\mathbb{1}"] & H_n(A_*\otimes G) \arrow[r] \arrow[d, "(\Phi\otimes\mathbb{1})_*"] & \operatorname{Tor}(H_{n-1}(A_*), G) \arrow[r] \arrow[d] & 0 \\
0 \arrow[r] & H_n(B_*)\otimes G \arrow[r, "h"] & H_n(B_*\otimes G) \arrow[r] & \operatorname{Tor}(H_{n-1}(B_*), G) \arrow[r] & 0
\end{tikzcd}$$

commutes, where $\operatorname{Tor}(H_{n-1}(A_*), G) \to \operatorname{Tor}(H_{n-1}(B_*), G)$ is the map induced by $\Phi_* : H_{n-1}(A_*) \to H_{n-1}(B_*)$ via the functoriality of Tor.

The splitting of the exact sequence means that there is always an R-module isomorphism

$$H_n(C_*\otimes G) \cong (H_n(C_*)\otimes G) \oplus \operatorname{Tor}(H_{n-1}(C_*), G),$$

so that is the sense in which $H_n(C_*)$ and $H_{n-1}(C_*)$ determine $H_n(C_*\otimes G)$. It should be emphasized however that this isomorphism depends in general on non-canonical choices, and does not fit into any nice commutative diagrams together with maps induced by chain maps—this is what is meant when we say that the splitting is "not natural". In any case, the formula is obviously not that useful unless one also has a means of computing $\operatorname{Tor}(H, G)$, at least for instance when $R = \mathbb{Z}$ and H, G are finitely-generated abelian groups. Such computations will turn out to be quite manageable, but they merit a slightly longer discussion, so we will come back to that subject in the next lecture and focus for now on filling in the gaps in both the statement and the proof of Theorem 42.2. A major portion of the proof was carried out already in the discussion above, but the following issues remain to be clarified:

(1) Why the sequence in (42.6) is exact at its last two nontrivial terms, though possibly not at the term that precedes them;
(2) The general definition of $\mathrm{Tor}(G, H)$, and why $\mathrm{Tor}(H_n(C_*), G)$ is the same thing as the kernel of the map $i_n \otimes \mathbb{1}$ in (42.6);
(3) Why the exact sequence in Theorem 42.2 splits;
(4) The naturality of the exact sequence.

As is often the case, the naturality of the sequence amounts to checking that certain diagrams that one would obviously hope should commute really do, and it will not be difficult to check this once all the other details are in place, so I will leave this part as an exercise. The reason why the sequence splits is somewhat less obvious, so let's address that right now.

Proof that the sequence in Theorem 42.2 splits. By Exercise 29.1, it will suffice to construct a left-inverse for the injective map $h : H_n(C_*) \otimes G \to H_n(C_* \otimes G)$. Recall first that the exact sequence $0 \to Z_n \hookrightarrow C_n \xrightarrow{\partial_n} B_{n-1} \to 0$ splits since B_{n-1} is free due to Proposition 42.1, thus there exists a left-inverse of the inclusion $Z_n \hookrightarrow C_n$, i.e. a projection homomorphism

$$p : C_n \to Z_n, \qquad p|_{Z_n} = \mathbb{1}_{Z_n}.$$

The composition of p with the quotient projection $\mathrm{pr} : Z_n \to Z_n/B_n = H_n(C_*)$ then defines a chain map

$$(C_*, \partial) \xrightarrow{\mathrm{pr} \circ p} (H_*(C_*), 0)$$

since $\mathrm{pr} \circ p \circ \partial = 0$, and this induces a chain map

$$(C_* \otimes G, \partial \otimes \mathbb{1}) \xrightarrow{(\mathrm{pr} \circ p) \otimes \mathbb{1}} (H_*(C_*) \otimes G, 0),$$

which then descends to a homomorphism on the homologies that we shall call

$$\pi := ((\mathrm{pr} \circ p) \otimes \mathbb{1})_* : H_n(C_* \otimes G) \to H_n(C_*) \otimes G$$

for each $n \in \mathbb{Z}$. One can now check that $\pi \circ h$ is the identity map on $H_n(C_*) \otimes G$. $\square$

The remainder of this lecture will be devoted to the first two points in the list above: the (non-)exactness of the sequence (42.6), and the definition of the Tor functor.

Remark 42.3. While it is conventional to call Tor a "functor," we will see from the definition that this is cheating a little bit, as one cannot simply feed a pair of modules G and H into Tor and extract a well-defined module $\mathrm{Tor}(G, H)$. In reality, the definition of $\mathrm{Tor}(G, H)$ requires some auxiliary choices beyond G and H, but we will see that for any two sets of these choices, there is a canonical isomorphism between our two definitions of $\mathrm{Tor}(G, H)$. In this sense, $\mathrm{Tor}(G, H)$ is well defined in the same sense that "the one-point space" is well defined: one can take any set of one element and label it "$\{\mathrm{pt}\}$," and any two spaces defined in this way are not technically the same, but there is a canonical homeomorphism between them.

Remark 42.4. For applications to singular or cellular homology, I cannot immediately think of a situation in which it is useful to make any choice other than $R = \mathbb{Z}$ in Theorem 42.2. However, we will soon prove a similar theorem involving product chain complexes, as well as a cohomological version of Theorem 42.2, and for each of these, having the freedom to choose R to be a field $\mathbb{K}$ is often quite useful. Moreover, much of the following digression into homological algebra seems most natural in the setting of modules over an arbitrary commutative ring, which need not even be a principal ideal domain in general—setting R equal to $\mathbb{Z}$ could potentially obscure some algebraic aspects of the story that are actually useful to understand.

Left- and right-exact functors. The argument of the previous subsection shows that the canonical map $h : H_n(C_*) \otimes G \xrightarrow{h} H_n(C_* \otimes G)$ appearing in the universal coefficient theorem is an isomorphism whenever the sequence in (42.6) is exact. The latter is obtained by feeding an actual exact sequence $0 \to B_n \to Z_n \to H_n(C_*) \to 0$ into the functor

$$\otimes G : \mathsf{Mod}^R \to \mathsf{Mod}^R : A \mapsto A \otimes G,$$

so the subtle aspects of the universal coefficient theorem arise from the fact that this particular functor does not always preserve the exactness of sequences. We will now introduce some notions from homological algebra to help understand this phenomenon.

REMARK 42.5. Several of the definitions in this section can be extended to the more general context of functors $\mathcal{F} : \mathcal{A} \to \mathcal{B}$ between a pair of arbitrary *abelian categories*. Philosophically, abelian categories are the correct setting in which to study exact sequences and diagram-chasing arguments, though for our present purposes, working in that general setting would impose an unnecesarily extreme level of abstraction. We will therefore keep things more concrete by working only in the specific abelian category Mod^R and considering only functors from Mod^R to itself.

DEFINITION 42.6. A covariant functor $\mathcal{F} : \mathsf{Mod}^R \to \mathsf{Mod}^R$ is called **additive** if for every pair of R-modules A and B, the map $\mathcal{F} : \operatorname{Hom}(A, B) \to \operatorname{Hom}(\mathcal{F}(A), \mathcal{F}(B))$ is a homomorphism of abelian groups.

Additive functors have two crucial properties that we will need to make use of. First, if $A \xrightarrow{f} B \xrightarrow{g} C$ are two homomorphisms whose composition $g \circ f : A \to C$ vanishes, then the fact that $\mathcal{F}$ defines a group homomorphism $\operatorname{Hom}(A, C) \to \operatorname{Hom}(\mathcal{F}(A), \mathcal{F}(C))$ implies

$$\mathcal{F}(g) \circ \mathcal{F}(f) = \mathcal{F}(g \circ f) = 0 \in \operatorname{Hom}(\mathcal{F}(A), \mathcal{F}(C)),$$

simply because group homomorphisms preserve the 0 element. It follows that additive functors preserve the property of being a chain complex, i.e. any chain complex

$$\ldots \longrightarrow A_{n+1} \xrightarrow{f_{n+1}} A_n \xrightarrow{f_n} A_{n-1} \longrightarrow \ldots$$

gives rise to a chain complex

$$\ldots \longrightarrow \mathcal{F}(A_{n+1}) \xrightarrow{\mathcal{F}(f_{n+1})} \mathcal{F}(A_n) \xrightarrow{\mathcal{F}(f_n)} \mathcal{F}(A_{n-1}) \longrightarrow \ldots .$$

The second important property is that additive functors preserve finite direct sums: in particular, for every pair of R-modules A and B, one obtains a natural isomorphism $\mathcal{F}(A \oplus B) \cong \mathcal{F}(A) \oplus \mathcal{F}(B)$. To understand why, we can write down a relation between certain maps that characterizes direct sums uniquely up to isomorphism. Indeed, consider the canonical inclusion and projection maps

(42.9)
$$\begin{tikzcd} A \arrow[dr, "i_A"] & & A \\ & A \oplus B \arrow[ur, "\pi_A"] \arrow[dr, "\pi_B"] & \\ B \arrow[ur, "i_B"] & & B \end{tikzcd}$$

i.e. $i_A(a) = (a, 0)$, $\pi_A(a, b) = a$ and so forth. Abbreviating $C := A \oplus B$, these maps satisfy the five relations

(42.10) $$\pi_A i_A = \mathbb{1}_A, \quad \pi_B i_B = \mathbb{1}_B, \quad \pi_A i_B = 0, \quad \pi_B i_A = 0, \quad i_A \pi_A + i_B \pi_B = \mathbb{1}_C.$$

Conversely, suppose C is some other R-module, and that we are given four homomorphisms

$$\begin{tikzcd} A \arrow[dr, "i_A"] & & A \\ & C \arrow[ur, "\pi_A"] \arrow[dr, "\pi_B"] & \\ B \arrow[ur, "i_B"] & & B \end{tikzcd}$$

that likewise satisfy the five relations in (42.10). It then follows that the maps

$$A \oplus B \xrightarrow{i_A \oplus i_B} C \qquad \text{and} \qquad C \xrightarrow{(\pi_A, \pi_B)} A \oplus B$$

are inverse to each other, and thus determine an isomorphism $C \cong A \oplus B$ that identifies the maps i_A, i_B and π_A, π_B with the canonical inclusions and projections respectively. The point is this: if $\mathcal{F}$ is an additive functor, then plugging the diagram (42.9) into $\mathcal{F}$ gives us four maps

$$\begin{tikzcd} \mathcal{F}(A) \arrow[dr, "\mathcal{F}(i_A)"] & & \mathcal{F}(A) \\ & \mathcal{F}(A \oplus B) \arrow[ur, "\mathcal{F}(\pi_A)"] \arrow[dr, "\mathcal{F}(\pi_B)"] & \\ \mathcal{F}(B) \arrow[ur, "\mathcal{F}(i_B)"] & & \mathcal{F}(B) \end{tikzcd}$$

that similarly satisfy the five relations

$$\begin{aligned} &\mathcal{F}(\pi_A)\mathcal{F}(i_A) = \mathbb{1}_{\mathcal{F}(A)}, \qquad \mathcal{F}(\pi_B)\mathcal{F}(i_B) = \mathbb{1}_{\mathcal{F}(B)}, \\ &\mathcal{F}(\pi_A)\mathcal{F}(i_B) = 0, \qquad \mathcal{F}(\pi_B)\mathcal{F}(i_A) = 0, \\ &\mathcal{F}(i_A)\mathcal{F}(\pi_A) + \mathcal{F}(i_B)\mathcal{F}(\pi_B) = \mathbb{1}_C, \end{aligned}$$

where C now denotes $\mathcal{F}(A \oplus B)$, and the last three relations all depend crucially on the fact that $\mathcal{F}$ preserves addition of morphisms. One therefore obtains an isomorphism

$$\mathcal{F}(A \oplus B) \cong \mathcal{F}(A) \oplus \mathcal{F}(B)$$

that identifies $\mathcal{F}(i_A), \mathcal{F}(i_B)$ and $\mathcal{F}(\pi_A), \mathcal{F}(\pi_B)$ with the canonical inclusions and projections respectively.

The main example of an additive functor you should have in mind at the moment is $\otimes G$ for any fixed R-module G, and there is indeed a natural isomorphism

$$(A \oplus B) \otimes G \cong (A \otimes G) \oplus (B \otimes G)$$

for any pair of R-modules A, B, which you can convince yourself is also characterized by inclusion and projection maps as described above. The word "natural" in this context can be given a precise meaning in terms of commuting diagrams; I will leave it as an exercise to spell out the details.

Our proof of the universal coefficient theorem made use of the fact that for every R-module G, the functor $\otimes G$ has the following property, which can now be understood as a consequence of the fact that additive functors preserve finite direct sums:

LEMMA 42.7. *For any split exact sequence of R-modules $0 \to A \to B \to C \to 0$ and any additive covariant functor $\mathcal{F} : \mathsf{Mod}^R \to \mathsf{Mod}^R$, the induced sequence $0 \to \mathcal{F}(A) \to \mathcal{F}(B) \to \mathcal{F}(C) \to 0$ is split exact.* □

This lemma can be applied in particular whenever the last term C in the original sequence is free, because one can then use a basis to write down a right-inverse $C \to B$ of the surjective map $B \to C$, and produce from this a splitting via Exercise 29.1. Our proof of the universal coefficient theorem used the lemma in precisely this way, taking advantage of the fact that $B_{n-1} \subset C_{n-1}$ was a submodule of a free module over a principal ideal domain, and therefore (by Prop. 42.1) also free.

It will be useful to observe that this trick also works under a somewhat weaker assumption than freeness. As you've seen in covering space theory, it is often useful to be able to recognize when "lifting" problems can be solved, and the next definition does something similar in algebraic settings.

Definition 42.8. An R-module G is called **projective** if for every surjective R-module homomorphism $\pi : B \to A$, every R-module homomorphism $\varphi : G \to A$ can be lifted to an R-module homomorphism $\tilde{\varphi} : G \to B$ so that the following diagram commutes:

$$\begin{tikzcd} & B \arrow[d, two heads, "\pi"] \\ G \arrow[ur, dashed, "\tilde{\varphi}"] \arrow[r, "\varphi"] & A \end{tikzcd}$$

Example 42.9. Every free R-module is also projective. Indeed, if G has a basis $\mathcal{B} \subset G$, then the required lift $\tilde{\varphi} : G \to B$ can be defined by choosing any $\tilde{\varphi}(e) \in \pi^{-1}(\varphi(e))$ for each $e \in \mathcal{B}$ and extending $\tilde{\varphi}$ to the unique homomorphism with these values on the basis elements.

Example 42.10. The abelian group $\mathbb{Z}_2$ is not a projective $\mathbb{Z}$-module. For instance, the lift in the diagram

$$\begin{tikzcd} & \mathbb{Z} \arrow[d, two heads, "\pi"] \\ \mathbb{Z}_2 \arrow[ur, dashed, "\tilde{\varphi}"] \arrow[r, "\mathbb{1}"] & \mathbb{Z}_2 \end{tikzcd}$$

can never exist since $\operatorname{Hom}(\mathbb{Z}_2, \mathbb{Z}) = 0$.

Now suppose $0 \to A \xrightarrow{i} B \xrightarrow{j} C \to 0$ is a short exact sequence of R-modules and, instead of assuming C is free, suppose it is projective. The identity map $C \to C$ then admits a lift $\varphi : C \to B$ satisfying $j \circ \varphi = \mathbb{1}$, so φ is a right-inverse of j. Using Exercise 29.1 and Lemma 42.7, we conclude:

Proposition 42.11. *If C is a projective R-module and $0 \to A \to B \to C \to 0$ is a short exact sequence, then for any additive functor $\mathcal{F} : \mathsf{Mod}^R \to \mathsf{Mod}^R$, the induced sequence $0 \to \mathcal{F}(A) \to \mathcal{F}(B) \to \mathcal{F}(C) \to 0$ is also split exact.* □

Now let's consider the question from a different angle: which additive functors can be shown to preserve exactness for *all* short exact sequences, even those that don't split? This would be the right moment to clarify a piece of terminology that we have not had occasion to use before in precisely this form: one sometimes needs to consider sequences of the form

$$A_0 \to \ldots \to A_n \quad \text{or} \quad A_0 \to A_1 \to A_2 \to \ldots \quad \text{or} \quad \ldots \to A_{-2} \to A_{-1} \to A_0,$$

which have a finite end point at one end or both. Calling such a sequence **exact** is meant to be a condition matching kernels and images at every term that has an arrow on *both sides*, while not imposing any condition on terms with only one arrow. So for instance, calling $0 \to A \xrightarrow{f} B$ exact means that $f : A \to B$ is injective (but not necessarily surjective), and calling $A \xrightarrow{f} B \xrightarrow{g} C$ exact means $\operatorname{im}(f) = \ker(g)$, but with no requirement that f should be injective or g surjective.

Definition 42.12. An additive covariant functor $\mathcal{F} : \mathsf{Mod}^R \to \mathsf{Mod}^R$ is called **right-exact** if for every exact sequence $A \xrightarrow{i} B \xrightarrow{j} C \to 0$, the sequence

$$\mathcal{F}(A) \xrightarrow{\mathcal{F}(i)} \mathcal{F}(B) \xrightarrow{\mathcal{F}(j)} \mathcal{F}(C) \longrightarrow 0$$

is also exact. Similarly, $\mathcal{F}$ is **left-exact** if for every exact sequence $0 \longrightarrow A \xrightarrow{i} B \xrightarrow{j} C$ the sequence

$$0 \longrightarrow \mathcal{F}(A) \xrightarrow{\mathcal{F}(i)} \mathcal{F}(B) \xrightarrow{\mathcal{F}(j)} \mathcal{F}(C)$$

is also exact, and $\mathcal{F}$ is **exact** if it is both left-exact and right-exact.

Exercise 42.13. Show that $\mathcal{F}$ is exact if and only if it preserves exactness for *all* (not just short) exact sequences. Equivalently: $\mathcal{F}$ is exact if and only if for every exact sequence $A \to B \to C$, the induced sequence $\mathcal{F}(A) \to \mathcal{F}(B) \to \mathcal{F}(C)$ is also exact.
Hint: First show that if $\mathcal{F}$ is exact, then it preserves injectivity or surjectivity of morphisms. Then replace any given exact sequence $A \xrightarrow{f} B \xrightarrow{g} C$ with one that takes the form $0 \to A/\ker(f) \to B \to \mathrm{im}(g) \to 0$.

The reason to distinguish between right- and left-exact functors is that there are important examples of functors that are one but not the other, and $\otimes G$ in particular is one of these. The following shows that it is not exact in general.

Example 42.14. Let $R := \mathbb{Z}$, so R-modules are abelian groups. Feeding the exact sequence

$$0 \longrightarrow \mathbb{Z} \xrightarrow{\cdot 2} \mathbb{Z} \xrightarrow{\mathrm{pr}} \mathbb{Z}_2 \to 0$$

into the functor $\otimes \mathbb{Z}_2$ then gives

$$0 \longrightarrow \mathbb{Z} \otimes \mathbb{Z}_2 = \mathbb{Z}_2 \xrightarrow{\cdot 2} \mathbb{Z} \otimes \mathbb{Z}_2 = \mathbb{Z}_2 \xrightarrow{\mathrm{pr} \otimes \mathbb{1}} \mathbb{Z}_2 \otimes \mathbb{Z}_2 \to 0,$$

which is not exact since multiplication by 2 is the trivial map on $\mathbb{Z} \otimes \mathbb{Z}_2 = \mathbb{Z}_2$.

On the other hand:

Proposition 42.15. *For any R-module G, the functor $\otimes G : \mathsf{Mod}^R \to \mathsf{Mod}^R$ is right-exact.*

Proof. Exactness of a sequence $A \xrightarrow{i} B \xrightarrow{j} C \to 0$ means that j is surjective and $\mathrm{im}\, i = \ker j \subset B$. Then for every $c \in C$ and $g \in G$, we can find $b \in B$ with $j(b) = c$ and write $(j \otimes \mathbb{1})(b \otimes g) = c \otimes g$, implying that $j \otimes \mathbb{1}$ is surjective. Clearly $(j \otimes \mathbb{1}) \circ (i \otimes \mathbb{1}) = (j \circ i) \otimes \mathbb{1} = 0$ since $j \circ i = 0$, so we only still need to show $\ker(j \otimes \mathbb{1}) \subset \mathrm{im}(i \otimes \mathbb{1})$, which is now equivalent to showing that the map

$$(B \otimes G)/\,\mathrm{im}(i \otimes \mathbb{1}) \to C \otimes G : [b \times g] \mapsto j(b) \otimes g \tag{42.11}$$

is injective (in which case it is an isomorphism). We can do this by constructing its inverse: define $\Phi : C \otimes G \to (B \otimes G)/\,\mathrm{im}(i \otimes \mathbb{1})$ by $\Phi(c \otimes g) := b \otimes g$ for any $b \in j^{-1}(c)$. This is well defined because for any two choices $b, b' \in j^{-1}(c)$, we have $b' - b \in \ker j = \mathrm{im}\, i$ and thus $b' - b = i(a)$ for some $a \in A$, thus $b' \otimes g - b \otimes g = i(a) \otimes g \in \mathrm{im}(i \otimes \mathbb{1})$. It is easy to check that Φ is an inverse for the map (42.11). □

Proposition 42.15 explains the reason for the isomorphism (42.7), which identifies the term $H_n(C_*) \otimes G$ appearing in the universal coefficient theorem with the cokernel of the map $i_n \otimes \mathbb{1} : B_n \otimes G \to Z_n \otimes G$ induced by the inclusion i_n of the group of n-boundaries B_n into the group of n-cycles $Z_n \subset C_n$. What we still need and do not yet have is a practical description of $\ker(i_{n-1} \otimes \mathbb{1})$, which appeared as the third term in a natural short exact sequence preceded by $0 \longrightarrow H_n(C_*) \otimes G \xrightarrow{h} H_n(C_* \otimes G)$. Since $\otimes G$ is a right-exact functor, $\ker(i_{n-1} \otimes \mathbb{1})$ can be interpreted as a measurement of its failure to be exact: it vanishes if and only if the sequence $0 \to B_{n-1} \to Z_{n-1} \to H_{n-1}(C_*) \to 0$ remains exact after feeding it into $\otimes G$.

DEFINITION 42.16. An R-module G is **flat** if for every pair of R-modules A, B with an injective R-module homomorphism $i : A \hookrightarrow B$, the homomorphism $i \otimes \mathbb{1} : A \otimes G \to B \otimes G$ is also injective. Equivalently: G is flat if the functor $\otimes G : \mathsf{Mod}^R \to \mathsf{Mod}^R$ is exact.

Example 42.14 shows that not all R-modules are flat; in particular, $\mathbb{Z}_2$ is not a flat $\mathbb{Z}$-module. Let us collect some simple tools for recognizing flatness.

LEMMA 42.17. *For any collection $\{G_\alpha\}_{\alpha \in J}$ of R-modules, $G := \bigoplus_{\alpha \in J} G_\alpha$ is flat if and only if G_α is flat for every $\alpha \in J$.*

PROOF. Given an injective homomorphism $i : A \to B$, tensoring with the direct sum G produces a "diagonal" homomorphism

$$A \otimes \left(\bigoplus_{\alpha \in J} G_\alpha \right) = \bigoplus_{\alpha \in J} (A \otimes G_\alpha) \xrightarrow{\bigoplus_\alpha i \otimes \mathbb{1}} \bigoplus_{\alpha \in J} (B \otimes G_\alpha) = B \otimes \left(\bigoplus_{\alpha \in J} G_\alpha \right)$$

which is injective if and only if all of its diagonal entries $i \otimes \mathbb{1} : A \otimes G_\alpha \to B \otimes G_\alpha$ are injective. $\square$

LEMMA 42.18. *Every free R-module is flat.*

PROOF. Free modules are isomorphic to direct sums of copies of R, so thanks to Lemma 42.17, this follows from the trivial observation that R itself is a flat R-module. $\square$

LEMMA 42.19. *Every projective R-module is a direct summand of a free R-module.*

PROOF. Given an R-module G, choose any free R-module F that admits a surjective homomorphism $\pi : F \to G$, e.g. one could take F to be the free R-module on the set of all elements of G, with π as the unique homomorphism determined by the inclusions of those elements. We now have a short exact sequence

$$0 \to \ker \pi \hookrightarrow F \xrightarrow{\pi} G \longrightarrow 0,$$

and if G is projective, then Proposition 42.11 gives a splitting of this sequence, and therefore an isomorphism $F \cong G \oplus \ker \pi$. $\square$

COROLLARY 42.20. *All projective R-modules are flat.* $\square$

REMARK 42.21. If R is a principal ideal domain, then Lemma 42.19 and Proposition 42.1 together imply that projective R-modules and free R-modules are the same thing. But we will not need to make any concrete use of this fact, as projectivity on its own is already a quite useful condition.

Projective resolutions. By now we have some motivation to believe that there is something special about R-modules that are projective: in particular, Proposition 42.11 and Corollary 42.20 tells us that in our proof of the universal coefficient theorem, the mysterious third term that prevents the injection $h : H_n(C_*) \otimes G \to H_n(C_* \otimes G)$ from being surjective would have vanished if we could assume that either $H_{n-1}(C_*)$ or G is projective. Unfortunately, plenty of interesting R-modules (such as the abelian group $\mathbb{Z}_2$) are not projective, so we still need a better understanding of what the third term in the sequence can be in such cases.

As an attempt at motivating the next definition, here is an unnecessarily verbose way of restating the observation that everything is fine if $H_{n-1}(C_*)$ is projective: everything is fine if there exists an exact sequence

$$0 \to P \to H_{n-1}(C_*) \to 0$$

in which all terms to the left of $H_{n-1}(C_*)$ are projective. This is clear: an exact sequence of this form is equivalent to an isomorphism $P \to H_{n-1}(C_*)$, and P is then projective if and only if $H_{n-1}(C_*)$ is projective. The point of framing the issue in these terms is that in homological algebra, exact sequences can be regarded as generalizations of isomorphisms, but they exist in many situations where actual isomorphisms do not.

DEFINITION 42.22. A **projective resolution** $A_* \xrightarrow{\alpha} A$ of an R-module A consists of an exact sequence

$$\ldots \longrightarrow A_2 \xrightarrow{\alpha_2} A_1 \xrightarrow{\alpha_1} A_0 \xrightarrow{\alpha} A \longrightarrow 0$$

such that all of the A_n for $n = 0, 1, 2, \ldots$ are projective R-modules and the final nontrivial term is A.

The notation $A_* \xrightarrow{\alpha} A$ for a projective resolution makes sense from the following perspective. Let A_* denote the chain complex

$$\ldots \longrightarrow A_2 \xrightarrow{\alpha_2} A_1 \xrightarrow{\alpha_1} A_0 \longrightarrow A_{-1} := 0 \longrightarrow A_{-2} := 0 \longrightarrow \ldots$$

obtained by truncating the exact sequence after A_0, and identify A with the trivial chain complex that has A itself in degree 0 and the trivial module in all other degrees. The fact that α is surjective and satisfies $\alpha \circ \alpha_1 = 0$ then allows us to regard α as a surjective chain map $A_* \to A$, also known as an *augmentation* of the chain complex A_*, for which the resulting augmented chain complex

$$\widetilde{A}_* := \bigoplus_{n=-1}^{\infty} \widetilde{A}_n \qquad \text{with} \qquad \widetilde{A}_{-1} := A \quad \text{and} \quad \widetilde{A}_n := A_n \text{ for } n \geqslant 0$$

is the original exact sequence in the projective resolution (cf. Remark 29.14). In these terms, a projective resolution of A is the same thing as a chain complex A_* of projective modules that is trivial in all negative degrees and has trivial homology in all positive degrees, together with an augmentation $\alpha : A_* \to A$ for which the reduced homology $\widetilde{H}_*(A_*) := H_*(\widetilde{A}_*)$ is trivial.

PROPOSITION 42.23. *Every R-module A admits a projective resolution $A_* \xrightarrow{\alpha} A$. Moreover, if R is a principal ideal domain, the resolution can be chosen such that A_n is the trivial R-module for every $n \geqslant 2$.*

PROOF. Pick any generating set $S_0 \subset A$, i.e. a set such that every element of A can be written (perhaps non-uniquely) as $\sum_{e \in S_0} r_e e$ for some coefficients $r_e \in R$, at most finitely-many of which are nonzero. This can always be done, since e.g. it would suffice to choose $S_0 = A$, though smaller subsets are also possible. We then set $A_0 := F_R^{\mathrm{mod}}(S_0)$ and define $\alpha : A_0 \to A$ as the unique R-module homomorphism that extends the inclusion $S_0 \hookrightarrow A$, noting that α is surjective by construction. Next, pick S_1 to be a generating subset of $\ker \alpha \subset A_0$, and define $A_1 := F_R^{\mathrm{mod}}(S_1)$ and $\alpha_1 : A_1 \to \ker \alpha$ analogously; this defines $\alpha_1 : A_1 \to A_0$ such that $\operatorname{im} \alpha_1 = \ker \alpha$. Now continue this process inductively: all of the modules A_n produced in this way are free, and therefore projective.

If R is a principal ideal domain, then after defining $A_0 \xrightarrow{\alpha} A$ as described above, we can exploit Proposition 42.1 to conclude that $\ker \alpha \subset A_0$ is also a free R-module, and thus simplify the construction by defining $A_1 := \ker \alpha$ with $A_1 \xrightarrow{\alpha_1} A_0$ as the inclusion, and then set $A_n := 0$ for all $n \geqslant 2$. $\square$

There seem to be quite a lot of arbitrary choices involved in constructing projective resolutions, but the next result shows that they are more unique than one might expect.

PROPOSITION 42.24. *Given an R-module homomorphism $\varphi : A \to B$ and any projective resolutions $A_* \xrightarrow{\alpha} A$ of A and $B_* \xrightarrow{\beta} B$ of B, there exists a sequence of R-module homomorphisms $\varphi_n : A_n \to B_n$ for $n \geqslant 0$ which, together with $\varphi : A \to B$, form a chain map $\varphi_* : \widetilde{A}_* \to \widetilde{B}_*$ between* *the corresponding augmented chain complexes, i.e. the diagram*

$$
\begin{tikzcd}
\dots \arrow[r] & A_2 \arrow[r, "\alpha_2"] \arrow[d, "\varphi_2"] & A_1 \arrow[r, "\alpha_1"] \arrow[d, "\varphi_1"] & A_0 \arrow[r, "\alpha"] \arrow[d, "\varphi_0"] & A \arrow[d, "\varphi"] \\
\dots \arrow[r] & B_2 \arrow[r, "\beta_2"] & B_1 \arrow[r, "\beta_1"] & B_0 \arrow[r, "\beta"] & B
\end{tikzcd}
\tag{42.12}
$$

commutes. Moreover, this chain map is unique up to chain homotopy.

COROLLARY 42.25. *For any two projective resolutions $A_* \xrightarrow{\alpha} A$ and $A'_* \xrightarrow{\alpha'} A$ of the same R-module A, the chain complexes A_* and A'_* admit a chain homotopy equivalence $\varphi_* : A_* \to A'_*$ whose restriction $\varphi_0 : A_0 \to A'_0$ to degree 0 satisfies $\alpha' \circ \varphi_0 = \alpha$.*

PROOF. Proposition 42.24 can be applied with $\varphi : A \to A$ as the identity map to produce chain maps between A_* and A'_* in both directions, and uniqueness of up to chain homotopy then implies that both of their compositions are chain homotopic to the identity. $\square$

PROOF OF PROPOSITION 42.24. For convenience, we can treat A and B as the degree -1 terms in augmented chain complexes and thus write $A_{-1} := A$, $B_{-1} := B$, $\varphi_{-1} := \varphi$, $\alpha_0 := \alpha$ and $\beta_0 := \beta$. Arguing by induction, assume for some integer $k \geqslant 0$ that the maps $\varphi_{-1}, \dots, \varphi_{k-1}$ in (42.12) have already been constructed so that all the relevant squares commute. We must then find a map $\varphi_k : A_k \to B_k$ such that $\beta_k \varphi_k = \varphi_{k-1}\alpha_k$. Notice that

$$\beta_{k-1}\varphi_{k-1}\alpha_k = \varphi_{k-2}\alpha_{k-1}\alpha_k = 0,$$

thus $\operatorname{im}(\varphi_{k-1}\alpha_k) \subset \ker \beta_{k-1} = \operatorname{im} \beta_k$, and we can therefore define φ_k to be any solution to the lifting problem

$$
\begin{tikzcd}
 & B_k \arrow[d, two heads, "\beta_k"] \\
A_k \arrow[ur, dashed, "\varphi_k"] \arrow[r, "\varphi_{k-1}\alpha_k"'] & \ker \beta_{k-1}
\end{tikzcd}
$$

A solution exists since A_k is projective. The existence of the complete chain map φ_* now follows by induction on k.

For uniqueness, suppose φ_* and ψ_* are two chain maps as above, and we want to define a chain homotopy between them, i.e. a sequence of maps $h_k : A_k \to B_{k+1}$ for $k \geqslant 0$ satisfying

$$\varphi_k - \psi_k = \beta_{k+1}h_k + h_{k-1}\alpha_k$$

for every k. For this to make sense when $k = 0$, we need also a map $h_{-1} : A \to B_0$, which we define as $h_{-1} := 0$. Assume for some $k \geqslant 0$ that $h_{-1}, \dots, h_{k-1}$ have already been constructed, so we now need to find a map $h_k : A_k \to B_{k+1}$ such that

$$\beta_{k+1}h_k = \varphi_k - \psi_k - h_{k-1}\alpha_k.$$

We observe that by commutativity and the chain homotopy relation for $k-1$,

$$
\begin{aligned}
\beta_k(\varphi_k - \psi_k - h_{k-1}\alpha_k) &= (\varphi_{k-1} - \psi_{k-1})\alpha_k - \beta_k h_{k-1}\alpha_k \\
&= (\beta_k h_{k-1} + h_{k-2}\alpha_{k-1} - \beta_k h_{k-1})\,\alpha_k = h_{k-2}\alpha_{k-1}\alpha_k = 0,
\end{aligned}
$$

so $\operatorname{im}(\varphi_k - \psi_k - h_{k-1}\alpha_k) \subset \ker \beta_k = \operatorname{im} \beta_{k+1}$, and h_k can now be defined as any solution to the lifting problem

$$
\begin{tikzcd}
 & B_{k+1} \arrow[d, two heads, "\beta_{k+1}"] \\
A_k \arrow[ur, dashed, "h_k"] \arrow[r, "\varphi_k - \psi_k - h_{k-1}\alpha_k"'] & \ker \beta_k
\end{tikzcd}
$$

The result now follows again by induction on k. $\square$

Left derived functors. Your first instinct when you see a chain map like $\varphi_* : A_* \to B_*$ as in Proposition 42.24 might be to look at the homomorphisms it induces between the homologies of the two chain complexes, but that is not very interesting in this situation since by exactness, those homologies can only be nontrivial in degree 0, where the original exact sequences of the projective resolutions have been truncated. Something much more interesting happens, however, if we now feed those exact sequences into an additive functor

$$\mathcal{F} : \mathsf{Mod}^R \to \mathsf{Mod}^R$$

that is right-exact but not necessarily exact. Given any projective resolution $A_* \xrightarrow{\alpha} A$ of an R-module A, we can let $\mathcal{F}$ operate on the chain complex A_* to obtain a chain complex $\mathcal{F}(A_*)$ with terms

$$\ldots \longrightarrow \mathcal{F}(A_2) \xrightarrow{\mathcal{F}(\alpha_2)} \mathcal{F}(A_1) \xrightarrow{\mathcal{F}(\alpha_1)} \mathcal{F}(A_0) \longrightarrow 0 \longrightarrow 0 \longrightarrow \ldots,$$

along with a homomorphism

$$\mathcal{F}(A_0) \xrightarrow{\mathcal{F}(\alpha)} \mathcal{F}(A).$$

If $\mathcal{F}$ is an exact functor, then by Exercise 42.13, the sequence

$$\ldots \longrightarrow \mathcal{F}(A_2) \xrightarrow{\mathcal{F}(\alpha_2)} \mathcal{F}(A_1) \xrightarrow{\mathcal{F}(\alpha_1)} \mathcal{F}(A_0) \xrightarrow{\mathcal{F}(\alpha)} \mathcal{F}(A) \longrightarrow 0$$

will also be exact, implying that the homologies $H_n(\mathcal{F}(A_*))$ of the chain complex $\mathcal{F}(A_*)$ are all trivial for $n > 0$. The only nontrivial homology here will be in degree 0, where the exactness of $\mathcal{F}$ forces $\mathcal{F}(\alpha) : \mathcal{F}(A_0) \to \mathcal{F}(A)$ to be a surjective map that descends to an isomorphism

$$H_0(\mathcal{F}(A_*)) = \mathcal{F}(A_0)\big/\operatorname{im}(\mathcal{F}(\alpha_1)) = \mathcal{F}(A_0)\big/\ker(\mathcal{F}(\alpha)) \xrightarrow{\cong} \mathcal{F}(A). \tag{42.13}$$

In fact, the latter is still true if $\mathcal{F}$ is not exact but only right-exact, because $\mathcal{F}(A_1) \xrightarrow{\mathcal{F}(\alpha_1)} \mathcal{F}(A_0) \xrightarrow{\mathcal{F}(\alpha)} \mathcal{F}(A) \longrightarrow 0$ is then still an exact sequence. But if $\mathcal{F}$ is not left-exact, then we no longer have any reason to expect $H_n(\mathcal{F}(A_*))$ to be trivial for $n > 0$, and in fact, these homologies can be regarded as measurements of the failure of $\mathcal{F}$ to be an exact functor. The crucial consequence of Proposition 42.24 and Corollary 42.25 is now that, up to isomorphism, these homologies are all *independent of the choice of projective resolution*! Indeed, one can operate with $\mathcal{F}$ on the entirety of the diagram (42.12) to produce a diagram

$$\begin{tikzcd}
\ldots \arrow[r] & \mathcal{F}(A_2) \arrow[r, "\mathcal{F}(\alpha_2)"] \arrow[d, "\mathcal{F}(\varphi_2)"] & \mathcal{F}(A_1) \arrow[r, "\mathcal{F}(\alpha_1)"] \arrow[d, "\mathcal{F}(\varphi_1)"] & \mathcal{F}(A_0) \arrow[r, "\mathcal{F}(\alpha)"] \arrow[d, "\mathcal{F}(\varphi_0)"] & \mathcal{F}(A) \arrow[d, "\mathcal{F}(\varphi)"] \\
\ldots \arrow[r] & \mathcal{F}(B_2) \arrow[r, "\mathcal{F}(\beta_2)"] & \mathcal{F}(B_1) \arrow[r, "\mathcal{F}(\beta_1)"] & \mathcal{F}(B_0) \arrow[r, "\mathcal{F}(\beta)"] & \mathcal{F}(B)
\end{tikzcd}$$

that represents a chain map $\mathcal{F}(\varphi_*) : \mathcal{F}(A_*) \to \mathcal{F}(B_*)$ associated to any homomorphism $\varphi : A \to B$, and applying $\mathcal{F}$ also to the chain homotopy relation shows that such chain maps are similarly unique up to chain homotopy. Letting the chain maps descend to homology now associates to each $\varphi : A \to B$ a sequence of R-module homomorphisms

$$H_n(\mathcal{F}(A_*)) \xrightarrow{\varphi_*} H_n(\mathcal{F}(B_*)) \qquad \text{for each } n \geqslant 0,$$

which are defined independently of the choices of chain maps in Proposition 42.24, and at the degree 0 level, the relation $\varphi \circ \alpha = \beta \circ \varphi_0$ implies that they match $\mathcal{F}(\varphi) : \mathcal{F}(A) \to \mathcal{F}(B)$ under the natural isomorphisms $H_0(\mathcal{F}(A_*)) \cong \mathcal{F}(A)$ and $H_0(\mathcal{F}(B_*)) \cong \mathcal{F}(B)$, i.e. the diagram

$$\begin{tikzcd}
H_0(\mathcal{F}(A_*)) \arrow[r, "\cong"] \arrow[d, "\varphi_*"] & \mathcal{F}(A) \arrow[d, "\varphi"] \\
H_0(\mathcal{F}(B_*)) \arrow[r, "\cong"] & \mathcal{F}(B)
\end{tikzcd}$$

commutes. If we now apply this construction with $\varphi : A \to B$ as the identity map $A \to A$ but with two different choices of projective resolution $A_* \xrightarrow{\alpha} A$ and $A'_* \xrightarrow{\alpha'} A$, then the fact that it can be done in both directions implies that the maps on homology are isomorphisms

$$H_n(\mathcal{F}(A_*)) \xrightarrow{\cong} H_n(\mathcal{F}(A'_*)) \qquad \text{for each } n \geqslant 0,$$

which for $n = 0$ fit into the diagram

$$\begin{tikzcd}
H_0(\mathcal{F}(A_*)) \arrow[r, "\cong"] \arrow[d, "\cong"] & \mathcal{F}(A) \arrow[d, equal] \\
H_0(\mathcal{F}(A'_*)) \arrow[r, "\cong"] & \mathcal{F}(A)
\end{tikzcd}$$

This discussion justifies the following definition.

Definition 42.26. Fix a choice of projective resolution $A_* \xrightarrow{\alpha} A$ for each R-module A. Given a covariant right-exact functor $\mathcal{F} : \mathsf{Mod}^R \to \mathsf{Mod}^R$, the associated **left derived functors**

$$L_n\mathcal{F} : \mathsf{Mod}^R \to \mathsf{Mod}^R$$

are defined for each integer $n \geqslant 0$ by associating to each R-module A the R-module

$$L_n\mathcal{F}(A) := H_n(\mathcal{F}(A_*)),$$

and to each homomorphism $\varphi : A \to B$ the homomorphism

$$L_n\mathcal{F}(A) = H_n(\mathcal{F}(A_*)) \xrightarrow{\varphi_*} H_n(\mathcal{F}(B_*)) = L_n\mathcal{F}(B)$$

determined by the unique chain homotopy class of chain maps $A_* \to B_*$ provided by Proposition 42.24.

You would be justified in finding something slightly unsatisfactory about the way Definition 42.26 is stated: on the one hand, making arbitrary choices of projective resolutions in advance for every conceivable R-module requires an unnecessarily broad invocation of the axiom of choice, and it seems to make the definition of the functors $L_n\mathcal{F} : \mathsf{Mod}^R \to \mathsf{Mod}^R$ rather far from unique or canonical. On the other hand, Corollary 42.25 ensures that this ambiguity is never really going to matter, because while making different choices of projective resolution for a single module A leads technically to two different definitions of the module $L_n\mathcal{F}(A)$ for each $n \geqslant 0$, these two modules come equipped with a canonical isomorphism between them. In practice, one does not actually make choices of projective resolutions in advance; one typically rather finds that in whatever application one is interested in, particular projective resolutions arise naturally, and are therefore the most convenient choices to use.

Definition 42.27. For any R-module G and each integer $n \geqslant 0$, the functor

$$\operatorname{Tor}_n(\cdot, G) : \mathsf{Mod}^R \to \mathsf{Mod}^R : A \mapsto \operatorname{Tor}_n(A, G)$$

is defined as the left derived functor $\operatorname{Tor}_n(\cdot, G) := L_n\mathcal{F}$ associated to the right-exact functor $\mathcal{F} := \otimes G$. More explicitly,

$$\operatorname{Tor}_n(A, G) := H_n(A_* \otimes G)$$

for any R-module A with a choice of projective resolution $A_* \xrightarrow{\alpha} A$. The case $n = 1$ is often denoted simply by

$$\operatorname{Tor}(A, G) := \operatorname{Tor}_1(A, G).$$

In situations where it is important to specify the ring R, we will occasionally make this more explicit by writing[69]

$$\mathrm{Tor}_n^R(A,G) := \mathrm{Tor}_n(A,G), \qquad \mathrm{Tor}^R(A,G) := \mathrm{Tor}_1^R(A,G) := \mathrm{Tor}(A,G).$$

Wrapping up the universal coefficient theorem. We can now fill in the last missing detail in the proof of Theorem 42.2. The third term following $h : H_n(C_*) \otimes G \to H_n(C_* \otimes G)$ in the short exact sequence was the kernel of the map $i_{n-1} \otimes \mathbb{1} : B_{n-1} \otimes G \to Z_{n-1} \otimes G$ induced by the first nontrivial map in the exact sequence

$$\ldots \longrightarrow 0 \longrightarrow 0 \longrightarrow B_{n-1} \xrightarrow{i_n} Z_{n-1} \xrightarrow{\mathrm{pr}} H_{n-1}(C_*) \longrightarrow 0. \tag{42.14}$$

Since C_{n-1} is free and R is a principal ideal domain, the submodules $B_{n-1} \subset Z_{n-1} \subset C_{n-1}$ are also free by Proposition 42.1, and in particular, they are projective, giving the exact sequence (42.14) a new interpretation as a projective resolution of $H_{n-1}(C_*)$, with B_{n-1} as the degree 1 term and trivial terms in all degrees greater than 1. Truncating the sequence at $H_{n-1}(C_*)$ and then feeding it into $\otimes G$ produces the chain complex

$$\ldots \longrightarrow 0 \longrightarrow 0 \longrightarrow B_{n-1} \otimes G \xrightarrow{i_n \otimes \mathbb{1}} Z_{n-1} \otimes G \longrightarrow 0 \longrightarrow 0 \longrightarrow \ldots,$$

whose homology in degree 1 gives

$$\mathrm{Tor}(H_{n-1}(C_*), G) := \ker(i_{n-1} \otimes \mathbb{1}) \subset B_{n-1} \otimes G,$$

thus explaining the appearance of the Tor functor in Theorem 42.2. Notice that if we similarly use $0 \to B_n \hookrightarrow Z_n \to H_n(C_*) \to 0$ as a projective resolution of $H_n(C_*)$, then the first nontrivial term in our original short exact sequence $0 \to \mathrm{coker}(i_n \otimes \mathbb{1}) \to H_n(C_* \otimes G) \to \ker(i_{n-1} \otimes \mathbb{1}) \to 0$ also appears in this picture as the result of a derived functor: we have

$$H_n(C_*) \otimes G \cong \mathrm{Tor}_0(H_n(C_*), G) := \mathrm{coker}(i_n \otimes \mathbb{1}).$$

It remains to examine how $\mathrm{Tor}(H_{n-1}(C_*), G)$ can actually be computed in practice, and we will address this in the next lecture.

43. Properties of the Tor functor (December 15, 2023)

The Tor functor was introduced at the end of the previous lecture as a particular example of a so-called *left derived* functor, one whose purpose is to measure the failure of short exact sequences of modules to remain exact after tensoring them with an auxiliary "coefficient" module. According to the universal coefficient theorem, it is the Tor functor that determines whether the canonical injection $H_*(X, A; \mathbb{Z}) \otimes G \to H_*(X, A; G)$ in either singular or cellular homology will also be a surjection: the general picture is that for each $n \geqslant 0$ and every pair (X, A), there is a split exact sequence

$$0 \to H_n(X, A; \mathbb{Z}) \otimes G \to H_n(X, A; G) \to \mathrm{Tor}^{\mathbb{Z}}(H_{n-1}(X, A; \mathbb{Z}), G) \to 0,$$

thus giving an isomorphism

$$H_n(X, A; G) \cong (H_n(X, A; \mathbb{Z}) \otimes G) \oplus \mathrm{Tor}^{\mathbb{Z}}(H_{n-1}(X, A; \mathbb{Z}), G).$$

In order to make use of this, one must of course be able to compute the groups $\mathrm{Tor}^{\mathbb{Z}}(H, G)$ for a sufficiently wide range of abelian groups H and G, and that is our first order of business for today.

As in the previous lecture, here we shall fix a commutative ring R and work in the category Mod^R of R-modules. It will mostly not be necessary to assume that R is a principal ideal domain, but we will have a few useful things to say about special properties that hold if it is.

[69] Be aware that some authors may prefer a different notational convention in which Tor by default means $\mathrm{Tor}^{\mathbb{Z}}$.

Some easy computations of Tor. Let's begin with some properties of the functors Tor_n that follow from more general facts about left derived functors.

PROPOSITION 43.1. *The left derived functors $L_n\mathcal{F} : \mathsf{Mod}^R \to \mathsf{Mod}^R$ associated to a right-exact functor $\mathcal{F} : \mathsf{Mod}^R \to \mathsf{Mod}^R$ have the following properties for every R-module A:*

(1) There is a natural isomorphism $L_0\mathcal{F}(A) \cong \mathcal{F}(A)$.
(2) If $\mathcal{F}$ is exact or A is projective, then $L_n\mathcal{F}(A) = 0$ for every $n \geqslant 1$.
(3) If R is a principal ideal domain, then $L_n\mathcal{F}(A) = 0$ for every $n \geqslant 2$.

PROOF. As was observed already in (42.13), the assumption that $\mathcal{F}$ is right-exact implies that for any projective resolution $A_* \xrightarrow{\alpha} A$, the map $\mathcal{F}(\alpha) : \mathcal{F}(A_0) \to \mathcal{F}(A)$ is surjective and descends to the quotient of its kernel $\ker \mathcal{F}(\alpha) = \operatorname{im} \mathcal{F}(\alpha_1)$ as an isomorphism

$$L_0\mathcal{F}(A) = \mathcal{F}(A_0)\big/ \operatorname{im} \mathcal{F}(\alpha_1) \xrightarrow{\cong} \mathcal{F}(A).$$

If $\mathcal{F}$ is an exact functor, then $H_n(\mathcal{F}(A_*)) = 0$ for all $n > 0$ because by Exercise 42.13, the exact sequence in the projective resolution remains exact after feeding it into $\mathcal{F}$. If $\mathcal{F}$ is not assumed exact but A is projective, then we can take

$$\ldots \to 0 \to A_0 := A \xrightarrow{\alpha:=\mathbb{1}} A$$

as a projective resolution of A, in which case $H_n(\mathcal{F}(A_*))$ trivially vanishes for all $n \geqslant 1$. Finally, if R is a principal ideal domain, then the freedom mentioned in Proposition 42.23 to choose projective resolutions $A_* \xrightarrow{\alpha} A$ with $A_n = 0$ for all $n \geqslant 2$ similarly causes $H_n(\mathcal{F}(A_*))$ to vanish for all $n \geqslant 2$. □

Applying this result to the functor $\mathcal{F} = \otimes G$ for any G-module G, we obtain natural isomorphisms

$$\mathrm{Tor}_0(A, G) \cong A \otimes G$$

for every pair of R-modules A and G, and

$$\mathrm{Tor}_n(A, G) = 0 \qquad \text{whenever } n \geqslant 2 \text{ if } R \text{ is a principal ideal domain,}$$

which is of course the reason to single out the case $n = 1$ with the special notation $\mathrm{Tor}(A, G) = \mathrm{Tor}_1(A, G)$. We also have

$$\mathrm{Tor}(A, G) = 0 \qquad \text{whenever } A \text{ is projective or } G \text{ is flat,}$$

which applies in particular whenever either A or G is free. This holds because all left derived functors in degree $n \geqslant 1$ vanish on projective modules, while taking G to be flat is synonymous with $\otimes G$ being an exact functor; recall that by Corollary 42.20, this also holds whenever G is projective.

Only a little bit more than this is needed if we want to compute $\mathrm{Tor}^{\mathbb{Z}}(A, B)$ for an arbitrary pair of finitely-generated abelian groups A and B, and even in many cases where they need not be finitely generated.

EXERCISE 43.2. Show that if A and B are both projective R-modules, then $A \oplus B$ is also projective.
Remark: If you're already a fan of universal properties, you may want to try proving this statement without using the concrete definition of the direct sum of two R-modules, but instead using the fact that it is a coproduct in the category Mod^R (cf. Exercise 39.24). Using this language essentially makes the result valid not just for R-modules but in arbitrary abelian categories.

Exercise 43.3. Use an intelligent choice of projective resolution to show that for any right-exact functor $\mathcal{F} : \mathsf{Mod}^R \to \mathsf{Mod}^R$, there are natural isomorphisms $L_n\mathcal{F}(A \oplus B) \cong L_n\mathcal{F}(A) \oplus L_n\mathcal{F}(B)$ for every integer $n \geqslant 0$ and any two R-modules A and B. In particular, this implies the formula

$$\operatorname{Tor}_n(A \oplus B, G) \cong \operatorname{Tor}_n(A, G) \oplus \operatorname{Tor}_n(B, G)$$

for any three R-modules A, B, G.

Exercise 43.4. Prove that there are also natural isomorphisms

$$\operatorname{Tor}_n(A, G \oplus H) \cong \operatorname{Tor}_n(A, G) \oplus \operatorname{Tor}_n(A, H)$$

for any three R-modules A, G, H.

Exercise 43.5. Assume $k \in \mathbb{N}$ is any number with the property that no nonzero element $x \in R$ satisfies $kx = 0$. Use an intelligent choice of projective resolution for the quotient module R/kR to show that for any right-exact functor $\mathcal{F} : \mathsf{Mod}^R \to \mathsf{Mod}^R$, $L_n\mathcal{F}(R/kR) = 0$ for all $n \geqslant 2$ and

$$L_1\mathcal{F}(R/kR) \cong \ker\left(\mathcal{F}(R) \to \mathcal{F}(R) : x \mapsto kx\right).$$

In particular, this implies the formula

$$\operatorname{Tor}(R/kR, G) \cong \ker\left(G \xrightarrow{\cdot k} G\right)$$

for any R-module G.
Hint: Every additive functor $\mathcal{F} : \mathsf{Mod}^R \to \mathsf{Mod}^R$ has the following property (why?). For every R-module A and every integer $k \in \mathbb{Z}$, $\mathcal{F}$ sends the morphism $A \to A : x \mapsto kx$ to the morphism $\mathcal{F}(A) \to \mathcal{F}(A) : y \mapsto ky$. Further hint: you already know this for $k = 1$, just because $\mathcal{F}$ is a functor.

Since every finitely-generated abelian group A is the direct sum of a free abelian group with finitely-many copies of $\mathbb{Z}_k = \mathbb{Z}/k\mathbb{Z}$ for various values of $k \in \mathbb{N}$, the exercises above provide us with an explicit formula for $\operatorname{Tor}^{\mathbb{Z}}(A, G)$ for any other abelian group G, so long as we are able to compute the kernels of the homomorphisms $G \to G : g \mapsto kg$. In particular, if A is finitely generated, then $\operatorname{Tor}^{\mathbb{Z}}(A, G)$ depends only on the torsion subgroups of both A and G, and vanishes if either of these is trivial.

In fact, that last statement is still true even if neither A nor G is finitely generated:

Exercise 43.6. Show that every torsion-free abelian group is a flat $\mathbb{Z}$-module.
Hint: Given abelian groups A, B, G and an injective homomorphism $i : A \hookrightarrow B$, show that any nontrivial element in the kernel of $i \otimes \mathbb{1} : A \otimes G \to B \otimes G$ is also in the kernel of the restriction of this map to $A \otimes G_0 \to B \otimes G_0$ for some finitely-generated subgroup $G_0 \subset G$. If G is torsion free, what does the classification of finitely-generated abelian groups tell you about G_0?

Corollary 43.7. *For any abelian groups A, G, $\operatorname{Tor}^{\mathbb{Z}}(A, G) = 0$ if G is torsion free.* □

Remark 43.8. We will see below that $\operatorname{Tor}(A, B)$ and $\operatorname{Tor}(B, A)$ are always isomorphic, so the result above implies that $\operatorname{Tor}^{\mathbb{Z}}(A, G)$ also vanishes whenever A is torsion free.

Exercise 43.9. In Lecture 40 we made use of the following convenient fact, which you can now easily prove using the universal coefficient theorem: for any chain complex C_* of free abelian groups and any field $\mathbb{K}$ of characteristic 0, the canonical map

$$H_*(C_*) \otimes \mathbb{K} \xrightarrow{h} H_*(C_* \otimes \mathbb{K}) : [c] \otimes k \mapsto [c \otimes k]$$

is an isomorphism.

$$\begin{tikzcd}
 & \vdots \arrow[d] & \vdots \arrow[d] & \vdots \arrow[d] & \\
0 \arrow[r] & A_2 \arrow[r, dashed, "f_2"] \arrow[d, "\alpha_2"] & B_2 \arrow[r, dashed, "g_2"] \arrow[d, dashed, "\beta_2"] & C_2 \arrow[r] \arrow[d, "\gamma_2"] & 0 \\
0 \arrow[r] & A_1 \arrow[r, dashed, "f_1"] \arrow[d, "\alpha_1"] & B_1 \arrow[r, dashed, "g_1"] \arrow[d, dashed, "\beta_1"] & C_1 \arrow[r] \arrow[d, "\gamma_1"] & 0 \\
0 \arrow[r] & A_0 \arrow[r, dashed, "f_0"] \arrow[d, "\alpha"] & B_0 \arrow[r, dashed, "g_0"] \arrow[d, dashed, "\beta"] & C_0 \arrow[r] \arrow[d, "\gamma"] & 0 \\
0 \arrow[r] & A \arrow[r, "f"] \arrow[d] & B \arrow[r, "g"] \arrow[d] & C \arrow[r] \arrow[d] & 0 \\
 & 0 & 0 & 0 &
\end{tikzcd}$$

FIGURE 24. The diagram constructed in the horseshoe lemma. The modules B_n and the dashed arrows need to be constructed; everything else is given.

The exact sequence of derived functors. The following provides a converse to the fact that the derived functors $L_n\mathcal{F}$ vanish for all $n \geqslant 1$ whenever the functor $\mathcal{F}$ is a exact. In the specific case $\mathcal{F} = \otimes G$, it will tell us exactly when and why a short exact sequence $0 \to A \to B \to C \to 0$ fails to remain exact after tensoring everything with G.

PROPOSITION 43.10 (horseshoe lemma). *Suppose $0 \longrightarrow A \xrightarrow{f} B \xrightarrow{g} C \longrightarrow 0$ is a short exact sequence of R-modules and $A_* \xrightarrow{\alpha} A$ and $C_* \xrightarrow{\gamma} C$ are projective resolutions of A and C respectively. Then there exists a projective resolution $B_* \xrightarrow{\beta} B$ of B and maps $A_n \xrightarrow{f_n} B_n$ and $B_n \xrightarrow{g_n} C_n$ for each $n \geqslant 0$ such that the diagram in Figure 24 commutes, and the rows $0 \to A_n \to B_n \to C_n \to 0$ for all $n \geqslant 0$ are split exact sequences.*

PROOF. Let us label the given exact sequence $0 \longrightarrow A_{-1} \xrightarrow{f_{-1}} B_{-1} \xrightarrow{g_{-1}} C_{-1} \longrightarrow 0$ and, arguing by induction, suppose for some $n \geqslant 0$ that rows $k = -1, \dots, n-1$ of the diagram above have already been constructed and have the desired properties. For convenience, label $\alpha_0 := \alpha$, $\beta_0 := \beta$ and $\gamma_0 := \gamma$, and define $\alpha_{-1}, \beta_{-1}, \gamma_{-1}$ to be the unique maps from A, B, C respectively to the trivial module. As a preliminary observation, we claim that the short exact sequence at row $n-1$ restricts to a short exact sequence

$$0 \longrightarrow \ker \alpha_{n-1} \xrightarrow{f_{n-1}} \ker \beta_{n-1} \xrightarrow{g_{n-1}} \ker \gamma_{n-1} \longrightarrow 0. \tag{43.1}$$

If $n = 0$ this just means that the original sequence $0 \to A \to B \to C \to 0$ is exact, and for $n \geqslant 1$, we can deduce this from a long exact sequence of homologies obtained as follows. Consider the short exact sequence of chain complexes obtained by adding to rows $-1, \dots, n-1$ of the diagram infinitely many rows that contain only trivial modules. The exactness of the original columns then forces the homology in degree $n-2$ to vanish, while the homologies in degree $n-1$ are just the kernels of the respective maps, and (43.1) thus arises as the portion of the long exact sequence corresponding to row $n-1$. Most importantly, we conclude from this that g_{n-1} maps $\ker \beta_{n-1}$ surjectively onto $\ker \gamma_{n-1}$.

Now we proceed with the inductive step. What makes the construction of row n of the diagram relatively straightforward is the expectation that the exact sequence $0 \to A_n \to B_n \to C_n \to 0$ we are looking for should split: indeed, any short exact sequence with C_n as its last nontrivial term will automatically split, since C_n is projective (cf. Prop. 42.11). We can therefore construct B_n with this in mind, i.e. we define $B_n := A_n \oplus C_n$, with $f_n : A_n \hookrightarrow A_n \oplus C_n$ and $g_n : A_n \oplus C_n \to C_n$ as the obvious inclusion and projection respectively. Note that by Exercise 43.2, $A_n \oplus C_n$ is projective. The remaining task is thus to find a map

$$A_n \oplus C_n \xrightarrow{\beta_n} B_{n-1}$$

that satisfies

$$\operatorname{im} \beta_n = \ker \beta_{n-1}, \qquad \gamma_n \circ g_n = g_{n-1} \circ \beta_n, \qquad \text{and} \qquad \beta_n \circ f_n = f_{n-1} \circ \alpha_n. \tag{43.2}$$

Denote the natural inclusions into the direct sum by

$$i_A : A_n \hookrightarrow A_n \oplus C_n, \qquad i_C : C_n \hookrightarrow A_n \oplus C_n,$$

so we have

$$i_A = f_n \qquad \text{and} \qquad g_n \circ i_C = \mathbb{1}_{C_n}.$$

Any pair of maps $\beta_A : A_n \to B_{n-1}$ and $\beta_C : C_n \to B_{n-1}$ will determine a unique $\beta_n : A_n \oplus C_n \to B_{n-1}$ via the relations

$$\beta_A := \beta_n \circ i_A : A_n \to B_{n-1} \qquad \beta_C := \beta_n \circ i_C : C_n \to B_{n-1},$$

and the relation $\beta_{n-1} \circ \beta_n = 0$ then holds trivially if $n = 0$, and for $n \geqslant 1$, it holds if and only if

$$\beta_{n-1} \circ \beta_A = 0 \qquad \text{and} \qquad \beta_{n-1} \circ \beta_C = 0.$$

Since $i_A = f_n$, (43.2) demands that β_A satisfy the relation

$$\beta_A = \beta_n \circ f_n = f_{n-1} \circ \alpha_n,$$

which uniquely determines it, i.e. we now define $\beta_A := f_{n-1} \circ \alpha_n$ and observe that for $n \geqslant 1$, it conveniently satisfies

$$\beta_{n-1} \circ \beta_A = \beta_{n-1} \circ f_{n-1} \circ \alpha_n = f_{n-2} \circ \alpha_{n-1} \circ \alpha_n = 0.$$

The relation $\beta_n \circ f_n = f_{n-1} \circ \alpha_n$ in (43.2) will now be satisfied regardless of how β_C is defined, and the additional conditions $\gamma_n \circ g_n = g_{n-1} \circ \beta_n$ and $\beta_{n-1} \circ \beta_n = 0$ will hold if and only if β_C is a map $C_n \to \ker \beta_{n-1}$ satisfying

$$g_{n-1} \circ \beta_C = \gamma_n \circ g_n \circ i_C = \gamma_n. \tag{43.3}$$

Since γ_n has image in $\ker \gamma_{n-1}$ and g_{n-1} maps $\ker \beta_{n-1}$ surjectively onto $\ker \gamma_{n-1}$, the fact that C_n is projective now guarantees the existence of a map $\beta_C : C_n \to \ker \beta_{n-1}$ that lifts $\gamma_n : C_n \to \ker \gamma_{n-1}$ via the surjection $g_{n-1} : \ker \beta_{n-1} \to \ker \gamma_{n-1}$ in precisely the sense of (43.3), so let us choose β_C to be any map with this property. The definition of row n in the diagram is now complete: by construction the diagram still commutes, and $\beta_{n-1} \circ \beta_n = 0$.

The only remaining question is whether $\operatorname{im} \beta_n = \ker \beta_{n-1}$, but in fact, this now follows for more-or-less formal reasons. Indeed, consider the short exact sequence of chain complexes obtained by replacing everything in the diagram except for rows $n-1$ and n with trivial modules, and replacing

$$\begin{tikzcd}
 & \vdots \arrow[d] & \vdots \arrow[d] & \vdots \arrow[d] & \\
0 \arrow[r] & 0 \arrow[r] \arrow[d] & 0 \arrow[r] \arrow[d] & 0 \arrow[r] \arrow[d] & 0 \\
0 \arrow[r] & A_n \arrow[r, "f_n"] \arrow[d, "\alpha_n"] & B_n \arrow[r, "g_n"] \arrow[d, "\beta_n"] & C_n \arrow[r] \arrow[d, "\gamma_n"] & 0 \\
0 \arrow[r] & \ker \alpha_{n-1} \arrow[r, "f_{n-1}"] \arrow[d] & \ker \beta_{n-1} \arrow[r, "g_{n-1}"] \arrow[d] & \ker \gamma_{n-1} \arrow[r] \arrow[d] & 0 \\
0 \arrow[r] & 0 \arrow[r] \arrow[d] & 0 \arrow[r] \arrow[d] & 0 \arrow[r] \arrow[d] & 0 \\
 & \vdots & \vdots & \vdots &
\end{tikzcd}$$

FIGURE 25. The short exact sequence of chain complexes that implies $\operatorname{im}\beta_n = \ker \beta_{n-1}$.

row $n-1$ with $0 \to \ker \alpha_{n-1} \to \ker \beta_{n-1} \to \gamma_{n-1} \to 0$, i.e. the diagram in Figure 25. The resulting long exact sequence of homologies[70] takes the form

$$\ldots \to 0 \to 0 \to \ker \alpha_n \to \ker \beta_n \to \ker \gamma_n \to \frac{\ker \alpha_{n-1}}{\operatorname{im} \alpha_n} \to \frac{\ker \beta_{n-1}}{\operatorname{im} \beta_n} \to \frac{\ker \gamma_{n-1}}{\operatorname{im} \gamma_n} \to 0 \to 0 \to \ldots,$$

and of the three quotients in this sequence, two of them are already assumed to be trivial, implying that the third must be as well. □

The horseshoe lemma allows us to choose suitable projective resolutions for any given short exact sequence $0 \longrightarrow A \xrightarrow{f} B \xrightarrow{g} C \longrightarrow 0$ so that we get a corresponding short exact sequence of chain complexes

$$0 \longrightarrow A_* \xrightarrow{f_*} B_* \xrightarrow{g_*} C_* \longrightarrow 0.$$

Even better, this sequence splits (due to the fact that its last nontrivial term is projective), so we can plug it into any additive functor $\mathcal{F}$ and obtain yet another short exact sequence of chain complexes

$$0 \longrightarrow \mathcal{F}(A_*) \xrightarrow{\mathcal{F}(f_*)} \mathcal{F}(B_*) \xrightarrow{\mathcal{F}(g_*)} \mathcal{F}(C_*) \longrightarrow 0.$$

I think you know what comes next: passing to homology now gives a long exact sequence of derived functors, and if we assume $\mathcal{F}$ is right-exact and use the natural isomorphisms to replace $L_0\mathcal{F}$ with $\mathcal{F}$, it takes the form

$$\begin{aligned} \ldots \to L_2\mathcal{F}(A) \to L_2\mathcal{F}(B) \to L_2\mathcal{F}(C) \to L_1\mathcal{F}(A) \to L_1\mathcal{F}(B) \to L_1\mathcal{F}(C) \\ \to \mathcal{F}(A) \to \mathcal{F}(B) \to \mathcal{F}(C) \to 0, \end{aligned} \tag{43.4}$$

where the maps for each arrow are all either induced by $f : A \to B$ or $g : B \to C$ via the functoriality of $L_n\mathcal{F}$ or are connecting homomorphisms constructed via the standard diagram

[70] What we are using here is a version of a popular result in homological algebra called the *snake lemma*. We are deducing it from the fact that short exact sequences of chain complexes give long exact sequences on homology, but it is also possible to do things the other way around, i.e. to prove the snake lemma directly via a diagram chase and then derive from it the usual result about short and long exact sequences.

chase. We now have a sharp criterion to determine which short exact sequences remain exact after feeding them into any given right-exact functor: in particular, the sequence $0 \to \mathcal{F}(A) \to \mathcal{F}(B) \to \mathcal{F}(C) \to 0$ will be exact whenever $L_1\mathcal{F}(C)$ vanishes!

Applying this construction to the specific right-exact functor $\mathcal{F} = \otimes G$ for any given R-module G gives

$$\begin{aligned} \ldots \to \operatorname{Tor}_2(A,G) \to \operatorname{Tor}_2(B,G) \to \operatorname{Tor}_2(C,G) \to \operatorname{Tor}_1(A,G) \to \operatorname{Tor}_1(B,G) \to \operatorname{Tor}_1(C,G) \\ \to A \otimes G \to B \otimes G \to C \otimes G \to 0, \end{aligned}$$

and if R is a principal ideal domain, this simplifies to an exact sequence with only six potentially nontrivial terms,

$$0 \to \operatorname{Tor}(A,G) \to \operatorname{Tor}(B,G) \to \operatorname{Tor}(C,G) \to A \otimes G \to B \otimes G \to C \otimes G \to 0.$$

Functoriality with respect to G. We have not yet discussed in what sense $\operatorname{Tor}_n(A,\cdot) : \mathsf{Mod}^R \to \mathsf{Mod}^R$ is also a functor for each $n \geqslant 0$, but this is not hard to understand. The first important observation is that the tensor product $A \otimes G$ can also be regarded in an obvious way as a functor $A\otimes : \mathsf{Mod}^R \to \mathsf{Mod}^R$ with respect to G for any fixed R-module A. In fact, this works just as well if one fixes not just a single module but a chain complex C_*, producing a functor

$$C_*\otimes : \mathsf{Mod}^R \to \mathsf{Chain}^R$$

that sends any R-module G to the chain complex $C_* \otimes G$ and any homomorphism $\varphi : G \to H$ to a chain map of the form

$$\mathbb{1} \otimes \varphi : C_* \otimes G \to C_* \otimes H,$$

thus inducing maps

$$H_n(C_* \otimes G) \xrightarrow{\varphi_*} H_n(C_* \otimes H)$$

for each n. Applying this in particular to the chain complex A_* arising from a projective resolution $A_* \xrightarrow{\alpha} A$ of any R-module A produces natural maps

$$\operatorname{Tor}_n(A,G) \xrightarrow{\varphi_*} \operatorname{Tor}_n(A,H)$$

induced by any homomorphism $\varphi : G \to H$, and one easily checks that for $n = 0$, the natural isomorphism identifies φ_* with $\mathbb{1} \otimes \varphi : A \otimes G \to A \otimes H$.

With that understood, consider what happens if we are given a short exact sequence of "coefficient" modules

$$0 \longrightarrow G \xrightarrow{\varphi} H \xrightarrow{\psi} K \longrightarrow 0.$$

If A is a projective module, then the functor $A\otimes$ is exact for the same reasons that $\otimes A$ is exact, thus producing another short exact sequence $0 \to A \otimes G \to A \otimes H \to A \otimes K \to 0$. The latter will not be exact in general if A is not assumed to be projective, but we can do the same trick with the modules A_n in any projective resolution $A_* \xrightarrow{\alpha} A$ of A, producing a short exact sequence of chain complexes

$$0 \longrightarrow A_* \otimes G \xrightarrow{\mathbb{1}\otimes\varphi} A_* \otimes H \xrightarrow{\mathbb{1}\otimes\psi} A_* \otimes K \longrightarrow 0,$$

and therefore a long exact sequence for any given R-module G gives

$$\begin{aligned} \ldots \to \operatorname{Tor}_2(A,G) \to \operatorname{Tor}_2(A,H) \to \operatorname{Tor}_2(A,K) \to \operatorname{Tor}_1(A,G) \to \operatorname{Tor}_1(A,H) \to \operatorname{Tor}_1(A,K) \\ \to A \otimes G \to A \otimes H \to A \otimes K \to 0. \end{aligned}$$

As usual, only six terms of this sequence can be nontrivial if R is a principal ideal domain.

$$
\begin{tikzcd}
 & \vdots \arrow[d] & \vdots \arrow[d] & \vdots \arrow[d] & \vdots \arrow[d] & \vdots \arrow[d] & \\
\cdots \arrow[r] & A_2 \otimes B_2 \arrow[r, "\mathbb{1}\otimes\beta_2"] \arrow[d, "\alpha_2\otimes\mathbb{1}"] & A_2 \otimes B_1 \arrow[r, "\mathbb{1}\otimes\beta_1"] \arrow[d, "\alpha_2\otimes\mathbb{1}"] & A_2 \otimes B_0 \arrow[r, "\mathbb{1}\otimes\beta"] \arrow[d, "\alpha_2\otimes\mathbb{1}"] & A_2 \otimes B \arrow[r] \arrow[d, "\alpha_2\otimes\mathbb{1}"] & 0 \arrow[r] \arrow[d] & \cdots \\
\cdots \arrow[r] & A_1 \otimes B_2 \arrow[r, "\mathbb{1}\otimes\beta_2"] \arrow[d, "\alpha_1\otimes\mathbb{1}"] & A_1 \otimes B_1 \arrow[r, "\mathbb{1}\otimes\beta_1"] \arrow[d, "\alpha_1\otimes\mathbb{1}"] & A_1 \otimes B_0 \arrow[r, "\mathbb{1}\otimes\beta"] \arrow[d, "\alpha_1\otimes\mathbb{1}"] & A_1 \otimes B \arrow[r] \arrow[d, "\alpha_1\otimes\mathbb{1}"] & 0 \arrow[r] \arrow[d] & \cdots \\
\cdots \arrow[r] & A_0 \otimes B_2 \arrow[r, "\mathbb{1}\otimes\beta_2"] \arrow[d, "\alpha\otimes\mathbb{1}"] & A_0 \otimes B_1 \arrow[r, "\mathbb{1}\otimes\beta_1"] \arrow[d, "\alpha\otimes\mathbb{1}"] & A_0 \otimes B_0 \arrow[r, "\mathbb{1}\otimes\beta"] \arrow[d, "\alpha\otimes\mathbb{1}"] & A_0 \otimes B \arrow[r] \arrow[d] & 0 \arrow[r] \arrow[d] & \cdots \\
\cdots \arrow[r] & A \otimes B_2 \arrow[r, "\mathbb{1}\otimes\beta_2"] \arrow[d] & A \otimes B_1 \arrow[r, "\mathbb{1}\otimes\beta_1"] \arrow[d] & A \otimes B_0 \arrow[r] \arrow[d] & 0 \arrow[r] \arrow[d] & 0 \arrow[r] \arrow[d] & \cdots \\
\cdots \arrow[r] & 0 \arrow[r] \arrow[d] & 0 \arrow[r] \arrow[d] & 0 \arrow[r] \arrow[d] & 0 \arrow[r] \arrow[d] & 0 \arrow[r] \arrow[d] & \cdots \\
 & \vdots & \vdots & \vdots & \vdots & \vdots &
\end{tikzcd}
$$

FIGURE 26. The double complex that implies $\mathrm{Tor}_n(A, B) \cong \mathrm{Tor}_n(B, A)$.

EXERCISE 43.11. Show that for an abelian group G with torsion subgroup $T(G) \subset G$ and any other abelian group A, the map

$$\mathrm{Tor}^{\mathbb{Z}}(A, T(G)) \to \mathrm{Tor}^{\mathbb{Z}}(A, G)$$

induced by the inclusion $T(G) \hookrightarrow G$ is an isomorphism.

Symmetry of the Tor functor. There is a natural isomorphism $A \otimes B \cong B \otimes A$ for every pair of R-modules, so the following result probably should not surprise you:

THEOREM 43.12. *There is a natural isomorphism $\mathrm{Tor}_n(A, B) \cong \mathrm{Tor}_n(B, A)$ for every $n \geqslant 0$ and every pair of R-modules A, B, and it matches the natural isomorphism $A \otimes B \cong B \otimes A$ in the case $n = 0$.*

Here is a nice way to see it. After choosing projective resolutions $A_* \xrightarrow{\alpha} A$ and $B_* \xrightarrow{\beta} B$, we can form the commutative diagram in Figure 26, which is called a **double complex**, because of each its rows and each of its columns is a chain complex. Let us number the rows and columns with integers so that $A_m \otimes B_n$ is in row m and column n. Since A_m and B_n are projective for each $m, n \geqslant 0$, Corollary 42.20 implies that the functors $A_m \otimes$ and $\otimes B_n$ are exact, which makes the mth row and nth column for each $m, n \geqslant 0$ into an exact sequence. The only potentially non-exact sequences we can see are therefore the chain complex $A_* \otimes B$ in column -1, whose homology is

$$H_m(A_* \otimes B) = \mathrm{Tor}_m(A, B),$$

and similarly the complex $A \otimes B_*$ in row -1, whose homology we can identify with

$$H_n(A \otimes B_*) \cong \mathrm{Tor}_n(B, A)$$

using the natural isomorphism $A \otimes B_* \cong B_* \otimes A$. The rest is diagram chasing:

$$\begin{tikzcd}
 & \vdots \arrow[d] & \vdots \arrow[d] & \vdots \arrow[d] & \vdots \arrow[d] & \vdots \arrow[d] & \\
\cdots \arrow[r] & C_2^2 \arrow[r, "d_2^2"] \arrow[d, "\partial_2^2"] & C_1^2 \arrow[r, "d_1^2"] \arrow[d, "\partial_1^2"] & C_0^2 \arrow[r, "d_0^2"] \arrow[d, "\partial_0^2"] & C^2 \arrow[r] \arrow[d, "\partial^2"] & 0 \arrow[r] \arrow[d] & \cdots \\
\cdots \arrow[r] & C_2^1 \arrow[r, "d_2^1"] \arrow[d, "\partial_2^1"] & C_1^1 \arrow[r, "d_1^1"] \arrow[d, "\partial_1^1"] & C_0^1 \arrow[r, "d_0^1"] \arrow[d, "\partial_0^1"] & C^1 \arrow[r] \arrow[d, "\partial^1"] & 0 \arrow[r] \arrow[d] & \cdots \\
\cdots \arrow[r] & C_2^0 \arrow[r, "d_2^0"] \arrow[d, "\partial_2^0"] & C_1^0 \arrow[r, "d_1^0"] \arrow[d, "\partial_1^0"] & C_0^0 \arrow[r, "d_0^0"] \arrow[d, "\partial_0^0"] & C^0 \arrow[r] \arrow[d, "\partial^0"] & 0 \arrow[r] \arrow[d] & \cdots \\
\cdots \arrow[r] & C_2 \arrow[r, "d_2"] \arrow[d] & C_1 \arrow[r, "d_1"] \arrow[d] & C_0 \arrow[r, "d_0"] \arrow[d] & C \arrow[r] \arrow[d] & 0 \arrow[r] \arrow[d] & \cdots \\
\cdots \arrow[r] & 0 \arrow[r] \arrow[d] & 0 \arrow[r] \arrow[d] & 0 \arrow[r] \arrow[d] & 0 \arrow[r] \arrow[d] & 0 \arrow[r] \arrow[d] & \cdots \\
 & \vdots & \vdots & \vdots & \vdots & \vdots &
\end{tikzcd}$$

FIGURE 27. The abstract double complex in Exercise 43.13.

EXERCISE 43.13. Suppose (C_*, d) and (C^*, ∂) are chain complexes with $C_{-1} = C^{-1} =: C$ which form row -1 and column -1 respectively of a double complex $\{C_j^i\}_{i,j\in\mathbb{Z}}$ as shown in Figure 27, with the property that all other rows and columns are exact, and $C_j^i = 0$ whenever $i < -1$ or $j < -1$. Find an isomorphism

$$H_n(C^*, \partial) \cong H_n(C_*, d)$$

for every $n \in \mathbb{Z}$. Then decide what it should mean to call this isomorphism *natural*, and convince yourself that it is.

Applying Exercise 43.13 to the specific double complex in Figure 26 yields the isomorphisms $\mathrm{Tor}_n(A, B) \cong \mathrm{Tor}_n(B, A)$ promised by Theorem 43.12.

Torsion on manifolds in codimension one. The exercise below gives a less straightforward application of the universal coefficient theorem in which information about $H_{n-1}(X; \mathbb{Z})$ gets extracted from information about $H_n(X; \mathbb{Z}_p)$ for prime numbers p. The result will acquire some importance when we study the homology of topological manifolds in general, as it serves as a step in the proof that every closed simply-connected 3-manifold is homotopy equivalent to S^3, a basic result in the background of the Poincaré conjecture.

For context, recall that if M is a connected topological n-manifold with empty boundary and an oriented triangulation, then its triangulation can be regarded as a cell decomposition, so that the cellular chain complex of M matches its oriented simplicial chain complex. If M is noncompact, then $H_n(M; G)$ clearly vanishes for all choices of coefficient group G, because the cellular n-chain formed out of some *but not all* n-simplices in a triangulation of a manifold can never be a cycle. If M is compact, then it admits an integral fundamental cycle which generates the group of cycles in $C_n^{\mathrm{CW}}(M; \mathbb{Z})$, and the resulting fundamental class thus generates $H_n^{\mathrm{CW}}(M; \mathbb{Z}) \cong \mathbb{Z}$, while for arbitrary choices of coefficient group G, one similarly obtains

$$H_n^{\mathrm{CW}}(M;G) \cong G \cong H_n^{\mathrm{CW}}(M;\mathbb{Z}) \otimes G.$$

We will see later that without any knowledge of triangulations, all of this remains true for the singular homology of orientable topological n-manifolds, due to the general construction of fundamental classes.

EXERCISE 43.14. Suppose C_* is a chain complex of free abelian groups such that for some $n \in \mathbb{N}$, $H_n(C_*)$ is finitely generated and satisfies

$$H_n(C_* \otimes \mathbb{Z}_p) \cong H_n(C_*) \otimes \mathbb{Z}_p$$

for every prime $p \in \mathbb{N}$. Prove that $H_{n-1}(C_*)$ is torsion free.

44. Product chain complexes (December 19, 2023)

The main question for the next two lectures is straightforward: if we understand $H_*(X)$ and $H_*(Y)$, how can we use them to compute $H_*(X \times Y)$?

If one starts by viewing the question in terms of cellular homology, one quickly discovers that there is a natural tensor product operation on chain complexes, such that if X and Y have cell decompositions, then $X \times Y$ inherits a cell decomposition for which $C_*^{\mathrm{CW}}(X \times Y) \cong C_*^{\mathrm{CW}}(X) \otimes C_*^{\mathrm{CW}}(Y)$. This reduces the problem to an essentially algebraic one: given two chain complexes A_* and B_* of the kind that appear in cellular or singular homology, how is $H_*(A_* \otimes B_*)$ determined by $H_*(A_*)$ and $H_*(B_*)$? The answer to that question is called the algebraic Künneth formula, and it will be the main result of the present lecture.

The product of two CW-complexes. Suppose X and Y are both compact CW-complexes, and consider the product $X \times Y$. This has a natural cell decomposition such that

$$(X \times Y)^n = \bigcup_{0 \leqslant k \leqslant n} X^k \times Y^{n-k}.$$

It is easiest to see this if we choose a homeomorphism of the disk $\mathbb{D}^n$ with the n-dimensional cube I^n and thus regard I^n as the domain of the characteristic maps of n-cells. Since $I^{k+\ell} = I^k \times I^\ell$, any pair consisting of a k-cell $e_\alpha^k \subset X$ and ℓ-cell $e_\beta^\ell \subset Y$ with characteristic maps $\Phi_\alpha : I^k \to X$ and $\Phi_\beta : I^\ell \to Y$ respectively gives rise to a $(k+\ell)$-cell

$$e_\alpha^k \times e_\beta^\ell \subset X \times Y$$

with characteristic map

$$\Phi_\alpha \times \Phi_\beta : I^{k+\ell} \to X \times Y : (s,t) \mapsto (\Phi_\alpha(s), \Phi_\beta(t)).$$

The bilinear operation

$$C_k^{\mathrm{CW}}(X;\mathbb{Z}) \times C_\ell^{\mathrm{CW}}(Y;\mathbb{Z}) \xrightarrow{\times} C_{k+\ell}^{\mathrm{CW}}(X \times Y;\mathbb{Z})$$

defined on the cellular chain complex by sending a pair of generators $(e_\alpha^k, e_\beta^\ell)$ to $e_\alpha^k \times e_\beta^\ell$ is called the (chain level) **cellular cross product**. The following formula for the boundary map arises from the geometric intuition that the boundary of a product of manifolds $M \times N$ consists of all points $(x,y) \in M \times N$ such that either $x \in \partial M$ or $y \in \partial N$; in particular, this description can be applied to the boundary of the cube $I^{k+\ell} = I^k \times I^\ell$, which is the domain of the characteristic map for a product cell $e_\alpha^k \times e_\beta^\ell$. One then has to think somewhat more carefully about orientations to get the signs right (see Exercise 44.10).[71]

[71] An explicit proof of the formula in Prop. 44.1 can also be found in [**Hat02**, Prop. 3B.1].

PROPOSITION 44.1. *For any pair of CW-complexes X and Y with a k-cell $e^k_\alpha \subset X$ and an ℓ-cell $e^\ell_\beta \subset Y$,*

$$\partial(e^k_\alpha \times e^\ell_\beta) = \partial e^k_\alpha \times e^\ell_\beta + (-1)^k e^k_\alpha \times \partial e^\ell_\beta \in C^{\mathrm{CW}}_{k+\ell-1}(X \times Y; \mathbb{Z}).$$

□

Being a bilinear map of abelian groups, the cross product is equivalent to a group homomorphism

$$C^{\mathrm{CW}}_k(X;\mathbb{Z}) \otimes C^{\mathrm{CW}}_\ell(Y;\mathbb{Z}) \to C^{\mathrm{CW}}_{k+\ell}(X \times Y;\mathbb{Z}),$$

namely the unique homomorphism that sends $e^k_\alpha \otimes e^\ell_\beta$ to $e^k_\alpha \times e^\ell_\beta$ for every pair of cells of the appropriate dimensions in X and Y. If we want to say something similar with an arbitrary coefficient group G in place of $\mathbb{Z}$, then in general we face a problem: there is no obviously canonical way to define $ge^k_\alpha \times he^\ell_\beta \in C^{\mathrm{CW}}_{k+\ell}(X \times Y; G)$ for all possible coefficients $g, h \in G$. If, on the other hand, G is not just an abelian group but also a commutative ring R, with a bilinear product structure $R \times R \to R : (g, h) \mapsto gh$, then it makes sense to define

$$(ge^k_\alpha) \times (he^\ell_\beta) := (gh)e^k_\alpha \times e^\ell_\beta \in C^{\mathrm{CW}}_{k+\ell}(X \times Y; R)$$

for any $g, h \in R$. The highbrow way to say this is that if our coefficient group is a ring R, then $C^{\mathrm{CW}}_*(X;R)$ and $C^{\mathrm{CW}}_*(Y;R)$ naturally become chain complexes of free R-modules, and the formula $e^k_\alpha \otimes e^\ell_\beta \mapsto e^k_\alpha \times e^\ell_\beta$ uniquely determines an R-module homomorphism

$$C^{\mathrm{CW}}_k(X;R) \otimes C^{\mathrm{CW}}_\ell(Y;R) \to C^{\mathrm{CW}}_{k+\ell}(X \times Y;R),$$

where the left hand side is to be understood as a tensor product of R-modules, not just of abelian groups. Taking the direct sum of these maps over all pairs of integers $k, \ell \geqslant 0$, the chain-level cellular cross product now determines an R-module homomorphism

$$C^{\mathrm{CW}}_*(X;R) \otimes C^{\mathrm{CW}}_*(Y;R) \xrightarrow{\times} C^{\mathrm{CW}}_*(X \times Y;R).$$

In fact, this map is an R-module *isomorphism*, due to the easy observation that both sides are free R-modules with canonical bases that are in bijective correspondence with each other.

The Künneth formula. The following purely algebraic definition should now hopefully seem quite natural. For the rest of this lecture, we fix a commutative ring R and assume (as in our presentation of the universal coefficient theorem) that all chain complexes, homologies and tensor products under discussion live in the category of R-modules. It will sometimes also be necessary to assume that R is a principal ideal domain, but this will be mentioned explicitly whenever it is needed.

DEFINITION 44.2. Given chain complexes (A_*, ∂^A) and (B_*, ∂^B) of R-modules, the **tensor product chain complex** $(A_* \otimes B_*, \partial)$ is defined by

$$(A_* \otimes B_*)_n = \bigoplus_{k+\ell=n} A_k \otimes B_\ell, \tag{44.1}$$

where the direct sum is understood to be over the set of all pairs $(k, \ell) \in \mathbb{Z}^2$ with $k + \ell = n$, and the boundary map is determined by the formula

$$\partial(a \otimes b) = \partial^A a \otimes b + (-1)^k a \otimes \partial^B b \quad \text{for } a \in A_k,\ b \in B_\ell. \tag{44.2}$$

You should take a moment to assure yourself that this really defines a chain complex: ∂^2 includes some terms that vanish because $(\partial^A)^2$ and $(\partial^B)^2$ both vanish, but also cross terms $\partial^A a \otimes \partial^B b$ that disappear due to sign cancelations. Here is another easy thing to check: given chain maps $f : A_* \to A'_*$ and $g : B_* \to B'_*$, there is a chain map

$$f \otimes g : A_* \otimes B_* \to A'_* \otimes B'_* : a \otimes b \mapsto f(a) \otimes g(b). \tag{44.3}$$

We can now rephrase Proposition 44.1 as follows:

PROPOSITION 44.3. *For any choice of coefficient ring R, the chain-level cellular cross product determines an isomorphism of chain complexes of R-modules*

$$C_*^{\mathrm{CW}}(X;R) \otimes C_*^{\mathrm{CW}}(Y;R) \to C_*^{\mathrm{CW}}(X \times Y;R) : a \otimes b \to a \times b.$$

$\square$

REMARK 44.4. The signs in formulas such as (44.2) can be deduced consistently from the *Koszul sign convention*, which made a previous appearance when we were constructing oriented triangulations of products of simplices (see Remark 31.20). The idea is to regard every element $a \in A_k$ in a chain complex A_* as *even* or *odd* depending on whether k is even or odd, while also regarding boundary maps such as $\partial : A_* \to A_*$ as having degree -1 (and thus odd) since they send A_k to A_{k-1}. The rule is then that a sign changes every time the order of two odd objects is interchanged. In other words, the sign in the formula $\partial(a \otimes b) = \partial a \otimes b + (-1)^k a \otimes \partial b$ comes from the fact that in the last term, we have interchanged the order of ∂ and a, which produces a sign if and only if a is odd (since ∂ is always odd). Similar sign conventions appear in many branches of mathematics, and they are often determined by the signs of permutations, e.g. a familiar example in differential geometry is the formula for the exterior derivative of a wedge product of differential forms.

With product cell complexes as motivation, it is important to be able to compute the homology of a tensor product chain complex, and it seems a good guess that the answer should be related to the tensor product of the individual homologies of the two complexes. As with the universal coefficient theorem, we can begin by observing that there is a canonical map: for any two chain complexes A_*, B_* and each $k, \ell \in \mathbb{Z}$, we can define

$$H_k(A_*) \otimes H_\ell(B_*) \to H_{k+\ell}(A_* \otimes B_*) : [a] \otimes [b] \mapsto [a \otimes b].$$

It is an easy exercise to check that this is a well-defined homomorphism, and taking the direct sum of these maps for all choices of $k, \ell \in \mathbb{Z}$ with a fixed sum produces a canonical map

$$\bigoplus_{k+\ell=n} H_k(A_*) \otimes H_\ell(B_*) \to H_n(A_* \otimes B_*) \tag{44.4}$$

for each $n \in \mathbb{Z}$. It seems reasonable to hope that this will at least sometimes be an isomorphism. What's actually true is in fact a direct generalization of the universal coefficient theorem.

THEOREM 44.5 (algebraic Künneth formula). *Assume R is a principal ideal domain, C_*, C'_* are chain complexes of R-modules, and C_* is free. Then the map (44.4) for every $n \in \mathbb{Z}$ fits into a natural short exact sequence*

$$0 \to \bigoplus_{k+\ell=n} H_k(C_*) \otimes H_\ell(C'_*) \to H_n(C_* \otimes C'_*) \to \bigoplus_{k+\ell=n-1} \mathrm{Tor}(H_k(C_*), H_\ell(C'_*)) \to 0,$$

and the sequence splits (but not naturally).

As usual, the word "natural" in this statement has a technical meaning in terms of natural transformations, so that for any two pairs of chain complexes A_*, A'_* and B_*, B'_* satisfying the hypotheses of the theorem, the maps in the two exact sequences will fit into commutative diagrams together with the maps induced by any pair of chain maps $A_* \to B_*$ and $A'_* \to B'_*$. The statement becomes a bit more concise if we define the operation $\otimes$ on the category $\mathsf{Mod}_{\mathbb{Z}}^R$ of $\mathbb{Z}$-graded R-modules via (44.1) and similarly define Tor as a functor $\mathsf{Mod}_{\mathbb{Z}}^R \times \mathsf{Mod}_{\mathbb{Z}}^R \to \mathsf{Mod}_{\mathbb{Z}}^R$ by

$$\big(\mathrm{Tor}(C_*, C'_*)\big)_n := \bigoplus_{k+\ell=n} \mathrm{Tor}(C_k, C'_\ell).$$

We will have some further comments below on why this is a sensible definition, but one immediate practical reason is that the exact sequence in Theorem 44.5 can now be written as

$$0 \to H_*(C_*) \otimes H_*(C'_*) \to H_*(C_* \otimes C'_*) \to \big(\mathrm{Tor}(H_*(C_*), H_*(C'_*))\big)_{*-1} \to 0,$$

where the subscript "$* - 1$" on the last term indicates the downward degree shift. The splitting gives rise to a (non-canonical) isomorphism

$$H_*(C_* \otimes C'_*) \cong \big(H_*(C_*) \otimes H_*(C'_*)\big) \oplus \big(\mathrm{Tor}(H_*(C_*), H_*(C'_*))\big)_{*-1},$$

which can be used in practice to compute the cellular homology of products. We will see in the next lecture how this can also be applied directly to singular homology, without needing to know that singular and cellular homologies are isomorphic.

Both the statement and the proof of the Künneth formula can be regarded as direct generalizations of the universal coefficient theorem if we extend the range of our definitions accordingly. Indeed, if one takes the chain complex C'_* to be a single R-module G in degree 0 and trivial in every other degree, then its homology is itself, and Theorem 44.5 reduces to precisely the universal coefficient theorem. Replacing G with a chain complex C'_* does not actually add much complication to the proof, if one adopts the right perspective.

Let us first clarify why it is sensible to extend Tor to a functor on $\mathsf{Mod}^R_{\mathbb{Z}}$ in the way that was described above. One can speak of additive, exact, left-exact and right-exact functors on $\mathsf{Mod}^R_{\mathbb{Z}}$ in exactly the same way as for functors on Mod^R: the direct sum $A_* \oplus B_*$ of two $\mathbb{Z}$-graded R-modules has an obvious $\mathbb{Z}$-grading with $(A_* \oplus B_*)_n := A_n \oplus B_n$, and exact sequences of $\mathbb{Z}$-graded R-modules are simply exact sequences of R-modules that each carry the grading as extra structure and thus require the morphisms in the sequence to preserve it. Given any $\mathbb{Z}$-graded R-module $G_* = \bigoplus_{n \in \mathbb{Z}} G_n$, the functor

$$\otimes G_* : \mathsf{Mod}^R_{\mathbb{Z}} \to \mathsf{Mod}^R_{\mathbb{Z}} : A_* \mapsto A_* \otimes G_*$$

is right-exact for the same reasons that $\otimes G : \mathsf{Mod}^R \to \mathsf{Mod}^R$ is right-exact, and one can similarly define left derived functors $\mathrm{Tor}_n(\cdot, G_*) := L_n(\otimes G_*) : \mathsf{Mod}^R_{\mathbb{Z}} \to \mathsf{Mod}^R_{\mathbb{Z}}$ as a way of measuring its failure to be left-exact. This requires choosing for any $\mathbb{Z}$-graded R-module A_* a projective resolution

$$\ldots \longrightarrow A_{*,2} \xrightarrow{\alpha_2} A_{*,1} \xrightarrow{\alpha_1} A_{*,0} \xrightarrow{\alpha} A_* \longrightarrow 0,$$

i.e. an exact sequence in which each $A_{*,n} = \bigoplus_{k \in \mathbb{Z}} A_{k,n}$ is a $\mathbb{Z}$-graded R-module that is projective, meaning that the lifting problem

$$\begin{tikzcd} & H_* \arrow[d, "\pi"] \\ A_{*,n} \arrow[r, "\varphi"] \arrow[ur, dashed, "\tilde{\varphi}"] & G_* \end{tikzcd}$$

can be solved in the category $\mathsf{Mod}^R_{\mathbb{Z}}$ whenever $\varphi : A_{*,n} \to G_*$ and $\pi : H_* \to G_*$ are homomorphisms that preserve gradings and π is surjective. It is straightforward to check that a $\mathbb{Z}$-graded R-module $G_* = \bigoplus_{n \in \mathbb{Z}} G_n$ is projective if and only if the individual R-modules G_n are all projective, and a projective resolution $A_{*,*} \xrightarrow{\alpha} A_*$ in $\mathsf{Mod}^R_{\mathbb{Z}}$ is thus equivalent to a collection of projective resolutions $A_{k,*} \xrightarrow{\alpha} A_k$ in Mod^R, one for each $k \in \mathbb{Z}$. For this reason, defining $\mathrm{Tor}_n(A_*, G_*)$ as a left derived functor in the category $\mathsf{Mod}^R_{\mathbb{Z}}$ gives the same result as the more naive definition, in which we simply regard $A_* = \bigoplus_k A_k$ and $G_* = \bigoplus_\ell G_\ell$ as R-modules and assign to the R-module $\mathrm{Tor}_n(A_*, G_*)$ the $\mathbb{Z}$-grading such that

$$\mathrm{Tor}_n(A_*, G_*)_m = \bigoplus_{k+\ell=m} \mathrm{Tor}_n(A_k, G_\ell)$$

for each $m \in \mathbb{Z}$.

With those formalities out of the way, let us go ahead and repeat the main details of the proof of the universal coefficient theorem in terms that are general enough to prove the Künneth formula as well.

Proof of Theorem 44.5. As in the proof of the universal coefficient theorem, we abbreviate the submodules of boundaries and cycles in C_n by $B_n \subset Z_n \subset C_n$, and think of $B_* := \bigoplus_n B_n$ and $Z_* := \bigoplus_n Z_n$ as chain complexes with trivial boundary maps, so their homologies are $H_n(Z_*) = Z_n$ and $H_n(B_*) = B_n$. We shall denote by B_{*-1} the chain complex that is the same as B_* but with all degrees shifted one step downward, meaning $(B_{*-1})_n = B_{n-1}$. Since R is a principal ideal domain and C_* (and therefore also its submodule $B_* \subset C_*$) is free, the exact sequence $0 \to Z_* \hookrightarrow C_* \xrightarrow{\partial} B_{*-1} \to 0$ splits, and so therefore does the sequence

$$0 \longrightarrow Z_* \otimes C'_* \longrightarrow C_* \otimes C'_* \to B_{*-1} \otimes C'_* \longrightarrow 0,$$

which is also a short exact sequence of chain complexes. One detail that is now different from the universal coefficient theorem is that the boundary operators on the first and third of these chain complexes may be nontrivial, though their homologies are still easy to write down. For instance, the submodules of homogeneous elements in $Z_* \otimes C'_*$ are

$$\left(Z_* \otimes C'_*\right)_n = \bigoplus_{k+\ell=n} Z_k \otimes C'_\ell,$$

and since Z_* is a free chain complex with trivial boundary, we can choose a basis $\mathcal{B}_k$ of each Z_k and thus obtain an isomorphism of this to

$$\left(Z_* \otimes C'_*\right)_n \cong \bigoplus_{k+\ell=n} \bigoplus_{e\in\mathcal{B}_k} R \otimes C'_\ell \cong \bigoplus_{k+\ell=n} \bigoplus_{e\in\mathcal{B}_k} C'_\ell,$$

so that $(Z_* \otimes C'_*)_n \xrightarrow{\partial} (Z_* \otimes C'_*)_{n-1}$ becomes the corresponding direct sum of the boundary maps $C'_\ell \to C'_{\ell-1}$. The homology of this complex is thus

$$H_n\left(Z_* \otimes C'_*\right) \cong \bigoplus_{k+\ell=n} \bigoplus_{e\in\mathcal{B}_k} H_\ell(C'_*) \cong \bigoplus_{k+\ell=n} Z_k \otimes H_\ell(C'_*),$$

a formula that can be written concisely as a natural isomorphism of $\mathbb{Z}$-graded R-modules

$$H_*\left(Z_* \otimes C'_*\right) \cong Z_* \otimes H_*(C'_*).$$

Since $B_{*-1} \subset Z_{*-1} \subset C_{*-1}$ is also free, the homology of $B_{*-1} \otimes C'_*$ admits a similar description, except that the degree shift gives

$$\left(B_{*-1} \otimes C'_*\right)_n = \bigoplus_{k+\ell=n} B_{k-1} \otimes C'_\ell = \bigoplus_{k+\ell=n-1} B_k \otimes C'_\ell$$

for each $n \in \mathbb{Z}$, and we therefore have natural isomorphisms

$$H_n(B_{*-1} \otimes C'_*) = H_{n-1}(B_* \otimes C'_*) \cong \bigoplus_{k+\ell=n-1} B_k \otimes H_\ell(C'_*),$$

or in concise form,

$$H_*(B_{*-1} \otimes C'_*) = H_{*-1}(B_* \otimes C'_*) \cong \left(B_* \otimes H_*(C'_*)\right)_{*-1}.$$

The short exact sequence of chain complexes now gives rise as usual to a long exact sequence of homologies

$$\ldots \to H_n(B_* \otimes C'_*) \xrightarrow{\Phi_n} H_n(Z_* \otimes C'_*) \to H_n(C_* \otimes C'_*) \to H_{n-1}(B_* \otimes C'_*) \xrightarrow{\Phi_{n-1}} H_{n-1}(Z_* \otimes C'_*) \to \ldots,$$

where the maps labeled Φ_n, Φ_{n-1} are the connecting homomorphisms, and we can then turn this into a short exact sequence centered around $H_n(C_* \otimes C'_*)$ in the usual way, namely

$$0 \to \operatorname{coker}\Phi_n \to H_n(C_* \otimes C'_*) \to \ker\Phi_{n-1} \to 0,$$

or if we take the direct sum over all $n \in \mathbb{Z}$, a short exact sequence of $\mathbb{Z}$-graded R-modules

$$0 \to \operatorname{coker}\Phi_* \to H_*(C_* \otimes C'_*) \to \ker\Phi_{*-1} \to 0. \tag{44.5}$$

Inspecting the diagram chase behind the long exact sequence reveals that the map $\Phi_* : H_*(B_* \otimes C'_*) \to H_*(Z_* \otimes C'_*)$ is the obvious thing: it is the map

$$H_*(B_* \otimes C'_*) = B_* \otimes H_*(C'_*) \xrightarrow{i_* \otimes \mathbb{1}} Z_* \otimes H_*(C'_*) = H_*(Z_* \otimes C'_*)$$

induced by the inclusions $B_* \overset{i_*}{\hookrightarrow} Z_*$. In order to understand the kernel and cokernel of $i_* \otimes \mathbb{1}$, one views the short exact sequence

$$0 \to B_* \overset{i_*}{\hookrightarrow} Z_* \xrightarrow{\text{pr}} H_*(C_*) \to 0$$

as a projective resolution of $H_*(C_*)$ in the category $\mathsf{Mod}^R_{\mathbb{Z}}$, so that the right-exactness of the functor $\otimes H_*(C'_*) : \mathsf{Mod}^R_{\mathbb{Z}} \to \mathsf{Mod}^R_{\mathbb{Z}}$ leads to a natural isomorphism

$$\operatorname{coker}(i_* \otimes \mathbb{1}) \cong H_*(C_*) \otimes H_*(C'_*),$$

and $\ker(i_* \otimes \mathbb{1})$ becomes a Tor module, namely

$$\ker(i_* \otimes \mathbb{1}) \cong \operatorname{Tor}(H_*(C_*), H_*(C'_*)).$$

The sequence we were looking for is now obtained by plugging these isomorphisms into (44.5), with attention to the downward degree shift in the third term.

The proofs of naturality and the splitting proceed as similar generalizations of the proof of the universal coefficient theorem, so we shall leave those steps as exercises. $\square$

It's worth taking special note of what the Künneth formula implies if we take R to be a field $\mathbb{K}$, so that all chain complexes in the discussion are vector spaces over $\mathbb{K}$. All such spaces are *free* $\mathbb{K}$-modules, since vector spaces always admit bases, thus

$$\operatorname{Tor}^{\mathbb{K}}(A, B) = 0 \qquad \text{for all vector spaces } A, B \text{ over } \mathbb{K},$$

and we therefore obtain:

Corollary 44.6. *For any field $\mathbb{K}$ and any two chain complexes C_* and C'_* of $\mathbb{K}$-vector spaces, the canonical map*

$$\bigoplus_{k+\ell=n} H_k(C_*) \otimes_{\mathbb{K}} H_\ell(C'_*) \to H_n(C_* \otimes_{\mathbb{K}} C'_*)$$

is a $\mathbb{K}$-linear isomorphism for every $n \in \mathbb{Z}$. $\square$

This result is one of the reasons why it is often easier to compute homology with field coefficients than over the integers.

Exercise 44.7. Using product cell complexes, describe a cell decomposition of the torus $\mathbb{T}^n$ for every $n \in \mathbb{N}$ such that the cellular boundary map vanishes. Use this to prove that for any axiomatic homology theory h_* with coefficient group G,

$$h_k(\mathbb{T}^n) \cong G^{\binom{n}{k}}$$

for all $n \in \mathbb{N}$ and $0 \leqslant k \leqslant n$.

The cross product on cellular homology. We've seen that if X and Y are CW-complexes and we assign the product cell decomposition to $X \times Y$, there is an obvious isomorphism of chain complexes of R-modules

$$C_*^{\mathrm{CW}}(X;R) \otimes_R C_*^{\mathrm{CW}}(Y;R) \to C_*^{\mathrm{CW}}(X \times Y;R) : a \otimes b \mapsto a \times b, \tag{44.6}$$

determined by the rule that for each pair of cells $e_\alpha^k \subset X$ and $e_\beta^\ell \subset Y$, $e_\alpha^k \otimes e_\beta^\ell$ is sent to the product $(k+\ell)$-cell $e_\alpha^k \times e_\beta^\ell \subset X \times Y$. The induced R-module isomorphism on homology then gives rise to an R-bilinear cross product

$$H_k^{\mathrm{CW}}(X;R) \otimes_R H_\ell^{\mathrm{CW}}(Y;R) \xrightarrow{\times} H_{k+\ell}^{\mathrm{CW}}(X \times Y;R).$$

If R is additionally a principal ideal domain, then the Künneth formula also holds, producing for each integer $n \geqslant 0$ a natural (and non-naturally split) short exact sequence of R-modules

$$0 \longrightarrow \bigoplus_{k+\ell=n} H_k^{\mathrm{CW}}(X;R) \otimes_R H_\ell^{\mathrm{CW}}(Y;R) \xrightarrow{\times} H_n^{\mathrm{CW}}(X \times Y;R)$$
$$\longrightarrow \bigoplus_{k+\ell=n-1} \mathrm{Tor}^R(H_k^{\mathrm{CW}}(X;R), H_\ell^{\mathrm{CW}}(Y;R)) \longrightarrow 0,$$

with the pleasing feature that the Tor term vanishes whenever R is taken to be a field. This exact sequence is the cellular version of the **topological Künneth formula**. We will discuss in the next lecture how to establish such an exact sequence directly in singular homology, without needing to assume that X and Y are CW-complexes.

REMARK 44.8. There is an annoying point that we've been glossing over so far in our discussion of product CW-complexes: if X and Y are two CW-complexes, then the product topology on $X \times Y$ might not always match the topology defined on $X \times Y$ via its product cell decomposition. The difference, however, is subtle: it turns out that both topologies are the same if X and Y are compact, and more generally, the two topologies always define the same notion of compact subsets in $X \times Y$, and their induced subspace topologies on any compact subset of $X \times Y$ are the same. In particular, this means that if our main concern is to determine when a map $K \to X \times Y$ from some compact space K is continuous, then both topologies give the same answer (see Exercise 44.9 below). Applying this observation for maps $\Delta^n \to X \times Y$, it follows that the singular homology of $X \times Y$ does not depend on whether we use the product topology or the CW-complex topology, hence the isomorphism $H_*(X \times Y;G) \cong H_*^{\mathrm{CW}}(X \times Y;G)$ holds as usual. With this in mind, we shall assume from now on that $X \times Y$ carries the product topology.

EXERCISE 44.9. Recall that the topology of a CW-complex X is defined normally as the strongest topology for which the characteristic maps of all cells $\Phi_\alpha : \mathbb{D}^k \to X$ are continuous. Given another CW-complex Y, let Z and Z' denote the set $X \times Y$ with two (potentially) different topologies: we assign to Z the product topology, and to Z' the topology of the product CW-complex induced by the cell decompositions of X and Y.

(a) Prove that every open set in Z is also an open set in Z', i.e. the identity map $Z' \to Z$ is continuous.
Remark: In general, the identity map $Z' \to Z$ might not be a homeomorphism![72]
(b) Prove that the identity map $Z' \to Z$ is a homeomorphism if X and Y are both compact.
(c) Prove that a subset $K \subset Z$ is compact if and only if it is compact in Z', and the two subspace topologies induced by Z and Z' on K are the same. Deduce from this that Z and Z' have the same singular homology groups.

[72]This is easily said, but writing down actual counterexamples is surprisingly difficult, e.g. it turns out that they must involve uncountable many cells. For more on such bizarre issues, see [**BT**].

EXERCISE 44.10. This problem is intended to elucidate in differential-geometric terms the intuitive reason behind the formula $\partial(e^k_\alpha \times e^\ell_\beta) = \partial e^k_\alpha \times e^\ell_\beta + (-1)^k e^k_\alpha \times \partial e^\ell_\beta$ stated in Proposition 44.1 for the boundary map on product CW-complexes.

Recall first that an **orientation** of a real n-dimensional vector space V means an equivalence class of bases, where two bases are equivalent if they are connected to each other by a continuous family of bases. The fact that the group $\mathrm{GL}(n,\mathbb{R})$ has two connected components (determined by whether the determinant is positive or negative) means that every real vector space of dimension $n > 0$ has exactly two choices of orientation.[73] On an oriented vector space, we call a basis **positive** whenever it belongs to the equivalence class determined by the orientation. A linear isomorphism $V \to W$ between two oriented vector spaces is called **orientation preserving** if it maps positive bases to positive bases, and is otherwise **orientation reversing**.

A smooth n-manifold M has a **tangent space** T_xM at every point x, which is an n-dimensional vector space. If you haven't seen this notion in differential geometry, then you should just picture M as a regular level-set $f^{-1}(0) \subset \mathbb{R}^k$ of some smooth function $f : \mathbb{R}^k \to \mathbb{R}^{k-n}$ for some $k \in \mathbb{N}$; a famous theorem of Whitney says that every smooth n-manifold can be described in this way if $k \geqslant 2n$. The tangent space T_xM at each point $x \in M$ is then the n-dimensional linear subspace $\ker df(x) \subset \mathbb{R}^k$. With this notion understood, an **orientation of** M means a choice of orientation for every tangent space T_xM such that the orientations vary continuously with x, i.e. every point $x_0 \in M$ has a neighborhood $\mathcal{U} \subset M$ admitting a continuous family of bases $\{(v_1(x), \dots, v_n(x))\}_{x \in \mathcal{U}}$ of the tangent spaces T_xM such that all of them are positive. If M and N are smooth manifolds of the same dimension, then any smooth map $f : M \to N$ has a derivative $df(x) : T_xM \to T_{f(x)}N$ at every point $x \in M$, and we call f an **immersion** if $df(x)$ is an isomorphism for every $x \in M$. If M and N are both oriented, then an immersion $f : M \to N$ is called **orientation preserving/reversing** if $df(x) : T_xM \to T_{f(x)}N$ is orientation preserving/reversing for every $x \in M$.

(a) Convince yourself that S^2 admits an orientation (i.e. it is **orientable**), but $\mathbb{RP}^2$ and the Klein bottle do not.

If V and W are both oriented vector spaces, we define the **product orientation** of $V \oplus W$ to be the one such that if $(v_1, \dots, v_n)$ and $(w_1, \dots, w_m)$ are positive bases of V and W respectively, then $(v_1, \dots, v_n, w_1, \dots, w_m)$ is a positive basis of $V \oplus W$. This notion carries over immediately to a product of manifolds M and N since for each $(x, y) \in M \times N$, $T_{(x,y)}(M \times N)$ can be naturally identified with $T_xM \oplus T_yN$, hence orientations of M and N give rise to a product orientation of $M \times N$.

(b) Show that if M and N are oriented manifolds of dimensions m and n respectively, then for the natural product orientations, the map $M \times N \to N \times M : (x, y) \mapsto (y, x)$ is orientation preserving if either m or n is even, and orientation reversing if both m and n are odd.

If M is an n-manifold with boundary, then its boundary ∂M is naturally an $(n-1)$-manifold, and for each $x \in \partial M$, the tangent space $T_x(\partial M)$ is naturally a codimension 1 linear subspace of T_xM. The set $T_xM \backslash T_x(\partial M)$ thus has two connected components, characterized as the tangent vectors in T_xM that point "outward" or "inward" with respect to the boundary. Now if M has an orientation, this induces on ∂M the so-called **boundary orientation**, defined such that for any choice of *outward* pointing vector $\nu \in T_xM$, a basis $(X_1, \dots, X_{n-1})$ of $T_x(\partial M)$ is positive (with respect to the orientation of ∂M) if and only if the basis $(\nu, X_1, \dots, X_{n-1})$ of T_xM is positive with respect to the orientation of M. Take a moment to convince yourself that this notion is well defined.

[73] Dimension zero must always be treated as a special case in orientation discussions. For this informal discussion we make our lives easier by assuming all dimensions are positive.

The simplest example is also the most relevant for our discussion of cell complexes: the closed n-disk $\mathbb{D}^n$ is a compact n-dimensional smooth manifold with boundary $\partial\mathbb{D}^n = S^{n-1}$. Since all the tangent spaces to $\mathbb{D}^n$ are canonically isomorphic to $\mathbb{R}^n$, $\mathbb{D}^n$ has a canonical orientation, and this determines a canonical orientation for S^{n-1}.

Finally, consider a product $M \times N$ of two smooth manifolds with boundary, with dimensions m and n respectively. This is a slightly more general object called a "smooth manifold with boundary and corners"; rather than defining this notion precisely, let us simply agree that in the complement of the "corner" $\partial M \times \partial N$, the object $M \times N$ is a smooth manifold whose boundary $\partial(M \times N)$ is the union of two smooth manifolds $\partial M \times N$ and $M \times \partial N$ of dimension $m + n - 1$. The question is: what orientations should these two pieces of $\partial(M \times N)$ carry?

(c) Assume M and N are both oriented, $M \times N$ is endowed with the resulting product orientation and ∂M and ∂N are each endowed with the boundary orientation. Show that the induced boundary orientation on $\partial(M \times N)$ always matches the product orientation of $\partial M \times N$, and that it matches the product orientation of $M \times \partial N$ if and only if m is even.

Remark: The result of part (c) can be summarized as follows. If M has an orientation and we denote the same manifold with the opposite orientation by $-M$, then for any two oriented manifolds M and N of dimensions m and n respectively,

$$\partial(M \times N) = (\partial M \times N) \cup (-1)^m (M \times \partial N).$$

If you apply this to the case $M = \mathbb{D}^m$ and $N = \mathbb{D}^n$ and consider that the degree of a map $S^k \to S^k$ changes sign if you compose it with an orientation-reversing homeomorphism, you may now be able to imagine the reason for the sign in the cellular boundary formula $\partial(e^k_\alpha \times e^\ell_\beta) = \partial e^k_\alpha \times e^\ell_\beta + (-1)^k e^k_\alpha \times \partial e^\ell_\beta$.

45. The singular cross product (December 22, 2023)

As in the previous lecture, let us fix a commutative coefficient ring R and assume by default that $\otimes$ means the tensor product of R-modules. Since cellular and singular homologies are isomorphic, the cellular cross product determines a homomorphism

$$H_k(X) \otimes H_\ell(Y) \xrightarrow{\times} H_{k+\ell}(X \times Y) \tag{45.1}$$

whenever X and Y come with cell decompositions. It is far from obvious at this stage whether $\times$ is independent of the choices of cell decompositions of X and Y. We shall deal with this by replacing the cellular cross product with an operation on singular homology that can be defined without reference to any cell decompositions. It should be emphasized that the construction we are about to give is distinctly for *singular* homology, i.e. it relies on the definition of H_* and not just on the Eilenberg-Steenrod axioms, so it does not give us anything for more general axiomatic homology theories. This does not mean that a cross product on other homology theories cannot be defined, but only that it must be defined for each theory separately, with the final step being to prove that it matches the cellular cross product when applied to CW-complexes.

There are good geometric reasons to expect that a product map (45.1) should exist. If you like to think about elements of $H_k(X;\mathbb{Z})$ as represented by closed oriented k-dimensional submanifolds $M \subset X$ as in Lecture 30, then since the product of two closed oriented manifolds is also a closed oriented manifold, it would make sense to define

$$[M] \times [N] := [M \times N] \in H_{k+\ell}(X \times Y;\mathbb{Z})$$

for a k-manifold $M \subset X$ and ℓ-manifold $N \subset Y$. It will be easy to see that the singular cross product has this property when $[M]$ and $[N]$ are defined via oriented triangulations, and we will be able to generalize this to a statement independent of triangulations once we have learned how to define fundamental classes on topological manifolds in general. But not every singular homology class can be represented by a submanifold, so the question remains: how should (45.1) be defined in general?

Since there is always a canonical homomorphism $H_*(X) \otimes H_*(Y) \to H_*(C_*(X) \otimes C_*(Y))$, we would obtain a map (45.1) if we had a chain map

$$C_*(X) \otimes C_*(Y) \xrightarrow{\times} C_*(X \times Y)$$

to play the role in singular homology that (44.6) plays in cellular homology. In order to write down such a map, we need to decide what $\sigma \times \tau \in C_{k+\ell}(X \times Y; \mathbb{Z})$ should mean if we are given a pair of singular simplices $\sigma : \Delta^k \to X$ and $\tau : \Delta^\ell \to Y$. Unfortunately, $\Delta^k \times \Delta^\ell$ is not a simplex in any canonical way, so we cannot simply write down the continuous map

$$\sigma \times \tau : \Delta^k \times \Delta^\ell \to X \times Y : (s, t) \mapsto (\sigma(s), \tau(t)) \tag{45.2}$$

and call it a generator of $C_{k+\ell}(X \times Y; \mathbb{Z})$. But we've dealt with this kind of thing before using subdivision: a natural approach is to fix a reasonable oriented triangulation of $\Delta^k \times \Delta^\ell$ for every pair of integers $k, \ell \geqslant 0$, giving rise to a relative fundamental cycle $c_{\Delta^k \times \Delta^\ell} \in C_{k+\ell}(\Delta^k \times \Delta^\ell; \mathbb{Z})$, and then use the continuous map (45.2) to push this fundamental cycle forward, defining

$$\sigma \times \tau := (\sigma \times \tau)_* c_{\Delta^k \times \Delta^\ell} \in C_{k+\ell}(X \times Y; \mathbb{Z}).$$

What this does in practice is make $\sigma \times \tau \in C_{k+\ell}(X \times Y; \mathbb{Z})$ a sum of singular simplices obtained by restricting the map (45.2) to the $(k+\ell)$-simplices in the triangulation, and if $\sigma \times \tau$ can be defined in this way for the chain complex with integer coefficients, then the definition extends immediately to coefficients in any commutative ring R. A good triangulation to use for this purpose was described in Lecture 31, and the formula in Exercise 31.19 for $\partial c_{\Delta^k \times \Delta^\ell}$ guarantees that the unique R-module homomorphism $C_*(X) \otimes C_*(Y) \xrightarrow{\times} C_*(X \times Y)$ obtained by defining $\sigma \times \tau$ in this way will be a chain map. This will serve as our first (but not last) definition of the chain-level singular cross product. The definition has a strong geometric advantage: if M and N are closed manifolds with oriented triangulations, then the cross product of their fundamental cycles by this definition will be the fundamental cycle for a triangulation of $M \times N$, thus justifying the formula $[M] \times [N] = [M \times N]$.

One may wonder, of course, whether defining cross products via the particular subdivision algorithm in Lecture 31 is the only sensible way to do things: there might be other good subdivision algorithms that produce different definitions of a chain-level cross product. This will be okay if it turns out that the dependence on choices disappears after descending from chain complexes to homology. In practice, the most convenient way to see how this works is to avoid mentioning triangulations at all, but instead employ an algebraic trick that accomplishes the same result: the trick is known as the method of **acyclic models**. One can find in various textbooks (e.g. [**Vic94**, **Spa95**]) a result called the *acyclic model theorem*, which is applicable to a wide variety of problems but difficult to digest, as it is typically expressed in heavily abstract category-theoretic language. We shall instead follow the approach of [**Bre93**] and demonstrate the method by example.

Two preparatory comments are in order before we continue. The first is that if we can define a cross product on singular homology with integer coefficients, then the definition and many (though not quite all) of its important properties will almost immediately extend to coefficients in an arbitrary commutative ring R. The main reason for this is a basic algebraic observation: if A_* and B_* are two chain complexes of abelian groups, then $A_* \otimes R$ and $B_* \otimes R$ can be regarded as chain complexes of R-modules by letting R act on the second factor in the tensor product, and one easily checks that the map

$$(A_* \otimes R) \otimes_R (B_* \otimes R) \to (A_* \otimes B_*) \otimes R : (a \otimes r) \otimes (b \otimes s) \mapsto (a \otimes b) \otimes (rs) \tag{45.3}$$

is then an isomorphism of chain complexes of R-modules. When applied to the singular chain complexes of two spaces X and Y, this means that the product chain complex of R-modules $C_*(X;R) \otimes_R C_*(Y;R)$ is completely equivalent to the complex constructed by first writing down the product complex of abelian groups $C_*(X;\mathbb{Z}) \otimes C_*(Y;\mathbb{Z})$ and then tensoring the latter with R.

The second preparatory comment involves a fact about reduced homology that was mentioned in Lecture 29 but has not really been used so far. Recall that the reduced singular homology $\widetilde{H}_*(X;G)$ of a space X with coefficient group G is also the homology of the so-called *augmented* chain complex $\widetilde{C}_*(X;G)$, defined by appending an extra nonzero term $\widetilde{C}_{-1}(X;G) := G$ to the end of the usual singular chain complex:

$$\ldots \longrightarrow C_2(X;G) \xrightarrow{\partial_2} C_1(X;G) \xrightarrow{\partial_1} C_0(X;G) \xrightarrow{\epsilon} G \longrightarrow 0 \longrightarrow 0 \longrightarrow \ldots.$$

The **augmentation** is the homomorphism $\epsilon : C_0(X;G) \to G$, which is determined by the property that on each of the generators $\sigma : \Delta^0 \to X$ of $C_0(X;\mathbb{Z})$, $\epsilon(\sigma) = 1$ (see Exercise 29.15). This makes ϵ a surjective map with $\epsilon \circ \partial_1 = 0$.

LEMMA 45.1. *One can assign to every tuple of topological spaces (X, Y) a chain map*

$$\Phi : C_*(X;\mathbb{Z}) \otimes C_*(Y;\mathbb{Z}) \to C_*(X \times Y;\mathbb{Z})$$

that satisfies $\Phi(x \otimes y) = (x, y)$ on 0-chains under the canonical identification of singular 0-simplices with points, and is natural in the sense that for any continuous maps $f : X \to X'$ and $g : Y \to Y'$, the diagram

$$\begin{array}{ccc} C_*(X;\mathbb{Z}) \otimes C_*(Y;\mathbb{Z}) & \xrightarrow{\Phi} & C_*(X \times Y;\mathbb{Z}) \\ \downarrow{\scriptstyle f_* \otimes g_*} & & \downarrow{\scriptstyle (f \times g)_*} \\ C_*(X';\mathbb{Z}) \otimes C_*(Y';\mathbb{Z}) & \xrightarrow{\Phi} & C_*(X' \times Y';\mathbb{Z}) \end{array}$$

commutes. Moreover, Φ with these properties is unique up to chain homotopy.

PROOF. In the following proof, the coefficient group is always assumed to be $\mathbb{Z}$ but will be omitted from the notation. We observe first that if $\Phi : C_0(X) \otimes C_0(Y) \to C_0(X \times Y)$ is defined as required, then it trivially satisfies the chain map relation $\Phi \circ \partial = \partial \circ \Phi$ on chains of degree 0 since they are all annihilated by the boundary maps, and it also satisfies the naturality condition

$$(f \times g)_* \Phi(x \otimes y) = (f(x), g(y)) = \Phi(f_* \otimes g_*)(x \otimes y)$$

for any maps $f : X \to X'$, $g : Y \to Y'$ and points $x \in X$, $y \in Y$ (regarded as singular 0-simplices). We shall now argue by induction and assume that maps $\Phi : C_k(X) \otimes C_\ell(Y) \to C_{k+\ell}(X \times Y)$ have been defined for all spaces X, Y and all integers $k, \ell \geqslant 0$ with $k + \ell \leqslant n - 1$ for some $n \geqslant 1$, such that the chain map and naturality conditions are satisfied on chains up to degree $n-1$. To extend this to chains of degree n, we start by defining Φ on a particular collection of models: for each integer $k \geqslant 0$, let $i_k : \Delta^k \to \Delta^k$ denote the identity map on the standard k-simplex, and regard this as a singular k-chain in the space Δ^k:

$$i_k \in C_k(\Delta^k).$$

Given integers $k, \ell \geqslant 0$ with $k + \ell = n$, let us consider $i_k \otimes i_\ell \in C_k(\Delta^k) \otimes C_\ell(\Delta^\ell)$ and try to define

$$\Phi(i_k \otimes i_\ell) \in C_n(\Delta^k \times \Delta^\ell).$$

To satisfy the chain map relation, $\Phi(i_k \otimes i_\ell)$ needs to have the property that

$$\partial\Phi(i_k \otimes i_\ell) = \Phi(\partial(i_k \otimes i_\ell)) \in C_{n-1}(\Delta^k \times \Delta^\ell), \tag{45.4}$$

where $\Phi(\partial(i_k \otimes i_\ell))$ is given by the inductive hypothesis since Φ has already been defined on chains up to degree $n-1$. Since it also satisfies the chain map relation up to degree $n-1$, we have

$$\partial\Phi\partial(i_k \otimes i_\ell) = \Phi\partial^2(i_k \otimes i_\ell) = 0, \tag{45.5}$$

so $\Phi\partial(i_k \otimes i_\ell)$ is a singular $(n-1)$-cycle in $\Delta^k \times \Delta^\ell$. This is a vacuous statement when $n = 1$, but in this case it can also be improved: letting $\epsilon : C_0(\Delta^k \times \Delta^\ell) \to \mathbb{Z}$ denote the augmentation in the augmented chain complex $\widetilde{C}_*(\Delta^k \times \Delta^\ell)$, we observe that if $k = 1$ and $\ell = 0$, then $\partial(i_1 \otimes i_0) = \partial i_1 \otimes i_0$ is a sum of two generators of $C_0(\Delta^1) \otimes C_0(\Delta^0)$ with coefficients 1 and -1 respectively, so $\Phi\partial(i_1 \otimes i_0)$ is similarly a sum of two generators with coefficients 1 and -1. The same holds in the case $k = 0$ and $\ell = 1$, proving that in either case,

$$\epsilon\Phi\partial(i_k \otimes i_\ell) = 0 \quad \text{when} \quad n = 1. \tag{45.6}$$

Now comes the crucial point: $\Delta^k \times \Delta^\ell$ is contractible, so its reduced singular homology is trivial. In light of (45.5) and (45.6), this means

$$[\Phi\partial(i_k \otimes i_\ell)] = 0 \in \widetilde{H}_{n-1}(\Delta^k \times \Delta^\ell),$$

implying $\Phi\partial(i_k \otimes i_\ell)$ is in the image of $C_n(\Delta^k \times \Delta^\ell) \xrightarrow{\partial} C_{n-1}(\Delta^k \times \Delta^\ell)$, hence the relation (45.4) has solutions, and we can define $\Phi(i_k \otimes i_\ell) \in C_n(\Delta^k \times \Delta^\ell)$ to be any element such that

$$\Phi(i_k \otimes i_\ell) \in \partial^{-1}\big(\Phi\partial(i_k \otimes i_\ell)\big). \tag{45.7}$$

This is an arbitrary choice, but such an element certainly exists.

Having chosen $\Phi(i_k \otimes i_\ell) \in C_n(\Delta^k \times \Delta^\ell)$ for every $k, \ell \geqslant 0$ with $k + \ell = n$, we claim that the general extension of $\Phi : C_*(X) \otimes C_*(Y) \to C_*(X \times Y)$ to all chains of degree n is uniquely determined by the naturality condition. Indeed, given any pair of spaces X and Y and singular simplices $\sigma : \Delta^k \to X$ and $\tau : \Delta^\ell \to Y$ with $k + \ell = n$, we have

$$\sigma = \sigma_* i_k \in C_k(X), \qquad \tau = \tau_* i_\ell \in C_\ell(Y),$$

so naturality requires $\Phi : C_k(X) \otimes C_\ell(Y) \to C_n(X \times Y)$ to have the property that

$$\Phi(\sigma \otimes \tau) = \Phi(\sigma_* \otimes \tau_*)(i_k \otimes i_\ell) = (\sigma \times \tau)_*\Phi(i_k \otimes i_\ell).$$

Let us take this as a *definition* of $\Phi(\sigma \otimes \tau)$, and verify that Φ now satisfies all the required properties on chains up to degree n. Keeping σ and τ as above, the fact that $\sigma_* : C_*(\Delta^k) \to C_*(X)$, $\tau_* : C_*(\Delta^\ell) \to C_*(Y)$ and $(\sigma \times \tau)_* : C_*(\Delta^k \times \Delta^\ell) \to C_*(X \times Y)$ are chain maps and the naturality of Φ up to degree $n-1$ implies

$$\begin{aligned}
\partial\Phi(\sigma \otimes \tau) &= \partial(\sigma \times \tau)_*\Phi(i_k \otimes i_\ell) = (\sigma \times \tau)_*\partial\Phi(i_k \otimes i_\ell) \\
&= (\sigma \times \tau)_*\Phi\partial(i_k \otimes i_\ell) = \Phi(\sigma_* \otimes \tau_*)\partial(i_k \otimes i_\ell) \\
&= \Phi\partial(\sigma_* \otimes \tau_*)(i_k \otimes i_\ell) = \Phi\partial(\sigma \otimes \tau),
\end{aligned}$$

where we have also used the fact that the tensor product of two chain maps induces a chain map on the tensor product chain complex (see (44.3)). This establishes the chain map property. To see that naturality also holds, consider two continuous maps $f : X \to X'$ and $g : Y \to Y'$: then

$$\begin{aligned}
\Phi(f_* \otimes g_*)(\sigma \otimes \tau) &= \Phi((f \circ \sigma)_* \otimes (g \circ \tau)_*)(i_k \otimes i_\ell) = \big((f \circ \sigma) \times (g \circ \tau)\big)_* \Phi(i_k \otimes i_\ell) \\
&= (f \times g)_*(\sigma \times \tau)_*\Phi(i_k \otimes i_\ell) = (f \times g)_*\Phi(\sigma \otimes \tau).
\end{aligned}$$

This completes the inductive step and thus proves the existence of the natural chain map Φ.

The same approach will establish uniqueness up to chain homotopy. Assuming Φ and Ψ are two natural chain maps as in the statement of the theorem, we would like to associate to each pair of spaces X and Y a collection of maps

$$h : C_k(X) \otimes C_\ell(Y) \to C_{n+1}(X \times Y)$$

for every pair of integers $k, \ell \geqslant 0$ and $n = k + \ell$, such that

$$\partial h + h\partial = \Phi - \Psi.$$

We claim that this can be done so that the obvious naturality property is also satisfied, i.e. so that the diagram

$$\begin{tikzcd} C_k(X) \otimes C_\ell(Y) \arrow[r, "h"] \arrow[d, "f_* \otimes g_*"] & C_{n+1}(X \times Y) \arrow[d, "(f \times g)_*"] \\ C_k(X') \otimes C_\ell(Y') \arrow[r, "h"] & C_{n+1}(X' \times Y') \end{tikzcd}$$

commutes for every pair of continuous maps $f : X \to X'$ and $g : Y \to Y'$.

Since Φ and Ψ match precisely on all 0-chains, we are free to define $h : C_0(X) \otimes C_0(Y) \to C_1(X \times Y)$ as the trivial map, and the naturality property is obviously also satisfied for this choice. Now by induction, assume h has been defined so as to satisfy both the chain map relation and naturality on all chains up to degree $n - 1$ for some $n \geqslant 1$. To extend this to degree n, we proceed as before by trying first to define h on the models $i_k \otimes i_\ell \in C_k(\Delta^k) \otimes C_\ell(\Delta^\ell)$ for $k + \ell = n$. We need $h(i_k \otimes i_\ell) \in C_{n+1}(\Delta^k \times \Delta^\ell)$ to satisfy

$$\partial h(i_k \otimes i_\ell) = (-h\partial + \Phi - \Psi)(i_k \otimes i_\ell),$$

where the right hand side is already determined since $\partial(i_k \otimes i_\ell)$ has degree $n - 1$. Applying ∂ to the right hand side, we use the chain homotopy relation in degree $n - 1$ and the fact that Φ and Ψ are chain maps to prove

$$\partial(-h\partial + \Phi - \Psi)(i_k \otimes i_\ell) = (-\partial h + \Phi - \Psi)\partial(i_k \otimes i_\ell) = (h\partial)\partial(i_k \otimes i_\ell) = 0,$$

hence $(-h\partial + \Phi - \Psi)(i_k \otimes i_\ell)$ is a cycle in $C_n(\Delta^k \times \Delta^\ell)$. It is therefore also a boundary since $H_n(\Delta^k \times \Delta^\ell) = 0$, so we can define $h(i_k \otimes i_\ell) \in C_{n+1}(\Delta^k \times \Delta^\ell)$ to be any element satisfying

$$h(i_k \otimes i_\ell) \in \partial^{-1}\big((-h\partial + \Phi - \Psi)(i_k \otimes i_\ell)\big).$$

Now we extend this definition to all possible $\sigma \otimes \tau \in C_k(X) \otimes C_\ell(Y)$ by requiring naturality, i.e. we define $h(\sigma \otimes \tau) \in C_{n+1}(X \times Y)$ by

$$h(\sigma \otimes \tau) = h(\sigma_* \otimes \tau_*)(i_k \otimes i_\ell) := (\sigma \times \tau)_* h(i_k \otimes i_\ell).$$

We must then check that the chain homotopy relation is satisfied on $\sigma \otimes \tau$, and indeed, we have

$$\begin{aligned} (\partial h + h\partial)(\sigma \otimes \tau) &= \partial(\sigma \times \tau)_* h(i_k \otimes i_\ell) + h\partial(\sigma_* \otimes \tau_*)(i_k \otimes i_\ell) \\ &= (\sigma \times \tau)_* \partial h(i_k \otimes i_\ell) + h(\sigma_* \otimes \tau_*)\partial(i_k \otimes i_\ell) \\ &= (\sigma \times \tau)_*(\partial h + h\partial)(i_k \otimes i_\ell) = (\sigma \times \tau)_*(\Phi - \Psi)(i_k \otimes i_\ell) \\ &= (\Phi - \Psi)(\sigma_* \otimes \tau_*)(i_k \otimes i_\ell) = (\Phi - \Psi)(\sigma \otimes \tau), \end{aligned}$$

where we've used the fact that $(\sigma \times \tau)_*$ and $\sigma_* \otimes \tau_*$ are chain maps, the naturality of h on $(n-1)$-chains, and the naturality of Φ and Ψ. Finally, we need to verify that our definition of h on n-chains satisfies naturality: given $f : X \to X'$ and $g : Y \to Y'$, we have

$$\begin{aligned} h(f_* \otimes g_*)(\sigma \otimes \tau) &= h\big((f \circ \sigma)_* \otimes (g \circ \tau)_*\big)(i_k \otimes i_\ell) = \big((f \circ \sigma) \times (g \circ \tau)\big)_* h(i_k \otimes i_\ell) \\ &= (f \times g)_*(\sigma \times \tau)_* h(i_k \otimes i_\ell) = (f \times g)_* h(\sigma \otimes \tau). \end{aligned}$$

This completes the inductive step and finishes the proof. □

The proof above was a bit long, but not conceptually difficult once the basic idea is understood, and we will need to make use of this idea several more times. The general pattern is always as follows. We want to define a chain map that is typically not unique or canonical, but should take a specific form on 0-chains and should also be "natural" in the sense of category theory; the latter is always a precise condition that can be expressed in terms of commutative diagrams. We then proceed by induction on the degree of the chains, where at each step in the induction, we start by trying to define the map on a specific set of "models," which are **acyclic** in the sense that their (reduced) homology vanishes. The latter makes it possible to define our map on the models so that the required conditions are satisfied, and the rest of the definition is then uniquely determined by naturality. Having extended the definition up by one degree in this way, we must then check that it still satisfies both the chain map and the naturality conditions. With this induction complete, one can then use the same approach again to prove that any two chain maps with the required properties are chain homotopic. I wanted to show you one example of this method with every step worked out in detail, but when I need to use this from now on, I will typically only tell you the main idea and leave the remaining details as exercises.

Thanks to the canonical isomorphism (45.3), the chain map $\Phi : C_*(X;\mathbb{Z}) \otimes C_*(Y;\mathbb{Z}) \to C_*(X \times Y;\mathbb{Z})$ from Lemma 45.1 uniquely determines a chain map of R-modules

$$C_*(X;R) \otimes_R C_*(Y;R) \to C_*(X \times Y;R)$$

for any commutative ring R, defined as $\Phi \otimes \mathbb{1}$ after identifying $C_*(X;R) \otimes_R C_*(Y;R)$ with $(C_*(X;\mathbb{Z}) \otimes C_*(Y;\mathbb{Z})) \otimes R$. For any two choices of the chain map Φ, a chain homotopy between them can similarly be tensored with $\mathbb{1} : R \to R$ to produce a chain homotopy between the resulting chain maps $C_*(X;R) \otimes_R C_*(Y;R) \to C_*(X \times Y;R)$. The induced map on homology is therefore independent of choices, and it can be composed with the canonical map from $H_*(X;R) \otimes_R H_*(Y;R)$ to define what we will henceforth call the **singular cross product**

$$\begin{tikzcd} H_*(X;R) \otimes_R H_*(Y;R) \arrow[r] \arrow[rr, bend right, "\times"'] & H_*\big(C_*(X;R) \otimes_R C_*(Y;R)\big) \arrow[r, "\Phi_*"] & H_*(X \times Y;R). \end{tikzcd}$$

This definition is not only independent of choices, but is also natural in the sense that there is a commutative diagram

$$\begin{tikzcd} H_*(X;R) \otimes_R H_*(Y;R) \arrow[r, "\times"] \arrow[d, "f_* \otimes g_*"] & H_*(X \times Y;R) \arrow[d, "(f \times g)_*"] \\ H_*(X';R) \otimes_R H_*(Y';R) \arrow[r, "\times"] & H_*(X' \times Y';R) \end{tikzcd}$$

for any pair of continuous maps $f : X \to X'$ and $g : Y \to Y'$.

Before we can feed this into the algebraic Künneth formula as we did with cellular homology, there is a missing ingredient. The cellular version of $\Phi : C_*(X;R) \otimes_R C_*(Y;R) \to C_*(X \times Y;R)$ was not only a chain map, but was in fact an *isomorphism* of chain complexes, which allowed us to replace the homology of a tensor product of chain complexes in the Künneth formula with the cellular homology of a product CW-complex. There is no obvious reason why Φ should be an isomorphism, except on 0-chains, for which it clearly is one; moreover, the cellular counterpart of Φ was canonically defined, whereas Φ itself depends on many choices and is canonical only up to chain homotopy. What we can therefore reasonably expect is for Φ to be a chain homotopy equivalence. This is where the method of acyclic models really demonstrates its power.

Lemma 45.2. *One can assign to every tuple of topological spaces (X, Y) chain maps*

$$C_*(X \times Y;\mathbb{Z}) \xrightarrow{\theta} C_*(X;\mathbb{Z}) \otimes C_*(Y;\mathbb{Z}),$$

$$C_*(X;\mathbb{Z}) \otimes C_*(Y;\mathbb{Z}) \xrightarrow{\alpha} C_*(X;\mathbb{Z}) \otimes C_*(Y;\mathbb{Z}),$$

$$C_*(X \times Y;\mathbb{Z}) \xrightarrow{\beta} C_*(X \times Y;\mathbb{Z}),$$

which are uniquely determined up to chain homotopy by a naturality condition and their definitions on 0*-chains,*

$$\theta(x,y) = x \otimes y, \qquad \alpha(x \otimes y) = x \otimes y, \qquad \beta(x,y) = (x,y).$$

Here, naturality of θ *means that there is a commutative diagram*

$$\begin{array}{ccc} C_*(X \times Y;\mathbb{Z}) & \xrightarrow{\theta} & C_*(X;\mathbb{Z}) \otimes C_*(Y;\mathbb{Z}) \\ \Big\downarrow{\scriptstyle (f\times g)_*} & & \Big\downarrow{\scriptstyle f_* \otimes g_*} \\ C_*(X' \times Y';\mathbb{Z}) & \xrightarrow{\theta} & C_*(X';\mathbb{Z}) \otimes C_*(Y';\mathbb{Z}) \end{array}$$

for any pair of continuous maps $f : X \to X'$ *and* $g : Y \to Y'$*, and naturality is defined similarly for* α *and* β*.*

Notice that for each of the last two maps, the identity is an example of a map satisfying the required conditions, and so are the compositions $\Phi \circ \theta$ and $\theta \circ \Phi$, thus the uniqueness up to chain homotopy implies that Φ and θ are chain homotopy inverses. This proves:

Corollary 45.3 (Eilenberg-Zilber theorem). *The natural chain maps*

$$C_*(X;\mathbb{Z}) \otimes C_*(Y;\mathbb{Z}) \underset{\theta}{\overset{\Phi}{\rightleftarrows}} C_*(X \times Y;\mathbb{Z})$$

are chain homotopy inverses, and are thus both chain homotopy equivalences. Moreover, for any commutative ring R*, the same holds for the chain maps of* R*-modules between* $C_*(X;R) \otimes_R C_*(Y;R)$ *and* $C_*(X \times Y;R)$ *obtained by tensoring* Φ *and* θ *with* $\mathbb{1} : R \to R$*.* □

Proof of Lemma 45.2. As before, we shall omit the coefficient group $\mathbb{Z}$ from the notation, but the fact that we are using this particular coefficient group will be relevant for the following reason: since $\mathbb{Z}$ is a principal ideal domain, the Künneth formula holds for chain complexes of abelian groups. The statement of the lemma uniquely specifies the definitions of the desired chain maps on 0-chains, and these clearly satisfy the naturality condition, so we use the method of acyclic models to extend the definition to chains of all degrees $n \geqslant 1$ by induction on n. For $\theta : C_*(X \times Y) \to C_*(X) \otimes C_*(Y)$, assume we already have a definition on $C_k(X \times Y)$ for all $k = 0, \dots, n-1$. We extend it to n-chains starting with the model

$$d_n : \Delta^n \to \Delta^n \times \Delta^n : t \mapsto (t,t),$$

interpreted as an element in $C_n(\Delta^n \times \Delta^n)$. The definition of $\theta(d_n) \in \bigoplus_{k+\ell=n} C_k(\Delta^n) \otimes C_\ell(\Delta^n)$ should be chosen to satisfy

$$\partial\theta(d_n) = \theta(\partial d_n) \in \bigoplus_{k+\ell=n-1} C_k(\Delta^n) \otimes C_\ell(\Delta^n),$$

where the right hand side is already determined since ∂d_n has degree $n-1$. To see if this is possible, we observe that since θ is a chain map up to degree $n-1$,

$$\partial(\theta\partial d_n) = \theta\partial^2(d_n) = 0,$$

so $\theta\partial d_n$ is an $(n-1)$-cycle in $C_*(\Delta^n) \otimes C_*(\Delta^n)$. Now observe that since Δ^n is contractible, the algebraic Künneth formula implies

$$H_m\big(C_*(\Delta^n) \otimes C_*(\Delta^n)\big) \cong \bigoplus_{k+\ell=m} H_k(\Delta^n) \otimes H_\ell(\Delta^n) \cong \begin{cases} \mathbb{Z} & \text{if } m = 0, \\ 0 & \text{otherwise,} \end{cases}$$

where all the Tor terms have vanished because every $H_k(\Delta^n)$ is a free abelian group. In particular this implies that the cycle $\theta\partial d_n$ is also a boundary if $n \geqslant 2$, and we can therefore choose $\theta(d_n)$ to satisfy

$$\theta(d_n) \in \partial^{-1}(\theta\partial d_n). \tag{45.8}$$

The case $n = 1$ is special since $H_0(\Delta^n) \otimes H_0(\Delta^n) = \mathbb{Z}$ is not trivial, but if we identify Δ^1 with the unit interval $I = [0,1]$, then it is easy to check that

$$\theta\partial(d_1) = \theta\big((1,1) - (0,0)\big) = 1 \otimes 1 - 0 \otimes 0$$

is a boundary, e.g. of $1 \otimes i_1 + i_1 \otimes 0$ if $i_1 \in C_1(\Delta^1)$ is the singular 1-simplex given by the identity map.[74] In either case, $\theta(d_n)$ can be defined so that (45.8) holds.

Now for an arbitrary singular n-simplex $\sigma : \Delta^n \to X \times Y$, we can use the projection maps $\pi_X : X \times Y \to X$ and $\pi_Y : X \times Y \to Y$ to write

$$\begin{tikzcd} \Delta^n \arrow[r, "d_n"] \arrow[rr, bend right, "\sigma"'] & \Delta^n \times \Delta^n \arrow[r, "(\pi_X\circ\sigma)\times(\pi_Y\circ\sigma)"] & X \times Y, \end{tikzcd}$$

so naturality requires that we define

$$\theta(\sigma) = \theta\left((\pi_X \circ \sigma) \times (\pi_Y \circ \sigma)\right)_* d_n := \left((\pi_X \circ \sigma)_* \otimes (\pi_Y \circ \sigma)_*\right) \theta(d_n).$$

It is then a straightforward matter to check that this extension of θ to all n-chains satisfies the chain map and naturality conditions, and one can use the same method to construct a chain homotopy between any two such natural chain maps. We leave these steps as exercises, along with the uniqueness up to chain homotopy of α and β, as none of these steps require any new ideas. □

REMARK 45.4. I will give you the same advice about acyclic models that I typically give about diagram chasing: the next time you find yourself bored on a long flight or train ride, finish the proof of Lemma 45.2. It's relaxing.

Corollary 45.3 implies that for any commutative ring R, the natural map

$$\Phi_* : H_*\big(C_*(X;R) \otimes_R C_*(Y;R)\big) \to H_*(X \times Y; R)$$

used in the definition of the singular cross product is an isomorphism, so we can now use it to replace the middle term in the algebraic Künneth formula, proving:

COROLLARY 45.5 (topological Künneth formula). *For any principal ideal domain R, any spaces X, Y and every integer $n \geqslant 0$, the singular cross product fits into a natural short exact sequence*

$$0 \longrightarrow \bigoplus_{k+\ell=n} H_k(X;R) \otimes_R H_\ell(Y;R) \xrightarrow{\times} H_n(X \times Y; R)$$
$$\longrightarrow \bigoplus_{k+\ell=n-1} \operatorname{Tor}(H_k(X;R), H_\ell(Y;R)) \longrightarrow 0,$$

and the sequence splits (but not naturally). □

In particular, we can always choose field coefficients to make the Tor terms vanish:

COROLLARY 45.6. *For any spaces X and Y and any field $\mathbb{K}$, the cross product on singular homology with coefficients in $\mathbb{K}$ defines natural $\mathbb{K}$-vector space isomorphisms*

$$\times : \bigoplus_{k+\ell=n} H_k(X;\mathbb{K}) \otimes_{\mathbb{K}} H_\ell(Y;\mathbb{K}) \xrightarrow{\cong} H_n(X \times Y; \mathbb{K}).$$

for every integer $n \geqslant 0$. □

[74] Equivalently, at this step one could introduce a natural augmentation on the complex $C_*(\Delta^1) \otimes C_*(\Delta^1)$ such that the resulting reduced homology vanishes and $\theta\partial(d_1)$ is in its kernel.

Exercise 45.7. The goal of this exercise is to prove the associativity of the cross product:

$$(A \times B) \times C = A \times (B \times C) \in H_*(X \times Y \times Z; R)$$

for all $A \in H_*(X; R)$, $B \in H_*(Y; R)$ and $C \in H_*(Z; R)$. Here R may be any commutative ring with unit.

(a) Use acyclic models to prove that for triples of spaces X, Y, Z, all natural chain maps

$$\Psi : C_*(X; \mathbb{Z}) \otimes C_*(Y; \mathbb{Z}) \otimes C_*(Z; \mathbb{Z}) \to C_*(X \times Y \times Z; \mathbb{Z})$$

that act on 0-chains by $\Psi(x \otimes y \otimes z) = (x, y, z)$ are chain homotopic.
Remark: The statement implicitly assumes that there is a well-defined notion of the tensor product of three chain complexes, which of course is true since there is a canonical chain isomorphism between $(C_*(X) \otimes C_*(Y)) \otimes C_*(Z)$ and $C_*(X) \otimes (C_*(Y) \otimes C_*(Z))$. Right?

(b) Given $A \in H_*(X; R)$, $B \in H_*(Y; R)$ and $C \in H_*(Z; R)$, show that the products $(A \times B) \times C$ and $A \times (B \times C) \in H_*(X \times Y \times Z; R)$ can each be expressed via natural chain maps as in part (a), and conclude that they are identical.

Exercise 45.8. Fix a commutative ring R with unit $1 \in R$, and recall that for any path-connected space Y, there is a canonical isomorphism $H_0(Y; R) = R$. Let $[\mathrm{pt}] \in H_0(Y; R)$ denote the homology class corresponding to the unit $1 \in R$ under this identification: more concretely, $[\mathrm{pt}]$ can be written as $[\sigma \otimes 1] \in H_0(C_*(Y; \mathbb{Z}) \otimes R) = H_0(C_*(Y; R))$ for any singular 0-simplex $\sigma : \Delta^0 \to Y$, and it serves as a canonical generator of the R-module $H_0(Y; R)$. Show that for any space X, the cross product of any $A \in H_n(X; R)$ with $[\mathrm{pt}] \in H_0(Y; R)$ is given by

$$A \times [\mathrm{pt}] = i_* A \in H_n(X \times Y; R)$$

for any inclusion map of the form $i : X \hookrightarrow X \times Y : x \mapsto (x, \mathrm{const})$. A similar formula holds for cross products with $[\mathrm{pt}] \in H_0(X; R)$ if X is path-connected. In particular, this means that the unit in $H_0(\{\mathrm{pt}\}; R) = R$ acts as a multiplicative identity element with respect to the cross product, under the obvious identifications $X \times \{\mathrm{pt}\} = X = \{\mathrm{pt}\} \times X$.
Hint: Remember that $\times$ is determined by a non-unique but natural chain map $\Phi : C_*(X; \mathbb{Z}) \otimes C_*(Y; \mathbb{Z}) \to C_*(X \times Y; \mathbb{Z})$. There are various choices one could make in defining $\Phi : C_n(X; \mathbb{Z}) \otimes C_0(Y; \mathbb{Z}) \to C_n(X \times Y; \mathbb{Z})$, but you may notice that if you make the "right" choice, the desired relation becomes obvious. Review the construction of Φ via acyclic models to show that this choice is always possible.

The alert reader may notice that there is at least one important question we have not addressed yet: if X and Y are CW-complexes, are the singular and cellular cross products the same? The answer is of course yes, but we will not discuss it at length, since we don't plan to carry out any serious applications of the cellular cross product—it is useful to have in mind for intuition and motivation, but the product on singular homology will play a much more important role in further developments. One other (and closely related) question we have not addressed is how to define the cross product on *relative* singular homology. We will come back to this when we introduce the cohomology cup product.

Exercise 45.9. Another nice application of acyclic models is the proof of Theorem 30.14, which states that for any simplicial pair (K, K') and any choice of coefficient group G, the canonical chain map

$$C^o_*(K, K'; G) \xrightarrow{\Phi} C^\Delta_*(K, K'; G)$$

determined by the formula $(v_0, \ldots, v_n) \mapsto [v_0, \ldots, v_n]$ is a chain homotopy equivalence. This is the reason why the ordered and oriented simplicial homology groups are naturally isomorphic. In practice, it will suffice to consider a single simplicial complex $K = (V, S)$, use coefficients $G := \mathbb{Z}$ and prove that $\Phi : C_*^o(K;\mathbb{Z}) \to C_*^\Delta(K;\mathbb{Z})$ admits a chain homotopy inverse; once this is done, the generalization to simplicial pairs and arbitrary coefficient groups follows almost immediately.

As a preliminary step, we need to introduce a *reduced* version of simplicial homology. The definition should seem familiar: assume P is a simplicial complex with only one vertex, let $\epsilon : K \to P$ denote the unique simplicial map, and define

$$\widetilde{H}_*^o(K) := \ker\left(H_*^o(K) \xrightarrow{\epsilon_*} H_*^o(P)\right), \qquad \widetilde{H}_*^\Delta(K) := \ker\left(H_*^\Delta(K) \xrightarrow{\epsilon_*} H_*^\Delta(P)\right).$$

(a) Prove that for any coefficient group G,
$$H_n^o(K;G) \cong \begin{cases} \widetilde{H}_n^o(K;G) \oplus G & \text{if } n = 0, \\ \widetilde{H}_n^o(K;G) & \text{if } n \neq 0, \end{cases}$$
and that the analogous relation between $\widetilde{H}_*^\Delta(K;G)$ and $H_*^\Delta(K;G)$ also holds.

(b) Show that $\widetilde{H}_*^o(K;G)$ is also the homology of an *augmented* chain complex $\widetilde{C}_*^o(K;G)$ of the form
$$\ldots \longrightarrow C_2^o(K;G) \xrightarrow{\partial} C_1^o(K;G) \xrightarrow{\partial} C_0^o(K;G) \xrightarrow{\epsilon} G \longrightarrow 0 \longrightarrow 0 \longrightarrow \ldots,$$
i.e. $\widetilde{C}_n^o(K;G) = C_n^o(K;G)$ for all $n \neq -1$ but $\widetilde{C}_{-1}^o(K;G) = G$. Describe the *augmentation* map $\epsilon : C_0^o(K;G) \to G$ in this complex explicitly, and show that the analogous statement also holds for $\widetilde{H}_*^\Delta(K;G)$.

We next define a simplicial analogue of the *cone* of a topological space. Let $CK = (CV, CS)$ denote the simplicial complex whose vertex set CV is the union of V with one extra element $v_\infty \notin V$, and whose simplices consist of all sets of the form $\sigma \cup \{v_\infty\}$ for $\sigma \in S$, plus all their subsets. It is not hard to show that the polyhedron $|CK|$ is then homeomorphic to the cone of $|K|$, thus it is contractible, and the isomorphism $H_*^\Delta(CK) \cong H_*(|CK|)$ implies $\widetilde{H}_*^\Delta(CK) = 0$. But this does not immediately imply $\widetilde{H}_*^o(CK) = 0$, since we haven't yet proved H_*^o and H_*^Δ are isomorphic.

(c) For integers $n \geqslant 0$, consider the homomorphism $h : C_n^o(CK) \to C_{n+1}^o(CK)$ determined by the formula $h(v_0, \ldots, v_n) := (v_\infty, v_0, \ldots, v_n)$. Find a definition of $h : \mathbb{Z} = \widetilde{C}_{-1}^o(CK;\mathbb{Z}) \to C_0^o(CK;\mathbb{Z})$ that makes $\widetilde{C}_*^o(CK;\mathbb{Z}) \xrightarrow{h} \widetilde{C}_{*+1}^o(CK;\mathbb{Z})$ into a chain homotopy between the chain maps $\mathbb{1} : \widetilde{C}_*^o(CK;\mathbb{Z}) \to \widetilde{C}_*^o(CK;\mathbb{Z})$ and $0 : \widetilde{C}_*^o(CK;\mathbb{Z}) \to \widetilde{C}_*^o(CK;\mathbb{Z})$, and deduce that $\widetilde{H}_*^o(CK) = 0$ for all choices of coefficient group.

Returning to the chain map $\Phi : C_*^o(K) \to C_*^\Delta(K)$, we notice that for any subcomplex $L \subset K$, $\Phi(C_*^o(L)) \subset C_*^\Delta(L)$. This can be interpreted as a form of *naturality* if we view C_*^o and C_*^Δ as functors on the category of subcomplexes of K, with morphisms $L \to L'$ defined by inclusion: indeed, any nested pair of subcomplexes $L \subset L' \subset K$ gives rise to a commutative diagram

$$\begin{tikzcd} C_*^o(L) \arrow[r, "\Phi"] \arrow[d] & C_*^\Delta(L) \arrow[d] \\ C_*^o(L') \arrow[r, "\Phi"] & C_*^\Delta(L'), \end{tikzcd}$$

where the two vertical maps are the chain maps induced by the inclusion $L \hookrightarrow L'$. Let us similarly say that a chain map $\Psi : C_*^\Delta(K) \to C_*^o(K)$ is *natural* if for every subcomplex $L \subset K$, Ψ sends $C_*^\Delta(L)$ into $C_*^o(L)$.

(d) Use the method of acyclic models to prove that there exists a natural chain map $\Psi : C_*^\Delta(K;\mathbb{Z}) \to C_*^o(K;\mathbb{Z})$ whose definition in degree 0 is determined by the formula $\Psi[v] := (v)$, and that natural chain maps with this property are unique up to chain homotopy.

> *Hint: Proceed by induction on the degree n, and construct Ψ first on "model" subcomplexes $L \subset K$ that consist of a single n-simplex and all its faces. Notice that subcomplexes of this type are cones.*

If you've gotten this far, then you can probably guess how the rest of the proof that $H_*^\Delta(K, K'; G) \cong H_*^o(K, K'; G)$ goes: one must similarly show the uniqueness up to chain homotopy of certain natural chain maps $C_*^o(K; \mathbb{Z}) \to C_*^\Delta(K; \mathbb{Z})$, $C_*^o(K; \mathbb{Z}) \to C_*^o(K; \mathbb{Z})$ and $C_*^\Delta(K; \mathbb{Z}) \to C_*^\Delta(K; \mathbb{Z})$, which would imply that $\Phi \circ \Psi$ and $\Psi \circ \Phi$ are each chain homotopic to the identity. I suggest you work out the remaining details the next time you get bored on a long train ride.

46. Čech homology and inverse limits (January 9, 2024)

By now I have said almost everything I can reasonably say about singular homology without bringing cohomology into the picture, so that will be the next major topic. But before that, I'd like to address a question that may have been nagging at you since we introduced the Eilenberg-Steenrod axioms many weeks ago: what other axiomatic homology theories are there?

Let me name a few theories that are *not* examples: first, cellular and simplicial homology are not axiomatic homology theories in the sense of Eilenberg-Steenrod [**ES52**], as they are not functors defined on $\mathsf{Top}_{\mathrm{rel}}$ or any subcategory of $\mathsf{Top}_{\mathrm{rel}}$. Both require auxiliary choices beyond a pair of spaces (X, A) for their definitions, e.g. one cannot define the group $H_*^{\mathrm{CW}}(X, A)$ without having chosen a cell decomposition for (X, A), thus H_*^{CW} is a functor $\mathsf{CW}_{\mathrm{rel}} \to \mathsf{Ab}_{\mathbb{Z}}$, and the fact that $H_*^{\mathrm{CW}}(X, A)$ up to isomorphism depends only on the topology of (X, A) is a deep theorem, but not an intrinsic feature of the definition. You may have heard of other theories that produce topological invariants in spite of auxiliary choices: for example, *Morse homology* (see [**AD14**] or [**Sch93**]) associates to any triple (X, f, g) consisting of a closed oriented smooth manifold X with a generic function $f : X \to \mathbb{R}$ and Riemannian metric g a graded abelian group $H_*^{\mathrm{Morse}}(X, f, g)$ that turns out (by a different deep theorem that has nothing to do with the axioms) to be isomorphic to $H_*(X)$. There is also a smooth variant of singular homology defined when X is a smooth manifold, where all singular simplices are required to be smooth maps. This is often used for proving de Rham's theorem, which relates $H_*(X; \mathbb{R})$ to the *de Rham cohomology* $H_{\mathrm{dR}}^*(X)$ of X, defined in terms of the exterior derivative on smooth differential forms. However, one can use smooth approximation results to find a canonical isomorphism from smooth singular homology to $H_*(X)$ whenever the former is defined, so it really is just another formulation of the same theory, restricted to a smaller category. Yet another variant is the *cubic* version of singular homology, which is used in place of $H_*(X)$ in a few books such as [**Mas91**]: here the point is to replace the standard n-simplex Δ^n with the standard n-cube I^n, which makes defining product structures easier, but various other things harder. In any case, cubic singular homology is indeed an axiomatic homology theory, but actually it is always isomorphic to $H_*(X)$, so it is not actually a different theory—which is why some authors feel free to use it as a substitute.

I want to describe a theory that satisfies the Eilenberg-Steenrod axioms and thus captures the same topological information as singular homology on nice spaces, but does not match it on all spaces and is based on a totally different idea. To explain the definition, we will also need to introduce inverse limits, the contravariant version of direct limits. For reasons that we will see, the *Čech homology* theory that I'm going to describe is not especially popular, partly because it requires some extra conditions in order to make it satisfy all the axioms. On the other hand, the closely related theory of *Čech cohomology* satisfies the analogous set of axioms without restriction, and is widely used in several branches of mathematics, especially in algebraic geometry. We will come back to that briefly after discussing the axioms for cohomology theories in a few lectures, but the present lecture is intended as a sketch of Čech *homology* in particular, and since it is only a sketch, we will leave several details unproved, but give suitable references wherever possible. The main source for most of this material is [ES52].

The idea behind Čech homology is to measure the topology of a space X in terms of the combinatorial data formed by the overlaps of open sets in an arbitrarily fine open covering of X. The starting point is the observation that for any given open covering, the overlaps can be encoded in the form of an abstract simplicial complex.

For a space X, let $\mathcal{O}(X)$ denote the set of open coverings of X, so each element $\mathfrak{U} \in \mathcal{O}(X)$ is a set whose elements are open subsets of X with the property that

$$\bigcup_{\mathcal{U} \in \mathfrak{U}} \mathcal{U} = X.$$

Similarly, for any pair of spaces (X, A), we define $\mathcal{O}(X, A)$ to be the set of all pairs $(\mathfrak{U}, \mathfrak{U}_A)$ such that

$$\mathfrak{U} \in \mathcal{O}(X), \quad \mathfrak{U}_A \subset \mathfrak{U} \quad \text{and} \quad A \subset \bigcup_{\mathcal{U} \in \mathfrak{U}_A} \mathcal{U}.$$

Definition 46.1. For each open covering $\mathfrak{U} \in \mathcal{O}(X)$ of a space X, the **nerve** of $\mathfrak{U}$ is the simplicial complex $\mathcal{N}(\mathfrak{U})$ whose set of vertices is $\mathfrak{U}$, and whose simplices are the finite subsets $\sigma \subset \mathfrak{U}$ such that

$$\bigcap_{\mathcal{U} \in \sigma} \mathcal{U} \neq \varnothing.$$

More generally, for each pair of spaces (X, A) and $(\mathfrak{U}, \mathfrak{U}_A) \in \mathcal{O}(X, A)$, the **nerve** of $(\mathfrak{U}, \mathfrak{U}_A)$ is the simplicial pair

$$\mathcal{N}(\mathfrak{U}, \mathfrak{U}_A) := (\mathcal{N}(\mathfrak{U}), \mathcal{N}(\mathfrak{U}_A)),$$

where $\mathcal{N}(\mathfrak{U}_A) \subset \mathcal{N}(\mathfrak{U})$ denotes the subcomplex whose set of vertices is $\mathfrak{U}_A$, and whose simplices are the finite subsets $\sigma \subset \mathfrak{U}_A$ such that

$$A \cap \bigcap_{\mathcal{U} \in \sigma} \mathcal{U} \neq \varnothing.$$

You should take a moment to contemplate why $\mathcal{N}(\mathfrak{U})$ satisfies all the conditions of a simplicial complex, with $\mathcal{N}(\mathfrak{U}_A)$ as a subcomplex; in particular, every subset of a simplex is also a simplex since the condition $\bigcap_{\mathcal{U} \in \sigma} \mathcal{U} \neq \varnothing$ clearly remains true after deleting some sets from the collection in σ. Notice also that $\mathcal{N}(\mathfrak{U}_A)$ is the nerve of the open covering of A formed by the sets $\{\mathcal{U} \cap A\}_{\mathcal{U} \in \mathfrak{U}_A}$. The Čech homology theory will be defined in terms of the (ordered) *simplicial* homology[75] of the nerves $\mathcal{N}(\mathfrak{U}, \mathfrak{U}_A)$ of open coverings of (X, A), denoted by

$$H_*^o(\mathcal{N}(\mathfrak{U}, \mathfrak{U}_A)) = H_*^o(\mathcal{N}(\mathfrak{U}), \mathcal{N}(\mathfrak{U}_A)) = H_*^o(\mathcal{N}(\mathfrak{U}), \mathcal{N}(\mathfrak{U}_A); G).$$

Here we are fixing an arbitrary choice of coefficient group G for simplicial homology, which will then also serve as the coefficient group of Čech homology. As usual, G will be omitted from the notation whenever nothing important depends on this choice.

Figure 28 shows some examples of open covers $\mathfrak{U}$ of S^1 and the polyhedra $|\mathcal{N}(\mathfrak{U})|$ that arise from their nerves. We see that in one case, $|\mathcal{N}(\mathfrak{U})|$ is homeomorphic to S^1; this is not a coincidence, and we'll come back to it shortly. In general, however, $|\mathcal{N}(\mathfrak{U})|$ need not be homeomorphic, nor even homotopy equivalent, to the space that is being covered, and in fact, the nerve of an open cover of X can easily be an infinite-dimensional simplicial complex, even when X is something as tame as a compact polyhedron. Thus we clearly cannot hope in general to use the nerve of a single covering

[75] We are using the *ordered* rather than *oriented* version of simplicial homology because, in the present context, the ordered variant is more convenient for bookkeeping purposes. Recall from Exercise 45.9 that they are in any case naturally isomorphic.

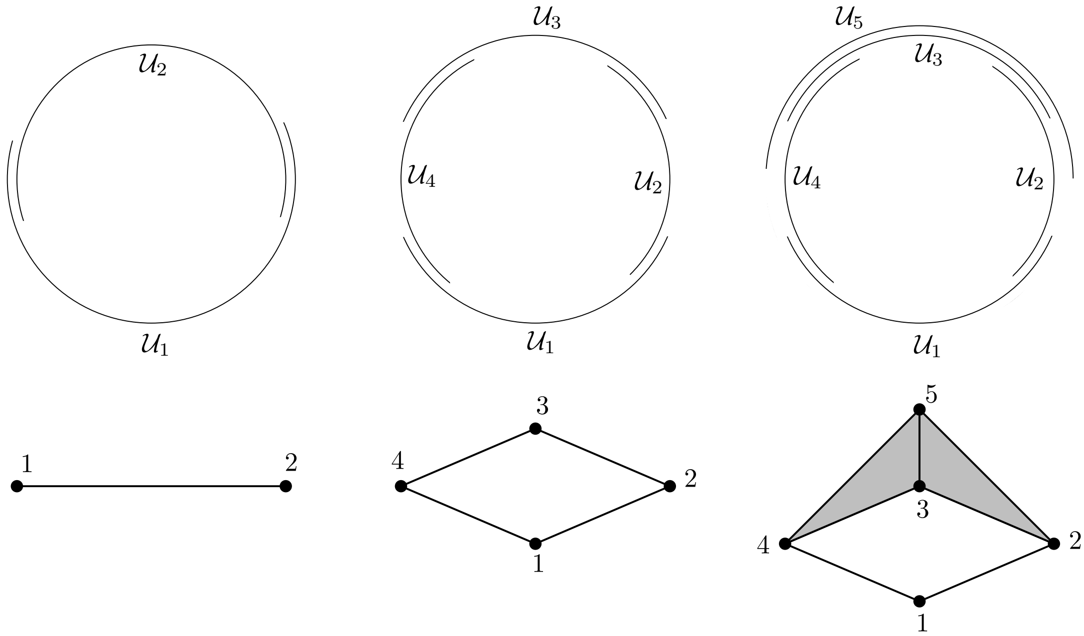

FIGURE 28. Three examples of open coverings of S^1 and their nerves, with vertices labeled $k \in \{1, 2, 3, 4, 5\}$ in correspondence with the open sets $\mathcal{U}_k \subset S^1$. The rightmost example includes two 2-simplices in addition to vertices and 1-simplices.

of X in order to define a topological invariant of X. What seems more promising, however, is to consider an open covering together with all of its possible refinements.

A **refinement** of an open cover $\mathfrak{U} \in \mathcal{O}(X)$ is another open covering $\mathfrak{U}' \in \mathcal{O}(X)$ such that every $\mathcal{U}' \in \mathfrak{U}'$ is a subset of some $\mathcal{U} \in \mathfrak{U}$. For pairs (X, A), we say similarly that a refinement of $(\mathfrak{U}, \mathfrak{U}_A) \in \mathcal{O}(X, A)$ is an element $(\mathfrak{U}', \mathfrak{U}'_A) \in \mathcal{O}(X, A)$ such that $\mathfrak{U}'$ is a refinement of $\mathfrak{U}$ and $\mathfrak{U}'_A$ is a refinement of $\mathfrak{U}_A$. The definition means that there exists a function

$$F : \mathfrak{U}' \to \mathfrak{U}, \qquad F(\mathfrak{U}'_A) \subset \mathfrak{U}_A$$

such that for every $\mathcal{U} \in \mathfrak{U}'$, $\mathcal{U} \subset F(\mathcal{U})$. It follows that if $\sigma \subset \mathfrak{U}'$ is a simplex of $\mathcal{N}(\mathfrak{U}')$, then

$$\bigcap_{\mathcal{U} \in \sigma} F(\mathcal{U}) \supset \bigcap_{\mathcal{U} \in \sigma} \mathcal{U} \neq \emptyset,$$

hence $F(\sigma) \subset \mathfrak{U}$ is a simplex of $\mathcal{N}(\mathfrak{U})$, and similarly, F maps simplices of $\mathcal{N}(\mathfrak{U}'_A)$ to simplices of $\mathcal{N}(\mathfrak{U}_A)$. In other words, F is a simplicial map from $\mathcal{N}(\mathfrak{U}')$ to $\mathcal{N}(\mathfrak{U})$, and in the relative case, a map of simplicial pairs:

$$F : \mathcal{N}(\mathfrak{U}', \mathfrak{U}'_A) \to \mathcal{N}(\mathfrak{U}, \mathfrak{U}_A).$$

It therefore induces a chain map between the corresponding ordered simplicial complexes

$$F_* : C^o_*(\mathcal{N}(\mathfrak{U}', \mathfrak{U}'_A)) \to C^o_*(\mathcal{N}(\mathfrak{U}, \mathfrak{U}_A)). \tag{46.1}$$

One obvious concern in this discussion is that F is not uniquely determined by the refinement, i.e. for each $\mathcal{U}' \in \mathfrak{U}'$, there may be more than one $\mathcal{U} \in \mathfrak{U}$ containing $\mathcal{U}'$. But the following result gives an enormous hint as to what we should do next:

PROPOSITION 46.2 ([ES52, Corollary IX.2.14]). *Given an open covering* $(\mathfrak{U}, \mathfrak{U}_A) \in \mathcal{O}(X, A)$ *and a refinement* $(\mathfrak{U}', \mathfrak{U}'_A)$ *of* $(\mathfrak{U}, \mathfrak{U}_A)$, *the chain homotopy class of the induced chain map* (46.1) *on ordered simplicial chain complexes is independent of choices.* □

It follows that we can associate to any refinement $\beta := (\mathfrak{U}', \mathfrak{U}'_A) \in \mathcal{O}(X, A)$ of an open covering $\alpha := (\mathfrak{U}, \mathfrak{U}_A) \in \mathcal{O}(X, A)$ a natural homomorphism of simplicial homology groups

$$\varphi_{\alpha\beta} : H_*^o(\mathcal{N}(\mathfrak{U}', \mathfrak{U}'_A)) \to H_*^o(\mathcal{N}(\mathfrak{U}, \mathfrak{U}_A))).$$

One can view this as defining something very similar to a direct system: indeed, let us define a pre-order on $\mathcal{O}(X, A)$ by writing

$$(\mathfrak{U}', \mathfrak{U}'_A) \succ (\mathfrak{U}, \mathfrak{U}_A)$$

whenever $(\mathfrak{U}', \mathfrak{U}'_A)$ is a refinement of $(\mathfrak{U}, \mathfrak{U}_A)$. (Note that it is not a partial order, as two open coverings can easily be refinements of each other without being identical.) Since any two open coverings have a common refinement, this makes $(\mathcal{O}(X, A), \prec)$ a directed set, and the result above associates a morphism of $\mathbb{Z}$-graded abelian groups to any pair $\alpha, \beta \in \mathcal{O}(X, A)$ with $\beta \succ \alpha$. The only trouble is that this morphism goes the wrong way: if the collection of graded abelian groups

$$\{H_*^o(\mathcal{N}(\alpha))\}_{\alpha \in \mathcal{O}(X,A)}$$

were to be viewed as a direct system, then we would have to have an associated morphism $\varphi_{\beta\alpha} : H_*^o(\mathcal{N}(\alpha)) \to H_*^o(\mathcal{N}(\beta))$ for every $\beta \succ \alpha$, but instead we have $\varphi_{\alpha\beta} : H_*^o(\mathcal{N}(\beta)) \to H_*^o(\mathcal{N}(\alpha))$. There is also a name for this.

DEFINITION 46.3. Given a category $\mathscr{C}$ and a directed set $(I, \prec)$, an **inverse system** *(projektives System)* $\{X_\alpha, \varphi_{\alpha\beta}\}$ in $\mathscr{C}$ over $(I, \prec)$ associates to each $\alpha \in I$ an object X_α of $\mathscr{C}$, along with morphisms

$$\varphi_{\alpha\beta} \in \operatorname{Mor}(X_\beta, X_\alpha) \quad \text{for each} \quad \alpha \prec \beta$$

such that

$$\varphi_{\alpha\alpha} = \operatorname{Id}_{X_\alpha}$$

and the diagram

$$\begin{tikzcd} X_\alpha & X_\beta \arrow[l, "\varphi_{\alpha\beta}"'] & X_\gamma \arrow[l, "\varphi_{\beta\gamma}"'] \arrow[ll, bend left, "\varphi_{\alpha\gamma}"] \end{tikzcd}$$

commutes for every triple $\alpha, \beta, \gamma \in I$ with $\alpha \prec \beta \prec \gamma$.

REMARK 46.4. In terms of the category $\mathscr{I}$ corresponding to the directed set $(I, \prec)$ as in Remark 39.2, an inverse system in $\mathscr{C}$ over $(I, \prec)$ is the same thing as a contravariant functor $\mathscr{I} \to \mathscr{C}$ (recall Definition 27.13).

Convergence of inverse systems is defined analogously to direct systems, the main difference being that most arrows go the other way.

DEFINITION 46.5. For an inverse system $\{X_\alpha, \varphi_{\alpha\beta}\}$ in $\mathscr{C}$ over $(I, \prec)$, a **target** $\{Y, f_\alpha\}$ of the system consists of an object Y of $\mathscr{C}$ together with associated morphisms $f_\alpha \in \operatorname{Mor}(Y, X_\alpha)$ for each $\alpha \in I$ such that the diagram

$$\begin{tikzcd} X_\alpha & & X_\beta \arrow[ll, "\varphi_{\alpha\beta}"'] \\ & Y \arrow[ul, "f_\alpha"] \arrow[ur, "f_\beta"'] & \end{tikzcd}$$

commutes for every pair $\alpha, \beta \in I$ with $\alpha \prec \beta$.

Definition 46.6. A target $\{X_\infty, \varphi_\alpha\}$ of the inverse system $\{X_\alpha, \varphi_{\alpha\beta}\}$ is called an **inverse limit**[76]*(projektiver Limes)* of the system and written as

$$X_\infty = \varprojlim\{X_\alpha\}$$

if it satisfies the following "universal" property: for all targets $\{Y, f_\alpha\}$ of $\{X_\alpha, \varphi_{\alpha\beta}\}$, there exists a unique morphism $f_\infty \in \mathrm{Mor}(Y, X_\infty)$ such that the diagram

$$\begin{tikzcd}
X_\alpha & X_\infty \arrow[l, "\varphi_\alpha"'] \\
& Y \arrow[ul, "f_\alpha"] \arrow[u, "f_\infty"']
\end{tikzcd}$$

commutes for every $\alpha \in I$.

The meaning of an inverse limit can be encoded in the diagram

$$\begin{tikzcd}
X_\alpha & X_\beta \arrow[l, "\varphi_{\alpha\beta}"'] & X_\gamma \arrow[l, "\varphi_{\beta\gamma}"'] & \ldots \arrow[l] & \varprojlim\{X_\alpha\} \arrow[l] \\
& & & & Y \arrow[ullll] \arrow[ulll] \arrow[ull] \arrow[u, dashed]
\end{tikzcd}$$

where we assume $\alpha \prec \beta \prec \gamma \prec \ldots \in I$, and the defining feature of $\varprojlim\{X_\alpha\}$ is that the morphism indicated by the dashed arrow must exist and be unique whenever all the other morphisms in the diagram are given.

As with direct limits, there is no guarantee from these definitions that an inverse limit must exist, but for the categories we are most interested in, its existence can be established by describing it more concretely. One should not confuse the statement that an inverse limit exists with any claim that it is nonempty—the empty set is also a topological space and can appear as the limit of an inverse system in Top (see Example 46.11 below).

Exercise 46.7. If $\{X_\alpha, \varphi_{\alpha\beta}\}$ is an inverse system in Top over $(I, \prec)$, show that its inverse limit is the space

$$\varprojlim\{X_\alpha\} = \left\{ \{x_\alpha\} \in \prod_{\alpha \in I} X_\alpha \;\middle|\; x_\alpha = \varphi_{\alpha\beta}(x_\beta) \text{ for all } \alpha, \beta \in I \text{ with } \alpha \prec \beta \right\},$$

with the associated morphisms $\varphi_\alpha : \varprojlim\{X_\beta\} \to X_\alpha$ defined via the obvious projections $\prod_{\beta \in I} X_\beta \to X_\alpha$ for each $\alpha \in I$. Conclude from this that the topology on $\varprojlim\{X_\alpha\}$ is the weakest for which the maps $\varphi_\alpha : \varprojlim\{X_\beta\} \to X_\alpha$ are all continuous.

Remark 46.8. The exercise extends in an obvious way to describe inverse limits in the category Set of sets (with morphisms defined as arbitrary maps).

Exercise 46.9. Consider an inverse system $\{X_\alpha, \varphi_{\alpha\beta}\}$ in Top for which the spaces X_α are all subspaces of some fixed topological space X, $\beta > \alpha$ holds if and only if $X_\beta \subset X_\alpha$, and the maps $\varphi_{\alpha\beta} : X_\beta \to X_\alpha$ are all inclusions. Show that $\varprojlim\{X_\alpha\} = \bigcap_{\alpha \in I} X_\alpha$, with the associated morphisms $\varphi_\alpha : \varprojlim\{X_\beta\} \to X_\alpha$ given by the obvious inclusions.
Comment: The obvious analogue of this exercise involving direct limits and unions is only sometimes true, e.g. it works for viewing any CW-complex as the direct limit of its skeleta, but Exercise 39.23 shows an example in which the direct limit and the union are the same set with different topologies. In this sense, inverse systems in the category Top *are somewhat better behaved than direct systems.*

[76]Inverse limits are also sometimes called **projective limits**, and they constitute a special case of the general category-theoretical notion of **limits** (as opposed to *colimits*, cf. Exercise 39.24).

Exercise 46.10. Prove that for any inverse system $\{X_\alpha, \varphi_{\alpha\beta}\}$ of topological spaces such that every X_α is nonempty, compact and Hausdorff, $\varprojlim\{X_\alpha\} \neq \varnothing$.
Hint: By Tychonoff's theorem, $\prod_\alpha X_\alpha$ is compact, which means that every net in $\prod_\alpha X_\alpha$ has a cluster point (see Lecture 5 from last semester). For every index β, one can choose an element $x^\beta = \{x^\beta_\alpha\} \in \prod_\alpha X_\alpha$ whose coordinates satisfy $x^\beta_\alpha = \varphi_{\alpha\beta}(x^\beta_\beta)$ for every $\alpha \prec \beta$ and are arbitrary for all other α. The collection $\{x^\beta \in \prod_\alpha X_\alpha\}_{\beta \in I}$ then defines a net in $\prod_\alpha X_\alpha$, which therefore has a cluster point. Prove that the cluster point belongs to $\varprojlim\{X_\alpha\}$. (For a slightly different argument that does not use nets, see [ES52, Theorem VIII.3.6]; it does still require Tychonoff's theorem.)

Example 46.11. Combining the previous two exercises produces the well-known fact that in any Hausdorff space, the intersection of any collection of nonempty compact subsets that all have nonempty pairwise intersections is nonempty. It is easy to see that the compactness condition cannot be dropped from this statement: for instance, taking the collection of intervals $\{(0, 1/n]\}_{n\in\mathbb{N}}$ as an inverse system in the sense of Exercise 46.9, the inverse limit is

$$\varprojlim\{(0, 1/n]\}_{n\in\mathbb{N}} = \bigcap_{n\in\mathbb{N}} (0, 1/n] = \varnothing.$$

Exercise 46.12. If $\{G_\alpha, \varphi_{\alpha\beta}\}$ is an inverse system in Ab over $(I, \prec)$, show that its inverse limit is a group of the form

$$\varprojlim\{G_\alpha\} = \left\{ \{g_\alpha\} \in \prod_{\alpha\in I} G_\alpha \;\middle|\; g_\alpha = \varphi_{\alpha\beta}(g_\beta) \text{ for all } \alpha, \beta \in I \text{ with } \alpha \prec \beta \right\},$$

with the associated homomorphisms $\varphi_\alpha : \varprojlim\{G_\beta\} \to G_\alpha$ defined via the projections $\prod_{\beta\in I} G_\beta \to G_\alpha$ all $\alpha \in I$.

Exercise 46.13. Prove the obvious analogues of the result in Exercise 46.12 for inverse systems in the categories $\mathsf{Ab}_\mathbb{Z}$ of $\mathbb{Z}$-graded abelian groups and Chain of chain complexes.

Exercise 46.14. Assume $\{X_\alpha, \varphi_{\alpha\beta}\}$ is an inverse system over $(I, \prec)$ in any category. A subset $I_0 \subset I$ is called a **cofinal set** if for every $\alpha \in I$ there exists some $\beta \in I_0$ such that $\beta \succ \alpha$. Suppose $I_0 \subset I$ is a cofinal set with the property that for every $\alpha, \beta \in I_0$ with $\alpha \prec \beta$, $\varphi_{\alpha\beta} \in \mathrm{Mor}(X_\beta, X_\alpha)$ is an isomorphism. Prove that $\varprojlim\{X_\alpha\}$ is then isomorphic to X_γ for any $\gamma \in I_0$, and describe the associated morphisms $\varprojlim\{X_\beta\} \xrightarrow{\varphi_\alpha} X_\alpha$ for every $\alpha \in I$.
Advice: This problem becomes a bit easier if you work in any of the categories Top, Ab *or* Chain *so that you can use the results of Exercises 46.7, 46.12 or 46.13 respectively. But it can also be done without that assumption, just by using the universal property and playing with commutative diagrams.*

We now have enough concepts in place to define the Čech homology groups.

Definition 46.15. The Čech homology of a pair of spaces (X, A) with coefficients in an abelian group G is defined as the $\mathbb{Z}$-graded abelian group

$$\check{H}_*(X, A) = \check{H}_*(X, A; G) := \varprojlim \left\{ H^o_*(\mathcal{N}(\mathfrak{U}, \mathfrak{U}_A); G) \right\}_{(\mathfrak{U}, \mathfrak{U}_A) \in \mathcal{O}(X,A)}.$$

It is slightly harder than for singular homology to see why this should define a functor $\mathsf{Top}_{\mathrm{rel}} \to \mathsf{Ab}_\mathbb{Z}$, but still not so hard. The main point is that whenever $f : (X, A) \to (Y, B)$ is a continuous map of pairs and $\alpha = (\mathfrak{U}, \mathfrak{U}_A) \in \mathcal{O}(Y, B)$ is an open covering of (Y, B), there is an induced open covering $f^*\alpha \in \mathcal{O}(X, A)$ of (X, A) consisting of the subsets $f^{-1}(\mathcal{U})$ for $\mathcal{U} \in \mathfrak{U}$, and whenever $\beta \in \mathcal{O}(Y, B)$ is a refinement of α, $f^*\beta \in \mathcal{O}(X, A)$ is clearly also a refinement of $f^*\alpha$. The obvious correspondence

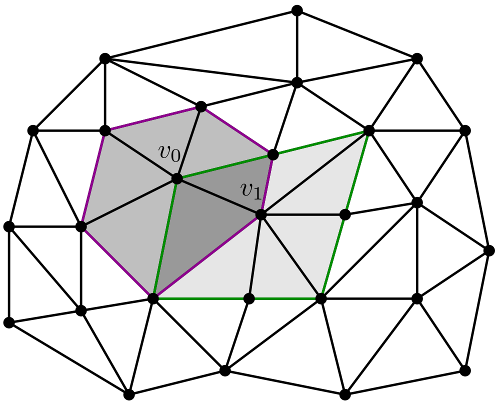

FIGURE 29. The open stars of two neighboring vertices v_0 and v_1 in a simplicial complex.

between the open sets in $f^*\alpha$ and those in α then defines a simplicial map $\mathcal{N}(f^*\alpha) \to \mathcal{N}(\alpha)$, giving a homomorphism

$$f_* : H_*^o\big(\mathcal{N}(f^*\alpha)\big) \to H_*^o\big(\mathcal{N}(\alpha)\big)$$

for every $\alpha \in \mathcal{O}(Y, B)$. Using the universal property of the inverse limit, one can derive from this a morphism

$$f_* : \check{H}_*(X, A) \to \check{H}_*(Y, B)$$

between the corresponding inverse limits, and prove that it satisfies the usual conditions for $\check{H}_*$ to be a functor. This implies in particular that homeomorphic pairs have the same Čech homology.

What is probably harder to see at this stage is why one should ever expect $\check{H}_*(X, A)$ to be the same as the singular homology $H_*(X, A)$. To this end, consider the case where X is the polyhedron of a finite simplicial complex $K = (V, S)$. We saw in Lecture 41 the notion of the *open star* of a vertex v in K, which defines an open set

$$\operatorname{st} v \subset X$$

containing all points that lie in simplices that have v as a vertex (see Figure 29). These sets define a distinguished open covering of X,

$$\mathfrak{U}_K := \{\operatorname{st} v \mid v \in V\},$$

and recall from Exercise 41.15 that for any finite collection of vertices $v_0, \ldots, v_n \in V$, we have

$$\bigcap_{k=0}^{n} \operatorname{st} v_k \neq \varnothing \quad \Leftrightarrow \quad \{v_0, \ldots, v_n\} \in S.$$

In other words, the nerve of $\mathfrak{U}_K$ is the complex K itself. Now if $\mathfrak{U} \in \mathcal{O}(X)$ is another open covering, since X is compact, we can always find a refinement of $\mathfrak{U}$ in the form $\mathfrak{U}_{K'}$ by applying barycentric subdivision to the simplices of K enough times, producing a new simplicial complex K' with more and smaller simplices but a homeomorphic polyhedron $|K'| = X$, and since barycentric subdivision induces chain homotopy equivalences, one can show that the induced map

$$H_*\big(\mathcal{N}(\mathfrak{U}_{K'})\big) \to H_*\big(\mathcal{N}(\mathfrak{U}_K)\big)$$

resulting from the fact that $\mathfrak{U}_{K'} > \mathfrak{U}_K$ is always an isomorphism. In other words, the open coverings that arise from successive barycentric subdivisions of K form a *cofinal set* in $\mathcal{O}(X)$ that satisfies the hypotheses of Exercise 46.14, and thus provides enough information to compute the inverse limit. The result is:

THEOREM 46.16. *For any compact polyhedron $X = |K|$ with underlying simplicial complex K, $\check{H}_*(X;G) \cong H_*^\Delta(K;G)$ for every coefficient group G.* ☐

It follows in particular that Čech homology is isomorphic to singular homology on compact polyhedra. Notice by the way that if we had not already constructed one example of an axiomatic homology theory, we could still use this argument to prove that simplicial homology is independent of the triangulation of a compact polyhedron—it follows now from the fact Čech homology is a topological invariant.

It is not *always* true however that $\check{H}_*(X) \cong H_*(X)$.

LEMMA 46.17. *If X is a connected space, then for every open cover $\mathfrak{U}$ of X, the nerve $\mathcal{N}(\mathfrak{U})$ is connected.*

PROOF. If $\mathcal{N}(\mathfrak{U})$ is not connected then it can be decomposed as a disjoint union of two nonempty subcomplexes $\mathcal{N}(\mathfrak{U}) \cong K_0 \amalg K_1$. Let $X_0 \subset X$ denote the union of all the sets $\mathcal{U} \in \mathfrak{U}$ that are vertices of K_0, and define $X_1 \subset X$ similarly via K_1. Then both are nonempty open sets, their union is X, and they are disjoint, since otherwise $\mathcal{N}(\mathfrak{U})$ would have to contain a 1-simplex with one vertex in K_0 and one in K_1. This proves that X is not connected. ☐

THEOREM 46.18. *For any connected space X and any coefficient group G, $\check{H}_0(X;G) \cong G$.*

PROOF. Lemma 46.17 implies that for every $\mathfrak{U} \in \mathcal{O}(X)$, $H_0^o(\mathcal{N}(\mathfrak{U});G) \cong H_0^\Delta(\mathcal{N}(\mathfrak{U});G) \cong G$. It is similarly easy to show that the canonical map $H_0^o(\mathcal{N}(\mathfrak{U}');G) \to H_0^o(\mathcal{N}(\mathfrak{U});G)$ for any refinement $\mathfrak{U}'$ of $\mathfrak{U}$ is an isomorphism, and that the inverse limit is therefore isomorphic to G. ☐

This result is different in general from singular homology in dimension 0, which splits over a direct sum of the *path-components* (not connected components) of each space. So, for instance, Figure 21 in Lecture 41 shows an example of compact space $X \subset \mathbb{R}^2$ with

$$H_0(X;\mathbb{Z}) \cong \mathbb{Z}^3 \quad \text{but} \quad \check{H}_0(X;\mathbb{Z}) \cong \mathbb{Z}.$$

Needless to say, that space is not a CW-complex, and one should expect better results in general for CW-complexes, as we saw with polyhedra in Theorem 46.16. At least $\check{H}_*$ and H_* will match on all CW-pairs if they have the same coefficient group and $\check{H}_*(\cdot;G)$ satisfies the Eilenberg-Steenrod axioms. So does it? The answer is a bit surprising.

THEOREM 46.19. *For every abelian group G, $\check{H}_*(\cdot;G) : \mathsf{Top}_{\mathrm{rel}} \to \mathsf{Ab}_{\mathbb{Z}}$ satisfies all of the Eilenberg-Steenrod axioms except for exactness, but it does not satisfy exactness in general.*

An actual counterexample to the exactness axiom is explained in [**ES52**, §X.4]. It would take at least a few lectures to either explain that counterexample or prove that the rest of the axioms are satisfied, so we'll mostly skip it since this lecture is meant to be only a brief digression away from the main topic of the course. But it's worth taking a closer look at how one would naturally try to prove the exactness axiom, and why it fails in general. It also succeeds in some cases, so the negative statement in Theorem 46.19 is not the end of the story.

The problem with exactness is traceable to a problem with the behavior of exact sequences under inverse limits. If (X,A) is a pair of spaces and $(\mathfrak{U},\mathfrak{U}_A) \in \mathcal{O}(X,A)$, then $\mathcal{N}(\mathfrak{U}_A) \subset \mathcal{N}(\mathfrak{U})$ is a subcomplex and the obvious short exact sequence of ordered simplicial chain complexes

$$0 \to C_*^o(\mathcal{N}(\mathfrak{U}_A)) \to C_*^o(\mathcal{N}(\mathfrak{U})) \to C_*^o\big(\mathcal{N}(\mathfrak{U},\mathfrak{U}_A)\big) = C_*^o(\mathcal{N}(\mathfrak{U}))/C_*^o(\mathcal{N}(\mathfrak{U}_A)) \to 0$$

gives rise to a long exact sequence of simplicial homology groups

$$\ldots \to H_n^o(\mathcal{N}(\mathfrak{U}_A)) \to H_n^o(\mathcal{N}(\mathfrak{U})) \to H_n^o\big(\mathcal{N}(\mathfrak{U},\mathfrak{U}_A)\big) \to H_{n-1}^o(\mathcal{N}(\mathfrak{U}_A)) \to \ldots \tag{46.2}$$

If $(\mathfrak{U}', \mathfrak{U}'_A) \in \mathcal{O}(X, A)$ is a refinement of $(\mathfrak{U}, \mathfrak{U}_A)$, it is not hard to show that the canonical maps in the inverse systems fit together with the long exact sequences for these two pairs into a commutative diagram

$$\begin{tikzcd}
\ldots \arrow[r] & H_n^o(\mathcal{N}(\mathfrak{U}_A)) \arrow[r] & H_n^o(\mathcal{N}(\mathfrak{U})) \arrow[r] & H_n^o\big(\mathcal{N}(\mathfrak{U}, \mathfrak{U}_A)\big) \arrow[r] & H_{n-1}^o(\mathcal{N}(\mathfrak{U}_A)) \arrow[r] & \ldots \\
\ldots \arrow[r] & H_n^o(\mathcal{N}(\mathfrak{U}'_A)) \arrow[r] \arrow[u] & H_n^o(\mathcal{N}(\mathfrak{U}')) \arrow[r] \arrow[u] & H_n^o\big(\mathcal{N}(\mathfrak{U}', \mathfrak{U}'_A)\big) \arrow[r] \arrow[u] & H_{n-1}^o(\mathcal{N}(\mathfrak{U}'_A)) \arrow[r] \arrow[u] & \ldots
\end{tikzcd}$$

An exact sequence is the same thing as a chain complex with trivial homology groups, so we can therefore regard this collection of exact sequences as an inverse system over $(\mathcal{O}(X, A), \prec)$ in the category Chain of chain complexes. By Exercise 46.13, this system will have an inverse limit, which will be a chain complex

$$\ldots \to \check{H}_n(A) \to \check{H}_n(X) \to \check{H}_n(X, A) \to \check{H}_{n-1}(A) \to \ldots \tag{46.3}$$

But it is not obvious whether this sequence of maps is exact. If this had been a direct limit, we could now appeal to Proposition 39.18, which would produce a natural isomorphism between the homology of the direct limit and the direct limit of the homology groups in the system; the latter are all zero since the sequences are all exact, so this would imply that the limit sequence is also exact. But we don't have an analogue of Proposition 39.18 for inverse limits. As a matter of fact, the result we would like to prove at this point is false:

Example 46.20. For every $n \in \mathbb{N}$, denote by $0 \to A_n \to B_n \to C_n \to 0$ the short exact sequence $0 \to \mathbb{Z} \xrightarrow{\cdot 2} \mathbb{Z} \xrightarrow{\text{pr}} \mathbb{Z}_2 \to 0$, and define homomorphisms $\varphi_{n-1,n}$ for each $n \geqslant 2$ by

$$A_n \xrightarrow{\cdot 3} A_{n-1}, \qquad B_n \xrightarrow{\cdot 3} B_{n-1}, \qquad C_n \xrightarrow{\mathbb{1}} C_{n-1}.$$

Then the resulting diagram

$$\begin{tikzcd}
\ldots \arrow[r] & 0 \arrow[r] & A_1 \arrow[r, "\cdot 2"] & B_1 \arrow[r, "\text{pr}"] & C_1 \arrow[r] & 0 \arrow[r] & \ldots \\
\ldots \arrow[r] & 0 \arrow[r] \arrow[u] & A_2 \arrow[r, "\cdot 2"] \arrow[u, "\cdot 3"] & B_2 \arrow[r, "\text{pr}"] \arrow[u, "\cdot 3"] & C_2 \arrow[r] \arrow[u, "\mathbb{1}"] & 0 \arrow[r] \arrow[u] & \ldots \\
\ldots \arrow[r] & 0 \arrow[r] \arrow[u] & A_3 \arrow[r, "\cdot 2"] \arrow[u, "\cdot 3"] & B_3 \arrow[r, "\text{pr}"] \arrow[u, "\cdot 3"] & C_3 \arrow[r] \arrow[u, "\mathbb{1}"] & 0 \arrow[r] \arrow[u] & \ldots \\
 & \vdots \arrow[u] & \vdots \arrow[u] & \vdots \arrow[u] & \vdots \arrow[u] & \vdots \arrow[u] &
\end{tikzcd}$$

commutes, meaning the $\varphi_{n-1,n}$ are all chain maps, and we can compose them to define further chain maps $\varphi_{m,n}$ for every $m < n$ and interpret this collection of data as an inverse system of chain complexes. By Exercises 46.12 and 46.13, the individual terms of the inverse limit complex are as follows: first,

$$\varprojlim\{A_n\} = \left\{ (a_1, a_2, a_3, \ldots) \in \prod_{n \in \mathbb{N}} A_n \;\middle|\; a_{n-1} = 3a_n \text{ for all } n \geqslant 2 \right\} = 0,$$

and $\varprojlim\{B_n\}$ similarly vanishes since no integer is divisible by arbitrarily large powers of 3. On the other hand,

$$\varprojlim\{C_n\} = \left\{ (c_1, c_2, c_3, \ldots) \in \prod_{n \in \mathbb{N}} C_n \;\middle|\; c_{n-1} = c_n \text{ for all } n \geqslant 2 \right\} \cong \mathbb{Z}_2,$$

so the full inverse limit chain complex is of the form

$$\ldots \longrightarrow 0 \longrightarrow 0 \longrightarrow 0 \longrightarrow \mathbb{Z}_2 \longrightarrow 0 \longrightarrow \ldots,$$

which is not an exact sequence.

REMARK 46.21. One can regard $\varprojlim$ as an additive functor $\mathsf{Inv}(\mathsf{Ab}) \to \mathsf{Ab}$, where $\mathsf{Inv}(\mathsf{Ab})$ is a category whose objects are inverse systems of abelian groups over a fixed directed set. In these terms, what Example 46.20 shows is that $\varprojlim$ is not an exact functor (cf. Definition 42.12). It is not hard to show however that it is a *left-exact* functor, and indeed, the failure of exactness after feeding the original short exact sequence in Example 46.20 into $\varprojlim$ occurs precisely at the third nontrivial term.

In general, suppose we are given an inverse system of chain complexes $\{C_*^\alpha, \varphi_{\alpha\beta}\}$ indexed by α in some directed set $(I, \prec)$, so we have a commuting diagram for every $\beta \succ \alpha$ in the form

$$\begin{tikzcd}
\ldots \arrow[r] & C_{n+1}^\alpha \arrow[r, "\partial^\alpha"] & C_n^\alpha \arrow[r, "\partial^\alpha"] & C_{n-1}^\alpha \arrow[r, "\partial^\alpha"] & \ldots \\
\ldots \arrow[r] & C_{n+1}^\beta \arrow[r, "\partial^\beta"] \arrow[u, "\varphi_{\alpha\beta}"] & C_n^\beta \arrow[r, "\partial^\beta"] \arrow[u, "\varphi_{\alpha\beta}"] & C_{n-1}^\beta \arrow[r, "\partial^\beta"] \arrow[u, "\varphi_{\alpha\beta}"] & \ldots
\end{tikzcd} \tag{46.4}$$

and assume moreover that the rows of these diagrams are always exact. The inverse limit is

$$C_*^\infty := \varprojlim\{C_*^\alpha\} = \left\{ \{x_\alpha\} \in \prod_{\alpha \in I} C_*^\alpha \;\middle|\; \varphi_{\alpha\beta}(x_\beta) = x_\alpha \text{ for all } \beta \succ \alpha \right\},$$

where the chain complex boundary map can be written as

$$\partial^\infty := \prod_{\alpha \in I} \partial^\alpha \Big|_{C_*^\infty} : C_*^\infty \to C_*^\infty,$$

the restriction to the subgroup $C_*^\infty \subset \prod_\alpha C_*^\alpha$ being well defined since $\varphi_{\alpha\beta}(\partial^\beta x_\beta) = \partial^\alpha \varphi_{\alpha\beta}(x_\beta) = \partial^\alpha x_\alpha$ for all $\beta \succ \alpha$ and $x_\beta \in C_*^\beta$. Given $x = \{x_\alpha\}_{\alpha \in I} \in C_n^\infty$ with $\partial^\infty \{x_\alpha\} = 0$, we have $\partial^\alpha x_\alpha = 0$ for all $\alpha \in I$ and thus $x_\alpha = \partial^\alpha y_\alpha$ for some $y_\alpha \in C_{n+1}^\alpha$. The trouble is that these elements y_α are not generally unique, and if they are chosen arbitrarily, then they need not satisfy

$$\varphi_{\alpha\beta}(y_\beta) = y_\alpha \quad \text{for all} \quad \beta \succ \alpha, \tag{46.5}$$

without which $\{y_\alpha\}_{\alpha \in I}$ will not be an element of C_{n+1}^∞.

To get a firmer handle on this problem, define for each $\alpha \in I$ the nonempty subset

$$K_\alpha := (\partial^\alpha)^{-1}(x_\alpha) \subset C_{n+1}^\alpha.$$

The chain map relation and the condition $\varphi_{\alpha\beta}(x_\beta) = x_\alpha$ then imply

$$\varphi_{\alpha\beta}(K_\beta) \subset K_\alpha \quad \text{for all } \beta \succ \alpha,$$

which makes the collection of sets $\{K_\alpha\}_{\alpha \in I}$ with maps $K_\beta \xrightarrow{\varphi_{\alpha\beta}} K_\alpha$ into an inverse system in Set over $(I, \prec)$. By Exercise 46.7 and Remark 46.8, $\varprojlim\{K_\alpha\}$ is then the set of all elements $\{y_\alpha\} \in \prod_{\alpha \in I} K_\alpha$ such that (46.5) is satisfied, in which case we then have $\{y_\alpha\} \in C_{n+1}^\infty$ with $\partial^\infty\{y_\alpha\} = \{x_\alpha\}$. The essential question thus boils down to this:

$$\text{Is } \varprojlim\{K_\alpha\} \text{ nonempty?}$$

Example 46.20 implies that the answer must sometimes be no, and indeed, we know from Example 46.11 that an inverse limit of nonempty sets or topological spaces can easily be the empty set.

To make progress, we need to add more assumptions. Suppose first of all that the individual groups C_n^α for each $n \in \mathbb{Z}$ and $\alpha \in I$ are finite. Then the sets K_α are also finite, and if we assign them the discrete topology, we can view them all as nonempty compact Haudroff spaces. In this case there is a positive result we can use: Exercise 46.10 implies that $\varprojlim\{K_\alpha\}$ will then always be nonempty, which fills the gap at the end of our proof that the limit sequence is exact!

I would like to point out that this trick for the case of finite groups is fairly abstract: hidden inside Exercise 46.10 is Tychonoff's theorem on the compactness of arbitrary products of compact spaces (cf. Lecture 6 from last semester), which depends on Zorn's lemma and thus the axiom of choice. As a consequence, we are guaranteed the existence of some $y \in (\partial^\infty)^{-1}(x)$ whenever $\partial^\infty x = 0$, but we cannot even begin to suggest how one might find y in practice. In the classic book of Eilenberg and Steenrod (see [**ES52**, Theorem 5.7 and Lemma 5.8 in Chapter VIII]), there is a linear-algebraic variation on this trick that also uses Zorn's lemma, and similarly solves the problem whenever the groups C_n^α are all assumed to be finite-dimensional vector spaces over a field $\mathbb{K}$, with ∂^α and $\varphi_{\alpha\beta}$ as $\mathbb{K}$-linear maps. These two scenarios are relevant to Čech homology under certain assumptions: in particular, suppose the coefficient group G is either finite or a finite-dimensional vector space over a field, and (X, A) is a *compact pair*, meaning X is a compact Hausdorff space and $A \subset X$ is closed. In this case, our open coverings of (X, A) always have finite refinements, whose nerves are then finite simplicial pairs, and the groups in the sequence (46.2) are therefore all either finite or are finite-dimensional vector spaces over a field $\mathbb{K}$. These conditions imply that exactness is preserved under the inverse limit, and we obtain:

THEOREM 46.22. *If G is either a finite abelian group or a finite-dimensional vector space over a field, then the restriction of Čech homology to the category $\mathsf{Cpct}_{\mathrm{rel}}$ of compact pairs defines an axiomatic homology theory*

$$\check{H}_*(\cdot\,; G) : \mathsf{Cpct}_{\mathrm{rel}} \to \mathsf{Ab}_{\mathbb{Z}}.$$

□

The restriction to compact pairs means that some details of the theory we have developed for axiomatic homology need to be handled with a bit more care: for instance, only the weaker form of the excision axiom (see Remark 32.4) makes sense in this category, so some of the excision tricks we used, e.g. for computing the homology of spheres and the isomorphism $h_*(X, A) \cong \widetilde{h}_*(X/A)$ for good pairs, need to be modified a bit when $h_* = \check{H}_*(\cdot\,; G)$. But this can be done, with the result that if G satisfies the conditions in the theorem above, then there is always a natural isomorphism

$$\check{H}_*(X, A; G) \cong H_*(X, A; G)$$

when (X, A) is a compact CW-pair.

Čech homology also has one nice property that singular homology does not: it is *continuous* with respect to inverse limits of spaces. The statement can be formulated for any axiomatic homology theory h_* as follows: suppose $\{(X_\alpha, A_\alpha), \varphi_{\alpha\beta}\}$ is an inverse system of pairs of spaces over some directed set $(I, \prec)$. The associated morphisms $\varphi_\alpha : \varprojlim\{(X_\beta, A_\beta)\} \to (X_\alpha, A_\alpha)$ then induce homomorphisms

$$\Phi_\alpha := (\varphi_\alpha)_* : h_*\big(\varprojlim\{(X_\beta, A_\beta)\}\big) \to h_*(X_\alpha, A_\alpha),$$

which make $\big\{h_*\big(\varprojlim\{(X_\beta, A_\beta)\}\big), \Phi_\alpha\big\}$ a target of the inverse system of $\mathbb{Z}$-graded abelian groups

$$\{h_*(X_\alpha, A_\alpha), (\varphi_{\alpha\beta})_*\}.$$

By the universal property of inverse limits, there is then a canonical limit morphism

$$\Phi_\infty : h_*\big(\varprojlim\{(X_\alpha, A_\alpha)\}\big) \to \varprojlim\{h_*(X_\alpha, A_\alpha)\}.$$

The following result is often quoted as the selling point of the Čech theory in comparison with singular homology. One can show in fact that every compact pair is the inverse limit of some inverse system of compact pairs that are homotopy equivalent to CW-pairs, thus the theorem can be used to understand the topology of very "wild" spaces for which singular homology cannot be expected to give a reasonable answer.

THEOREM 46.23 (continuity in Čech homology; see [ES52, Chapter X]). *For any inverse system of compact pairs $\{(X_\alpha, A_\alpha), \varphi_{\alpha\beta}\}$ and any abelian coefficient group G, the canonical map*

$$\check{H}_*\big(\varprojlim\{(X_\alpha, A_\alpha)\}; G\big) \to \varprojlim\Big\{\check{H}_*(X_\alpha, A_\alpha; G)\Big\}$$

is an isomorphism. □

EXERCISE 46.24. Find an example of a compact space X that is connected but not path-connected and is the inverse limit of a system $\{X_\alpha\}$ of path-connected spaces. Conclude that for this example,

$$H_*\big(\varprojlim\{X_\alpha\}\big) \not\cong \varprojlim\{H_*(X_\alpha)\}.$$

Hint: Use Exercise 46.9.

EXERCISE 46.25. Find an example of a path-connected space X for which $\check{H}_1(X; \mathbb{Z}_2) = 0$ but $H_1(X; \mathbb{Z}_2) \neq 0$. Can you also describe a specific nontrivial element of $\pi_1(X)$?
Hint: Take the suspension of something that is connected but not path-connected.

47. Singular cohomology (January 12, 2024)

Motivation. Singular cohomology assigns to each topological space X and each abelian group G a $\mathbb{Z}$-graded abelian group denoted by

$$H^*(X; G) = \bigoplus_{n\in\mathbb{Z}} H^n(X; G), \qquad \text{or more succinctly,} \qquad H^*(X) = \bigoplus_{n\in\mathbb{Z}} H^n(X).$$

It is closely related to singular homology, and in many (though not all) cases is isomorphic to it, but it has a slightly different structure. The most obvious difference is that as a functor from Top to $\mathsf{Ab}_{\mathbb{Z}}$, it is *contravariant*, meaning that continuous maps $f : X \to Y$ induce homomorphisms

$$f^* : H^n(Y) \to H^n(X),$$

going the opposite direction from homology. You may at this stage rightfully question what is to be gained from this cosmetic difference: as we will see, the most significant advantage is that if we choose the coefficient group G to be a ring R, then $H^*(X; R)$ has a natural product structure, called the **cup product**

$$H^k(X; R) \otimes_R H^\ell(X; R) \xrightarrow{\cup} H^{k+\ell}(X; R).$$

It is closely related to the homology cross product, but the latter is something that we use to relate the homologies of two spaces X and Y to that of their product $X \times Y$, whereas $\cup$ produces extra algebraic structure on $H^*(X; R)$ itself. This can be extremely useful in computations. Moreover, we will see that in the special case where X is a closed oriented n-manifold, $\cup$ gives rise to a product structure on homology that has deep geometric meaning, the **intersection product**

$$H_{n-k}(X; \mathbb{Z}) \otimes H_{n-\ell}(X; \mathbb{Z}) \to H_{n-k-\ell}(X; \mathbb{Z}) : [M] \otimes [N] \mapsto [M] \cdot [N] := [M \cap N].$$

This expression assumes that M and N are closed oriented submanifolds of codimension k and ℓ respectively in X, and the right hand side should be taken with a grain of salt at the moment since extra conditions are required in order for it to make sense, i.e. in order for the intersection $M \cap N \subset$

X to be a submanifold of the correct dimension and thus represent a homology class. Before explaining this, we will need to introduce **Poincaré duality**, which gives natural isomorphisms

$$H^k(X;\mathbb{Z}) \xrightarrow{\cong} H_{n-k}(X;\mathbb{Z})$$

whenever X is a closed oriented n-manifold, and thus implies various unexpected relations among the numerical invariants that one can define out of homology, e.g. the fact that every closed odd-dimensional manifold has Euler characteristic zero. These relations can be motivated geometrically in terms of triangulations, thus they were at least partially understood long before the development of cohomology theory, but the proper formulation of the isomorphism requires that we first define $H^*(X)$.

As further motivation, I would like to start by explaining a concrete topological application to a familiar problem, but one that cannot be solved using homology alone. The proof below is complete modulo a few major technical details that we will have to work through over the next several lectures, so you may consider this as motivation for the effort that will go into those details. We recall from Exercise 38.6 the complex projective space $\mathbb{CP}^n$, defined as the space of all complex lines through the origin in $\mathbb{C}^{n+1}$, meaning literally the quotient space

$$\mathbb{CP}^n = (\mathbb{C}^{n+1}\backslash\{0\})/\mathbb{C}^*,$$

where the multiplicative group $\mathbb{C}^* := \mathbb{C}\backslash\{0\}$ is understood to act on $\mathbb{C}^{n+1}\backslash\{0\}$ by scalar multiplication.

THEOREM 47.1. *For every even $n \geqslant 0$, every continuous map $f : \mathbb{CP}^n \to \mathbb{CP}^n$ has a fixed point.*

PROOF (MODULO TECHNICAL DETAILS). We saw in Exercise 38.6 that $\mathbb{CP}^n$ has a cell decomposition of the form $e^0 \cup e^2 \cup \ldots \cup e^{2n}$, i.e. it has a single k-cell for each even k from 0 to $2n$, which makes its cellular homology trivial to compute since the boundary map is necessarily zero. We will see that its singular cohomology can be computed in the same way via this cell decomposition, and gives the same answer:

$$H^k(\mathbb{CP}^n;\mathbb{Z}) \cong H_k(\mathbb{CP}^n;\mathbb{Z}) \cong \begin{cases} \mathbb{Z} & \text{for } k = 0, 2, 4, \ldots, 2n, \\ 0 & \text{for all other } k. \end{cases}$$

We will also see that there is a universal coefficient theorem expressing $H^k(X;G)$ up to isomorphism in terms of $H_k(X;\mathbb{Z})$, $H_{k-1}(X;\mathbb{Z})$ and G, and it implies moreover that the Lefschetz number $L(f) \in \mathbb{Z}$ of a map $f : X \to X$ can be computed equally well using homology or cohomology. Thus for a map $f : \mathbb{CP}^n \to \mathbb{CP}^n$, we can write

$$L(f) = \sum_{k\in\mathbb{Z}} (-1)^k \operatorname{tr}\left(H^k(\mathbb{CP}^n;\mathbb{Z}) \xrightarrow{f^*} H^k(\mathbb{CP}^n;\mathbb{Z})\right) = \sum_{k=0}^{n} \operatorname{tr}\left(H^{2k}(\mathbb{CP}^n;\mathbb{Z}) \xrightarrow{f^*} H^{2k}(\mathbb{CP}^n;\mathbb{Z})\right).$$

Now we take advantage of the cup product on $H^*(\mathbb{CP}^n)$, which has the following properties:

- It is *natural*, i.e. for all $\alpha, \beta \in H^*(\mathbb{CP}^n)$, $f^*(\alpha \cup \beta) = f^*\alpha \cup f^*\beta$. (This is a general property of the cup product with respect to continuous maps between arbitrary spaces.)
- If $\alpha \in H^2(\mathbb{CP}^n;\mathbb{Z}) \cong \mathbb{Z}$ is a generator, then for each $k = 0, 1, \ldots, n$,

  $$\alpha^k := \underbrace{\alpha \cup \ldots \cup \alpha}_{k} \in H^{2k}(\mathbb{CP}^n;\mathbb{Z}) \cong \mathbb{Z}$$

  is also a generator. We will prove this as a corollary of Poincaré duality, which holds since $\mathbb{CP}^n$ is a closed and oriented manifold.

Now fixing a generator $\alpha \in H^2(\mathbb{CP}^n;\mathbb{Z})$, every continuous map $f : \mathbb{CP}^n \to \mathbb{CP}^n$ gives rise to a unique integer $m \in \mathbb{Z}$ such that

$$f^*\alpha = m\alpha$$

since $H^2(\mathbb{CP}^n;\mathbb{Z}) \cong \mathbb{Z}$. It follows via the two properties above that for each $k = 0,\dots,n$, the generator $\alpha^k \in H^{2k}(\mathbb{CP}^n;\mathbb{Z})$ satisfies

$$f^*(\alpha^k) = f^*(\alpha \cup \ldots \cup \alpha) = f^*\alpha \cup \ldots \cup f^*\alpha = m^k\alpha^k,$$

and the Lefschetz number of f is therefore

$$L(f) = 1 + m + \ldots + m^n \in \mathbb{Z}.$$

This is clearly not equal to 0 if $m = 1$. On the other hand, if $m \neq 1$, then we can rewrite it as

$$L(f) = \frac{1 - m^{n+1}}{1 - m},$$

which is zero if and only if $m^{n+1} = 1$. Since m is an integer and we have already excluded the case $m = 1$, this can only happen if $m = -1$, and then only if n is odd. The result thus follows from the Lefschetz fixed point theorem. $\square$

The functor $\mathrm{Hom}(\cdot, G)$ and cochains. Let's talk about algebra. Given a chain complex (C_*, ∂) of abelian groups, we obtain its homology by applying the functor $H_* : \mathsf{Chain} \to \mathsf{Ab}_{\mathbb{Z}}$, which discards some of the information in (C_*, ∂) in the hope of obtaining something computable. For a little more flexibility, we can also choose an abelian coefficient group G and "pre-process" the chain complex via the functor $\otimes G$, producing the composition of functors

$$\mathsf{Chain} \xrightarrow{\otimes G} \mathsf{Chain} \xrightarrow{H_*} \mathsf{Ab}_{\mathbb{Z}}, \tag{47.1}$$

which sends the chain complex C_* to $H_*(C_* \otimes G)$.

The idea of cohomology is to pre-process the chain complex in a different way: instead of applying $\otimes G$, we apply the functor $\mathrm{Hom}(\cdot, G)$ and thus *dualize it*. You are certainly already familiar with the notion of the dual space of a vector space; more generally, the **dual** of a module A over a commutative ring R is defined as the module of R-module homomorphisms to R,

$$\mathrm{Hom}_R(A, R) := \{\lambda \in \mathrm{Hom}(A, R) \mid \lambda(ra) = r\lambda(a) \text{ for all } r \in R,\ a \in A\},$$

which reproduces the definition familiar from linear algebra if R is a field. Restricting to the case $R = \mathbb{Z}$ defines the dual of an abelian group A to be

$$A^* := \mathrm{Hom}(A, \mathbb{Z}).$$

More generally, we can fix an arbitrary abelian group G and consider the functor

$$\mathsf{Ab} \to \mathsf{Ab} : A \mapsto \mathrm{Hom}(A, G).$$

This is perhaps the simplest example of a **contravariant** functor, as one can naturally associate to each homomorphism $\Phi : A \to B$ a homomorphism in the other direction

$$\Phi^* : \mathrm{Hom}(B, G) \to \mathrm{Hom}(A, G)$$

defined by

$$\Phi^*(\lambda) := \lambda \circ \Phi \in \mathrm{Hom}(A, G) \qquad \text{for } \lambda \in \mathrm{Hom}(B, G).$$

You should take a moment to convince yourself that this satisfies the relations characteristic of a contravariant functor (see Definition 27.13): the identity map $\mathbb{1} : A \to A$ induces the identity map $\mathbb{1}^* : \mathrm{Hom}(A, G) \to \mathrm{Hom}(A, G)$, and $(\Phi\Psi)^* = \Psi^*\Phi^*$ whenever Φ and Ψ can be composed.

We next define what $\mathrm{Hom}(C_*, G)$ should mean when C_* is a chain complex with boundary map $\partial : C_* \to C_{*-1}$. Since C_* is a $\mathbb{Z}$-graded abelian group, we would like $\mathrm{Hom}(C_*, G)$ to be another $\mathbb{Z}$-graded abelian group: the obvious definition is then

$$\mathrm{Hom}(C_*, G) := \bigoplus_{n\in\mathbb{Z}} \mathrm{Hom}(C_n, G),$$

so that $\mathrm{Hom}(C_n, G)$ is the subgroup of elements with degree n in $\mathrm{Hom}(C_*, G)$.[77] Now we can dualize the map $\partial : C_* \to C_*$ to obtain a map

$$\partial^* : \mathrm{Hom}(C_*, G) \to \mathrm{Hom}(C_*, G) : \alpha \mapsto \alpha \circ \partial,$$

which sends $\mathrm{Hom}(C_n, G)$ to $\mathrm{Hom}(C_{n+1}, G)$ for each $n \in \mathbb{Z}$ and clearly satisfies $(\partial^*)^2 = 0$. For reasons that are best not to worry about right now (but see Remark 47.3), we're going to introduce an extra sign and define

$$\delta : \mathrm{Hom}(C_*, G) \to \mathrm{Hom}(C_*, G) : \alpha \mapsto (-1)^{|\alpha|+1}\partial^*\alpha, \tag{47.2}$$

where $\alpha \in \mathrm{Hom}(C_*, G)$ here is assumed to be a homogeneous element of degree $|\alpha|$, i.e. it belongs to $\mathrm{Hom}(C_n, G)$ for $n = |\alpha|$. This clearly also satisfies the relation

$$\delta^2 = 0,$$

and it is a map of degree $+1$, meaning it sends $\mathrm{Hom}(C_n, G)$ to $\mathrm{Hom}(C_{n+1}, G)$ for every $n \in \mathbb{Z}$.

We shall refer to any $\mathbb{Z}$-graded abelian group A_* endowed with a homomorphism $\delta : A_* \to A_*$ of degree $+1$ satisfying $\delta^2 = 0$ as a **cochain complex**. Up to a minor matter of bookkeeping, this is the same thing as a chain complex, and the notions of **chain map** and **chain homotopy** carry over in obvious ways: in particular, a chain homotopy between two chain maps $\varphi, \psi : A_* \to B_*$ of cochain complexes (A_*, δ_A) and (B_*, δ_B) is a homomorphism $h : A_* \to B_*$ of degree -1 that satisfies the usual chain homotopy relation

$$\varphi - \psi = h\delta_A + \delta_B h.$$

The **homology** of a cochain complex (A_*, δ) is the $\mathbb{Z}$-graded abelian group

$$H_*(A_*, \delta) = \ker\delta / \operatorname{im}\delta,$$

so in other words $H_n(A_*, \delta) = \ker\delta_n / \operatorname{im}\delta_{n-1}$ if $A_n \xrightarrow{\delta_n} A_{n+1}$ denotes the restriction of δ for each $n \in \mathbb{Z}$. With these notions in place, we can associate to any chain complex (C_*, ∂) its **cohomology** with coefficients in G: this is the $\mathbb{Z}$-graded abelian group

$$H^*(C_*, \partial; G) := H_*(\mathrm{Hom}(C_*, G), \delta).$$

Much like (47.1), the functor that replaces a chain complex with its cohomology can be expressed as the composition of two functors:

$$\mathsf{Chain} \xrightarrow{\mathrm{Hom}(\cdot,G)} \mathsf{Cochain} \xrightarrow{H_*} \mathsf{Ab}_{\mathbb{Z}}.$$

Here $\mathsf{Cochain}$ denotes the category whose objects are cochain complexes, with morphisms defined as chain maps, and $H_* : \mathsf{Cochain} \to \mathsf{Ab}_{\mathbb{Z}}$ is a covariant functor sending each cochain complex to its homology and chain maps to the induced homomorphisms on homology. The functor

$$\mathrm{Hom}(\cdot, G) : \mathsf{Chain} \to \mathsf{Cochain}$$

---

[77] If you know enough algebra and are paying close attention, you might now notice an incongruity in our notation: unless C_* happens to be nonzero in only finitely many degrees, $\mathrm{Hom}(C_*, G)$ as we've defined it is not literally the group of all homomorphisms $C_* \to G$. That would be $\prod_{n\in\mathbb{Z}} \mathrm{Hom}(C_n, G)$, as dualizing infinite direct sums generally gives rise to direct products. This should not be a cause for concern, you just need to keep in mind that the notation $\mathrm{Hom}(C_*, G)$ is not to be interpreted too literally.

replaces a chain complex (C_*, ∂) with the cochain complex $(\mathrm{Hom}(C_*, G), \delta)$ as defined above, and it is contravariant: it associates to each chain map $\varphi : (A_*, \partial_A) \to (B_*, \partial_B)$ the dual map

$$\varphi^* : (\mathrm{Hom}(B_*, G), \delta_A) \to (\mathrm{Hom}(A_*, G), \delta_B),$$

which is a chain map since for $\beta \in \mathrm{Hom}(B_n, G)$,

$$\begin{aligned}\varphi^* \delta_B \beta &= \varphi^* \left((-1)^{n+1} \partial_B^* \beta\right) = (-1)^{n+1} \varphi^* \partial_B^* \beta = (-1)^{n+1} (\partial_B \varphi)^* \beta = (-1)^{n+1} (\varphi \partial_A)^* \beta \\ &= (-1)^{n+1} \partial_A^* \varphi^* \beta = \delta_A \varphi^* \beta.\end{aligned}$$

As a consequence, the composition functor $H^*(\cdot; G) : \mathsf{Chain} \to \mathsf{Ab}_{\mathbb{Z}}$ is also contravariant: it associates to each chain map $\varphi : (A_*, \partial_A) \to (B_*, \partial_B)$ the homomorphism $H_*(\mathrm{Hom}(B_*, G), \delta_B) \to H_*(\mathrm{Hom}(A_*, G), \delta_A)$ induced by the chain map φ^*, and we shall also denote the induced morphism of $\mathbb{Z}$-graded abelian groups by

$$\varphi^* : H^*(B_*, \partial_B; G) \to H^*(A_*, \partial_A; G).$$

Two further algebraic observations are worth recording before we go back to topology.

PROPOSITION 47.2. *If $\varphi, \psi : (A_*, \partial_A) \to (B_*, \partial_B)$ are chain maps between chain complexes and $h : A_* \to B_{*+1}$ is a chain homotopy between φ and ψ, then the map $\eta : \mathrm{Hom}(B_*, G) \to \mathrm{Hom}(A_{*-1}, G)$ defined for each $n \in \mathbb{Z}$ by*

$$\mathrm{Hom}(B_n, G) \xrightarrow{\eta} \mathrm{Hom}(A_{n-1}, G) : \beta \mapsto (-1)^n h^* \beta$$

is a chain homotopy between φ^* and ψ^*.

PROOF. We have $\varphi^* - \psi^* = (\varphi - \psi)^* = (h\partial_A + \partial_B h)^* = \partial_A^* h^* + h^* \partial_B^*$, thus for any $\beta \in \mathrm{Hom}(B_n, G)$,

$$(\delta_A \eta + \eta \delta_B)\beta = \partial_A^* h^* \beta + h^* \partial_B^* \beta = (\varphi^* - \psi^*)\beta.$$

□

In category-theoretic terms, the proposition means that $\mathrm{Hom}(\cdot, G)$ descends to a well-defined functor

$$\mathrm{Hom}(\cdot, G) : \mathsf{Chain}^h \to \mathsf{Cochain}^h,$$

where $\mathsf{Cochain}^h$ is the category with cochain complexes as objects and chain homotopy classes of chain maps as morphisms. As a consequence, $H^*(\cdot; G)$ likewise descends to a functor

$$H^*(\cdot; G) : \mathsf{Chain}^h \to \mathsf{Ab}_{\mathbb{Z}}.$$

The second observation is that for any chain complex (C_*, ∂), the canonical pairing

$$\mathrm{Hom}(C_n, G) \times C_n \to G : (\alpha, c) \mapsto \alpha(c) \tag{47.3}$$

descends to homology to give a well-defined pairing

$$H^n(C_*, \partial; G) \times H_n(C_*, \partial) \to G : ([\alpha], [c]) \mapsto \langle [\alpha], [c] \rangle := \alpha(c). \tag{47.4}$$

To see that this is well defined, we observe that if $\delta\alpha$ and ∂c are both assumed to be zero, then in the case $c = \partial a$ for some $a \in C_{n+1}$, we have

$$\alpha(\partial a) = (\partial^* \alpha)(a) = \pm(\delta\alpha)(a) = 0,$$

and similarly if $\alpha = \delta\beta$ for some $\beta \in \mathrm{Hom}(C_{n-1}, G)$,

$$(\delta\beta)(c) = \pm(\partial^*\beta)(c) = \pm\beta(\partial c) = 0.$$

We will often refer to (47.4) as the **evaluation** of cohomology classes on homology classes.

REMARK 47.3. The reason for the sign in (47.2) can be understood in terms of the "chain-level" evaluation map (47.3). Since it is bilinear, it can be expressed as a homomorphism

$$\operatorname{Hom}(C_n, G) \otimes C_n \to G,$$

which extends in a trivial way to all degrees as a homomorphism

$$\operatorname{Hom}(C_*, G) \otimes C_* \to G \tag{47.5}$$

if we define $\alpha(c) := 0$ whenever $\alpha \in \operatorname{Hom}(C_k, G)$ and $c \in C_\ell$ for $k \neq \ell$. With a little care, we can then rephrase the fact that (47.4) is well defined as a corollary of the fact that (47.5) is a chain map. For this we need to make sense of $\operatorname{Hom}(C_*, G) \otimes C_*$ as a tensor product chain complex, even though $\operatorname{Hom}(C_*, G)$ strictly speaking is not a chain complex but a *cochain* complex: however, any cochain complex becomes a chain complex if we simply reverse the degrees by a sign, so let us write

$$\operatorname{Hom}(C_*, G)_n := \operatorname{Hom}(C_{-n}, G)$$

and think of δ as a homomorphism that sends $\operatorname{Hom}(C_*, G)_n$ to $\operatorname{Hom}(C_*, G)_{n-1}$. The fact that $\alpha(c) = 0$ whenever $\alpha \in \operatorname{Hom}(C_k, G)$ and $c \in C_\ell$ with $k \neq \ell$ then means that the map (47.5) vanishes on all elements of degree nonzero in the tensor product chain complex, so it becomes natural to understand the right hand side as a chain complex that has the group G in degree 0 and the trivial group in all other degrees. With this convention in place, the boundary map on the right hand side is zero, so the chain map condition demands that for all $\alpha \in \operatorname{Hom}(C_k, G)$ and $c \in C_\ell$,

$$\partial(\alpha \otimes c) = \delta\alpha \otimes c + (-1)^k \alpha \otimes \partial c \mapsto (\delta\alpha)(c) + (-1)^k \alpha(\partial c) = 0,$$

leading in the case $k = \ell - 1 = n$ to the formula

$$(\delta\alpha)(c) = -(-1)^n \alpha(\partial c) = (-1)^{n+1}(\partial^*\alpha)(c).$$

The sign in (47.2) is therefore necessary in order to make the evaluation $\operatorname{Hom}(C_*, G) \otimes C_* \to G$ a chain map in this sense.

It is not strictly necessary to adopt this sign convention, and many textbooks do not; you will notice of course that the definition of $H^*(C_*, \partial; G)$ does not care whether the sign is included since it does not change $\ker \delta$ or $\operatorname{im} \delta$. But if we don't include the sign here, we will be forced to insert a different unwanted sign somewhere later in the development of the theory. I am trying to stay consistent with the conventions in [**Bre93**].

REMARK 47.4. The entirety of this discussion admits a straightforward generalization in which (C_*, ∂) is a chain complex of modules over a commutative ring R with unit, G is also an R-module, and the functor $\operatorname{Hom}(\cdot, G)$ is replaced by $\operatorname{Hom}_R(\cdot, G)$, which transforms a chain complex of R-modules into a cochain complex of R-modules. The homology $H_*(C_*, \partial)$ and cohomology $H^*(C_*, \partial; G)$ then both also have natural R-module structures, and the evaluation (47.4) becomes an R-module homomorphism

$$H^*(C_*, \partial; G) \otimes_R H_*(C_*, \partial) \to G,$$

which just means that (47.4) is R-bilinear. This situation arises naturally if we start with (C_*, ∂) as a chain complex of abelian groups but then introduce R as a coefficient ring by replacing it with $(C_* \otimes R, \partial \otimes \mathbb{1})$, which has a natural R-module structure. The following exercise shows that if we now take the cohomology of $(C_* \otimes R, \partial \otimes \mathbb{1})$ with coefficients in an R-module G by applying $\operatorname{Hom}_R(\cdot, G)$ and then H_*, we obtain exactly the same result as the cohomology of (C_*, ∂) with coefficients in G, except that instead of just the pairing of abelian groups that is defined in (47.4), we have a bilinear R-module pairing

$$H^*(C_*, \partial; G) \times H_*(C_* \otimes R, \partial \otimes \mathbb{1}) \to G.$$

EXERCISE 47.5. Assume R is a commutative ring with unit, A is an abelian group and G is an R-module.

(a) Show that the abelian groups $A \otimes G$ and $\mathrm{Hom}(A, G)$ each have natural R-module structures defined via the relations
$$r(a \otimes g) := a \otimes (rg) \quad \text{for } r \in R,\ a \in A,\ g \in G,$$
$$(r\Phi)(a) := r\,(\Phi(a)) \quad \text{for } r \in R,\ a \in A,\ \Phi \in \mathrm{Hom}(A, G),$$
and that $\otimes G$ and $\mathrm{Hom}(\cdot, G)$ can each be understood as functors (covariant and contravariant respectively) from the category of abelian groups to the category of R-modules.

(b) Show that there is a canonical R-module isomorphism between $\mathrm{Hom}(A, G)$ and $\mathrm{Hom}_R(A \otimes R, G)$, which defines a natural transformation between the functors
$$\mathsf{Ab} \to \mathsf{Mod}^R : A \mapsto \mathrm{Hom}(A, G) \quad \text{and} \quad \mathsf{Ab} \to \mathsf{Mod}^R : A \mapsto \mathrm{Hom}_R(A \otimes R, G).$$

The singular cochain complex. The **singular cohomology** of a pair (X, A) with coefficients in an abelian group G is now defined by applying the algebraic processing described above to the singular chain complex with integer coefficients: that is,
$$H^*(X, A) = H^*(X, A; G) := H^*\big(C_*(X, A; \mathbb{Z}); G\big) = H_*\big(\mathrm{Hom}(C_*(X, A; \mathbb{Z}), G)\big).$$
As we do with homology, we shall follow the practice of omitting the coefficient group G from the notation for cohomology in most situations where the choice of coefficients is unimportant. It is standard to abbreviate the cochain complex $\mathrm{Hom}(C_*(X, A; \mathbb{Z}), G)$ by
$$C^*(X, A) = C^*(X, A; G) := \mathrm{Hom}\left(C_*(X, A; \mathbb{Z}), G\right)$$
and refer to elements of $C^*(X, A)$ as **singular cochains** with coefficients in G. Elements of $\ker \delta \subset C^*(X, A)$ and $\mathrm{im}\, \delta \subset C^*(X, A)$ are likewise called (singular) **cocycles** and **coboundaries** respectively. Each element $\varphi \in C^n(X)$ is a homomorphism $\varphi : C_n(X; \mathbb{Z}) \to G$, and since $C_n(X; \mathbb{Z})$ is a free abelian group, all such homomorphisms can be described uniquely via their values on the generators, i.e. the singular n-simplices in X. We thus have a canonical identification
$$C^n(X) = G^{\mathcal{K}_n(X)} = \prod_{\sigma \in \mathcal{K}_n(X)} G = \{\text{functions } \varphi : \mathcal{K}_n(X) \to G\},$$
where $\mathcal{K}_n(X)$ again denotes the set of all singular n-simplices in X. We will often use this identification to regard cochains $\varphi \in C^n(X)$ simply as functions $\varphi : \mathcal{K}_n(X) \to G$. With this understood, we plug in (47.2) and the usual formula for the boundary operator $\partial : C_{n+1}(X; \mathbb{Z}) \to C_n(X; \mathbb{Z})$ to find a corresponding formula for the **coboundary operator** $\delta : C^n(X) \to C^{n+1}(X)$, in the form
$$(\delta\varphi)(\sigma) = (-1)^{n+1} \sum_{k=0}^{n+1} (-1)^k \varphi\big(\sigma|_{\partial_{(k)}\Delta^{n+1}}\big) \quad \text{for} \quad \varphi : \mathcal{K}_n(X) \to G, \quad \sigma \in \mathcal{K}_{n+1}(X). \tag{47.6}$$
In the relative case, we can think of a homomorphism $\varphi : C_n(X, A; \mathbb{Z}) = C_n(X; \mathbb{Z})/C_n(A; \mathbb{Z}) \to G$ as equivalent to a homomorphism $\varphi : C_n(X; \mathbb{Z}) \to G$ that vanishes on the subgroup $C_n(A; \mathbb{Z}) \subset C_n(X; \mathbb{Z})$, so this is the same thing as a function $\mathcal{K}_n(X) \to G$ that vanishes on the subset $\mathcal{K}_n(A) \subset \mathcal{K}_n(X)$:
$$C^n(X, A) = \left\{ \varphi : \mathcal{K}_n(X) \to G \;\middle|\; \varphi|_{\mathcal{K}_n(A)} = 0 \right\}.$$
The formula (47.6) then gives the correct homomorphism $\delta : C^n(X, A) \to C^{n+1}(X, A)$ by restriction.

As a functor, $H^* = H^*(\cdot; G) : \mathsf{Top}_{\mathrm{rel}} \to \mathsf{Ab}_{\mathbb{Z}}$ is the composition of three functors,
$$\mathsf{Top}_{\mathrm{rel}} \xrightarrow{C_*(\cdot;\mathbb{Z})} \mathsf{Chain} \xrightarrow{\mathrm{Hom}(\cdot,G)} \mathsf{Cochain} \xrightarrow{H_*} \mathsf{Ab}_{\mathbb{Z}},$$

one of which is contravariant, thus H^* is also contravariant. Concretely, this means that continuous maps of pairs $f : (X, A) \to (Y, B)$ induce "pullback" homomorphisms

$$f^* : H^n(Y, B) \to H^n(X, A)$$

for every $n \in \mathbb{Z}$. These maps are induced by the chain map $f^* : C^*(Y, B) \to C^*(X, A)$ defined by

$$(f^*\varphi)(c) := \varphi(f_*c) \quad \text{for} \quad \varphi \in C^n(Y, B),\ c \in C_n(X, A; \mathbb{Z}).$$

By the previous algebraic discussion, there is a natural pairing

$$H^*(X, A; G) \otimes H_*(X, A; \mathbb{Z}) \to G : [\varphi] \otimes [c] \mapsto \langle [\varphi], [c] \rangle := \varphi(c),$$

which we call the **evaluation** of the cohomology class $[\varphi]$ on the homology class $[c]$, and it satisfies

$$(47.7)\quad \langle f^*[\varphi], [c] \rangle = \langle [\varphi], f_*[c] \rangle \quad \text{for} \quad [\varphi] \in H^*(Y, B; G),\ [c] \in H_*(X, A; \mathbb{Z}),\ (X, A) \xrightarrow{f} (Y, B).$$

More generally, if G is a module over a commutative ring R with unit, one can define an R-bilinear pairing

$$H^*(X, A; G) \otimes_R H_*(X, A; R) \to G$$

as in Remark 47.4.

Let us conclude this lecture with two straightforward but revealing computations of $H^n(X)$ for particular values of n. We start with the case $n = 0$.

For any space X, $C^{-1}(X) = 0$, thus $H^0(X)$ is simply the kernel of the map $C^0(X) \xrightarrow{\delta} C^1(X)$, also known as the group of 0-cocycles. Under the usual identification of $\mathcal{K}_0(X)$ with X and $\mathcal{K}_1(X)$ with the set of paths $\gamma : I \to X$, (47.6) gives

$$(\delta\varphi)(\gamma) = \pm [\varphi(\gamma(1)) - \varphi(\gamma(0))] \quad \text{for} \quad \varphi : X \to G, \quad \gamma : I \to X,$$

which vanishes for all paths γ if and only if $\varphi(x) = \varphi(y)$ for every pair of points $x, y \in X$ that are in the same path-component of X. A function $\varphi : X \to G$ is therefore a 0-cocycle if and only if it is constant on path-components, meaning it is equivalent to a function $\pi_0(X) \to G$. We've proved:

THEOREM 47.6. *For any space X and abelian group G, there is a canonical isomorphism*

$$H^0(X; G) \cong \prod_{\pi_0(X)} G.$$

□

REMARK 47.7. This proves that $H^0(X; G) \cong H_0(X; G)$ if X has only finitely-many path-components, but otherwise $H^0(X; G)$ is larger than $H_0(X; G)$. Indeed, for any collection of groups $\{G_\alpha\}_{\alpha \in I}$, the direct sum $\bigoplus_{\alpha \in I} G_\alpha$ can be identified with the subgroup of the direct product $\prod_{\alpha \in I} G_\alpha$ consisting of tuples $\{g_\alpha\}_{\alpha \in I}$ that have at most finitely-many nonzero coordinates. For example, if the index set I is $\mathbb{N}$ and $G_\alpha = \mathbb{Z}_2$ for every $\alpha \in I$, then $\bigoplus_{\alpha \in I} G_\alpha$ is countably infinite but $\prod_{\alpha \in I} G_\alpha$ is uncountable.

The second computation relates $H^1(X; G)$ to $\pi_1(X)$; we shall give a brief sketch and leave the details as exercises. Assume X is a path-connected space, and identify Δ^1 with $I = [0, 1]$ as usual so that singular 1-cochains $\varphi \in C^1(X; G)$ can be interpreted as functions from the set of paths $\{\gamma : I \to X\}$ to G.

EXERCISE 47.8. Show that a singular 1-cochain $\varphi \in C^1(X; G)$ is a cocycle if and only if it satisfies both of the following:

(i) For all paths $\gamma : I \to X$, $\varphi(\gamma) \in G$ depends only on the homotopy class of γ with fixed end points;
(ii) For every pair of paths $\alpha, \beta : I \to X$ with $\alpha(1) = \beta(0)$, $\varphi(\alpha \cdot \beta) = \varphi(\alpha) + \varphi(\beta)$.

Hint: If $\sigma : \Delta^2 \to X$ is a singular 2-simplex, one can identify its three boundary faces with paths $\alpha, \beta, \gamma : I \to X$ such that $\alpha \cdot \beta$ is homotopic to γ with fixed end points.

EXERCISE 47.9. Show that a singular 1-cochain $\varphi \in C^1(X; G)$ is a coboundary if and only if there exists a function[78] $\psi : X \to G$ such that for all paths $\gamma : I \to X$, $\varphi(\gamma) = \psi(\gamma(1)) - \psi(\gamma(0))$.

EXERCISE 47.10. Prove that for any $x \in X$, there is a well-defined homomorphism

$$\Psi : H^1(X; G) \to \mathrm{Hom}(\pi_1(X, x), G) : [\varphi] \mapsto \Psi_\varphi$$

such that for each 1-cocycle $\varphi \in C^1(X; G)$, $\Psi_\varphi : \pi_1(X, x) \to G$ is given by

$$\Psi_\varphi([\gamma]) = \varphi(\gamma) \quad \text{for} \quad x \overset{\gamma}{\rightsquigarrow} x.$$

Then prove that Ψ is injective and surjective.
Hint: For injectivity, you need to show that if $\varphi(\gamma) = 0$ for all loops γ then φ satisfies the condition in Exercise 47.9. For surjectivity, it might help to observe that since $H_1(X; \mathbb{Z})$ is the abelianization of $\pi_1(X, x)$ and G is abelian, $\mathrm{Hom}(\pi_1(X, x), G) = \mathrm{Hom}(H_1(X; \mathbb{Z}), G)$, so the map $\Psi : H^1(X; G) \to \mathrm{Hom}(\pi_1(X, x), G)$ can then be identified with

$$H^1(X; G) \to \mathrm{Hom}(H_1(X; \mathbb{Z}), G) : [\varphi] \mapsto \langle [\varphi], \cdot \rangle.$$

You then need to show that every homomorphism to G from the group Z_1 of 1-cycles that vanishes on the subgroup $B_1 \subset Z_1$ of boundaries can be extended to a homomorphism $C_1(X; \mathbb{Z}) \to G$. Use the fact that $0 \to Z_1 \hookrightarrow C_1(X; \mathbb{Z}) \xrightarrow{\partial} B_0 \to 0$ is a split exact sequence. (Why?)

48. Axioms for cohomology (January 16, 2024)

Eilenberg-Steenrod revisited. Each of the Eilenberg-Steenrod axioms for homology theories has an analogue that is satisfied by singular cohomology, thus giving rise to the notion of *axiomatic* cohomology theories. The proof that $H^*(\cdot; G)$ satisfies the axioms is at this point quite easy; it is mostly a matter of reusing the same lemmas that were used for proving properties of $H_*(\cdot; G)$, but with most of the arrows reversed.

DEFINITION 48.1. An **axiomatic cohomology theory** h^* is a contravariant functor

$$\mathsf{Top}_{\mathrm{rel}} \to \mathsf{Ab}_{\mathbb{Z}} : (X, A) \mapsto h^*(X, A) = \bigoplus_{n \in \mathbb{Z}} h^n(X, A),$$

together with a natural transformation δ^* from the functor $\mathsf{Top}_{\mathrm{rel}} \to \mathsf{Ab} : (X, A) \mapsto h^n(A)$ to the functor $\mathsf{Top}_{\mathrm{rel}} \to \mathsf{Ab} : (X, A) \mapsto h^{n+1}(X, A)$ for each $n \in \mathbb{Z}$ such that the following axioms are satisfied:

- (EXACTNESS) For all pairs (X, A) with inclusion maps $i : A \hookrightarrow X$ and $j : (X, \varnothing) \hookrightarrow (X, A)$, the sequence

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

  is exact.
- (HOMOTOPY) For any two homotopic maps $f, g : (X, A) \to (Y, B)$, the induced morphisms $f^*, g^* : h^*(Y, B) \to h^*(X, A)$ are identical.
- (EXCISION) For any pair (X, A) and any subset $B \subset X$ with closure in the interior of A, the inclusion $(X \backslash B, A \backslash B) \hookrightarrow (X, A)$ induces an isomorphism

  $$h^*(X, A) \xrightarrow{\cong} h^*(X \backslash B, A \backslash B).$$

[78]Note that since neither G nor the set of singular 0-simplices $\mathcal{K}_0(X) \cong X$ in this discussion is understood to be endowed with a topology, there is no continuity assumption on the function $\psi : X \to G$.

- (DIMENSION) For any space $\{\text{pt}\}$ containing only one point, $h^n(\{\text{pt}\}) = 0$ for all $n \neq 0$.
- (ADDITIVITY) For any collection of spaces $\{X_\alpha\}_{\alpha\in J}$ with inclusion maps $i_\alpha : X_\alpha \hookrightarrow \coprod_{\beta\in J} X_\beta$, the induced homomorphisms $i_\alpha^* : h^*\left(\coprod_{\beta\in J} X_\beta\right) \to h^*(X_\alpha)$ determine an isomorphism
$$\prod_{\alpha\in J} i_\alpha^* : h^*\left(\coprod_{\beta\in J} X_\beta\right) \xrightarrow{\cong} \prod_{\alpha\in J} h^*(X_\alpha).$$

The group $h^0(\{\text{pt}\})$ is called the **coefficient group** of the theory.

THEOREM 48.2. *For any abelian group G, the singular cohomology $H^*(\cdot;G)$ is an axiomatic cohomology theory with coefficient group G.*

PROOF. Exactness follows from the fact that if we dualize the usual short exact sequence of singular chain complexes $0 \to C_*(A;\mathbb{Z}) \xrightarrow{i_*} C_*(X;\mathbb{Z}) \xrightarrow{j_*} C_*(X,A;\mathbb{Z}) \to 0$, then the resulting sequence of chain maps

$$0 \longleftarrow C^*(A;G) \xleftarrow{i^*} C^*(X;G) \xleftarrow{j^*} C^*(X,A;G) \longleftarrow 0 \tag{48.1}$$

is also exact. Indeed, under the canonical identifications of these groups with sets of functions $\mathcal{K}_n(X) \to G$ or $\mathcal{K}_n(A) \to G$, j^* becomes the obvious inclusion

$$j^* : \left\{\varphi : \mathcal{K}_n(X) \to G \;\middle|\; \varphi|_{\mathcal{K}_n(A)} = 0\right\} \hookrightarrow \{\varphi : \mathcal{K}_n(X) \to G\},$$

and i^* becomes the restriction map

$$i^* : \{\varphi : \mathcal{K}_n(X) \to G\} \to \{\varphi : \mathcal{K}_n(A) \to G\} : \varphi \mapsto \varphi|_{\mathcal{K}_n(A)},$$

which is manifestly surjective and has kernel equal to $\operatorname{im} j^*$. I should caution you against thinking that the exactness of this dualized sequence follows automatically from abstract nonsense—we will see in the next lecture that not every short exact sequence of abelian groups remains exact after it is dualized. But this one does. As a result, (48.1) is what we may sensibly call a short exact sequence of *cochain* complexes, which is the same thing as a short exact sequence of chain complexes except that the coboundary operator raises degrees instead of lowering them. The usual diagram-chasing argument therefore produces from this a long exact sequence of the homology groups of the complexes, with a connecting homomorphism that raises the degree by 1.

The main reason for the homotopy axiom is Proposition 47.2 in the previous lecture, which implies that if the two chain maps $f_*, g_* : C_*(X,A;\mathbb{Z}) \to C_*(Y,B;\mathbb{Z})$ are chain homotopic, then so are the two chain maps $f^*, g^* : C^*(Y,B;G) \to C^*(X,A;G)$.

For excision, recall from Theorem 29.20 that if $B \subset \bar{B} \subset \mathring{A} \subset A \subset X$, then the inclusion $i : (X\backslash B, A\backslash B) \hookrightarrow (X,A)$ induces a chain homotopy equivalence $i_* : C_*(X\backslash B, A\backslash B;\mathbb{Z}) \to C_*(X,A;\mathbb{Z})$, meaning in particular that there is a chain map $\rho_* : C_*(X,A;\mathbb{Z}) \to C_*(X\backslash B, A\backslash B;\mathbb{Z})$ such that $\rho_* i_*$ and $i_*\rho_*$ are each chain homotopy equivalent to the identity. Dualizing both i_* and ρ_* then produces chain maps $i^* : C^*(X,A;G) \to C^*(X\backslash B, A\backslash B;G)$ and $\rho^* : C^*(X\backslash B, A\backslash B;G) \to C^*(X,A;G)$ such that by Proposition 47.2, $i^*\rho^*$ and $\rho^* i^*$ are also chain homotopic to the identity, hence

$$i^* : C^*(X,A;G) \to C^*(X\backslash B, A\backslash B;G)$$

is a chain homotopy equivalence and induces an isomorphism $H^*(X,A;G) \to H^*(X\backslash B, A\backslash B;G)$.

The dimension axiom and the computation of the coefficient group are straightforward since there is only one singular n-simplex $\sigma_n \in \mathcal{K}_n(\{\text{pt}\})$ for each $n \geqslant 0$, giving canonical isomorphisms

$$C^n(\{\text{pt}\};G) \xrightarrow{\cong} G : \varphi \mapsto \varphi(\sigma_n).$$

The map $\delta : C^n(\{\mathrm{pt}\}; G) \to C^{n+1}(\{\mathrm{pt}\}; G)$ then becomes

$$\delta_n : G \to G : g \mapsto (-1)^{n+1} \sum_{k=0}^{n+1} (-1)^k g = \begin{cases} 0 & \text{if } n \text{ is even,} \\ (-1)^{n+1} g & \text{if } n \text{ is odd.} \end{cases}$$

For $n > 0$ even, this means $\ker \delta_n = \operatorname{im} \delta_{n-1}$ and thus $H^n(\{\mathrm{pt}\}; G) = 0$. For $n > 0$ odd, we instead have $\ker \delta_n = 0$ and thus $H^n(\{\mathrm{pt}\}; G) = 0$. The only special case is $n = 0$, for which $H^0(\{\mathrm{pt}\}; G) = \ker \delta_0 = G$.

The additivity axiom is a straightforward consequence of the fact that since no individual singular simplex can have image in more than one component of a disjoint union, the chain complex $C_*\big(\coprod_\beta X_\beta; \mathbb{Z}\big)$ splits naturally into a direct sum of chain complexes $\bigoplus_\beta C_*(X_\beta; \mathbb{Z})$. Dualizing then changes the direct sum to a direct product as we saw in the computation of $H^0(X; G)$ in the previous lecture. We leave the details as an exercise. $\square$

Exercise 48.3. Describe a cohomological version of the "braid" diagram in Lecture 32 and use it to prove that for every triple of spaces (X, A, B) with $B \subset A \subset X$ and every axiomatic cohomology theory h^*, the maps induced by the inclusions $i : (A, B) \hookrightarrow (X, B)$ and $j : (X, B) \hookrightarrow (X, A)$ fit into a long exact sequence

$$\ldots \longleftarrow h^{n+1}(X, A) \xleftarrow{\delta^*} h^n(A, B) \xleftarrow{i^*} h^n(X, B) \xleftarrow{j^*} h^n(X, A) \xleftarrow{\delta^*} h^{n-1}(A, B) \longleftarrow \ldots.$$

Give also an alternative proof of this for singular cohomology using a short exact sequence of cochain complexes.

Reduced cohomology. Every cohomology theory h^* also has a **reduced** version, which is again defined in terms of the unique map

$$\epsilon : X \to \{\mathrm{pt}\}.$$

Choosing any embedding $i : \{\mathrm{pt}\} \to X$, the fact that $\epsilon \circ i$ is the identity map implies that

$$(\epsilon \circ i)^* = i^* \epsilon^* : h^*(\{\mathrm{pt}\}) \to h^*(\{\mathrm{pt}\})$$

is also the identity, so $\epsilon^* : h^*(\{\mathrm{pt}\}) \to h^*(X)$ is injective and has i^* as a left-inverse. We then define

$$\widetilde{h}^*(X) := \operatorname{coker} \epsilon^* = h^*(X) / \operatorname{im} \epsilon^*,$$

so that the quotient projection $h^*(X) \to \widetilde{h}^*(X)$ fits into a split exact sequence

$$0 \longrightarrow h^*(\{\mathrm{pt}\}) \xrightarrow{\epsilon^*} h^*(X) \longrightarrow \widetilde{h}^*(X) \longrightarrow 0,$$

implying that if h^* has coefficient group G,

$$h^n(X) \cong \begin{cases} \widetilde{h}^n(X) \oplus G & \text{for } n = 0, \\ \widetilde{h}^n(X) & \text{for } n \neq 0. \end{cases}$$

If X is contractible, then ϵ is a homotopy equivalence and $\epsilon^* : h^*(\{\mathrm{pt}\}) \to h^*(X)$ is thus an isomorphism, so its cokernel is trivial:

Theorem 48.4. *For any axiomatic cohomology theory h^*, if X is contractible, $\widetilde{h}^*(X) = 0$.* $\square$

As with homology, this result is mainly useful because of the role that trivial homology groups play in exact sequences. We showed in Lecture 29 via diagram-chasing arguments that the homology long exact sequence of a pair (X, A) is also exact if all homology groups are replaced by their reduced versions, where the reduced homology of a pair (X, A) with $A \neq \varnothing$ is defined to match the ordinary homology. We can do the same thing here: if we define

$$\widetilde{h}^*(X, A) := h^*(X, A) \quad \text{if} \quad A \neq \varnothing,$$

then repeating the arguments of Lecture 29 with reversed arrows gives:

THEOREM 48.5. *For any pair* (X, A) *and any axiomatic cohomology theory, the sequence*

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

is also well defined and exact. □

EXERCISE 48.6. Show that $\widetilde{H}^*(X; G)$ is also the cohomology (with coefficient group G) of the *augmented* chain complex

$$\ldots \longrightarrow C_2(X; \mathbb{Z}) \xrightarrow{\partial} C_1(X; \mathbb{Z}) \xrightarrow{\partial} C_0(X; \mathbb{Z}) \xrightarrow{\epsilon} \widetilde{C}_{-1}(X; \mathbb{Z}) := \mathbb{Z} \longrightarrow 0 \longrightarrow 0 \longrightarrow \ldots$$

described in Remark 29.14.

EXERCISE 48.7. Adapt the proof of Theorem 32.14 to prove that for any axiomatic cohomology theory h^* and any space X, there is a natural isomorphism $\widetilde{h}^n(X) \to \widetilde{h}^{n+1}(SX)$ for every $n \in \mathbb{Z}$.

EXERCISE 48.8. For any axiomatic cohomology theory h^* and two spaces X and Y with maps $\epsilon_X : X \to \{\mathrm{pt}\}$ and $\epsilon_Y : Y \to \{\mathrm{pt}\}$, show that the isomorphism $h^*(X \amalg Y) \cong h^*(X) \times h^*(Y)$ given by the additivity axiom identifies $\widetilde{h}_*(X \amalg Y)$ with the cokernel of the map

$$(\epsilon_X^*, \epsilon_Y^*) : h^*(\{\mathrm{pt}\}) \to h^*(X) \times h^*(Y).$$

Then apply this in the case $X = Y = \{\mathrm{pt}\}$ to identify $\widetilde{h}^0(\{\mathrm{pt}\} \amalg \{\mathrm{pt}\})$ with the cokernel of the diagonal map $G \to G \times G$, where $G = h^0(\{\mathrm{pt}\})$. Conclude in particular

$$\widetilde{h}^n(S^0) \cong \begin{cases} G & \text{if } n = 0, \\ 0 & \text{if } n \neq 0. \end{cases}$$

EXERCISE 48.9. Combine the previous two exercises to prove by induction on $n \in \mathbb{N}$ that for any axiomatic cohomology theory h^* with coefficient group G,

$$h^k(S^n) \cong \begin{cases} G & \text{if } k = 0 \text{ or } k = n, \\ 0 & \text{otherwise.} \end{cases}$$

EXERCISE 48.10. Adapt the proof of Theorem 32.23 to prove that for any axiomatic cohomology theory h^* and any good pair (X, A), there is a natural isomorphism

$$h^*(X, A) \cong \widetilde{h}^*(X/A).$$

REMARK 48.11. You may by now be getting the impression that cohomology is always isomorphic to homology, especially in light of the computation above for S^n. There is a grain of truth in this, but the whole story is more complicated: e.g. we will see in the next lecture that $H^*(X; G)$ is fully determined up to isomorphism by $H_*(X; \mathbb{Z})$ and G, but it is not always isomorphic to $H_*(X; G)$, especially if $H_*(X; \mathbb{Z})$ has torsion. It also deserves to be emphasized that for arbitrary axiomatic theories, the premise does not always make sense: in contrast to the obvious "duality" between $H^*(\cdot; G)$ and $H_*(\cdot; \mathbb{Z})$, not every axiomatic cohomology theory h^* has a corresponding axiomatic homology theory h_* (cf. Remark 48.19 at the end of ths lecture).

The Mayer-Vietoris sequence. One can use a diagram-chase as in Exercise 33.3 to derive from the axioms a Mayer-Vietoris sequence for any axiomatic cohomology theory, but for singular cohomology it also can be seen more directly. Indeed, suppose $A, B \subset X$ are subsets whose interiors cover X, let

$$a : A \cap B \hookrightarrow A, \qquad b : A \cap B \hookrightarrow B$$

denote the obvious continuous inclusions of spaces, and

$$\alpha : C_*(A;\mathbb{Z}) \hookrightarrow C_*(A;\mathbb{Z}) + C_*(B;\mathbb{Z}), \qquad \beta : C_*(B;\mathbb{Z}) \hookrightarrow C_*(A;\mathbb{Z}) + C_*(B;\mathbb{Z})$$

the obvious inclusions of subgroups of $C_*(X;\mathbb{Z})$. The Mayer-Vietoris sequence in singular homology was derived in Lecture 33 from a short exact sequence of chain complexes in the form

$$0 \longrightarrow C_*(A \cap B;\mathbb{Z}) \xrightarrow{(a_*,-b_*)} C_*(A;\mathbb{Z}) \oplus C_*(B;\mathbb{Z}) \xrightarrow{\alpha\oplus\beta} C_*(A;\mathbb{Z}) + C_*(B;\mathbb{Z}) \longrightarrow 0.$$

Applying to this the functor $\mathrm{Hom}(\cdot;G)$ and using the natural isomorphism

$$\mathrm{Hom}\left(C_*(A;\mathbb{Z}),G\right) \oplus \mathrm{Hom}\left(C_*(B;\mathbb{Z}),G\right) \xrightarrow{\cong} \mathrm{Hom}\left(C_*(A;\mathbb{Z}) \oplus C_*(B;\mathbb{Z}),G\right)$$
$$(\varphi,\psi) \mapsto \varphi \oplus \psi$$

transforms it into the sequence

$$0 \longleftarrow C^*(A \cap B;G) \xleftarrow{a^*\oplus(-b^*)} C^*(A;G) \oplus C^*(B;G) \xleftarrow{(\alpha^*,\beta^*)} C^*(A+B;G) \longleftarrow 0,$$

where we are abbreviating

$$C^*(A+B;G) := \mathrm{Hom}\left(C_*(A;\mathbb{Z}) + C_*(B;\mathbb{Z}),G\right).$$

The dual maps

$$a^* : C^*(A;G) \to C^*(A \cap B;G), \qquad b^* : C^*(B;G) \to C^*(A \cap B;G),$$
$$\alpha^* : C^*(A+B;G) \to C^*(A;G), \qquad \beta^* : C^*(A+B;G) \to C^*(B;G)$$

are all canonical restriction maps, e.g. α^* replaces a homomorphism $\varphi : C_*(A;\mathbb{Z}) + C_*(B;\mathbb{Z}) \to G$ with its restriction to the subgroup $C_*(A;\mathbb{Z})$. It is now an easy exercise to check that the dualized sequence is also exact. To make use of this, we need to identify the homology of the cochain complex $C^*(A+B;G)$ with something more familiar. By Lemma 29.21, the condition $X = \mathring{A} \cup \mathring{B}$ guarantees that the inclusion $j : C_*(A;\mathbb{Z}) + C_*(B;\mathbb{Z}) \hookrightarrow C_*(X;\mathbb{Z})$ is a chain homotopy equivalence, so by Proposition 47.2, the dual map

$$j^* : C^*(X;G) \to C^*(A+B;G)$$

is also a chain homotopy equivalence and therefore induces an isomorphism

$$H^*(X;G) \xrightarrow{\cong} H_*\left(C^*(A+B;G)\right).$$

Combining this with the usual diagram-chasing result gives:

Theorem 48.12 (Mayer-Vietoris sequence for cohomology). *If $A, B \subset X$ are subsets such that $X = \mathring{A} \cup \mathring{B}$ and*

$$i_A : A \cap B \hookrightarrow A, \quad i_B : A \cap B \hookrightarrow B, \quad j_A : A \hookrightarrow X, \quad j_B : B \hookrightarrow X,$$

denote the obvious inclusions, then there exist connecting homomorphisms $\delta^* : H^n(A \cap B;G) \to H^{n+1}(X;G)$ for every $n \in \mathbb{Z}$ such that the sequence

$$\ldots \longleftarrow H^{n+1}(X;G) \xleftarrow{\delta^*} H^n(A \cap B;G) \xleftarrow{i_A^*\oplus(-i_B^*)} H^n(A;G) \oplus H^n(B;G)$$
$$\xleftarrow{(j_A^*,j_B^*)} H^n(X;G) \xleftarrow{\delta^*} H^{n-1}(A \cap B;G) \longleftarrow \ldots$$

is exact, and this sequence is also natural with respect to maps $f : X \to X' = \mathring{A}' \cup \mathring{B}'$ satisfying $f(A) \subset A'$ and $f(B) \subset B'$. □

Exercise 48.13. Adapt the diagram-chasing arguments in Lecture 33 to show that every axiomatic cohomology theory h^* admits a Mayer-Vietoris sequence under the same hypotheses on $X = A \cup B$, and that it also works if h^* is replaced by $\widetilde{h}^*$.

Cellular cohomology. The cellular cohomology of a CW-pair (X, A) with coefficients in G is defined as the cohomology of the cellular chain complex, or equivalently,

$$H^*_{\mathrm{CW}}(X, A; G) := H_*\big(C^*_{\mathrm{CW}}(X, A; G)\big),$$

where we define the **cellular cochain complex**

$$C^*_{\mathrm{CW}}(X, A; G) := \mathrm{Hom}\left(C_*^{\mathrm{CW}}(X, A; \mathbb{Z}), G\right).$$

This gives a contravariant functor $H^*_{\mathrm{CW}} : \mathsf{CW}_{\mathrm{rel}} \to \mathsf{Ab}_{\mathbb{Z}}$ that is typically not very hard to compute. The coboundary map $\delta : C^n_{\mathrm{CW}}(X, A; G) \to C^{n+1}_{\mathrm{CW}}(X, A; G)$ can be expressed in terms of the same incidence numbers that describe the cellular boundary map: indeed, for each n-cell $e^n_\alpha \subset X$, define its dual cochain

$$\varphi^n_\alpha \in C^n_{\mathrm{CW}}(X; \mathbb{Z}), \qquad \varphi^n_\alpha(e^n_\beta) := \begin{cases} 1 & \text{if } \beta = \alpha, \\ 0 & \text{otherwise.} \end{cases}$$

These generators form a basis of $C^n_{\mathrm{CW}}(X; \mathbb{Z})$ if there are only finitely many n-cells, and any element of $C^n_{\mathrm{CW}}(X; G)$ can then similarly be described as a linear combination of the φ^n_α with coefficients in G. Formally, the latter remains true when there are infinitely-many n-cells, so long as we adopt a suitable interpretation of infinite sums of the form $\sum_{e^n_\alpha \subset X} g_\alpha \varphi^n_\alpha$, which have well-defined evaluations on $C_n^{\mathrm{CW}}(X; \mathbb{Z})$ even if infinitely-many of the coefficients $g_\alpha \in G$ are nonzero. A more precise way to say this is that dualizing direct sums gives direct products, so in light of the isomorphism $\mathrm{Hom}(\mathbb{Z}, G) \cong G$ defined by evaluation on the generator $1 \in \mathbb{Z}$, there are canonical isomorphisms

$$C^n_{\mathrm{CW}}(X; G) = \mathrm{Hom}(C_n^{\mathrm{CW}}(X; \mathbb{Z}), G) = \mathrm{Hom}\left(\bigoplus_{e^n_\alpha \subset X} \mathbb{Z}, G\right) \cong \prod_{e^n_\alpha \subset X} \mathrm{Hom}(\mathbb{Z}, G) \cong \prod_{e^n_\alpha \subset X} G,$$

and $\sum_{e^n_\alpha \subset X} g_\alpha \varphi^n_\alpha \in C^n_{\mathrm{CW}}(X; G)$ is to be interpreted as the element that corresponds canonically with $\{g_\alpha\} \in \prod_{e^n_\alpha \subset X} G$. With this understood, we shall abuse terminology by calling the elements φ^n_α a "basis" of $C^n_{\mathrm{CW}}(X; \mathbb{Z})$, and observe that the coboundary operator $\delta : C^n_{\mathrm{CW}}(X; G) \to C^{n+1}_{\mathrm{CW}}(X; G)$ is uniquely determined if we write down a formula for $\delta\varphi^n_\alpha \in C^{n+1}(X; \mathbb{Z})$ for each n-cell $e^n_\alpha \subset X$. For any $(n+1)$-cell $e^{n+1}_\beta \subset X$, we have

$$(\delta\varphi^n_\alpha)(e^{n+1}_\beta) = (-1)^{n+1}\varphi^n_\alpha(\partial e^{n+1}_\beta) = (-1)^{n+1} \sum_{e^n_\gamma \subset X} \varphi^n_\alpha\big([e^n_\gamma : e^{n+1}_\beta] e^n_\gamma\big) = (-1)^{n+1}[e^n_\alpha : e^{n+1}_\beta],$$

thus the required formula is

$$\delta\varphi^n_\alpha = (-1)^{n+1} \sum_{e^{n+1}_\beta \subset X} [e^n_\alpha : e^{n+1}_\beta] \varphi^{n+1}_\beta.$$

As with homology, cellular cohomology provides a powerful tool for computing arbirary axiomatic cohomology theories on spaces that have cell decompositions:

Theorem 48.14. *For any axiomatic cohomology theory h^* with coefficient group G and every CW-pair (X, A), there exists an isomorphism $H^*_{\mathrm{CW}}(X, A; G) \to h^*(X, A)$, which is natural in the sense that every cellular map $f : (X, A) \to (Y, B)$ gives rise to a commutative diagram*

$$\begin{array}{ccc} H^*_{\mathrm{CW}}(X, A; G) & \xrightarrow{\cong} & h^*(X, A) \\ {\scriptstyle f^*}\uparrow & & {\scriptstyle f^*}\uparrow \\ H^*_{\mathrm{CW}}(Y, B; G) & \xrightarrow{\cong} & h^*(Y, B) \end{array}$$

For finite-dimensional complexes, this theorem can be proved in a way that closely parallels the corresponding argument for cellular homology carried out in Lectures 37 and 38. One starts by deriving from the axioms (in particular the correspondence $h^*(X, A) \cong \widetilde{h}^*(X/A)$ for good pairs) a natural isomorphism

$$h^k(X^n, X^{n-1}) \cong \prod_{e^n_\alpha \subset X} h^k(\mathbb{D}^n, \partial\mathbb{D}^n)$$

for every k and n; here the direct product takes on the role formerly played by the direct sum, due to its appearance in the cohomological version of the additivity axiom. One then uses the long exact sequence of $(\mathbb{D}^n, \partial\mathbb{D}^n)$ in cohomology to prove that the right hand side is zero for all $k \neq n$ but (since $h^n(\mathbb{D}^n, \partial\mathbb{D}^n) \cong \widetilde{h}^{n-1}(S^{n-1}) \cong G$) is identical to the cellular n-cochain group $C^n_{\mathrm{CW}}(X; G) \cong \prod_{e^n_\alpha \subset X} G$ when $k = n$. Putting $h^n(X^n, X^{n-1})$ for each $n \geqslant 0$ in the role of $C^n_{\mathrm{CW}}(X; G)$, one then assembles the long exact sequences of (X^{n+1}, X^n) and (X^n, X^{n-1}) into the diagram

$$\begin{tikzcd}
0 \arrow[d, equal] & & & h^n(X^{n-1}) \arrow[r, equal] & 0 & \\
h^{n+1}(X^n) & h^{n+1}(X^{n+1}) \arrow[l] & h^{n+1}(X^{n+1}, X^n) \arrow[l, "j^*_{n+1}"'] & h^n(X^n) \arrow[l, "\delta^*_n"'] \arrow[u] & h^n(X^{n+1}) \arrow[l, "i^*_n"'] & h^n(X^{n+1}, X^n) \arrow[l] \arrow[d, equal] \\
& & & h^n(X^n, X^{n-1}) \arrow[u, "j^*_n"] \arrow[ul, "\gamma_n"] & & 0 \\
& & & h^{n-1}(X^{n-1}) \arrow[u, "\delta^*_{n-1}"] & &
\end{tikzcd}$$

in which the diagonal arrow defines maps γ_n so that the sequence

$$h^0(X^0) \xrightarrow{\gamma_0} h^1(X^1, X^0) \xrightarrow{\gamma_1} h^2(X^2, X^1) \longrightarrow \dots$$

becomes a cochain complex. One can check that γ_n is equivalent to the cellular coboundary map $C^n_{\mathrm{CW}}(X; G) \xrightarrow{\delta} C^{\mathrm{CW}}_{n+1}(X; G)$ under the natural isomorphisms $h^n(X^n, X^{n+1}) \cong C^n_{\mathrm{CW}}(X; G)$. The diagram then allows us to deduce that the map $i^*_n : h^n(X^{n+1}) \to h^n(X^n)$ is injective and j^*_n descends to an isomorphism

$$\ker \gamma_n / \operatorname{im} \gamma_{n-1} \xrightarrow{j^*_n} \operatorname{im} i^*_n \cong h^n(X^{n+1}) \cong h^n(X^{n+2}) \cong \dots,$$

thus giving an isomorphism $H^n_{\mathrm{CW}}(X; G) \cong h^n(X)$ if $X = X^N$ for some $N \in \mathbb{N}$ sufficiently large. To handle CW-pairs (X, A) with $A \neq \varnothing$, one carries out this same argument with $h^n(X^n, X^{n-1})$ replaced by $h^n(X^n \cup A, X^{n-1} \cup A)$, and the long exact sequence of (X^n, X^{n-1}) replaced by the sequence of the triple $(X^n \cup A, X^{n-1} \cup A, A)$.

Exercise 48.15. Work out the further details of the proof of Theorem 48.14 for finite-dimensional CW-pairs.

Some additional arguments are needed if the CW-pair (X, A) is infinite dimensional. In Lecture 39, we handled this for singular homology by viewing any CW-complex as the direct limit of its skeleta and establishing sufficient conditions for singular homology to behave continuously under direct limits. You may recall that this issue was a little bit subtle: some topological condition is required on a direct system of spaces in order for the singular chain complex functor to behave continuously under direct limits, but fortunately, all CW-complexes satisfy the condition (cf. Prop. 39.19). For cohomology, the situation is worse: since $H^*(\cdot;G)$ is contravariant, the cohomology groups of the skeleta $\{H^*(X^n;G)\}_{n\geqslant 0}$ define an *inverse* system of graded abelian groups, and we've seen in Lecture 46 that in contrast to direct systems, the algebraic functor taking chain complexes to their homology groups does not always behave continuously under inverse limits. This does not mean that the situation is hopeless, but it does make things more complicated: it means that one must introduce another "derived" functor (analogous to Tor) to account for the nonexactness of the functor that takes each inverse system to its inverse limit. This is done in [**Mil62**] in the more general context of axiomatic cohomology theories, but we will not discuss it any further here since we do not need that level of generality. For our purposes, it will suffice if we can prove the general case of the isomorphism $H^*_{\mathrm{CW}}(X,A;G) \cong h^*(X,A)$ when h^* is taken to be singular cohomology $H^*(\cdot;G)$, and in the next lecture, we will see a cheap way of deducing this from facts already proven about $H_*(X,A;\mathbb{Z})$, via a universal coefficient theorem.

Other cohomology theories. Finally, I would like to give brief sketches of two axiomatic cohomology theories other than singular cohomology. They will demonstrate in particular that the two properties of H^* we discussed at the end of the previous lecture,

$$H^0(X;G) \cong \prod_{\pi_0(X)} G \quad \text{and} \quad H^1(X;G) \cong \operatorname{Hom}(\pi_1(X),G),$$

do not follow from the axioms, but are distinctive to the singular theory.

We begin with Čech cohomology. Recall from Lecture 46 that for any open cover $(\mathfrak{U},\mathfrak{U}_A) \in \mathcal{O}(X,A)$ of a pair of spaces (X,A), one can define its *nerve* $\mathcal{N}(\mathfrak{U},\mathfrak{U}_A)$, which consists of a simplicial complex $\mathcal{N}(\mathfrak{U})$ and subcomplex $\mathcal{N}(\mathfrak{U}_A) \subset \mathcal{N}(\mathfrak{U})$. The *ordered simplicial cohomology* of the nerve with coefficients in G is defined in the obvious way as the cohomology of the ordered simplicial chain complex

$$H^*_o(\mathcal{N}(\mathfrak{U},\mathfrak{U}_A);G) := H^*\big(C^o_*(\mathfrak{U},\mathfrak{U}_A;\mathbb{Z});G\big).$$

Recall moreover that the set $\mathcal{O}(X,A)$ of all open coverings of (X,A) is a directed set with respect to refinement, and any refinement $(\mathfrak{U}',\mathfrak{U}'_A)$ of $(\mathfrak{U},\mathfrak{U}_A)$ gives rise to a simplicial map $F : \mathcal{N}(\mathfrak{U}',\mathfrak{U}'_A) \to \mathcal{N}(\mathfrak{U},\mathfrak{U}_A)$ whose induced chain map

$$F_* : C^o_*(\mathcal{N}(\mathfrak{U}',\mathfrak{U}'_A);\mathbb{Z}) \to C^o_*(\mathcal{N}(\mathfrak{U},\mathfrak{U}_A);\mathbb{Z})$$

is unique up to chain homotopy. It follows via Proposition 47.2 that dualizing this map produces a map of cochain complexes that is also unique up to chain homotopy, producing a canonically defined morphism

$$F^* : H^*_o(\mathcal{N}(\mathfrak{U},\mathfrak{U}_A);G) \to H^*_o(\mathcal{N}(\mathfrak{U}',\mathfrak{U}'_A);G).$$

Notice what has happened as a result of dualization: the collection of simplicial homology groups $\{H^o_*(\mathcal{N}(\mathfrak{U},\mathfrak{U}_A);G)\}_{(\mathfrak{U},\mathfrak{U}_A)\in\mathcal{O}(X,A)}$ in Lecture 46 was an inverse system, but the reversal of arrows now means that the corresponding cohomology groups

$$\{H^*_o(\mathcal{N}(\mathfrak{U},\mathfrak{U}_A);G)\}_{(\mathfrak{U},\mathfrak{U}_A)\in\mathcal{O}(X,A)}$$

form a **direct** system, and we define the **Čech cohomology** of (X,A) with coefficients in G to be the direct limit

$$\check{H}^*(X,A) = \check{H}^*(X,A;G) := \varinjlim \{H^*_o(\mathcal{N}(\mathfrak{U},\mathfrak{U}_A);G)\}_{(\mathfrak{U},\mathfrak{U}_A)\in\mathcal{O}(X,A)}.$$

There is a huge technical advantage in the fact that $\check{H}^*(X,A)$ is defined via a direct limit instead of an inverse limit: exactness of sequences is preserved under direct limits in the category of abelian groups (cf. Prop. 39.18), and one can use this to prove that unlike $\check{H}_*$, the cohomology $\check{H}^*$ satisfies the exactness axiom without any restrictions. It also satisfies all the other axioms:

THEOREM 48.16 (see [ES52, Spa95]). *For any abelian group G, the Čech cohomology $\check{H}^* = \check{H}^*(\cdot; G) : \mathsf{Top}_{\mathrm{rel}} \to \mathsf{Ab}_{\mathbb{Z}}$ is an axiomatic cohomology theory with coefficient group G.* ☐

It follows that $\check{H}^*(X;G)$ and $H^*(X;G)$ are isomorphic whenever X is a CW-complex. To find examples in which $\check{H}^*(X;G)$ and $H^*(X;G)$ differ, it suffices again to consider a space X that is connected but not path-connected. Recall from Lemma 46.17 that whenever $\mathfrak{U} \in \mathcal{O}(X)$ is an open covering of a connected space X, the nerve $\mathcal{N}(\mathfrak{U})$ is also connected, thus $H^0(\mathcal{N}(\mathfrak{U}); G) \cong G$. One can deduce from this that if X is connected, then $\check{H}^0(X;G) \cong G$, and the reduced Čech cohomology of X in degree zero vanishes. Exercise 48.7 then implies $\check{H}^1(SX;G) = 0$. But if X has more than one path-component, then $\widetilde{H}^0(X;G)$ and $H^1(SX;G)$ are both nontrivial; the latter is isomorphic to $\operatorname{Hom}(\pi_1(SX), G)$ since the suspension SX is always path-connected, thus SX is an example of a space for which $\check{H}^1(SX;G) \not\cong \operatorname{Hom}(\pi_1(SX), G)$.

The **Alexander-Spanier cohomology** is yet another theory that satisfies all of the Eilenberg-Steenrod axioms, but is based on a different idea of how to detect topological information. Let us describe the absolute version.

For integers $n \geqslant 0$ and a fixed choice of abelian coefficient group G, let $\bar{C}^n(X) = \bar{C}^n(X;G)$ denote the additive abelian group of equivalence classes of functions

$$\varphi : X^{n+1} = \underbrace{X \times \ldots \times X}_{n+1} \to G,$$

where we say $\varphi \sim \psi$ whenever φ and ψ are identical on some neighborhood of the **diagonal**

$$\Delta := \{(x, \ldots, x) \in X^{n+1} \mid x \in X\}.$$

The group operation on $\bar{C}^n(X)$ is defined via pointwise addition, so for two equivalence classes $[\varphi], [\psi] \in \bar{C}^n(X)$, $[\varphi] + [\psi] \in \bar{C}^n(X)$ is represented by the function $\varphi + \psi : X^{n+1} \to G$ defined by

$$(\varphi + \psi)(x_0, \ldots, x_n) := \varphi(x_0, \ldots, x_n) + \psi(x_0, \ldots, x_n).$$

You should take a moment to assure yourself that the equivalence class of $\varphi + \psi$ is independent of the choice of representatives $\varphi \in [\varphi]$ and $\psi \in [\psi]$. Note that since the group G is not assumed to have a topology, there is no continuity condition on the functions $X^{n+1} \to G$ representing elements of $\bar{C}^n(X)$. Instead, this group detects the topology of X via the notion of "neighborhoods of $\Delta \subset X^{n+1}$" that is used to define the equivalence relation.

To make the collection of groups $\bar{C}^n(X)$ for $n \geqslant 0$ into a cochain complex, we associate to each function $\varphi : X^{n+1} \to G$ the function $\delta\varphi : X^{n+2} \to G$ defined by

$$(\delta\varphi)(x_0, \ldots, x_{n+1}) := \sum_{k=0}^{n+1} (-1)^k \varphi(x_0, \ldots, x_{k-1}, x_{k+1}, \ldots, x_{n+1}).$$

This defines a homomorphism from the group of $(n+1)$-functions to the group of $(n+2)$-functions such that $\delta^2 = 0$, and it preserves the subgroup of functions that vanish near the diagonal, thus it descends to a coboundary homomorphism

$$\delta : \bar{C}^n(X) \to \bar{C}^{n+1}(X).$$

Extending this to all $n \in \mathbb{Z}$ by defining $\bar{C}^n(X) = 0$ for $n < 0$, we obtain a cochain complex $(\bar{C}^*(X), \delta)$, and its homology is the Alexander-Spanier cohomology of X with coefficients in G, denoted by

$$\bar{H}^*(X) = \bar{H}^*(X;G) := H_*\big(\bar{C}^*(X;G)\big).$$

It is not hard to give $\bar{H}^*$ the structure of a contravariant functor: given a continuous map $f : X \to Y$, one defines a chain map

$$f^* : \bar{C}^*(Y) \to \bar{C}^*(X) : \varphi \mapsto \varphi \circ (f \times \ldots \times f),$$

thus inducing homomorphisms $f^* : \bar{H}^*(Y) \to \bar{H}^*(X)$. With some more effort, one can also define relative groups $\bar{H}^*(X, A)$ and prove that $\bar{H}^*$ satisfies all of the Eilenberg-Steenrod axioms for a cohomology theory. A good exposition of the details can be found e.g. in [**Spa95**, §6.4–6.5].

It is instructive to unpack more explicitly the conditions that define cocycles and coboundaries in $\bar{C}^0(X)$ and $\bar{C}^1(X)$. Elements $\varphi \in \bar{C}^0(X)$ are simply functions $\varphi : X \to G$; here the equivalence relation is trivial since the diagonal in $X^1 = X$ is the whole space. Acting on such a function with δ gives the equivalence class of functions $X \times X \to G$ represented by

$$(\delta\varphi)(x, y) = \varphi(y) - \varphi(x),$$

hence $\delta\varphi = 0 \in \bar{C}^1(X)$ means that $\varphi(y) = \varphi(x)$ for all $(x, y) \in X \times X$ in some neighborhood of the diagonal. In other words, the 0-cocycles in $\bar{C}^*(X)$ are precisely the *locally constant* functions $X \to G$, i.e. those which are constant on the connected components of X, and $\bar{H}^0(X; G)$ is therefore naturally isomorphic to the group of locally constant functions $X \to G$, or equivalently, the direct product of copies of G over the set of connected components of X. This, of course, often matches the description of the singular cohomology group $H^0(X; G)$ that we gave in Theorem 47.6, but not always: $\bar{H}^0(X)$ and $H^0(X)$ differ on spaces whose connected components and path-components do not match. In this respect, $\bar{H}^0(X)$ matches the Čech cohomology $\check{H}^0(X)$ rather than $H^0(X)$.

Elements of $\bar{C}^1(X)$ are represented by functions $\varphi : X \times X \to G$, and such a function represents a coboundary if and only if there exists a function $\psi : X \to G$ such that

$$\varphi(x, y) = \psi(y) - \psi(x)$$

for all $x, y \in X$ sufficiently close together, i.e. for $(x, y) \in X \times X$ in some neighborhood of the diagonal. More generally, φ represents a 1-cocycle if and only if it satisfies

$$(\delta\varphi)(x, y, z) = \varphi(y, z) - \varphi(x, z) + \varphi(x, y) = 0,$$

or equivalently

$$\varphi(x, z) = \varphi(x, y) + \varphi(y, z),$$

for all triples of points $x, y, z \in X$ that are sufficiently close to each other. Looking at special cases with $z = x$, this relation forces $\varphi(x, y)$ to be an antisymmetric function of x and y in some neighborhood of the diagonal. I recommend thinking through the following exercise in order to gain some intuition on what this cocycle condition means and what $\bar{H}^1(X)$ actually measures.

EXERCISE 48.17. Show that for any path-connected space X and any choice of base point $p \in X$, there is a well-defined and injective homomorphism

$$\bar{H}^1(X; G) \to \operatorname{Hom}(\pi(X, p), G) : [\varphi] \mapsto \Psi_{[\varphi]}$$

such that for any representative cocycle $\varphi \in \bar{C}^1(X; G)$ and loop $\gamma : [0, 1] \to X$ from p to itself, one can choose a sufficiently fine partition $0 =: t_0 < t_1 < \ldots < t_{N-1} < t_N := 1$ of $[0, 1]$ to compute

$$\Psi_{[\varphi]}([\gamma]) = \sum_{j=1}^{N} \varphi(\gamma(t_j), \gamma(t_{j-1})).$$

Show moreover that this map is an isomorphism if X is S^1 or $\mathbb{R}$.

Both $\check{H}^*$ and $\bar{H}^*$ satisfy an "extra" axiom that singular cohomology does not, the so-called **continuity** axiom (cf. Theorem 46.23). Since both are contravariant functors, any inverse system of spaces $\{X_\alpha, \varphi_{\beta\alpha}\}$ gives rise to direct systems of cohomology groups, e.g. in the Alexander-Spanier theory, we obtain the system $\{\bar{H}^*(X_\alpha; G), \varphi^*_{\beta\alpha}\}$. It turns out that whenever the spaces X_α are all compact and Hausdorff, there is an isomorphism

$$\bar{H}^*(\varprojlim\{X_\alpha\}; G) \cong \varinjlim\{\bar{H}^*(X_\alpha; G)\},$$

and the same is true for Čech cohomology. It is not hard to find examples (e.g. involving spaces that are connected but not path-connected) for which this is not true in singular cohomology, cf. Exercise 46.24.

Remark 48.18. One can show that every compact Hausdorff space is an inverse limit of some inverse system of compact Hausdorff spaces homotopy equivalent to CW-complexes. It follows that up to isomorphism, there is only one cohomology theory on compact Hausdorff spaces that satisfies all of the Eilenberg-Steenrod axioms plus continuity. In particular, $\check{H}^*(X;G) \cong \bar{H}^*(X;G)$ whenever X is compact and Hausdorff, though both may be different from $H^*(X;G)$. (This result can be generalized beyond compact spaces using sheaf cohomology; details are carried out in [**Spa95**, Chapter 6].)

Remark 48.19. It is interesting to note that $\bar{C}^*(X;G)$ is not in any obvious way the dual complex of a chain complex, thus it is far from obvious at this stage what the definition of "Alexander-Spanier homology" might be. A corresponding homology theory was defined in an appendix of [**Spa48**], but its definition is much more complicated, requiring inverse limits, and as a result it suffers from the same drawbacks as Čech homology, i.e. it fails to satisfy the exactness axiom of Eilenberg-Steenrod.

49. Universal coefficients and the Ext functor (January 19, 2024)

The main goal of this lecture is to understand how $H^*(X;G)$ for an arbitrary choice of coefficient group G is determined in general by $H_*(X;\mathbb{Z})$, and in particular, under what circumstances the natural homomorphism

$$h : H^n(X;G) \to \mathrm{Hom}(H_n(X;\mathbb{Z}), G) : \varphi \mapsto \langle \varphi, \cdot \rangle \tag{49.1}$$

is an isomorphism. The answer to that question is the cohomological analogue of the universal coefficient theorem for homology (see Lecture 42), and just as in that setting, the question at hand is an algebraic one, whose complete answer requires the introduction of a derived functor defined in terms of projective resolutions. It will also be useful to frame the question somewhat more generally by fixing a commutative ring R and allowing G to be an R-module rather than just an abelian group, so that $H_*(X;R)$ and $H^*(X;G)$ have natural R-module structures, and (49.1) becomes an R-module homomorphism

$$H^n(X;G) \xrightarrow{h} \mathrm{Hom}_R(H_n(X;R), G). \tag{49.2}$$

This level of generality is especially useful when R is taken to be a field $\mathbb{K}$, so that the groups in question are all vector spaces over $\mathbb{K}$. Since you've already seen what happens with the Tor functor, it will probably not surprise you to learn that h in that setting is *always* an isomorphism, thus in particular, the vector space $H^n(X;\mathbb{K})$ is naturally isomorphic to the dual of the vector space $H_n(X;\mathbb{K})$.

Throughout this lecture, the commutative ring R and the R-module G are considered fixed pieces of data, and all (co-)chain complexes, tensor products and homomorphisms should be understood in the context of R-modules.

Preliminary remarks. Let's first discuss two nitpicky details that can cause confusion if we aren't careful.

First, there are *two* equivalent algebraic prescriptions for the construction of $H^*(X,A;G)$ for any pair of spaces (X,A) if G is an R-module. Both of them start with the chain complex of free abelian groups $C_* := C_*(X,A;\mathbb{Z})$, and then process this further via certain algebraic functors. In the standard presentation, one feeds C_* into the contravariant functor $\mathrm{Hom}(\cdot, G) : \mathsf{Chain} \to \mathsf{Cochain}$ and then applies the algebraic homology functor $H_* : \mathsf{Cochain} \to \mathsf{Ab}_{\mathbb{Z}}$. This defines $H^*(X,A;G)$

as a $\mathbb{Z}$-graded abelian group, but since G is an R-module, $H^*(X, A; G)$ also inherits a natural R-module structure, because there is a natural pairing of R with $\mathrm{Hom}(C, G)$ for any abelian group C, allowing us to view $\mathrm{Hom}(\cdot, G)$ as a functor $\mathsf{Chain} \to \mathsf{Cochain}^R$ and replace $H_* : \mathsf{Cochain} \to \mathsf{Ab}_{\mathbb{Z}}$ with $H_* : \mathsf{Cochain}^R \to \mathsf{Mod}^R_{\mathbb{Z}}$. A slightly different way to produce the same result is as follows: first feed C_* into $\otimes R$ in order produce $C_*(X, A; R)$, which has the extra structure of an R-module, and thus lives in the category Chain^R, then feed the latter successively into $\mathrm{Hom}_R(\cdot, G) : \mathsf{Chain}^R \to \mathsf{Cochain}^R$ and $H_* : \mathsf{Cochain}^R \to \mathsf{Mod}^R_{\mathbb{Z}}$. Exercise 47.5 shows that the two versions of the $\mathbb{Z}$-graded R-module $H^*(X, A; G)$ produced by these two procedures are canonically isomorphic, so for most purposes, it doesn't matter which one we use. The only real difference between the two constructions is that they lead to slightly different versions of the canonical evaluation pairing between cohomology and homology. Indeed, from the first perspective, where $H^*(X, A; G)$ is obtained by feeding $C_*(X, A; \mathbb{Z})$ into $\mathrm{Hom}(\cdot, G)$, the pairing $\langle\ ,\ \rangle$ becomes a homomorphism of abelian groups

$$\langle\ ,\ \rangle : H^*(X, A; G) \otimes H_*(X, A; \mathbb{Z}) \to G,$$

and the fact that $H^*(X, A; G)$ and G have R-module structures isn't really relevant since $H_*(X, A; \mathbb{Z})$ does not. From the second perspective, one gets $H^*(X, A; G)$ by feeding $C_*(X, A; R)$ into $\mathrm{Hom}_R(\cdot, G)$ and thus obtains an R-module homomorphism

$$\langle\ ,\ \rangle : H^*(X, A; G) \otimes_R H_*(X, A; R) \to G.$$

This second map is not quite the same thing as the first one, because a tensor product of R-modules is not generally the same set as the tensor product of the underlying abelian groups, but of course, the second perspective produces something more general than the first, since we can always recover the first version of the pairing by setting $R := \mathbb{Z}$. For this reason, the most useful perspective algebraically is to work with R-modules the entire time, and that is what we will do: the subscript R on $\otimes$ and Hom will thus be omitted henceforth from the notation whenever we have no need to emphasize the module structure, but one should always keep in mind that that structure is there, and must be respected.

The second detail concerns the sign $(-1)^{n+1}$ that we inserted into the definition of the dualized coboundary operator δ, which determines the functor $\mathrm{Hom}(\cdot, G) : \mathsf{Chain}^R \to \mathsf{Cochain}^R$ (see (47.2) and Remark 47.3). There are contexts in which the inclusion of that sign makes life less inconvenient (I will not go so far as to say "more convenient") than it otherwise would be, but the present lecture is not one of them. We are therefore going to pretend in this lecture that we hadn't included that sign in the definition, and that the way to produce a cochain complex $(\mathrm{Hom}(C_*, G), \delta)$ from a chain complex (C_*, ∂) is simply to dualize the boundary operator and define

$$\delta := \partial^* : \mathrm{Hom}(C_n, G) \to \mathrm{Hom}(C_{n+1}, G) : \varphi \mapsto \varphi \circ \partial.$$

This changes the definition of the functor $\mathrm{Hom}(\cdot, G) : \mathsf{Chain}^R \to \mathsf{Cochain}^R$, but the reason we can get away with it is that it does *not* change either of the functors $H^*(\cdot; G) : \mathsf{Chain}^R \to \mathsf{Mod}^R_{\mathbb{Z}}$ or $H^*(\cdot; G) : \mathsf{Top}_{\mathrm{rel}} \to \mathsf{Mod}^R_{\mathbb{Z}}$, so our cohomology groups are exactly the same thing that they were before, and have the same structure. In fact, the definition of the functor $\mathrm{Hom}(\cdot, G) : \mathsf{Chain}^R \to \mathsf{Cochain}^R$ does not play an important role in the universal coefficient theorem; what does play a crucial role is its simpler cousin $\mathrm{Hom}(\cdot, G) : \mathsf{Mod}^R \to \mathsf{Mod}^R$, and we have not changed that.

Review of the covariant case. Let's recall briefly how the universal coefficient theorem for homology was proved, but frame the discussion a bit more generally. The given data is a chain complex C_* of free R-modules, where R is assumed to be a principal ideal domain, plus a covariant functor

$$\mathcal{F} : \mathsf{Mod}^R \to \mathsf{Mod}^R$$

that is additive and right-exact. In Lecture 42, $\mathcal{F}$ was assumed to be $\otimes G$ for a fixed R-module G, but the main argument does not require that information, so we will not assume it. Writing $C_n \xrightarrow{\partial} C_{n-1}$ for the boundary operator on C_*, the condition that $\mathcal{F}$ is additive implies that we can feed C_* into it and obtain another chain complex $\mathcal{F}(C_*)$, given by

$$\ldots \longrightarrow \mathcal{F}(C_2) \xrightarrow{\mathcal{F}(\partial)} \mathcal{F}(C_1) \xrightarrow{\mathcal{F}(\partial)} \mathcal{F}(C_0) \longrightarrow \ldots.$$

The natural question then becomes: what relation is there between the R-modules $H_n(\mathcal{F}(C_*))$ and $\mathcal{F}(H_n(C_*))$ for each n? Note that without having a more concrete description of the functor $\mathcal{F}$, it is no longer so straightforward to write down a canonical map $\mathcal{F}(H_n(C_*)) \to H_n(\mathcal{F}(C_*))$, which would generalize the map $H_n(C_*) \otimes G \to H_n(C_* \otimes G) : [c] \otimes g \mapsto [c \otimes g]$ that we have when $\mathcal{F} = \otimes G$. We will see however that a natural map $\mathcal{F}(H_n(C_*)) \to H_n(\mathcal{F}(C_*))$ does exist, and as long as $\mathcal{F}$ is right-exact, it will also be injective, with surjectivity hinging on the vanishing of a derived functor that measures the failure of $\mathcal{F}$ to be left-exact.

Let B_n and Z_n denote the submodules of n-boundaries and n-cycles respectively in C_n for each $n \in \mathbb{Z}$, with inclusion maps

$$B_n \overset{i_n}{\hookrightarrow} Z_n \overset{j_n}{\hookrightarrow} C_n,$$

and recall that $B_* := \bigoplus_n B_n$ and $Z_* := \bigoplus_n Z_n$ can each be regarded as chain complexes with trivial boundary operators, so that the inclusions

$$B_* \overset{i}{\hookrightarrow} Z_* \overset{j}{\hookrightarrow} C_*$$

are chain maps. The boundary operator $C_n \xrightarrow{\partial_n} C_{n-1}$ has image $B_{n-1} \subset C_{n-1}$ by definition, so let us denote by

$$C_* \xrightarrow{\partial'} B_{*-1}$$

the resulting surjective chain map, and write its restriction to the degree n level as

$$C_n \xrightarrow{\partial'_n} B_{n-1}.$$

This notation may seem pedantic at first, because $C_n \xrightarrow{\partial_n} C_{n-1}$ and $C_n \xrightarrow{\partial'_n} B_{n-1}$ are the same map, just with different understandings of what their targets are. But it will be useful to have this notational distinction when we feed them into the functor $\mathcal{F}$, because $\mathcal{F}(B_{n-1})$ cannot necessarily be understood as a submodule of $\mathcal{F}(C_{n-1})$: indeed, applying $\mathcal{F}$ to the inclusion $j_{n-1}i_{n-1} : B_{n-1} \hookrightarrow C_{n-1}$ will produce a homomorphism

$$\mathcal{F}(B_{n-1}) \xrightarrow{\mathcal{F}(j_{n-1}i_{n-1})} \mathcal{F}(C_{n-1})$$

that might not be injective if $\mathcal{F}$ is not left-exact! What we can say instead is the following: the obvious commutative diagram

$$\begin{tikzcd} C_n \arrow[r, "\partial'_n"] \arrow[d, "\partial_n"'] & B_{n-1} \arrow[dl, "j_{n-1}i_{n-1}"] \\ C_{n-1} & \end{tikzcd}$$

remains commutative after applying $\mathcal{F}$, thus producing

$$\mathcal{F}(\partial_n) = \mathcal{F}(j_{n-1}i_{n-1})\mathcal{F}(\partial'_n) = \mathcal{F}(j_{n-1})\mathcal{F}(i_{n-1})\mathcal{F}(\partial'_n) \tag{49.3}$$

as the essential relation between $\mathcal{F}(C_n) \xrightarrow{\mathcal{F}(\partial_n)} \mathcal{F}(C_{n-1})$ and $\mathcal{F}(C_n) \xrightarrow{\mathcal{F}(\partial'_n)} \mathcal{F}(B_{n-1})$.

With that out of the way, the first step in the proof of the universal coefficient theorem for homology was to observe that

$$0 \longrightarrow Z_* \xrightarrow{j} C_* \xrightarrow{\partial'} B_{*-1} \longrightarrow 0$$

is a short exact sequence of chain maps, and it splits since B_{*-1} is free, being a submodule of a free module over a principal ideal domain. Since $\mathcal{F}$ is an additive functor,

$$0 \longrightarrow \mathcal{F}(Z_*) \xrightarrow{\mathcal{F}(j)} \mathcal{F}(C_*) \xrightarrow{\mathcal{F}(\partial')} \mathcal{F}(B_{*-1}) \longrightarrow 0$$

is then also a short exact sequence of chain maps, and the boundary operators on the chain complexes $\mathcal{F}(Z_*)$ and $\mathcal{F}(B_{*-1})$ are still trivial, so the resulting long exact sequence of homologies takes the form

$$\ldots \longrightarrow \mathcal{F}(B_n) \xrightarrow{\Phi_n} \mathcal{F}(Z_n) \xrightarrow{\mathcal{F}(j)_*} H_n(\mathcal{F}(C_*)) \xrightarrow{\mathcal{F}(\partial')_*} \mathcal{F}(B_{n-1}) \xrightarrow{\Phi_{n-1}} \mathcal{F}(Z_{n-1}) \longrightarrow \ldots,$$

where Φ_n and Φ_{n-1} are connecting homomorphisms. We brushed this detail under the rug in Lecture 42, but let's take a closer look at what the map $\Phi_n : \mathcal{F}(B_n) \to \mathcal{F}(Z_n)$ actually is. If you recall how connecting homomorphisms were constructed in Proposition 28.18, the diagram that we need to chase is:

$$\begin{tikzcd}
 & \vdots \arrow[d] & \vdots \arrow[d] & \vdots \arrow[d] & \\
0 \arrow[r] & \mathcal{F}(Z_{n+1}) \arrow[r, "\mathcal{F}(j_{n+1})"] \arrow[d, "0"] & \mathcal{F}(C_{n+1}) \arrow[r, "\mathcal{F}(\partial'_{n+1})"] \arrow[d, "\mathcal{F}(\partial_{n+1})"] & \mathcal{F}(B_n) \arrow[r] \arrow[d, "0"] & 0 \\
0 \arrow[r] & \mathcal{F}(Z_n) \arrow[r, "\mathcal{F}(j_n)"] \arrow[d] & \mathcal{F}(C_n) \arrow[r, "\mathcal{F}(\partial'_n)"] \arrow[d, "\mathcal{F}(\partial_n)"] & \mathcal{F}(B_{n-1}) \arrow[r] \arrow[d] & 0 \\
 & \vdots & \vdots & \vdots &
\end{tikzcd}$$

Every element $b \in \mathcal{F}(B_n)$ is a cycle (of degree $n+1$) in the chain complex $\mathcal{F}(B_{*-1})$ and is also $\mathcal{F}(\partial'_{n+1})c$ for some $c \in \mathcal{F}(C_{n+1})$, and $\Phi_n(b) \in \mathcal{F}(Z_n)$ will then be determined by the condition that for a suitable choice of $c \in \mathcal{F}(C_{n+1})$,

$$\mathcal{F}(j_n)\Phi_n(b) = \mathcal{F}(\partial_{n+1})c.$$

Now notice: according to (49.3), $\mathcal{F}(\partial_{n+1})c = \mathcal{F}(j_n)\mathcal{F}(i_n)\mathcal{F}(\partial'_{n+1})c = \mathcal{F}(j_n)\mathcal{F}(i_n)b$, thus the required condition is satisfied by $\Phi_n(b) := \mathcal{F}(i_n)b$, and our long exact sequence therefore becomes

$$\ldots \longrightarrow \mathcal{F}(B_n) \xrightarrow{\mathcal{F}(i_n)} \mathcal{F}(Z_n) \xrightarrow{\mathcal{F}(j)_*} H_n(\mathcal{F}(C_*)) \xrightarrow{\mathcal{F}(\partial')_*} \mathcal{F}(B_{n-1}) \xrightarrow{\mathcal{F}(i_{n-1})} \mathcal{F}(Z_{n-1}) \longrightarrow \ldots.$$

Letting $\mathcal{F}(j)_*$ descend to the quotient by its kernel and rewriting the image of $\mathcal{F}(\partial')_*$ as a kernel gives the short exact sequence

$$0 \longrightarrow \operatorname{coker} \mathcal{F}(i_n) \xrightarrow{\mathcal{F}(j)_*} H_n(\mathcal{F}(C_*)) \xrightarrow{\mathcal{F}(\partial')_*} \ker \mathcal{F}(i_{n-1}) \longrightarrow 0.$$

So far in this argument, $\mathcal{F}$ was only required to be an additive functor; the assumption that it is right-exact only becomes important when we want to understand more precisely what $\operatorname{coker} \mathcal{F}(i_n)$ and $\ker \mathcal{F}(i_{n-1})$ are. This understanding comes from the short exact sequence

$$0 \to B_n \xrightarrow{i_n} Z_n \xrightarrow{q_n} H_n(C_*) \to 0,$$

in which q_n denotes the quotient projection. We used this in Lecture 42 by observing that since B_n and Z_n are free, this sequence can be interpreted as a projective resolution of $H_n(C_*)$, whose homology after acting on it with $\mathcal{F}$ thus produces the left derived functors of $\mathcal{F}$. Actually, what we need to know can be deduced without mentioning projective resolutions, so long as we understand the *formal* properties of the left derived functors $L_k\mathcal{F}$, namely: $L_k\mathcal{F}(A) = 0$ for $k > 0$ whenever A is projective (Proposition 43.1), and every short exact sequence $0 \to A \to B \to C \to 0$ naturally determines a long exact sequence (Proposition 43.10 and Equation 43.4)

$$\ldots \to L_2\mathcal{F}(C) \to L_1\mathcal{F}(A) \to L_1\mathcal{F}(B) \to L_1\mathcal{F}(C) \to \mathcal{F}(A) \to \mathcal{F}(B) \to \mathcal{F}(C) \to 0.$$

Indeed, feeding the specific short exact sequence above into $\mathcal{F}$ and remembering that Z_n is free and therefore projective, we now obtain an exact sequence of the form

$$0 \longrightarrow L_1\mathcal{F}(H_n(C_*)) \longrightarrow \mathcal{F}(B_n) \xrightarrow{\mathcal{F}(i_n)} \mathcal{F}(Z_n) \xrightarrow{\mathcal{F}(q_n)} \mathcal{F}(H_n(C_*)) \longrightarrow 0,$$

giving rise to natural isomorphisms

$$L_1\mathcal{F}(H_n(C_*)) \xrightarrow{\cong} \ker \mathcal{F}(i_n), \qquad \text{and} \qquad \operatorname{coker} \mathcal{F}(i_n) \xrightarrow{\cong} \mathcal{F}(H_n(C_*)).$$

Using these to replace the first and last nontrivial terms in our previous short exact sequence gives

$$0 \longrightarrow \mathcal{F}(H_n(C_*)) \longrightarrow H_n(\mathcal{F}(C_*)) \longrightarrow L_1\mathcal{F}(H_{n-1}(C_*)) \longrightarrow 0, \tag{49.4}$$

which is precisely the sequence in the universal coefficient theorem if we take $\mathcal{F}$ to be $\otimes G$ and thus write $\mathcal{F}(H_n(C_*)) = H_n(C_*) \otimes G$ and $L_1\mathcal{F} = \operatorname{Tor}(\cdot, G)$.

One drawback of the argument leading to (49.4) is that while we can now see the existence of a natural map

$$\mathcal{F}(H_n(C_*)) \xrightarrow{h} H_n(\mathcal{F}(C_*))$$

and prove that it is injective, it seems less than obvious what exactly the map h is. At this level of generality, one cannot quite write down an explicit formula for it, but the argument above leads to the following characterization: h is the unique map such that

$$h(x) = [\mathcal{F}(j_n)z] \in H_n(\mathcal{F}(C_*)) \qquad \text{for every } z \in \mathcal{F}(Z_n) \text{ with } \mathcal{F}(q_n) = x \in \mathcal{F}(H_n(C_*)). \tag{49.5}$$

Indeed, the right-exactness of $\mathcal{F}$ implies that $\mathcal{F}(q_n) : \mathcal{F}(Z_n) \to \mathcal{F}(H_n(C_*))$ is surjective, so one can choose a suitable $z \in \mathcal{F}(Z_n)$ for any given $x \in \mathcal{F}(H_n(C_*))$, and its equivalence class in the quotient $\mathcal{F}(Z_n)/\operatorname{im}\mathcal{F}(i_n) = \operatorname{coker}\mathcal{F}(i_n)$ is unique, reflecting the isomorphism $\operatorname{coker}\mathcal{F}(i_n) \cong \mathcal{F}(H_n(C_*))$. The map $\operatorname{coker}\mathcal{F}(i_n) \to H_n(\mathcal{F}(C_*))$ in the exact sequence then arises by letting the chain map $\mathcal{F}(j) : \mathcal{F}(Z_*) \to \mathcal{F}(C_*)$ descend to homology as the map $\mathcal{F}(Z_n) \to H_n(\mathcal{F}(C_*)) : z \mapsto [\mathcal{F}(j_n)z]$, and then letting the latter descend to the quotient of $\mathcal{F}(Z_n)$ by $\operatorname{im}\mathcal{F}(i_n)$. The result is precisely the characterization in (49.5). For the special case $\mathcal{F} = \otimes G$, one easily deduces from this that h is the usual canonical map $H_n(C_*) \otimes G \to H_n(C_* \otimes G)$.

Now a final detail before we do something new: the sequence (49.4) splits, though of course, not naturally. One sees this by choosing left-inverses $p_n : C_n \to Z_n$ for the inclusions $j_n : Z_n \hookrightarrow C_n$ for every $n \in \mathbb{Z}$, and using them to write down a left-inverse for the injection $h : \mathcal{F}(H_n(C_*)) \to H_n(\mathcal{F}(C_*))$. Indeed, the existence of left-inverses $p_n : C_n \to Z_n$ is guaranteed because the exact sequence $0 \to Z_n \hookrightarrow C_n \to B_{n-1} \to 0$ splits, due to the fact that B_{n-1} is free. Put these left-inverses together for all n to define a left-inverse of the chain map $j : Z_* \to C_*$, denoted by $p : C_* \to Z_*$. Unfortunately, p is not typically a chain map, for fairly obvious reasons: the boundary operator on Z_* is trivial, but $p_{n-1}\partial_n : C_n \to Z_{n-1}$ sends C_n onto B_{n-1} and then simply includes it into Z_{n-1}, giving a nontrivial map. But if we compose p_n with the quotient projection $q_n : Z_n \to H_n(C_*)$ and regard $H_*(C_*)$ as a chain complex with trivial boundary operator, then the composition $q_{n-1}p_{n-1}\partial_n$ vanishes, giving rise to a chain map

$$C_* \xrightarrow{qp} H_*(C_*),$$

which at degree n is the composition $q_n \circ p_n : C_n \to H_n(C_*)$. Feeding this into $\mathcal{F}$ and regarding the $\mathbb{Z}$-graded module $\mathcal{F}(H_*(C_*))$ similarly as a chain complex with trivial boundary operators, we obtain a chain map

$$\mathcal{F}(C_*) \xrightarrow{\mathcal{F}(qp)} \mathcal{F}(H_*(C_*)),$$

and the map that this induces on homology at degree n takes the form

$$\mathcal{F}(qp)_* : H_n(\mathcal{F}(C_*)) \to \mathcal{F}(H_n(C_*)) : [y] \mapsto \mathcal{F}(q_n p_n) y.$$

We claim that this map is a left-inverse of $h : \mathcal{F}(H_n(C_*)) \to H_n(\mathcal{F}(C_*))$. Indeed, for any $x \in \mathcal{F}(H_n(C_*))$, choosing $z \in \mathcal{F}(Z_n)$ with $\mathcal{F}(q_n)z = x$ and applying (49.5), we find

$$\mathcal{F}(qp)_* h(x) = \mathcal{F}(qp)_* ([\mathcal{F}(j_n)z]) = \mathcal{F}(q_n p_n)\mathcal{F}(j_n)z = \mathcal{F}(q_n)\mathcal{F}(p_n j_n)z = \mathcal{F}(q_n)z = x,$$

due to the fact that p_n is a left-inverse of j_n.

Contravariant and left-exact. Before we can replace $\mathcal{F}$ in the discussion above with the contravariant functor $\mathrm{Hom}(\cdot, G)$, reverse all the arrows and prove something useful, we need to be clear on what happens when a contravariant functor $\mathcal{F} : \mathsf{Mod}^R \to \mathsf{Mod}^R$ fails to be exact. Let's start with the definitions.

The notion of an **additive** functor $\mathcal{F} : \mathsf{Mod}^R \to \mathsf{Mod}^R$ admits a straightforward adaptation to the contravariant case: the only difference is that for each pair of R-modules A, B, $\mathcal{F}$ defines a group homomorphism from $\mathrm{Hom}(A, B)$ to $\mathrm{Hom}(\mathcal{F}(B), \mathcal{F}(A))$ instead of $\mathrm{Hom}(\mathcal{F}(A), \mathcal{F}(B))$. Feeding a chain complex A_* with boundary maps $A_n \xrightarrow{f_n} A_{n-1}$ into $\mathcal{F}$ then produces a cochain complex $\mathcal{F}(A_*)$ with coboundary maps $\mathcal{F}(A_n) \xrightarrow{\mathcal{F}(f_{n+1})} \mathcal{F}(A_{n+1})$, and similarly, $\mathcal{F}$ transforms cochain complexes into chain complexes. One also obtains natural isomorphisms $\mathcal{F}(A \oplus B) \cong \mathcal{F}(A) \oplus \mathcal{F}(B)$ whenever $\mathcal{F}$ is additive, but here there is an important difference to be aware of. Recall from Lecture 42: the reason why additive covariant functors preserve direct sums is that one can characterize direct sums uniquely up to isomorphism in terms of the natural inclusion and projection maps

$$\begin{tikzcd}
A \arrow[dr, "i_A"] & & A \\
& A \oplus B \arrow[ur, "\pi_A"] \arrow[dr, "\pi_B"] & \\
B \arrow[ur, "i_B"] & & B
\end{tikzcd}$$

which satisfy the relations

$$\pi_A i_A = \mathbb{1}_A, \quad \pi_B i_B = \mathbb{1}_B, \quad \pi_A i_B = 0, \quad \pi_B i_A = 0, \quad i_A \pi_A + i_B \pi_B = \mathbb{1}_C$$

if we abbreviate $C := A \oplus B$. Plugging these maps into an additive *contravariant* functor $\mathcal{F}$ reverses both the arrows and the order of composition, so we obtain four homomorphisms

$$\begin{tikzcd}
\mathcal{F}(A) \arrow[dr, "\mathcal{F}(\pi_A)"] & & \mathcal{F}(A) \\
& \mathcal{F}(A \oplus B) \arrow[ur, "\mathcal{F}(i_A)"] \arrow[dr, "\mathcal{F}(i_B)"] & \\
\mathcal{F}(B) \arrow[ur, "\mathcal{F}(\pi_B)"] & & \mathcal{F}(B)
\end{tikzcd}$$

that satisfy the relations

$$\mathcal{F}(i_A)\mathcal{F}(\pi_A) = \mathbb{1}_{\mathcal{F}(A)}, \qquad \mathcal{F}(i_B)\mathcal{F}(\pi_B) = \mathbb{1}_{\mathcal{F}(B)},$$
$$\mathcal{F}(i_B)\mathcal{F}(\pi_A) = 0, \qquad \mathcal{F}(i_A)\mathcal{F}(\pi_B) = 0,$$
$$\mathcal{F}(\pi_A)\mathcal{F}(i_A) + \mathcal{F}(\pi_B)\mathcal{F}(i_B) = \mathbb{1}_{\mathcal{F}(C)}$$

for $C := \mathcal{F}(A \oplus B)$. These relations thus determine a natural isomorphism

$$\mathcal{F}(A \oplus B) \cong \mathcal{F}(A) \oplus \mathcal{F}(B),$$

but the roles of the inclusions and projections have been reversed: the isomorphism identifies $\mathcal{F}(\pi_A) : \mathcal{F}(A) \to \mathcal{F}(A \oplus B)$ with the inclusion $\mathcal{F}(A) \to \mathcal{F}(A) \oplus \mathcal{F}(B)$, $\mathcal{F}(i_A) : \mathcal{F}(A \oplus B) \to \mathcal{F}(A)$ with the projection $\mathcal{F}(A) \oplus \mathcal{F}(B) \to \mathcal{F}(A)$, and so forth. If you think about it, this reversal is exactly what we need, because it means that applying $\mathcal{F}$ to the split exact sequence $0 \longrightarrow A \xrightarrow{i_A} A \oplus B \xrightarrow{\pi_B} B \longrightarrow 0$ produces a sequence

$$0 \longrightarrow \mathcal{F}(B) \xrightarrow{\mathcal{F}(\pi_B)} \mathcal{F}(A \oplus B) \xrightarrow{\mathcal{F}(i_A)} \mathcal{F}(A) \longrightarrow 0$$

that is naturally isomorphic to the obvious split exact sequence $0 \to \mathcal{F}(B) \to \mathcal{F}(A) \oplus \mathcal{F}(B) \to \mathcal{F}(A) \to 0$. This proves:

PROPOSITION 49.1. *For any additive contravariant functor $\mathcal{F}$ and any split exact sequence $0 \to A \xrightarrow{f} B \xrightarrow{g} C \to 0$, the sequence*

$$0 \longrightarrow \mathcal{F}(C) \xrightarrow{\mathcal{F}(g)} \mathcal{F}(B) \xrightarrow{\mathcal{F}(f)} \mathcal{F}(A) \longrightarrow 0$$

is also split exact. □

An exact sequence of the form $A \to B \to C \to 0$, as appears in the definition of right-exactness for covariant functors, gets transformed by an additive contravariant functor $\mathcal{F}$ into a chain complex of the form $0 \to \mathcal{F}(C) \to \mathcal{F}(B) \to \mathcal{F}(A)$, so we call $\mathcal{F}$ **left-exact** if the latter is always guaranteed to be an exact sequence. Similarly, $\mathcal{F}$ is **right-exact** if it transforms all exact sequences of the form $0 \to A \to B \to C$ into exact sequences $\mathcal{F}(C) \to \mathcal{F}(B) \to \mathcal{F}(A) \to 0$, and $\mathcal{F}$ is simply **exact** if it is both left- and right-exact. One can adapt Exercise 42.13 to show that in the contravariant case just as in the covariant case, being exact means that $\mathcal{F}$ transforms *all* exact sequences (of arbitrary lengths) into exact sequences.

The main example of interest in this lecture turns out to be left-exact, so let's be more explicit about what that means when $\mathcal{F}$ is contravariant. The exactness of a sequence $A \xrightarrow{f} B \xrightarrow{g} C \to 0$ means firstly that g is surjective, and secondly that $\operatorname{im} f = \ker g$; both conditions together are equivalent to the statement that the map $g : B \to C$ descends to an isomorphism

$$\operatorname{coker} f = B/\operatorname{im}(f) \xrightarrow{g} C.$$

Exactness for $0 \to \mathcal{F}(C) \xrightarrow{\mathcal{F}(g)} \mathcal{F}(B) \xrightarrow{\mathcal{F}(f)} \mathcal{F}(A)$ will then mean two conditions that are essentially dual to these: first that $\mathcal{F}(g) : \mathcal{F}(C) \to \mathcal{F}(B)$ is injective, and second that its image is the kernel of $\mathcal{F}(f)$, meaning in total that we have an isomorphism

$$\mathcal{F}(C) \xrightarrow{\mathcal{F}(g)} \ker \mathcal{F}(f) \subset \mathcal{F}(B).$$

If this always holds, then $\mathcal{F}$ will be exact if and only if it satisfies one additional condition: for every injective homomorphism $f : A \to B$, the homomorphism $\mathcal{F}(f) : \mathcal{F}(B) \to \mathcal{F}(A)$ is surjective.

We already know how to measure the non-exactness of a covariant right-exact functor $\mathcal{F}$ in terms of the left derived functors $L_n\mathcal{F}$ for $n \geqslant 1$. If $\mathcal{F}$ is instead contravariant and left-exact, then more-or-less the same prescription works: one only has to reverse all the arrows. Indeed, fix a choice of projective resolution $A_* \xrightarrow{\alpha} A$ for each R-module A. Feeding the chain complex A_* into

$\mathcal{F}$ then gives a cochain complex $\mathcal{F}(A_*)$, whose homologies[79] we define to be the **right derived functors**

$$R_n\mathcal{F}(A) := H_n(\mathcal{F}(A_*)).$$

These are contravariant functors due to Proposition 42.24, which associates to any homorphism $\varphi : A \to B$ a chain map $\varphi_* : A_* \to B_*$, giving rise to a chain map $\mathcal{F}(\varphi_*) : \mathcal{F}(B_*) \to \mathcal{F}(A_*)$ that induces maps $R_n\mathcal{F}(B) \to R_n\mathcal{F}(A)$ for each $n \geqslant 0$. The chain map $\varphi_* : A_* \to B_*$ is of course not unique, but any other such map $\psi_* : A_* \to B_*$ is related to it by a chain homotopy $h_* : A_* \to B_{*+1}$, which can similarly be fed into $\mathcal{F}$ to produce a chain homotopy $\mathcal{F}(h_*) : \mathcal{F}(B_*) \to \mathcal{F}(A_{*-1})$ between $\mathcal{F}(\varphi_*)$ and $\mathcal{F}(\psi_*)$, showing that the induced map $R_n\mathcal{F}(B) \to R_n\mathcal{F}(A)$ is independent of choices. Applying the same argument to $\mathbb{1} : A \to A$ with two different choices of projective resolution gives canonical isomorphisms between the two versions of $R_n\mathcal{F}(A)$ defined via these choices.

The left-exactness of $\mathcal{F}$ implies that the sequence

$$0 \longrightarrow \mathcal{F}(A) \xrightarrow{\mathcal{F}(\alpha)} \mathcal{F}(A_0) \xrightarrow{\mathcal{F}(\alpha_1)} \mathcal{F}(A_1) \xrightarrow{\mathcal{F}(\alpha_2)} \dots$$

is exact at the first two nontrivial terms, meaning that $\mathcal{F}(\alpha)$ defines an isomorphism of $\mathcal{F}(A)$ onto $\ker \mathcal{F}(\alpha_1) \subset \mathcal{F}(A_0)$. Since $\mathcal{F}(A_0)$ is the first nontrivial term in the cochain complex $\mathcal{F}(A_*)$, this kernel is just the zeroth homology of that complex, and we therefore have a natural isomorphism

$$\mathcal{F}(A) \xrightarrow{\mathcal{F}(\alpha)} R_0\mathcal{F}(A).$$

The ability to make intelligent choices of projective resolution in certain circumstances now gives us the following analogues of the results stated for covariant right-exact functors in Lecture 43:

Proposition 49.2. *The right derived functors $R_n\mathcal{F} : \mathsf{Mod}^R \to \mathsf{Mod}^R$ associated to a contravariant left-exact functor $\mathcal{F} : \mathsf{Mod}^R \to \mathsf{Mod}^R$ have the following properties for every R-module A:*

(1) There is a natural isomorphism $R_0\mathcal{F}(A) \cong \mathcal{F}(A)$.
(2) If $\mathcal{F}$ is exact or A is projective, then $R_n\mathcal{F}(A) = 0$ for every $n \geqslant 1$.
(3) If R is a principal ideal domain, then $R_n\mathcal{F}(A) = 0$ for every $n \geqslant 2$.
(4) There are natural isomorphisms $R_n\mathcal{F}(A \oplus B) \cong R_n\mathcal{F}(A) \oplus R_n\mathcal{F}(B)$ for every $n \geqslant 0$ and pair of R-modules A, B.
(5) If $k \in \mathbb{N}$ has the property that $kx \neq 0$ for all nonzero $x \in R$, then $R_n\mathcal{F}(R/kR) = 0$ for all $n \geqslant 2$, and

$$R_1\mathcal{F}(R/kR) \cong \operatorname{coker}\left(\mathcal{F}(R) \to \mathcal{F}(R) : x \mapsto kx\right) = \mathcal{F}(R)\Big/k\mathcal{F}(R).$$

□

Most importantly, of course, there are long exact sequences of right derived functors. One obtains them from the horseshoe lemma (Proposition 43.10), just as in the covariant case: the only

[79] I am trying to be consistent in the way that I distinguish between the words "homology" and "cohomology" as algebraic notions, but it seems to cause confusion sometimes, so let's clarify. The algebraic homology functor can be applied to either chain complexes or cochain complexes; it is a covariant functor in both cases, and I denote it in both cases by H_*, not H^*. Given a chain complex C_*, one can apply a contravariant dualization functor such as $\operatorname{Hom}(\cdot, G)$ to produce a cochain complex $\operatorname{Hom}(C_*, G)$, whose homology $H_*(\operatorname{Hom}(C_*, G))$ I would also call the *cohomology* of C_*, written as $H^*(C_*; G)$; this makes $H^*(\cdot; G)$ a contravariant functor on the category of chain complexes. I try to avoid calling $H^*(C_*; G)$ the "cohomology of $\operatorname{Hom}(C_*, G)$" or writing $H^*(\operatorname{Hom}(C_*, G))$, because the functor that takes $\operatorname{Hom}(C_*, G)$ to the cohomology of C_* is covariant, i.e. it is the usual homology functor H_*. Similarly, if A^* is a cochain complex that is not given as the dualization of any chain complex, then we can define its homology $H_*(A^*)$ in the usual way, but I would not call this the "cohomology of A^*" because I do not see any actual contravariant functor in the picture. I think this convention makes sense, but not everyone agrees; some prefer to use the word "cohomology" whenever there is a cochain complex (i.e. with boundary operators of degree $+1$ instead of -1) involved.

difference is that feeding the diagram of Figure 24 into $\mathcal{F}$ reverses all the arrows, so that for any given short exact sequence $0 \to A \xrightarrow{f} B \xrightarrow{g} C \to 0$ and a suitable choice of projective resolutions for its three terms, we obtain a short exact sequence of cochain complexes

$$0 \longrightarrow \mathcal{F}(C_*) \xrightarrow{\mathcal{F}(g_*)} \mathcal{F}(B_*) \xrightarrow{\mathcal{F}(f_*)} \mathcal{F}(A_*) \longrightarrow 0,$$

and the resulting long exact sequence therefore takes the form

$$\begin{aligned}(49.6)\qquad 0 \to \mathcal{F}(C) \to \mathcal{F}(B) \to \mathcal{F}(A) \to R_1\mathcal{F}(C) \to R_1\mathcal{F}(B) \to R_1\mathcal{F}(A) \\ \to R_2\mathcal{F}(C) \to R_2\mathcal{F}(B) \to R_2\mathcal{F}(A) \to \dots .\end{aligned}$$

The Ext **functors.** We now consider the example $\mathcal{F} = \mathrm{Hom}(\cdot, G)$.

Proposition 49.3. *For any fixed R-module G, the contravariant functor $\mathrm{Hom}(\cdot, G) : \mathsf{Mod}^R \to \mathsf{Mod}^R$ is left-exact.*

Proof. If $g : B \to C$ is surjective, then its dualization $g^* : \mathrm{Hom}(C, G) \to \mathrm{Hom}(B, G) : \varphi \mapsto \varphi \circ g$ is clearly injective. Assuming also that g descends to an isomorphism $\mathrm{coker}(f) \to C$ for some map $f : A \to B$, we need to show that everything in the kernel of $f^* : \mathrm{Hom}(B, G) \to \mathrm{Hom}(A, G)$ is also in the image of g^*. An element $\varphi \in \mathrm{Hom}(B, G)$ is in $\ker(f^*)$ if and only if $\varphi(f(a)) = 0$ for every $a \in A$, so the claim is that if the latter holds, then $\varphi = g^*\psi = \psi \circ g$ for some $\psi \in \mathrm{Hom}(C, G)$. Define ψ as follows: every $c \in C$ can be written as $g(b)$ for some $b \in B$, and we claim that

$$\psi(c) := \varphi(b)$$

is independent of the choice of $b \in g^{-1}(c)$. Indeed, for any other choice $b' \in g^{-1}(c)$, $b' - b \in \ker(g) = \mathrm{im}(f)$ implies $b' = b + f(a)$ for some $a \in A$, thus $\varphi(b') = \varphi(b) + \varphi(f(a)) = \varphi(b)$ since $f^*\varphi = 0$. By construction, $g^*\psi(b) = \psi(g(b)) = \varphi(b)$ for all $b \in B$. □

In order for $\mathrm{Hom}(\cdot, G)$ to be a fully exact functor, we would also need to know that for any injective homomorphism $f : A \to B$, the dualization $f^* : \mathrm{Hom}(B, G) \to \mathrm{Hom}(A, G)$ is surjective. In practice, this would mean that for every $f : A \to B$ injective and every $\varphi : A \to G$, a third homomorphism $\tilde{\varphi} : B \to G$ can always be found so that the diagram

$$\begin{tikzcd} B \arrow[dr, dashed, "\tilde{\varphi}"] & \\ A \arrow[u, hook, "f"] \arrow[r, "\varphi"] & G \end{tikzcd}$$

commutes. This is an extension problem: since f is injective, an equivalent condition would be that for every module B and submodule $A \subset B$, every homomorphism $A \to G$ admits an extension to a homomorphism $B \to G$. We call G an **injective** module if it has this property, and the discussion above can then be summarized as follows:

Proposition 49.4. *For an R-module G, the contravariant left-exact functor $\mathrm{Hom}(\cdot, G) : \mathsf{Mod}^R \to \mathsf{Mod}^R$ is exact if and only if G is injective.* □

It is pretty easy to think of modules that are *not* injective, e.g. the abelian group $\mathbb{Z}$ is not an injective $\mathbb{Z}$-module, since the isomorphism $2\mathbb{Z} \to \mathbb{Z} : m \mapsto m/2$ cannot be extended from the subgroup $2\mathbb{Z} \subset \mathbb{Z}$ to a homomorphism $\mathbb{Z} \to \mathbb{Z}$. In this regard, the obvious problem with $\mathbb{Z}$ is that it is not *divisible*, meaning it contains elements $m \in \mathbb{Z}$ that cannot be written as $m = k\ell$ for some $\ell \in \mathbb{Z}$ and arbitrary natural numbers k. Any abelian group G that is not divisble will fail to be injective due to examples such as $\varphi : k\mathbb{Z} \to G$ for $k \geqslant 2$ as described above. A somewhat less obvious fact is that the converse also holds: every divisible abelian group is an injective $\mathbb{Z}$-module (see e.g. [**Bre93**, Proposition V.6.2]), thus producing simple examples of injective $\mathbb{Z}$-modules such as the rational numbers $\mathbb{Q}$. With this knowledge, in fact, it is not so difficult to show that every abelian group is isomorphic to a subgroup of one that is injective, and there are also relatively simple ways of extending that result to the context of modules over an arbitrary commutative ring (see [**Bae40**, **ES53**]). I do not intend to either prove or make essential use of such a fact, but let us record it here for future reference, since it forms an important component of the big picture:

LEMMA 49.5. *Every R-module is isomorphic to a submodule of one that is injective.* □

There is an obvious analogy between injective modules and projective modules, and the roles that they play can be understood as dual to each other. The following exercise makes this statement slightly more precise.

EXERCISE 49.6. Show that for an R-module G, the covariant functor $\mathrm{Hom}(G,\cdot) : \mathsf{Mod}^R \to \mathsf{Mod}^R$ is left-exact, and is exact if and only if G is projective.

REMARK 49.7. The fact that every module is isomorphic to a submodule of an injective one is relatively easy to prove once one understands that divisible abelian groups are injective, but standard proofs of the latter require Zorn's lemma, which (to my mind at least) makes the result seem a bit unintuitive. The analogous fact about projective modules is that every module is isomorphic to a quotient of one, which seems obvious because we know how to construct a free module with a surjection onto any given module. But if we're being fully honest about it, our knowledge that free modules are projective also depends, in an essential way, on a slightly sneaky invocation of the axiom of choice. For those who find this kind of thing amusing: it is known (see [**Bla79**]) that each of the statements "free $\Rightarrow$ projective" and "divisible $\Rightarrow$ injective" for abelian groups is *equivalent* to the axiom of choice.

DEFINITION 49.8. For any R-module G and each integer $n \geqslant 0$, the contravariant functor

$$\mathrm{Ext}_n(\cdot, G) : \mathsf{Mod}^R \to \mathsf{Mod}^R : A \mapsto \mathrm{Ext}_n(A, G)$$

is defined as the right derived functor $R_n\mathcal{F}$ associated to the left-exact contravariant functor $\mathcal{F} := \mathrm{Hom}(\cdot, G)$. More explicitly, $\mathrm{Ext}_n(A, G)$ is the nth cohomology with coefficients in G of the chain complex A_* formed by a choice of projective resolution $A_* \xrightarrow{\alpha} A$, i.e.

$$\mathrm{Ext}_n(A, G) := H^n(A_*; G) := H_n(\mathrm{Hom}(A_*, G)).$$

Since the case $n = 1$ arises most often (and the cases $n \geqslant 2$ all vanish if R is a principal ideal domain), this case is often denoted simply by

$$\mathrm{Ext}(A, G) := \mathrm{Ext}_1(A, G),$$

and in situations where it is important to specify the ring R, we occasionally write[80]

$$\mathrm{Ext}^R_n(A, G) := \mathrm{Ext}_n(A, G), \qquad \mathrm{Ext}^R(A, G) := \mathrm{Ext}^R_1(A, G) := \mathrm{Ext}(A, G).$$

Applying $\mathrm{Hom}(\cdot, G)$ to any short exact sequence $0 \to A \to B \to C \to 0$ now produces via (49.6) a long exact sequence of the form

$$\begin{aligned} 0 \to \mathrm{Hom}(C, G) \to \mathrm{Hom}(B, G) \to \mathrm{Hom}(A, G) \to \mathrm{Ext}(C, G) \to \mathrm{Ext}(B, G) \\ \to \mathrm{Ext}(A, G) \to \mathrm{Ext}_2(C, G) \to \mathrm{Ext}_2(B, G) \to \mathrm{Ext}_2(A, G) \to \dots, \end{aligned} \tag{49.7}$$

[80]Same caveat here as when we introduced notation for Tor: some authors prefer to regard $\mathbb{Z}$ as the default choice of ring and thus write $\mathrm{Ext} := \mathrm{Ext}^{\mathbb{Z}}$, which necessitates writing Ext^R whenever any other ring is to be used. In practice this issue seldom arises, because the most important other choices of R are fields, for which Ext always vanishes because vector spaces are projective.

in which only the terms up to $\mathrm{Ext}(A, G)$ can be nontrivial if R is a principal ideal domain. Proposition 49.2 also gives

$$\mathrm{Ext}_n(A, G) = 0 \qquad \text{for every } n \geqslant 1 \text{ if } A \text{ is projective or } G \text{ is injective,}$$

where the statement follows in the projective case by choosing the simplest possible projective resolution $\ldots \to 0 \to A \xrightarrow{\mathbb{1}} A \to 0$, and in the injective case because $\mathrm{Hom}(\cdot, G)$ is then an exact functor. The latter is the cohomological analogue of the fact that $\mathrm{Tor}_n(A, G)$ vanishes for $n \geqslant 1$ whenever G is a flat R-module. Proposition 49.2 also gives natural isomorphisms $\mathrm{Ext}_0(A, G) \cong \mathrm{Hom}(A, G)$, $\mathrm{Ext}_n(A \oplus B, G) \cong \mathrm{Ext}_n(A, G) \oplus \mathrm{Ext}_n(B, G)$, and the formula

$$\mathrm{Ext}(R/kR, G) \cong \mathrm{Hom}(R, G)\Big/ k\,\mathrm{Hom}(R, G)$$

for any $k \in \mathbb{N}$ such that multiplication by k is an injective map $R \to R$.

Let us describe a few properties of the functors Ext_n that are not already covered by the general properties of right derived functors. The first is that $\mathrm{Ext}_n(A, G)$ is not only a contravariant functor with respect to A, but also a *covariant* functor with respect to G. Indeed, Hom has this property, since for any fixed A and homomorphism $\varphi : G \to H$, there is an induced homomorphism

$$\mathrm{Hom}(A, G) \xrightarrow{\varphi_*} \mathrm{Hom}(A, H)$$

defined by composing homomorphisms $A \to G$ with φ. Given a projective resolution $A_* \xrightarrow{\alpha} A$ of A, the homomorphisms $\mathrm{Hom}(A_n, G) \xrightarrow{\varphi_*} \mathrm{Hom}(A_n, H)$ for every $n \geqslant 0$ then define a chain map between cochain complexes

$$\mathrm{Hom}(A_*, G) \xrightarrow{\varphi_*} \mathrm{Hom}(A_*, H),$$

so that the induced map on the homology in degree n defines a homomorphism

$$\mathrm{Ext}_n(A, G) \xrightarrow{\varphi_*} \mathrm{Ext}_n(A, H).$$

One easily checks that for $n = 0$, the natural isomorphisms identify this map with the map $\mathrm{Hom}(A, G) \to \mathrm{Hom}(A, H)$ described above.

Now recall from Exercise 49.6: if A is projective, then the covariant functor $\mathrm{Hom}(A, \cdot)$ is exact. If follows that for any (not necessarily projective) module A and any short exact sequence

$$0 \longrightarrow G \xrightarrow{\varphi} H \xrightarrow{\psi} K \longrightarrow 0,$$

we can apply the functors $\mathrm{Hom}(A_n, \cdot)$ arising from a projective resolution $A_* \xrightarrow{\alpha} A$ of A to produce a short exact sequence of cochain complexes

$$0 \longrightarrow \mathrm{Hom}(A_*, G) \xrightarrow{\varphi_*} \mathrm{Hom}(A_*, H) \xrightarrow{\psi_*} \mathrm{Hom}(A_*, K) \longrightarrow 0.$$

The resulting long exact sequence of cohomologies has the form

$$\begin{aligned} 0 \to \mathrm{Hom}(A, G) \to \mathrm{Hom}(A, H) \to \mathrm{Hom}(A, K) \to \mathrm{Ext}(A, G) \to \mathrm{Ext}(A, H) \\ \to \mathrm{Ext}(A, K) \to \mathrm{Ext}_2(A, G) \to \mathrm{Ext}_2(A, H) \to \mathrm{Ext}_2(A, K) \to \ldots, \end{aligned} \tag{49.8}$$

in which, as usual, there are only six potentially nontrivial terms if R is a principal ideal domain. Notice that the direction of this sequence is different from (49.7), a consequence of the fact that $\mathrm{Hom}(A, \cdot)$ is covariant while $\mathrm{Hom}(\cdot, G)$ is contravariant. On the other hand, both sequences have in common that their connecting homomorphisms adjust the degree upward instead of downward: this is so because the dualized projective resolutions are *cochain* complexes, not chain complexes.

The properties stated so far furnish us with a workable recipe for computing $\mathrm{Ext}^{\mathbb{Z}}(A, G)$ whenever A is a finitely-generated abelian group. We have, for instance:

PROPOSITION 49.9. *For any finitely generated abelian group A, $\mathrm{Ext}^{\mathbb{Z}}(A, \mathbb{Z})$ is isomorphic to the torsion subgroup of A.*

Proof. By the classification of finitely generated abelian groups, A is isomorphic to the direct sum of some free abelian group F with the torsion subgroup $T \subset A$, and T in turn is a finite direct sum of finite cyclic groups $\mathbb{Z}_{k_1}, \dots, \mathbb{Z}_{k_N}$. Since F is projective and $\operatorname{Hom}(\mathbb{Z}, \mathbb{Z}) \cong \mathbb{Z}$, Proposition 49.2 then gives

$$\operatorname{Ext}^{\mathbb{Z}}(A, \mathbb{Z}) \cong \operatorname{Ext}^{\mathbb{Z}}(F, \mathbb{Z}) \oplus \left(\bigoplus_{j=1}^{N} \operatorname{Ext}^{\mathbb{Z}}(\mathbb{Z}_{k_j}, \mathbb{Z}) \right) \cong \bigoplus_{j=1}^{N} \mathbb{Z}_{k_j} \cong T.$$

□

The other definition of Ext. We will not need the following detail in subsequent developments, but it would seem criminal not to mention it. Recall from Lecture 43 that the Tor_n functors are symmetric, meaning there are always natural isomorphisms $\operatorname{Tor}_n(A, B) \cong \operatorname{Tor}_n(B, A)$. We proved this via a diagram chase involving a double complex, based on the fact that there are always natural isomorphisms $A \otimes B \cong B \otimes A$. You may be wondering what the analogous property of the functors Ext_n is.

Symmetry is not the answer, and the obvious reason why not is that Ext_n is contravariant in one of its variables but covariant in the other. The key instead is to fix an R-module G and develop derived functors for the covariant functor $\operatorname{Hom}(G, \cdot)$.

According to Exercise 49.6, $\operatorname{Hom}(G, \cdot)$ is left-exact, and it is exact if and only if G is projective. The theory of left derived functors developed in Lectures 42 and 43 is not appropriate for left-exact covariant functors, but one can develop an analogous theory, in which most of the arrows are reversed. The main idea is to replace projective resolutions $A_* \xrightarrow{\alpha} A$ of a module A by **injective resolutions** $A \xrightarrow{\alpha} A^*$, which are exact sequences

$$0 \longrightarrow A \xrightarrow{\alpha} A^0 \xrightarrow{\alpha_0} A^1 \xrightarrow{\alpha_1} A^2 \longrightarrow \dots$$

in which the modules A^n for all $n \geqslant 0$ are injective. Injective resolutions always exist, and you will easily come up with a proof of this fact if you already believe Lemma 49.5, which can be used to define $\alpha : A \to A^0$ as the inclusion of A into a larger injective module A^0, whose quotient by $\operatorname{im}(\alpha) \subset A^0$ can then be included into another injective module A^1, and so forth. An injective resolution of A yields a cochain complex

$$\dots \longrightarrow 0 \longrightarrow A^0 \xrightarrow{\alpha_0} A^1 \xrightarrow{\alpha_1} A^2 \longrightarrow \dots,$$

which we shall abbreviate by A^*, and plugging this into an additive covariant functor $\mathcal{F}$ then yields another cochain complex $\mathcal{F}(A^*)$. If $\mathcal{F}$ is left-exact, then the right derived functors of $\mathcal{F}$ are defined for each $n \geqslant 0$ by

$$R_n\mathcal{F}(A) := H_n(\mathcal{F}(A^*)).$$

It is straightforward to prove an injective analogue of Proposition 42.24, so that any homomorphism $\varphi : A \to B$ gives rise to a unique chain homotopy class of chain maps $\varphi_* : A^* \to B^*$ between the corresponding injective resolutions. Feeding such a chain map into $\mathcal{F}$ gives a chain map $\mathcal{F}(A^*) \to \mathcal{F}(B^*)$, thus inducing a natural homomorphism $R_n\mathcal{F}(A) \to R_n\mathcal{F}(B)$ that is independent of choices, and proving at the same time that $R_n\mathcal{F}$ is (up to natural isomorphisms) independent of the choices of injective resolutions. There is also an injective analogue of the horseshoe lemma (Proposition 43.10), turning any short exact sequence $0 \to A \to B \to C \to 0$ into a short exact sequence of cochain complexes $0 \to \mathcal{F}(A^*) \to \mathcal{F}(B^*) \to \mathcal{F}(C^*) \to 0$; the exactness of the latter follows from the observation that a short exact sequence $0 \to A^n \to B^n \to C^n \to 0$ splits whenever A^n is injective, because injectivity guarantees the existence of a left-inverse for the injective map $A^n \to B^n$. The result is that for a left-exact covariant functor $\mathcal{F}$, any short exact sequence

$0 \to A \to B \to C \to 0$ gives rise to a long exact sequence

$$0 \to \mathcal{F}(A) \to \mathcal{F}(B) \to \mathcal{F}(C) \to R_1\mathcal{F}(A) \to R_1\mathcal{F}(B) \to R_1\mathcal{F}(C) \\ \to R_2\mathcal{F}(A) \to R_2\mathcal{F}(B) \to R_2\mathcal{F}(C) \to \dots.$$

Exploiting left-exactness and the freedom to make intelligent choices of injective resolutions, one proves the following properties analogous to Proposition 49.2:

(1) There is a natural isomorphism $\mathcal{F}(A) \cong R_0\mathcal{F}(A)$.
(2) If $\mathcal{F}$ is exact or A is injective, then $R_n\mathcal{F}(A) = 0$ for every $n \geqslant 1$.
(3) There are natural isomorphisms $R_n\mathcal{F}(A \oplus B) \cong R_n\mathcal{F}(A) \oplus R_n\mathcal{F}(B)$ for every $n \geqslant 0$.

Since $\mathrm{Hom}(A, \cdot)$ is left-exact for any A, we can apply this machinery to define a second variant of the sequence Ext functors, which we shall denote for now by

$$\widehat{\mathrm{Ext}}_n(A, \cdot) := R_n(\mathrm{Hom}(A, \cdot)), \qquad \widehat{\mathrm{Ext}} := \widehat{\mathrm{Ext}}_1,$$

so explicitly,

$$\widehat{\mathrm{Ext}}_n(A, B) = H_n(\mathrm{Hom}(A, B^*))$$

for the cochain complex $\mathrm{Hom}(A, B^*)$ arising from any injective resolution $B \xrightarrow{\beta} B^*$ of B. We can then observe that $\widehat{\mathrm{Ext}}_n$ has several properties matching those of Ext_n, though seemingly for different reasons. Indeed,

$$\widehat{\mathrm{Ext}}_n(A, B) = 0 \qquad \text{for every } n \geqslant 1 \text{ if } A \text{ is projective or } B \text{ is injective},$$

the reason being that $\mathrm{Hom}(A, \cdot)$ is exact if A is projective, and $0 \to B \xrightarrow{\mathbb{1}} B \to 0 \to \dots$ is an injective resolution if B is injective. Similarly, a short exact sequence $0 \to G \to H \to K \to 0$ produces a long exact sequence

$$0 \to \mathrm{Hom}(A, G) \to \mathrm{Hom}(A, H) \to \mathrm{Hom}(A, K) \to \widehat{\mathrm{Ext}}(A, G) \to \widehat{\mathrm{Ext}}(A, H) \\ \to \widehat{\mathrm{Ext}}(A, K) \to \widehat{\mathrm{Ext}}_2(A, G) \to \widehat{\mathrm{Ext}}_2(A, H) \to \widehat{\mathrm{Ext}}_2(A, K) \to \dots$$

due to the general properties of right derived functors, and this sequence has the same form as (49.8), which was obtained via a more direct argument (with no need of the horseshoe lemma) using the covariant functoriality of Ext_n in the second variable. The analogous result for $\widehat{\mathrm{Ext}}_n$ shows that it is likewise a contravariant functor in its first variable, and that short exact sequences $0 \to A \to B \to C \to 0$ give rise to long exact sequences

$$0 \to \mathrm{Hom}(C, G) \to \mathrm{Hom}(B, G) \to \mathrm{Hom}(A, G) \to \widehat{\mathrm{Ext}}(C, G) \to \widehat{\mathrm{Ext}}(B, G) \\ \to \widehat{\mathrm{Ext}}(A, G) \to \widehat{\mathrm{Ext}}_2(C, G) \to \widehat{\mathrm{Ext}}_2(B, G) \to \widehat{\mathrm{Ext}}_2(A, G) \to \dots,$$

thus matching the form of the sequence (49.7) for Ext_n obtained from the general theory of contravariant right derived functors. All this provides strong evidence that Ext_n and $\widehat{\mathrm{Ext}}_n$ are, secretly, the same thing.

We will leave the details as an exercise, but an explicit (and natural) isomorphism $\mathrm{Ext}_n(A, B) \cong \widehat{\mathrm{Ext}}_n(A, B)$ can be obtained from the double complex in Figure 30. Here, $A_* \xrightarrow{\alpha} A$ is a projective resolution of A, $B \xrightarrow{\beta} B^*$ is an injective resolution of B, and the vertical and horizontal maps are all obtained via the contravariant and covariant functoriality of Hom in its first and second variables respectively. The nth homology of the cochain complex $\mathrm{Hom}(A_*, B)$ in the leftmost nontrivial column is $\mathrm{Ext}_n(A, B)$, the nth homology of the complex $\mathrm{Hom}(A, B^*)$ in the lowest nontrivial row is $\widehat{\mathrm{Ext}}_n(A, B)$, and all the other rows and columns are exact, due to the fact that each A_n is projective and each B^n is injective, making $\mathrm{Hom}(A_n, \cdot)$ and $\mathrm{Hom}(\cdot, B_n)$ exact functors. The rest, as they say, is diagram chasing (cf. Exercise 43.13).

$$\begin{tikzcd}
 & \vdots & \vdots & \vdots & \vdots & \vdots & \\
\cdots \arrow[r] & 0 \arrow[r] \arrow[u] & \mathrm{Hom}(A_2,B) \arrow[r] \arrow[u] & \mathrm{Hom}(A_2,B^0) \arrow[r] \arrow[u] & \mathrm{Hom}(A_2,B^1) \arrow[r] \arrow[u] & \mathrm{Hom}(A_2,B^2) \arrow[r] \arrow[u] & \cdots \\
\cdots \arrow[r] & 0 \arrow[r] \arrow[u] & \mathrm{Hom}(A_1,B) \arrow[r] \arrow[u] & \mathrm{Hom}(A_1,B^0) \arrow[r] \arrow[u] & \mathrm{Hom}(A_1,B^1) \arrow[r] \arrow[u] & \mathrm{Hom}(A_1,B^2) \arrow[r] \arrow[u] & \cdots \\
\cdots \arrow[r] & 0 \arrow[r] \arrow[u] & \mathrm{Hom}(A_0,B) \arrow[r] \arrow[u] & \mathrm{Hom}(A_0,B^0) \arrow[r] \arrow[u] & \mathrm{Hom}(A_0,B^1) \arrow[r] \arrow[u] & \mathrm{Hom}(A_0,B^2) \arrow[r] \arrow[u] & \cdots \\
\cdots \arrow[r] & 0 \arrow[r] \arrow[u] & 0 \arrow[r] \arrow[u] & \mathrm{Hom}(A,B^0) \arrow[r] \arrow[u] & \mathrm{Hom}(A,B^1) \arrow[r] \arrow[u] & \mathrm{Hom}(A,B^2) \arrow[r] \arrow[u] & \cdots \\
\cdots \arrow[r] & 0 \arrow[r] \arrow[u] & 0 \arrow[r] \arrow[u] & 0 \arrow[r] \arrow[u] & 0 \arrow[r] \arrow[u] & 0 \arrow[r] \arrow[u] & \cdots \\
 & \vdots \arrow[u] & \vdots \arrow[u] & \vdots \arrow[u] & \vdots \arrow[u] & \vdots \arrow[u] &
\end{tikzcd}$$

FIGURE 30. The double complex that implies $\mathrm{Ext}_n(A,B) \cong \widehat{\mathrm{Ext}}_n(A,B)$.

The universal coefficient theorem. With the necessary machinery in place, we now assume again that R is a principal ideal domain and C_* is a chain complex of free R-modules, and investigate the natural relation between $H_n(\mathcal{F}(C_*))$ and $\mathcal{F}(H_n(C_*))$ when $\mathcal{F} : \mathsf{Mod}^R \to \mathsf{Mod}^R$ is a contravariant left-exact functor. The argument is a straightforward adaptation of the covariant right-exact case.

As before, we begin with the short exact sequence of chain complexes

$$0 \longrightarrow Z_* \stackrel{j}{\longrightarrow} C_* \stackrel{\partial'}{\longrightarrow} B_{*-1} \longrightarrow 0,$$

which splits, so plugging it into $\mathcal{F}$ gives a short exact sequence of cochain complexes

$$0 \longleftarrow \mathcal{F}(Z_*) \stackrel{\mathcal{F}(j)}{\longleftarrow} \mathcal{F}(C_*) \stackrel{\mathcal{F}(\partial')}{\longleftarrow} \mathcal{F}(B_{*-1}) \longleftarrow 0,$$

and therefore a long exact sequence

$$\ldots \longleftarrow \mathcal{F}(B_n) \stackrel{\Phi_n}{\longleftarrow} \mathcal{F}(Z_n) \stackrel{\mathcal{F}(j)_*}{\longleftarrow} H_n(\mathcal{F}(C_*)) \stackrel{\mathcal{F}(\partial')_*}{\longleftarrow} \mathcal{F}(B_{n-1}) \stackrel{\Phi_{n-1}}{\longleftarrow} \mathcal{F}(Z_{n-1}) \longleftarrow \ldots.$$

One needs to do a quick diagram-chasing exercise to figure out what the connecting homomorphisms Φ_n are, but the answer will not surprise you: as in the covariant case, one finds

$$\Phi_n = \mathcal{F}(i_n) : \mathcal{F}(Z_n) \to \mathcal{F}(B_n),$$

and the proof depends on the relation obtained by applying $\mathcal{F}$ to the equation $\partial_{n+1} = j_n i_n \partial'_{n+1}$. Turning the long exact sequence into a short exact sequence with $H_n(\mathcal{F}(C_*))$ as the middle term thus gives

$$0 \to \operatorname{coker} \mathcal{F}(i_{n-1}) \to H_n(\mathcal{F}(C_*)) \to \ker \mathcal{F}(i_n) \to 0.$$

The left-exactness of $\mathcal{F}$ now becomes relevant because it turns the short exact sequence $0 \to B_n \stackrel{i_n}{\to} Z_n \stackrel{q_n}{\to} H_n(C_*) \to 0$ into the (not very) long exact sequence

$$0 \longrightarrow \mathcal{F}(H_n(C_*)) \stackrel{\mathcal{F}(q_n)}{\longrightarrow} \mathcal{F}(Z_n) \stackrel{\mathcal{F}(i_n)}{\longrightarrow} \mathcal{F}(B_n) \longrightarrow R_1\mathcal{F}(H_n(C_*)) \longrightarrow R_1\mathcal{F}(Z_n) = 0 \longrightarrow \ldots,$$

where the term $R_1\mathcal{F}(Z_n)$ vanishes because Z_n is free (and therefore projective). This exact sequence gives us natural isomorphisms

$$\mathcal{F}(H_n(C_*)) \stackrel{\cong}{\longrightarrow} \ker \mathcal{F}(i_n) \subset \mathcal{F}(Z_n), \qquad \text{and} \qquad \operatorname{coker} \mathcal{F}(i_n) \stackrel{\cong}{\longrightarrow} R_1\mathcal{F}(H_n(C_*)),$$

in which the first is the map $\mathcal{F}(q_n)$, and the second is the result of letting the connecting homomorphism $\mathcal{F}(B_n) \to R_1\mathcal{F}(H_n(C_*))$ descend to the quotient. Our previous long exact sequence thus becomes

$$0 \longrightarrow R_1\mathcal{F}(H_{n-1}(C_*)) \longrightarrow H_n(\mathcal{F}(C_*)) \xrightarrow{h} \mathcal{F}(H_n(C_*)) \longrightarrow 0. \tag{49.9}$$

Inspecting where the maps in this sequence came from, one finds that $h : H_n(\mathcal{F}(C_*)) \to \mathcal{F}(H_n(C_*))$ is characterized uniquely by the relation

$$\mathcal{F}(q_n)h([\varphi]) = \mathcal{F}(j_n)\varphi \qquad \text{for all } \varphi \in \ker \mathcal{F}(\partial_{n+1}) \subset \mathcal{F}(C_n), \tag{49.10}$$

due to the fact that $\mathcal{F}(q_n)$ maps $\mathcal{F}(H_n(C_*))$ isomorphically to $\ker \mathcal{F}(i_n)$.

Does the sequence split? Of course it does: the proof of this in the covariant case proceeded by composing left-inverses $p_n : C_n \to Z_n$ of the inclusions $j_n : Z_n \hookrightarrow C_n$ with the quotient projections $q_n : Z_n \to H_n(C_*)$ to produce a chain map $qp : C_* \to H_*(C_*)$, for which the induced chain map $\mathcal{F}(qp) : \mathcal{F}(C_*) \to \mathcal{F}(H_*(C_*))$ determined maps on homology $H_n(\mathcal{F}(C_*)) \to \mathcal{F}(H_n(C_*))$ that were left-inverses of the injection $\mathcal{F}(H_n(C_*)) \to H_n(\mathcal{F}(C_*))$ in the exact sequence. In the contravariant case, feeding $qp : C_* \to H_*(C_*)$ into $\mathcal{F}$ produces a chain map

$$\mathcal{F}(qp) : \mathcal{F}(H_*(C_*)) \to \mathcal{F}(C_*)$$

between two cochain complexes, with the $\mathbb{Z}$-graded R-module $\mathcal{F}(H_*(C_*))$ regarded as a cochain complex with trivial coboundary operator. The induced map on homology in degree n is then a homomorphism

$$\mathcal{F}(qp)_* : \mathcal{F}(H_n(C_*)) \to H_n(\mathcal{F}(C_*)),$$

which we claim is now a right-inverse of the surjective map $h : H_n(\mathcal{F}(C_*)) \to \mathcal{F}(H_n(C_*))$. Indeed, for any $\varphi \in \mathcal{F}(H_n(C_*))$, $\mathcal{F}(qp)_*\varphi \in H_n(\mathcal{F}(C_*))$ is represented by the cycle $\mathcal{F}(q_n p_n)\varphi \in \mathcal{F}(C_n)$, and since $p_n j_n$ is the identity map on Z_n, we have

$$\mathcal{F}(j_n)\mathcal{F}(q_n p_n)\varphi = \mathcal{F}(q_n p_n j_n)\varphi = \mathcal{F}(q_n)\varphi,$$

implying $h\mathcal{F}(qp)_*\varphi = \varphi$ according to the characterization of h in (49.10).

Finally, fix an R-module G and set $\mathcal{F} := \operatorname{Hom}(\cdot, G)$. The right derived functor appearing in (49.9) now becomes $\operatorname{Ext}(H_{n-1}(C_*), G)$, while the other two terms are $H_n(\operatorname{Hom}(C_*, G)) = H^n(C_*; G)$ and $\operatorname{Hom}(H_n(C_*), G)$. According to (49.10), the map $h : H^n(C_*; G) \to \operatorname{Hom}(H_n(C_*), G)$ is uniquely characterized by the condition that for any n-cocycle $\varphi \in \operatorname{Hom}(C_n, G)$ representing a class $[\varphi] \in H^n(C_*; G)$,

$$q_n^* h([\varphi]) = j_n^* \varphi \in \operatorname{Hom}(Z_n, G)$$

for the inclusion $j_n : Z_n \hookrightarrow C_n$ and quotient projection $q_n : Z_n \to H_n(C_*)$, which concretely means that for every n-cycle $c \in Z_n$, we have

$$h([\varphi])\big([c]\big) = \varphi(c).$$

This means that $h : H^n(C_*; G) \to \operatorname{Hom}(H_n(C_*), G)$ is exactly what we hoped it would be: the canonical map defined via the evaluation pairing. We can now state the main theorem, and the proof is complete except for the detail about naturality, which I will leave as an exercise. For extra clarity, the notation in the following statement specifies explicitly that Hom and Ext are to be understood in the category of R-modules.

THEOREM 49.10 (universal coefficient theorem). *For any chain complex C_* of free modules over a principle ideal domain R, a fixed R-module G and $n \in \mathbb{Z}$, the canonical map $h : H^n(C_*; G) \to$ $\operatorname{Hom}_R(H_n(C_*), G)$ defined as in (49.2) via the evaluation pairing fits into a split exact sequence*

$$0 \longrightarrow \operatorname{Ext}^R\big(H_{n-1}(C_*), G\big) \longrightarrow H^n(C_*; G) \xrightarrow{h} \operatorname{Hom}_R\big(H_n(C_*), G\big) \longrightarrow 0.$$

Moreover, the sequence (though not its splitting) is natural in the sense that for any chain map $\varphi : A_* \to B_*$ between two chain complexes of free R-modules, the exact sequences for both fit into a commutative diagram

$$\begin{array}{ccccccccc} 0 & \longrightarrow & \mathrm{Ext}^R\big(H_{n-1}(A_*), G\big) & \longrightarrow & H^n(A_*; G) & \xrightarrow{h} & \mathrm{Hom}_R\big(H_n(A_*), G\big) & \longrightarrow & 0 \\ & & {\scriptstyle\varphi^*}\big\uparrow & & {\scriptstyle\varphi^*}\big\uparrow & & {\scriptstyle\varphi^*}\big\uparrow & & \\ 0 & \longrightarrow & \mathrm{Ext}^R\big(H_{n-1}(B_*), G\big) & \longrightarrow & H^n(B_*; G) & \xrightarrow{h} & \mathrm{Hom}_R\big(H_n(B_*), G\big) & \longrightarrow & 0, \end{array}$$

where the map $\varphi^* : \mathrm{Ext}^R(H_{n-1}(B_*), G) \to \mathrm{Ext}^R(H_{n-1}(A_*), G)$ arises from the functoriality of Ext^R, and $\varphi^* : \mathrm{Hom}_R(H_n(B_*), G) \to \mathrm{Hom}_R(H_n(A_*), G)$ is defined by dualizing $\varphi_* : H_n(A_*) \to H_n(B_*)$. ☐

Remark 49.11. The proof of Theorem 49.10 carried out above did not require any specific knowledge of the definition of Ext, but relied instead on two of the properties that it satisfies: namely that $\mathrm{Ext}(A, B)$ vanishes whenever A is projective, and that short exact sequences $0 \to A \to B \to C \to 0$ give rise to natural long exact sequences $0 \to \mathrm{Hom}(C, G) \to \mathrm{Hom}(B, G) \to \mathrm{Hom}(A, G) \to \mathrm{Ext}(C, G) \to \ldots$. It therefore would have been possible to prove the theorem without mentioning projective resolutions at all; one could instead use injective resolutions to give the alternative definition of $\mathrm{Ext}(A, \cdot)$ as a derived functor for $\mathrm{Hom}(A, \cdot)$, prove that it vanishes when A is projective because $\mathrm{Hom}(A, \cdot)$ is exact, and construct the required long exact sequence using its contravariant functoriality with respect to A (which does not require the horseshoe lemma). This is the approach taken in [**Bre93**, §V.6]. Of course, one cannot get around talking about projective resolutions at some point (or alternatively *free* resolutions, as is done in [**Hat02**]), as they are necessary for the definition of Tor, and thus for the homological version of the universal coefficient theorem, among other things.

Applications. When applied to the singular chain complex of a pair (X, A), the splitting of the sequence in Theorem 49.10 gives an isomorphism

$$H^n(X, A; G) \cong \mathrm{Hom}\big(H_n(X, A; \mathbb{Z}), G\big) \oplus \mathrm{Ext}^{\mathbb{Z}}\big(H_{n-1}(X, A; \mathbb{Z}), G\big)$$

for any coefficient group G, revealing that $H^n(X, A; G)$ is determined up to isomorphism by $H_n(X, A; \mathbb{Z})$, $H_{n-1}(X, A; \mathbb{Z})$ and G. More generally, if G is a module over a principal ideal domain R, then we can view $H_*(X, A; R)$ as the homology of a chain complex of free R-modules $C_*(X, A; R) \cong C_*(X, A; \mathbb{Z}) \otimes R$, and $H^*(X, A; G)$ is the cohomology of the same chain complex due to Exercise 47.5, so the theorem also gives an R-module isomorphism

$$H^n(X, A; G) \cong \mathrm{Hom}_R\big(H_n(X, A; R), G\big) \oplus \mathrm{Ext}^R\big(H_{n-1}(X, A; R), G\big),$$

which reduces to the same statement again if $R = \mathbb{Z}$. All of this applies equally well to cellular or simplicial chain complexes, since these are also freely generated.

There is a particularly appealing corollary whenever R and G are both chosen to be a field $\mathbb{K}$. All modules in the picture are then vector spaces over $\mathbb{K}$, which are automatically free and thus projective, so the vector space $\mathrm{Ext}^{\mathbb{K}}(H_{n-1}(X, A; \mathbb{K}), \mathbb{K})$ is always trivial, and we conclude that cohomology is just the dual vector space of homology:

Corollary 49.12. *For any field $\mathbb{K}$ and any pair of spaces (X, A), the natural map*

$$H^n(X, A; \mathbb{K}) \to \mathrm{Hom}_{\mathbb{K}}(H_n(X, A; \mathbb{K}), \mathbb{K}) : [\varphi] \mapsto \langle [\varphi], \cdot \rangle$$

is an isomorphism. ☐

Here is another situation in which the Ext term vanishes automatically: since $H_0(X;\mathbb{Z})$ is the free abelian group generated by $\pi_0(X)$, Proposition 49.9 implies $\operatorname{Ext}^{\mathbb{Z}}(H_0(X;\mathbb{Z}),G) = 0$ for every X and G.

COROLLARY 49.13. *The natural map $H^1(X;G) \to \operatorname{Hom}(H_1(X;\mathbb{Z}),G)$ is an isomorphism for all spaces X and abelian groups G.* □

REMARK 49.14. Note that if X is path-connected, then $\operatorname{Hom}(H_1(X;\mathbb{Z}),G)$ is canonically isomorphic to $\operatorname{Hom}(\pi_1(X),G)$; indeed, since G is abelian, every homomorphism $\pi_1(X) \to G$ vanishes on the commutator subgroup and thus descends to the abelianization, which is $H_1(X;\mathbb{Z})$. Corollary 49.13 thus confirms the result of Exercise 47.10.

If we take $R = \mathbb{Z}$, then the Ext term need not vanish in general, but one still obtains something revealing whenever the homology groups are finitely generated. Assume C_* is a chain complex of free abelian groups, abbreviate $H_n := H_n(C_*)$ and $H^n := H^n(C_*) := H^n(C_*;\mathbb{Z})$, and let $T_n \subset H_n$ and $T^n \subset H^n$ denote their respective torsion subgroups. If H_n and H^n are finitely generated, we can define their **free parts** as the quotients

$$H_n^{\text{free}}(C_*) = H_n^{\text{free}} := H_n/T_n, \qquad H^n_{\text{free}}(C_*) = H^n_{\text{free}} := H^n/T^n,$$

and these are finitely-generated free abelian groups such that there exist isomorphisms

$$H_n \cong H_n^{\text{free}} \oplus T_n, \qquad H^n \cong H^n_{\text{free}} \oplus T^n.$$

Applying the universal coefficient theorem with $G = \mathbb{Z}$ now produces the formula

$$H^n \cong \operatorname{Hom}(H_n,\mathbb{Z}) \oplus \operatorname{Ext}^{\mathbb{Z}}(H_{n-1},\mathbb{Z}).$$

Since all homomorphisms $H_n \to \mathbb{Z}$ kill torsion elements and $\operatorname{Hom}(\mathbb{Z}^m,\mathbb{Z}) \cong \mathbb{Z}^m$ for each $m \in \mathbb{N}$, we have $\operatorname{Hom}(H_n,\mathbb{Z}) \cong \operatorname{Hom}(H_n^{\text{free}} \oplus T_n,\mathbb{Z}) \cong \operatorname{Hom}(H_n^{\text{free}},\mathbb{Z}) \cong H_n^{\text{free}}$, and if H_{n-1} is also finitely generated, Proposition 49.9 gives an isomorphism $\operatorname{Ext}^{\mathbb{Z}}(H_{n-1},\mathbb{Z}) \cong T_{n-1}$, resulting in the formula

$$H^n \cong H_n^{\text{free}} \oplus T_{n-1}. \tag{49.11}$$

This implies $H^n_{\text{free}} \cong H_n^{\text{free}}$ and $T^n \cong T_{n-1}$. The first isomorphism can also be understood as follows. According to Theorem 49.10, the natural map $h : H^n \to \operatorname{Hom}(H_n,\mathbb{Z})$ descends to an isomorphism

$$H^n/T \xrightarrow{h} \operatorname{Hom}(H_n,\mathbb{Z}), \tag{49.12}$$

where $T \subset H^n$ is the image of the injective map $\operatorname{Ext}^{\mathbb{Z}}(H_{n-1},\mathbb{Z}) \to H^n$ in the long exact sequence. Composing homomorphisms $H_n \to \mathbb{Z}$ with the quotient projection $H_n \to H_n^{\text{free}}$ gives a natural isomorphism $\operatorname{Hom}(H_n^{\text{free}},\mathbb{Z}) = \operatorname{Hom}(H_n,\mathbb{Z})$ since $\operatorname{Hom}(T_n,\mathbb{Z}) = 0$, and $\operatorname{Hom}(H_n^{\text{free}},\mathbb{Z}) \cong H_n^{\text{free}}$ is a finitely generated free abelian group. Since $\operatorname{Ext}^{\mathbb{Z}}(H_{n-1},\mathbb{Z}) \cong T_{n-1}$ is a torsion group, the subgroup $T \subset H^n$ is necessarily contained in the torsion subgroup, but it must in fact be all of it since the isomorphism (49.12) implies that H^n/T is free. Combining all this with the naturality in Theorem 49.10 gives:

COROLLARY 49.15. *For any chain complex C_* of free abelian groups and any $n \in \mathbb{Z}$ such that $H_n(C_*)$ and $H_{n-1}(C_*)$ are both finitely generated, the natural map $h : H^n(C_*;\mathbb{Z}) \to \operatorname{Hom}(H_n(C_*),\mathbb{Z}) = \operatorname{Hom}(H_n^{\text{free}}(C_*),\mathbb{Z})$ descends to a natural isomorphism*

$$H^n_{\text{free}}(C_*) \xrightarrow{h} \operatorname{Hom}(H_n^{\text{free}}(C_*),\mathbb{Z}),$$

where naturality means that for any two chain complexes A_*, B_* satisfying the above conditions and a chain map $\varphi_* : A_* \to B_*$, there is a commutative diagram

$$\begin{tikzcd} H^n_{\text{free}}(A_*) \arrow[r, "h", "\cong"'] & \operatorname{Hom}(H_n^{\text{free}}(A_*), \mathbb{Z}) \\ H^n_{\text{free}}(B_*) \arrow[r, "h", "\cong"'] \arrow[u, "\varphi^*"] & \operatorname{Hom}(H_n^{\text{free}}(B_*), \mathbb{Z}). \arrow[u, "\varphi^*"'] \end{tikzcd}$$

In particular, $H^n_{\text{free}}(C_*) \cong H_n^{\text{free}}(C_*)$. Moreover, the torsion of $H^n(C_*; \mathbb{Z})$ is isomorphic to the torsion of $H_{n-1}(C_*)$. □

Exercise 49.16. Use cellular cohomology to compute $H^*(\mathbb{RP}^2; \mathbb{Z})$, and compare the result with $H_*(\mathbb{RP}^2; \mathbb{Z})$ in light of Corollary 49.15.

Recall the numerical invariants defined in Lecture 40, namely the Betti numbers $b_n(X) = \operatorname{rank} H_n(X; \mathbb{Z})$ and the Euler characteristic $\chi(X) = \sum_{n=0}^{\infty}(-1)^n b_n(X)$. The universal coefficient theorem now gives us the freedom to compute these in terms of cohomology instead of homology:

Corollary 49.17. *For any space X such that $H_*(X; \mathbb{Z})$ is finitely generated,*

$$b_n(X) = \operatorname{rank} H^n(X; \mathbb{Z}) = \dim_{\mathbb{K}} H^n(X; \mathbb{K})$$

for every integer $n \geqslant 0$ and any field $\mathbb{K}$ of characteristic zero. Moreover, if X is a compact CW-complex, then the formula

$$\chi(X) = \sum_{n=0}^{\infty} (-1)^n \dim_{\mathbb{K}} H^n(X; \mathbb{K})$$

holds for any field $\mathbb{K}$.

Proof. Since $b_n(X)$ depends only on the free part of $H_n(X; \mathbb{Z})$, it matches $\operatorname{rank} H^n(X; \mathbb{Z})$ due to Corollary 49.15. For any field $\mathbb{K}$ of characteristic zero, the universal coefficient theorem for homology then implies $\operatorname{rank} H_n(X; \mathbb{Z}) = \dim_{\mathbb{K}} H_n(X; \mathbb{K})$, and the latter matches $\dim_{\mathbb{K}} H^n(X; \mathbb{K})$ by Corollary 49.12. Finally if X is a finite cell complex, then we already know by applying Proposition 40.9 to the cellular chain complex that $\chi(X) = \sum_{n=0}^{\infty}(-1)^n \dim_{\mathbb{K}} H_n(X; \mathbb{K})$ holds for every field $\mathbb{K}$, and Corollary 49.12 enables us to replace $H_n(X; \mathbb{K})$ by $H^n(X; \mathbb{K})$ in this expression. □

We can also use these kinds of tricks to compute the Lefschetz number of a map $f : X \to X$. Here is one of the technical results that was used without proof in Lecture 47, when we sketched an argument showing that $L(f) \neq 0$ for every map $f : \mathbb{CP}^n \to \mathbb{CP}^n$ if n is even (cf. Theorem 47.1).

Corollary 49.18. *Assume X is a space such that $H_*(X; \mathbb{Z})$ is finitely generated. Then for any map $f : X \to X$ and any field $\mathbb{K}$,*

$$L_{\mathbb{K}}(f) = \sum_{n=0}^{\infty} (-1)^n \operatorname{tr}\left(H^n(X; \mathbb{K}) \xrightarrow{f^*} H^n(X; \mathbb{K}) \right),$$

while for the special case $\mathbb{K} = \mathbb{Q}$,

$$L(f) = \sum_{n=0}^{\infty} (-1)^n \operatorname{tr}\left(H^n_{\text{free}}(X) \xrightarrow{f^*} H^n_{\text{free}}(X) \right),$$

where $H^n_{\text{free}}(X)$ denotes the free part of $H^n(X; \mathbb{Z})$.

PROOF. Under the isomorphism of $H^n(X;\mathbb{K})$ with the dual space of $H_n(X;\mathbb{K})$ given by Corollary 49.12, the map $f^* : H^n(X;\mathbb{K}) \to H^n(X;\mathbb{K})$ is the transpose of f_*, as we have

$$\langle f^*[\varphi], [c]\rangle = (f^*\varphi)(c) = \varphi(f_*c) = \langle[\varphi], f_*[c]\rangle$$

for all $[\varphi] \in H^n(X;\mathbb{K})$ and $[c] \in H_n(X;\mathbb{K})$. This implies the formula above for $L_{\mathbb{K}}(f)$ since every linear map has the same trace as its transpose. The formula for $L(f)$ follows similarly from the naturality statement in Corollary 49.15, which gives a commutative diagram

$$\begin{tikzcd}
H^n_{\text{free}}(X) \arrow[r, "\cong"] & \operatorname{Hom}(H_n^{\text{free}}(X), \mathbb{Z}) \\
H^n_{\text{free}}(X) \arrow[r, "\cong"] \arrow[u, "\varphi^*"] & \operatorname{Hom}(H_n^{\text{free}}(X), \mathbb{Z}) \arrow[u, "\varphi^*"]
\end{tikzcd}$$

and thus identifies $f^* : H^n_{\text{free}}(X) \to H^n_{\text{free}}(X)$ with the transpose of $f_* : H_n^{\text{free}}(X) \to H_n^{\text{free}}(X)$. ☐

We still have a loose end to tie up regarding cellular cohomology: the argument of the previous lecture only proves $H^*_{\text{CW}}(X;G) \cong H^*(X;G)$ when the CW-complex X is finite dimensional. Without this assumption, what it proves is that

$$H^n_{\text{CW}}(X;G) \cong H^n(X^N;G)$$

for every $n \geqslant 0$ and $N \geqslant n+1$. Thus it will suffice to prove:

LEMMA 49.19. *For any CW-complex X, the inclusion $X^{n+2} \hookrightarrow X$ induces an isomorphism $H^n(X;G) \to H^n(X^{n+2};G)$ for every integer $n \geqslant 0$.*

PROOF. The direct limit approach in Lecture 39 proves that the inclusion $X^{n+2} \hookrightarrow X$ induces isomorphisms $H_k(X^{n+2};\mathbb{Z}) \xrightarrow{\cong} H_k(X;\mathbb{Z})$ for every $k \leqslant n+1$, and the long exact sequence

$$\ldots \longrightarrow H_k(X^{n+2};\mathbb{Z}) \xrightarrow{\cong} H_k(X;\mathbb{Z}) \longrightarrow H_k(X, X^{n+2};\mathbb{Z}) \longrightarrow H_{k-1}(X^{n+2};\mathbb{Z}) \xrightarrow{\cong} H_{k-1}(X;\mathbb{Z}) \longrightarrow \ldots$$

then implies $H_k(X, X^{n+2};\mathbb{Z}) = 0$. Plugging the relative singular chain complex $C_*(X, X^{n+2};\mathbb{Z})$ into the universal coefficient theorem now gives a short exact sequence

$$0 \longrightarrow \operatorname{Ext}^{\mathbb{Z}}\left(H_{k-1}(X, X^{n+2};\mathbb{Z}), G\right) \longrightarrow H^k(X, X^{n+2};G) \longrightarrow \operatorname{Hom}\left(H_k(X, X^{n+2};\mathbb{Z}), G\right) \longrightarrow 0,$$

in which the first and last terms both vanish, therefore so does $H^k(X, X^{n+2};G)$. The long exact sequence of (X, X^{n+2}) in cohomology then has a segment of the form

$$0 = H^n(X, X^{n+2};G) \to H^n(X;G) \to H^n(X^{n+2};G) \to H^{n+1}(X, X^{n+2};G) = 0,$$

implying that $H^n(X;G) \to H^n(X^{n+2};G)$ is an isomorphism. ☐

COROLLARY 49.20. *The isomorphism $H^*_{\text{CW}}(X;G) \cong H^*(X;G)$ holds for all (not just finite-dimensional) CW-complexes X.* ☐

EXERCISE 49.21. Extend this discussion to prove that for all CW-pairs (X,A), the obvious inclusions of pairs induce isomorphisms

$$H_n(X^{n+1} \cup A, A;G) \xrightarrow{\cong} H_n(X,A;G),$$
$$H^n(X,A;G) \xrightarrow{\cong} H^n(X^{n+2} \cup A, A;G)$$

for all $n \geqslant 0$, and conclude from this that the isomorphisms $H_*^{\text{CW}}(X,A;G) \cong H_*(X,A;G)$ and $H^*_{\text{CW}}(X,A;G) \cong H^*(X,A;G)$ hold for all (possibly infinite-dimensional) CW-pairs (X,A).
Hint: For homology, you need to extend the direct limit discussion in Lecture 39 to accommodate direct limits in $\mathsf{Top}_{\text{rel}}$. You can then derive the cohomological statement from this by plugging the singular chain complex of $(X, X^{n+2} \cup A)$ into the universal coefficient theorem and using the exact sequence of the triple $(X, X^{n+1} \cup A, A)$.

50. The cup product (January 23, 2024)

In this lecture, we shall define the cup product

$$\cup : H^k(X;R) \otimes_R H^\ell(X;R) \to H^{k+\ell}(X;R)$$

on absolute singular cohomology with coefficients in a commutative ring R, thus turning $H^*(X;R)$ into a (not quite commutative) ring. The initial definition is based on the cross product, which we defined for homology in Lecture 45: we will first have to adapt that definition to define a cohomological version of the cross product, and then use the contravariance of H^* to derive the cup product from this and prove its basic properties. Once this definition is in place, we will see that there are also ways to characterize $\cup$ without mentioning $\times$, which will come in handy when we generalize everything to relative cohomology in the next lecture.

We assume throughout this lecture that R is a commutative ring with unit; it will not need to be a principal ideal domain. All chain complexes, homology and cohomology groups should be assumed to have R as coefficient group unless otherwise specified, so in particular, they have natural R-module structures.

Cross product. The cross product in singular homology was based on the existence and uniqueness (up to chain homotopy) of natural chain maps

$$\Phi_{(X,Y)} : C_*(X) \otimes C_*(Y) \to C_*(X \times Y). \tag{50.1}$$

Note that by our stated notational convention, the singular chain complexes on both sides of this expression use coefficients in the ring R, and $\otimes$ denotes a tensor product of R-modules, but one should keep in mind that both $\Phi_{(X,Y)}$ and the chain homotopies between different choices of $\Phi_{(X,Y)}$ were initially defined with $\mathbb{Z}$ coefficients; these determined corresponding R-module homomorphisms by applying to them the functor $\otimes R$. The same remark applies to the natural chain maps

$$\theta_{(X,Y)} : C_*(X \times Y) \to C_*(X) \otimes C_*(Y),$$

which we saw define chain homotopy inverses of $\Phi_{(X,Y)}$, as a consequence of uniqueness up to chain homotopy. The homological cross product $\times : H_*(X) \otimes H_*(Y) \to H_*(X \times Y)$ was defined as the composition of the map induced by $\Phi_{(X,Y)}$ on homology with the canonical map $H_*(X) \otimes H_*(Y) \to H_*\big(C_*(X) \otimes C_*(Y)\big)$, the latter being a purely algebraic construction that can be defined for any pair of chain complexes of R-modules. Exercise 45.7 outlined an argument via acyclic models to show that $\times$ is associative, and we saw in Exercise 45.8 that the canonical generator of $H_0(\{\mathrm{pt}\};R) = R$ acts as a multiplicative identity element.

The question of commutativity is a bit subtler: in the first place, the relation $A \times B = B \times A$ would not make sense in general since for $A \in H_*(X)$ and $B \in H_*(Y)$, $A \times B$ and $B \times A$ are homology classes on different spaces, i.e. $H_*(X \times Y)$ is strictly speaking not the same thing as $H_*(Y \times X)$. Of course we can choose the obvious homeomorphism

$$\tau : X \times Y \to Y \times X : (x,y) \mapsto (y,x),$$

and use it to identify these two homologies via the isomorphism

$$H_*(X \times Y) \xrightarrow{\tau_*} H_*(Y \times X),$$

but even then, things are not quite so simple. Instead of strict commutativity, we run into the Koszul sign convention (cf. Remark 31.20), which requires a sign change whenever the order of two elements with odd degree is interchanged:

Proposition 50.1. *For any $A \in H_k(X)$ and $B \in H_\ell(Y)$, $\tau_*(A \times B) = (-1)^{k\ell}(B \times A)$.*

PROOF. The general case will follow easily from the case $R := \mathbb{Z}$, so let's first assume the latter. Define $\Psi : C_*(X) \otimes C_*(Y) \to C_*(Y) \otimes C_*(X)$ by

$$\Psi(a \otimes b) := (-1)^{k\ell}(b \otimes a) \quad \text{for} \quad a \in C_k(X),\ b \in C_\ell(Y).$$

This is a chain map, since for $a \in C_k(X)$ and $b \in C_\ell(Y)$,

$$\Psi\partial(a \otimes b) = \Psi(\partial a \otimes b + (-1)^k a \otimes \partial b) = (-1)^{(k-1)\ell} b \otimes \partial a + (-1)^k(-1)^{k(\ell-1)}\partial b \otimes a,$$

which is the same as

$$\partial\Psi(a \otimes b) = (-1)^{k\ell}\partial(b \otimes a) = (-1)^{k\ell}(\partial b \otimes a + (-1)^\ell b \otimes \partial a).$$

It also satisfies a naturality property: if $f : X \to X'$ and $g : Y \to Y'$ are continuous maps, then we have a commutative diagram

$$\begin{tikzcd} C_*(X) \otimes C_*(Y) \arrow[r, "\Psi"] \arrow[d, "f_* \otimes g_*"] & C_*(Y) \otimes C_*(X) \arrow[d, "g_* \otimes f_*"] \\ C_*(X') \otimes C_*(Y') \arrow[r, "\Psi"] & C_*(Y) \otimes C_*(X). \end{tikzcd}$$

Similarly, $\tau_*^{-1} : C_*(Y \times X) \to C_*(X \times Y)$ is a natural chain map in the sense that for any f and g as above, the diagram

$$\begin{tikzcd} C_*(Y \times X) \arrow[r, "\tau_*^{-1}"] \arrow[d, "(g \times f)_*"] & C_*(X \times Y) \arrow[d, "(f \times g)_*"] \\ C_*(Y' \times X') \arrow[r, "\tau_*^{-1}"] & C_*(X' \times Y') \end{tikzcd}$$

also commutes. We can therefore compose three natural chain maps

$$\begin{tikzcd} C_*(X) \otimes C_*(Y) \arrow[r, "\Psi"] \arrow[rrr, bend right] & C_*(Y) \otimes C_*(X) \arrow[r, "\Phi_{(Y,X)}"] & C_*(Y \times X) \arrow[r, "\tau_*^{-1}"] & C_*(X \times Y), \end{tikzcd}$$

obtaining a natural chain map $C_*(X) \otimes C_*(Y) \to C_*(X \times Y)$, which acts on 0-chains in the obvious way $x \otimes y \mapsto (x, y)$ under the usual identification of singular 0-simplices with points. By the uniqueness statement for natural chain maps that we proved in Lemma 45.1, it follows that $\tau_*^{-1} \circ \Phi_{(Y,X)} \circ \Psi$ is chain homotopic to $\Phi_{(X,Y)} : C_*(X) \otimes C_*(Y) \to C_*(X \times Y)$, so these two chain maps induce the same map on homology, which proves $A \times B = (-1)^{k\ell}\tau_*^{-1}(B \times A)$ for $A \in H_k(X; \mathbb{Z})$ and $B \in H_\ell(Y; \mathbb{Z})$. The case of a general coefficient ring R is now obtained by applying $\otimes R$ to put all of these chain maps (and the related chain homotopies) into the category of R-modules. $\square$

EXERCISE 50.2. Adapt Proposition 50.1 to prove the analogous statement about the cellular cross product $\times : H_k^{\mathrm{CW}}(X) \otimes H_\ell^{\mathrm{CW}}(Y) \to H_{k+\ell}^{\mathrm{CW}}(X \times Y)$, where X and Y are assumed to be CW-complexes and $X \times Y$ and $Y \times X$ carry the resulting product cell decompositions.
Remark: The tricky part here is that one must compute the relevant incidence numbers for the cellular map $\tau : X \times Y \to Y \times X$; this is where the strange sign change will come from.

We would now like to define a similar product for the singular cohomology of two spaces: for each pair of integers $k, \ell \geqslant 0$, this should define a homomorphism

$$H^k(X) \otimes H^\ell(Y) \xrightarrow{\times} H^{k+\ell}(X \times Y) \tag{50.2}$$

that is dual to the homology cross product in the sense of the natural pairing of H^* with H_*. The first step is an easy algebraic observation: cohomology groups are just the homology groups of dualized chain complexes, which are in this case cochain complexes of R-modules, so in light of the canonical map

$$H_*(A_*) \otimes H_*(B_*) \to H_*\big(A_* \otimes B_*\big) : [a] \otimes [b] \mapsto [a \otimes b]$$

that exists for any pair of chain complexes of R-modules, plugging in $A_* := C^*(X) := \mathrm{Hom}(C_*(X), R)$ and $B_* := C^*(Y) := \mathrm{Hom}(C_*(Y), R)$ gives a canonical map

$$H^k(X) \otimes H^\ell(Y) \to H_{k+\ell}\big(C^*(X) \otimes C^*(Y)\big). \tag{50.3}$$

The complex at the right of this expression is to be understood as a *tensor product cochain complex*, meaning its degree n group is the direct sum of all $C^k(X) \otimes C^\ell(Y)$ for $k+\ell = n$, and its coboundary map is determined by the usual "graded Leibniz rule"

$$\delta(\varphi \otimes \psi) = \delta\varphi \otimes \psi + (-1)^k \varphi \otimes \delta\psi \quad \text{ for } \quad \varphi \in C^k(X),\ \psi \in C^\ell(Y).$$

The obvious way we should try to get from here to (50.2) is by finding a natural chain map $C^*(X) \otimes C^*(Y) \to C^*(X \times Y)$ and composing the induced map on homology with (50.3). A suitable map for this purpose is found by dualizing the chain homotopy inverse

$$\theta_{(X,Y)} : C_*(X \times Y) \to C_*(X) \otimes C_*(Y)$$

of $\Phi_{(X,Y)}$, which produces a chain map

$$\theta^*_{(X,Y)} : \mathrm{Hom}\big(C_*(X) \otimes C_*(Y), R\big) \to \mathrm{Hom}(C_*(X \times Y), R) = C^*(X \times Y),$$

and the desired chain map $C^*(X) \otimes C^*(Y) \to C^*(X \times Y)$ is then defined as the composition

$$\begin{tikzcd} C^*(X) \otimes C^*(Y) \arrow[r, "F"] \arrow[rr, bend right, "\times"'] & \mathrm{Hom}\big(C_*(X) \otimes C_*(Y), R\big) \arrow[r, "\theta^*_{(X,Y)}"] & C^*(X \times Y), \end{tikzcd} \tag{50.4}$$

where the map at the left arises canonically from the following purely algebraic construction:

EXERCISE 50.3. Show that for any two chain complexes of R-modules A_*, B_*, the canonical map

$$F : \mathrm{Hom}(A_*, R) \otimes \mathrm{Hom}(B_*, R) \to \mathrm{Hom}(A_* \otimes B_*, R)$$

defined for $\alpha \in \mathrm{Hom}(A_k, R)$ and $\beta \in \mathrm{Hom}(B_\ell, R)$ by

$$F(\alpha \otimes \beta)(a \otimes b) = \begin{cases} (-1)^{k\ell}\alpha(a)\beta(b) & \text{if } a \in A_k \text{ and } b \in B_\ell, \\ 0 & \text{otherwise} \end{cases}$$

is a chain map of cochain complexes.
Hint: To evaluate each of $\delta F(\alpha \otimes \beta)$ and $F\delta(\alpha \otimes \beta)$ on some element $a \otimes b$, you need to distinguish two cases, depending on the individual degrees of a and b. Getting all the signs right is a bit tricky.
Remark: The sign convention in the definition of F would need to change if we had defined coboundary operators simply by $\delta = \partial^*$, instead of with the extra sign that appears in (47.2).

In the following, we shall remove the symbol F from the notation and simply regard elements of $C^*(X) \otimes C^*(Y)$ as homomorphisms $C_*(X) \otimes C_*(Y) \to R$ via the formula that results from Exercise 50.3, namely

$$(\varphi \otimes \psi)(a \otimes b) := \begin{cases} (-1)^{|\psi||a|}\varphi(a)\psi(b) & \text{if } |\varphi| = |a| \text{ and } |\psi| = |b|, \\ 0 & \text{otherwise.} \end{cases}$$

The image of $\varphi \otimes \psi \in C^k(X) \otimes C^\ell(Y)$ under the map (50.4) will be denoted by $\varphi \times \psi \in C^{k+\ell}(X \times Y)$, which can now be written as a homomorphism $C_{k+\ell}(X \times Y) \to R$ in the form

$$\varphi \times \psi = (\varphi \otimes \psi) \circ \theta_{(X,Y)} : C_{k+\ell}(X \times Y) \to R,$$

and the fact that (50.4) is a chain map means that it satisfies the Leibniz rule

$$\delta(\varphi \times \psi) = \delta\varphi \times \psi + (-1)^{|\varphi|}\varphi \times \delta\psi.$$

It follows that $\times$ descends to the singular cohomology groups, and we can define an R-bilinear product

$$H^k(X) \otimes H^\ell(Y) \to H^{k+\ell}(X \times Y) : [\varphi] \otimes [\psi] \mapsto [\varphi] \times [\psi] := [\varphi \times \psi].$$

Note that the cochain-level cross product $\varphi \times \psi$ depends in general on an arbitrary choice, namely the chain map $\theta_{(X,Y)} : C_*(X \times Y) \to C_*(X) \otimes C_*(Y)$, but since the latter is unique up to chain homotopy, its dual $\theta^*_{(X,Y)}$ is similarly unique up to chain homotopy, and so therefore is the resulting chain map $C^*(X) \otimes C^*(Y) \to C^*(X \times Y)$, proving that $[\varphi] \times [\psi]$ does not depend on any choices.

REMARK 50.4. After Lecture 45, you may have been tempted to try constructing the chain map $C^*(X) \otimes C^*(Y) \to C^*(X \times Y)$ used above via the method of acyclic models. This idea runs into the following difficulty: a cochain group such as $C^k(X; \mathbb{Z}) = \mathrm{Hom}(C_k(X), \mathbb{Z})$ is not generally free, but is a direct *product* rather than a direct sum, thus it does not admit a basis. It is therefore quite unclear what models one might try to use for defining a map on $C^*(X) \otimes C^*(Y)$. The solution is to do the same thing we have done all along in our treatment of singular cohomology: reuse results that we've already proved about homology, but dualize them where appropriate.

PROPOSITION 50.5. *The cross product on singular cohomology with coefficients in a ring R has the following properties:*

(1) It is associative: $(\varphi \times \psi) \times \eta = \varphi \times (\psi \times \eta) \in H^*(X \times Y \times Z)$ for any $\varphi \in H^*(X)$, $\psi \in H^*(Y)$ and $\eta \in H^*(Z)$.
(2) It is graded commutative: $\varphi \times \psi = (-1)^{k\ell}\tau^*(\psi \times \varphi)$ for any $\varphi \in H^k(X)$, $\psi \in H^\ell(Y)$ and the homeomorphism $\tau : X \times Y \to Y \times X : (x, y) \mapsto (y, x)$.
(3) It is dual to the homology cross product: we have

$$\langle \varphi \times \psi, A \times B\rangle = (-1)^{k\ell}\langle \varphi, A\rangle\langle \psi, B\rangle \tag{50.5}$$

for all $\varphi \in H^k(X)$, $A \in H_k(X)$, $\psi \in H^\ell(Y)$ and $B \in H_\ell(Y)$.

PROOF. We shall leave the first two properties as exercises; they can be solved as in Exercise 45.7 and Proposition 50.1 using the uniqueness up to chain homotopy of certain natural chain maps. For associativity, for instance, it is possible to express each of $(\varphi \times \psi) \times \eta$ and $\varphi \times (\psi \times \eta)$ as the composition of $\varphi \otimes \psi \otimes \eta$ with some natural chain map.

Let's quickly check the third property: given cycles $a \in C_k(X)$, $b \in C_\ell(Y)$ and cocycles $\varphi \in C^k(X)$, $\psi \in C^\ell(Y)$, we have by definition

$$\begin{aligned}
\langle [\varphi] \times [\psi], [a] \times [b]\rangle &= \langle [\varphi \times \psi], [\Phi_{(X,Y)}(a \otimes b)]\rangle = (\varphi \times \psi)\big(\Phi_{(X,Y)}(a \otimes b)\big) \\
&= (\varphi \otimes \psi) \circ \theta_{(X,Y)} \circ \Phi_{(X,Y)}(a \otimes b) = (\varphi \otimes \psi) \circ (\mathbb{1} + \partial h + h\partial)(a \otimes b) \\
&= (\varphi \otimes \psi)(a \otimes b) = (-1)^{k\ell}\varphi(a)\psi(b) = (-1)^{k\ell}\langle [\varphi], [a]\rangle\langle [\psi], [b]\rangle,
\end{aligned}$$

where in the second line we've used the existence of a chain homotopy h between $\theta_{(X,Y)} \circ \Phi_{(X,Y)}$ and the identity map, plus the fact that $\partial(a \otimes b)$ and $(\varphi \otimes \psi) \circ \partial$ both vanish since $\partial a = \partial b = 0$ and $\delta\varphi = \delta\psi = 0$. □

EXERCISE 50.6. The cross product on cellular cohomology is defined in the same way as above, but with the simplifying feature that since there is a canonical isomorphism of chain complexes $\Phi : C^{\mathrm{CW}}_*(X) \otimes C^{\mathrm{CW}}_*(Y) \to C^{\mathrm{CW}}_*(X \times Y)$, the required map $\theta : C^{\mathrm{CW}}_*(X \times Y) \to C^{\mathrm{CW}}_*(X) \otimes C^{\mathrm{CW}}_*(Y)$ is simply its inverse, and is likewise canonical. Prove that the resulting product $H^*_{\mathrm{CW}}(X) \otimes H^*_{\mathrm{CW}}(Y) \to H^*_{\mathrm{CW}}(X \times Y)$ satisfies the same (or analogous) properties as in Proposition 50.5.

EXERCISE 50.7. Prove that the naturality formula

$$(f \times g)^*(\varphi \times \psi) = f^*\varphi \times g^*\psi$$

holds for all $\varphi \in H^*(X)$ and $\psi \in H^*(Y)$ given continuous maps $f : X' \to X$ and $g : Y' \to Y$.

EXERCISE 50.8. The goal of this problem is to show that the canonical generator $1 \in R = H^0(\{\mathrm{pt}\}; R)$ acts as a multiplicative identity element for the cross product on cohomology with coefficients in R.

(a) Suppose Ψ associates to every space X a chain map $\Psi : C_*(X; \mathbb{Z}) \to C_*(X; \mathbb{Z})$, and call this a **natural chain map** $C_*(X; \mathbb{Z}) \to C_*(X; \mathbb{Z})$ if it acts as the identity map on 0-chains and satisfies $\Psi \circ f_* = f_* \circ \Psi$ for every continuous map $f : X \to Y$. Use the method of acyclic models to show that any two choices of natural chain maps in this sense are chain homotopic for all X.

(b) Identify the chain complex $C_*(X \times \{\mathrm{pt}\}; \mathbb{Z})$ with $C_*(X; \mathbb{Z})$ via the obvious canonical isomorphism between them, and consider the following two maps:

$$C_*(X \times \{\mathrm{pt}\}; \mathbb{Z}) \xrightarrow{\theta} C_*(X; \mathbb{Z}) \otimes C_*(\{\mathrm{pt}\}; \mathbb{Z}) \xrightarrow{\mathbb{1}\otimes\epsilon} C_*(X; \mathbb{Z}) \otimes \mathbb{Z} = C_*(X; \mathbb{Z}),$$

$$C_*(X \times \{\mathrm{pt}\}; \mathbb{Z}) \xrightarrow{(\pi_X)_*} C_*(X; \mathbb{Z}),$$

where $\pi_X : X \times \{\mathrm{pt}\} \to X$ is the canonical projection, θ is any natural chain homotopy inverse for the natural chain map $\Phi : C_*(X; \mathbb{Z}) \otimes C_*(\{\mathrm{pt}\}; \mathbb{Z}) \to C_*(X \times \{\mathrm{pt}\}; \mathbb{Z})$ as used in the construction of the cross product, and $\epsilon : C_*(\{\mathrm{pt}\}; \mathbb{Z}) \to \mathbb{Z}$ is the *augmentation* map, which vanishes on $C_n(\{\mathrm{pt}\}; \mathbb{Z})$ for $n \neq 0$ and sends the unique generator $\sigma \in C_0(\{\mathrm{pt}\}; \mathbb{Z})$ to 1. Verify that both of these define natural chain maps, hence by part (a), they are chain homotopic.

(c) Deduce from part (b) that for any space X and any coefficient ring R, the cross product of $\varphi \in H^*(X)$ with $1 \in R \subset H^0(\{\mathrm{pt}\})$ satisfies $\varphi \times 1 = \pi_X^*\varphi$.

Cup product. Thus far the development of cohomology looks quite similar to that of homology, i.e. every theorem or construction for homology has had a cohomological analogue. But we can now introduce something in cohomology that has no homological analogue: it is possible due to the fact that cohomology is *contravariant*, so in particular, the diagonal map $d : X \to X \times X$ induces a natural map $H^*(X \times X) \to H^*(X)$.

DEFINITION 50.9. The **cup product** on $H^*(X) = H^*(X; R)$ is an R-bilinear map $\cup : H^k(X) \oplus H^\ell(X) \to H^{k+\ell}(X)$ for each $k, \ell \geqslant 0$ defined by

$$\varphi \cup \psi := d^*(\varphi \times \psi),$$

where $d : X \to X \times X$ is the diagonal map $x \mapsto (x, x)$.

On the cochain level, one can write $[\varphi] \cup [\psi] = [\varphi \cup \psi]$ if

$$\cup : C^k(X) \otimes C^\ell(X) \to C^{k+\ell}(X) : \varphi \otimes \psi \mapsto \varphi \cup \psi$$

is defined by

$$\varphi \cup \psi := d^*(\varphi \times \psi) = (\varphi \otimes \psi) \circ \theta_{(X,X)} \circ d_*. \tag{50.6}$$

This really is just the composition of our previous chain map $C^*(X) \otimes C^*(X) \to C^*(X \times X)$ with the chain map $d^* : C^*(X \times X) \to C^*(X)$ induced by the diagonal map, thus it is also a chain map and therefore satisfies the Leibniz rule

$$\delta(\varphi \cup \psi) = \delta\varphi \cup \psi + (-1)^{|\varphi|}\varphi \cup \delta\psi.$$

The presence of the mysterious object $\theta_{(X,X)} : C_*(X \times X) \to C_*(X) \otimes C_*(X)$ in (50.6) prevents it from being a useful formula on its own; we should also keep in mind that it is not uniquely defined since $\theta_{(X,X)}$ itself is not unique. We will see however that one can exploit this freedom to produce useful formulas.

DEFINITION 50.10. A **diagonal approximation** is an assignment to every space X of a chain map

$$\Psi_X : C_*(X;\mathbb{Z}) \to C_*(X;\mathbb{Z}) \otimes C_*(X;\mathbb{Z})$$

that is defined on the degree 0 level by $\Psi_X(x) = x \otimes x$ under the usual identification of singular 0-simplices $\Delta^0 \to X$ with points $x \in X$, and is natural in the sense that every continuous map $f : X \to Y$ gives rise to a commutative diagram

$$\begin{tikzcd} C_*(X;\mathbb{Z}) \arrow[r, "\Psi_X"] \arrow[d, "f_*"] & C_*(X;\mathbb{Z}) \otimes C_*(X;\mathbb{Z}) \arrow[d, "f_* \otimes f_*"] \\ C_*(Y;\mathbb{Z}) \arrow[r, "\Psi_Y"] & C_*(Y;\mathbb{Z}) \otimes C_*(Y;\mathbb{Z}). \end{tikzcd}$$

Applying the functor $\otimes R$ to any diagonal approximation $\Psi_X : C_*(X;\mathbb{Z}) \to C_*(X;\mathbb{Z}) \otimes C_*(X;\mathbb{Z})$ determines a similar chain map between chain complexes of R-modules, which we shall also denote simply by

$$\Psi_X : C_*(X) \to C_*(X) \otimes C_*(X)$$

when there is no need to be explicit about the coefficient ring R. We will sometimes abuse terminology by referring to chain maps of this form as diagonal approximations, but one should keep in mind that, technically, they are always assumed to be determined by natural chain maps defined over integer coefficients.

EXERCISE 50.11. Show that the maps $\theta_{(X,X)} \circ d_*$ appearing in (50.6) define a diagonal approximation.

EXERCISE 50.12. Prove via acyclic models that all diagonal approximations are chain homotopic, and deduce that for any two cocycles $\varphi, \psi \in C^*(X)$, the cohomology class of $(\varphi \otimes \psi) \circ \Psi_X \in C^*(X)$ is independent of the choice of diagonal approximation Ψ_X.

Exercise 50.12 reveals an alternative but equivalent definition of the cup product that we could have taken if we had wanted to leave the cross product out of the discussion: $\cup$ on cohomology is induced by any cochain-level product $\cup : C^k(X) \otimes C^\ell(X) \to C^{k+\ell}(X)$ of the form

$$\varphi \cup \psi := (\varphi \otimes \psi) \circ \Psi_X,$$

where Ψ_X is an arbitrary choice of diagonal approximation.

The most popular diagonal approximation in the literature is called the **Alexander-Whitney** diagonal approximation, and is defined as follows. Number the vertices of the standard n-simplex $\Delta^n \subset \mathbb{R}^{n+1}$ as $0, \dots, n$, and given any integers $0 \leqslant j_0 < j_1 < \dots < j_k \leqslant n$, let

$$[j_0, \dots, j_k] \subset \Delta^n$$

denote the k-simplex spanned by the vertices $j_0, \dots, j_k$, which is identified naturally with the standard k-simplex. For instance, in this notation, the jth boundary face of Δ^n is $\partial_{(j)}\Delta^n = [0, \dots, j-1, j+1, \dots, n]$ for each $j = 0, \dots, n$. Now define $\Psi_X : C_*(X;\mathbb{Z}) \to C_*(X;\mathbb{Z}) \otimes C_*(X;\mathbb{Z})$ on each singular n-simplex $\sigma : \Delta^n \to X$ by

$$\Psi_X(\sigma) := \sum_{k+\ell=n} \left(\sigma|_{[0,\dots,k]}\right) \otimes \left(\sigma|_{[k,\dots,n]}\right)$$

EXERCISE 50.13. Verify that Ψ_X as defined above is a diagonal approximation.

Plugging the Alexander-Whitney approximation into $\varphi \cup \psi = (\varphi \otimes \psi) \circ \Psi_X$ gives the following formula for the cup product of cochains: for any singular n-simplex $\sigma : \Delta^n \to X$ with $n = k + \ell$,[81]

$$(\varphi \cup \psi)(\sigma) = (-1)^{k\ell} \varphi(\sigma|_{[0,\ldots,k]}) \psi(\sigma|_{[k,\ldots,n]}).$$

On its own, this formula is seldom very useful since explicit computations with singular cochains are almost never practical. What is slightly more reasonable, however, is to use the same formula for computing the cup product in the *simplicial* cohomology of a simplicial complex, which of course is a special case of cellular cohomology and is therefore isomorphic to its singular cohomology. This trick is sometimes used for explicit computations of singular cohomology rings; see for instance [**Hat02**, Examples 3.7 and 3.8], or [**Bre93**, Example VI.4.6]. I will avoid computations like that in these notes, essentially for two reasons: first, they depend on a nontrivial fact we have not proved about the natural product structures on singular and simplicial cohomology being the same; second, they are ugly. We will see there are more elegant ways to carry out all the computations we need.

To that end, let us now establish some properties of the cup product that will be essential in further developments. To understand the second property below, we need to be aware that there is always a canonical inclusion

$$R \subset H^0(X) = H^0(X; R)$$

that makes the coefficient ring a submodule of $H^0(X)$: namely, each $r \in R$ is identified with the cohomology class of the cocycle $\varphi_r : C_0(X; \mathbb{Z}) \to R$ that has the value $\varphi_r(\sigma) = r$ on every singular 0-simplex $\sigma : \Delta^0 \to X$, regarded as a generator of $C_0(X; \mathbb{Z})$.

THEOREM 50.14. *The cup product $\cup : H^*(X) \otimes H^*(X) \to H^*(X)$ on cohomology with coefficients in the ring R has the following properties.*

1. *It is* natural*: for all continuous maps $f : Y \to X$ and $\varphi, \psi \in H^*(X)$, $f^*(\varphi \cup \psi) = f^*\varphi \cup f^*\psi$.*
2. *It has a* unit*: under the canonical inclusion $R \subset H^0(X)$, $1 \in R$ satisfies $1 \cup \varphi = \varphi \cup 1 = \varphi$ for all $\varphi \in H^*(X)$.*
3. *It is* associative*: $(\varphi \cup \psi) \cup \eta = \varphi \cup (\psi \cup \eta)$ for all $\varphi, \psi, \eta \in H^*(X)$.*
4. *It is* graded commutative*: $\varphi \cup \psi = (-1)^{k\ell} \psi \cup \varphi$ for all $\varphi \in H^k(X)$ and $\psi \in H^\ell(X)$.*
5. *It is related to the cross product by*
$$\varphi \times \psi = \pi_X^* \varphi \cup \pi_Y^* \psi \quad \textit{for} \quad \varphi \in H^*(X),\ \psi \in H^*(Y),$$
where $\pi_X : X \times Y \to X$ and $\pi_Y : X \times Y \to Y$ are the natural projections.

PROOF. Most of these properties are relatively straightforward exercises using some combination of the cross product properties proved earlier and acyclic model arguments. Let's go quickly through the list.

Naturality is an easy consequence of the similar formula $(f \times g)^*(\varphi \times \psi) = f^*\varphi \times g^*\psi$ from Exercise 50.7. The unit property follows similarly from Exercise 50.8, while associativity and graded commutativity are easy consequences of the corresponding properties for the cross product.

For the last property, we observe that for any cocycles $\varphi \in C^*(X)$, $\psi \in C^*(Y)$ and any choice of diagonal approximation $\Psi_{(X \times Y)} : C_*(X \times Y) \to C_*(X \times Y) \otimes C_*(X \times Y)$, the resulting expression for the cocycle $\pi_X^* \varphi \cup \pi_Y^* \psi : C_*(X \times Y) \to R$ is the composition

$$C_*(X \times Y) \xrightarrow{\Psi_{(X \times Y)}} C_*(X \times Y) \otimes C_*(X \times Y) \xrightarrow{(\pi_X)_* \otimes (\pi_Y)_*} C_*(X) \otimes C_*(Y) \xrightarrow{\varphi \otimes \psi} R.$$

[81] The formula we have derived here for the cochain $\varphi \cup \psi$ matches a formula in [**Bre93**] but differs from [**Hat02**] by a sign if k and ℓ are both odd. This is due to the sign convention in (47.2) for the definition of coboundary maps.

As can easily be checked, this is the composition of $\varphi \otimes \psi$ with a natural chain map $C_*(X \times Y) \to C_*(X) \otimes C_*(Y)$ that is defined in the canonical way on 0-chains, thus the chain map is chain homotopic to $\theta_{(X,Y)}$ and we end up with the usual formula for $\varphi \times \psi$. $\square$

The cohomology rings of tori. In order to access the full range of applications of cohomology, one often needs to compute not just the individual groups $H^k(X)$ for $k \geqslant 0$, but also the ring structure that $H^*(X)$ inherits from the cup product. We saw an example of this with $X = \mathbb{CP}^n$ in Theorem 47.1. At this point we do not yet have enough machinery to compute the ring $H^*(\mathbb{CP}^n)$, but the computation of cup products on $H^*(\mathbb{T}^n)$ for each $n \geqslant 0$ is possible based on the results established so far. Let's write down the details.[82]

The homology of $\mathbb{T}^n$ is fairly easy to compute because $\mathbb{T}^n = S^1 \times \ldots \times S^1$ has a natural structure as a product cell complex (see Exercise 44.7). Without mentioning cell complexes, we can also use an inductive argument based on the Künneth formula. Indeed, the case $n = 1$ is trivial since $\mathbb{T}^1 = S^1$, so in particular, $H_*(S^1;\mathbb{Z})$ is a finitely generated free abelian group. Let's call its canonical generators

$$[\mathrm{pt}] \in H_0(S^1;\mathbb{Z}), \qquad [S^1] \in H_1(S^1;\mathbb{Z}),$$

i.e. $[\mathrm{pt}]$ is the homology class represented by any singular 0-simplex $\Delta^0 \to S^1$, and $[S^1]$ is the class represented by the identity map $S^1 \to S^1$ under the isomorphism $H_1(S^1;\mathbb{Z}) \cong \pi_1(S^1)$. Now suppose we assume for a given integer $n \geqslant 2$ that $H_*(\mathbb{T}^q;\mathbb{Z})$ is finitely generated and free for every $q \leqslant n-1$. Then the Künneth formula gives for every $m \geqslant 0$ an isomorphism

$$\times : \bigoplus_{k+\ell=m} H_k(\mathbb{T}^{n-1};\mathbb{Z}) \otimes H_\ell(S^1;\mathbb{Z}) \xrightarrow{\cong} H_m(\mathbb{T}^n;\mathbb{Z})$$

since all Tor terms vanish, while $H_\ell(S^1;\mathbb{Z})$ on the left hand side is only nontrivial for $\ell = 0, 1$ and is then $\mathbb{Z}$, giving

$$H_m(\mathbb{T}^n;\mathbb{Z}) \cong H_{m-1}(\mathbb{T}^{n-1};\mathbb{Z}) \oplus H_m(\mathbb{T}^{n-1};\mathbb{Z}).$$

This proves that $H_m(\mathbb{T}^n;\mathbb{Z})$ is also a finitely-generated free abelian group, and its rank is an entry in Pascal's triangle,

$$\operatorname{rank} H_m(\mathbb{T}^n;\mathbb{Z}) = \binom{n}{m}.$$

Moreover, the cross product provides a canonical set of generators of $H_m(\mathbb{T}^n;\mathbb{Z})$: for each choice of integers $1 \leqslant j_1 < \ldots < j_m \leqslant n$, we define

$$e_{j_1,\ldots,j_m} := A_1 \times \ldots \times A_n \in H_m(\mathbb{T}^n;\mathbb{Z})$$

by setting $A_{j_i} := [S^1]$ for each $i = 1, \ldots, m$ and $A_j := [\mathrm{pt}]$ for all other $j = 1, \ldots, n$. So far so good.

It will be useful to have an alternative description of the degree 1 generators $e_j \in H_1(\mathbb{T}^n;\mathbb{Z})$. Pick a base point $t_0 \in S^1$ and consider the embedding

$$i_j : S^1 \hookrightarrow \mathbb{T}^n : x \mapsto (\underbrace{t_0 \times \ldots \times t_0}_{j-1}, x, \underbrace{t_0 \times \ldots \times t_0}_{n-j}). \tag{50.7}$$

Note that different choices of the base point $t_0 \in S^1$ give homotopic maps $i_j : S^1 \to \mathbb{T}^n$, thus the induced map $(i_j)_* : H_*(S^1;\mathbb{Z}) \to H_*(\mathbb{T}^n;\mathbb{Z})$ is independent of this choice.

LEMMA 50.15. *For each $j = 1, \ldots, n$, $(i_j)_*[S^1] = e_j$.*

[82] The contents of this section were not covered in any lecture, but were instead outsourced to a problem session.

Proof. Consider first the case $j = n$. Under the obvious identification of S^1 with $\{\mathrm{pt}\} \times S^1$, we can then write $i_n = c \times \mathrm{Id} : \{\mathrm{pt}\} \times S^1 \to \mathbb{T}^{n-1} \times S^1$, where $c : \{\mathrm{pt}\} \to \mathbb{T}^{n-1}$ denotes the constant map with value $(t_0, \ldots, t_0)$. The naturality of the cross product now gives a commutative diagram

$$\begin{tikzcd} H_0(\{\mathrm{pt}\}; \mathbb{Z}) \otimes H_1(S^1; \mathbb{Z}) \arrow[r, "\times"] \arrow[d, "c_* \otimes \mathbb{1}"] & H_1(S^1; \mathbb{Z}) \arrow[d, "(i_n)_*"] \\ H_0(\mathbb{T}^{n-1}; \mathbb{Z}) \otimes H_1(S^1; \mathbb{Z}) \arrow[r, "\times"] & H_1(\mathbb{T}^n; \mathbb{Z}). \end{tikzcd}$$

In light of Exercise 45.8, this proves $(i_n)_*[S^1] = (i_n)_*\left([\mathrm{pt}] \times [S^1]\right) = c_*[\mathrm{pt}] \times [S^1] = [\mathrm{pt}] \times [S^1] = e_n$ since $c_* : H_0(\{\mathrm{pt}\}; \mathbb{Z}) \to H_0(\mathbb{T}^{n-1}; \mathbb{Z})$ is an isomorphism relating the canonical generators $[\mathrm{pt}]$. The general case $j \in \{1, \ldots, n\}$ follows from this same argument after permuting the coordinates. □

Before we turn to cohomology, we note that the computation of $H_*(\mathbb{T}^n; \mathbb{Z})$ is easily upgraded to a computation of $H_*(\mathbb{T}^n; G)$ for any coefficient group G, since the freeness of $H_*(\mathbb{T}^n; \mathbb{Z})$ forces all Tor terms to vanish, and the universal coefficient theorem thus gives a natural isomorphism

$$H_*(\mathbb{T}^n; \mathbb{Z}) \otimes G \xrightarrow{\cong} H_*(\mathbb{T}^n; G) : [c] \otimes g \mapsto [c \otimes g].$$

If G is chosen to be a commutative ring R with unit element $1 \in R$, then $H_m(\mathbb{T}^n; R)$ is now a free R-module with a basis that is in bijective correspondence with the basis of $H_m(\mathbb{T}^n; \mathbb{Z})$ described above, so we shall abuse notation and denote by

$$e_{j_1,\ldots,j_m} \in H_m(\mathbb{T}^n; R)$$

the generator obtained from $e_{j_1,\ldots,j_m} \in H_m(\mathbb{T}^n; \mathbb{Z})$ by feeding $e_{j_1,\ldots,j_m} \otimes 1 \in H_m(\mathbb{T}^n; \mathbb{Z}) \otimes R$ into the canonical isomorphism to $H_m(\mathbb{T}^n; R)$. In particular, one obtains in this manner canonical generators

$$[\mathrm{pt}] \in H_0(S^1; R), \qquad [S^1] \in H_1(S^1; R),$$

whose cross products in homology with coefficients in R produce the basis elements of $H_*(\mathbb{T}^n; R)$.

Henceforth, we fix R as the coefficient ring for both $H_*(\mathbb{T}^n)$ and $H^*(\mathbb{T}^n)$ and omit it from the notation wherever possible. The computation of $H^m(\mathbb{T}^n)$ for each m is a similarly easy application of the universal coefficient theorem for cohomology: the Ext terms vanish since $H_*(\mathbb{T}^n)$ is a free R-module, implying that the canonical map

$$H^m(\mathbb{T}^n) \to \mathrm{Hom}(H_m(\mathbb{T}^n), R) : \varphi \mapsto \langle \varphi, \cdot \rangle$$

is an isomorphism. Since $H_m(\mathbb{T}^n) \cong R^{\binom{n}{m}}$, this means $H^m(\mathbb{T}^n) \cong R^{\binom{n}{m}}$, and we can write down a canonical set of generators as follows. For $n = 1$, define

$$\lambda \in H^1(S^1)$$

to be the unique cohomology class such that

$$\langle \lambda, [S^1] \rangle = 1 \in R.$$

Now for each choice of integers $1 \leqslant j_1 < \ldots < j_m \leqslant n$, define

$$\lambda_{j_1,\ldots,j_m} := \alpha_1 \times \ldots \times \alpha_n \in H^m(\mathbb{T}^n),$$

where we choose $\alpha_{j_i} := \lambda$ for each $i = 1, \ldots, m$ and $\alpha_j = 1 \in H^0(S^1)$ for all other $j = 1, \ldots, n$. By (50.5), we have

$$\begin{aligned} \langle \lambda_{j_1,\ldots,j_m}, e_{k_1,\ldots,k_m} \rangle &= \langle \alpha_1 \times \ldots \times \alpha_n, A_1 \times \ldots \times A_n \rangle = \pm \langle \alpha_1, A_1 \rangle \ldots \langle \alpha_n, A_n \rangle \\ &= \begin{cases} \pm 1 & \text{if } j_i = k_i \text{ for all } i = 1, \ldots, m, \\ 0 & \text{otherwise,} \end{cases} \end{aligned}$$

proving that the collection of classes $\lambda_{j_1,\dots,j_m}$ for all choices $1 \leqslant j_1 < \dots < j_m \leqslant n$ is a basis for $H^*(\mathbb{T}^n)$ as a free R-module.

To describe $H^*(\mathbb{T}^n)$ as a ring, we now need to compute each product of the form $\lambda_{j_1,\dots,j_m} \cup \lambda_{k_1,\dots,k_q} \in H^{m+q}(\mathbb{T}^n)$. We start with an observation about the 1-dimensional classes $\lambda_j \in H^1(\mathbb{T}^n)$. Consider for each $j = 1,\dots,n$ the projection map

$$\pi_j : \mathbb{T}^n \to S^1 : (x_1,\dots,x_n) \mapsto x_j,$$

which is related to the inclusions $i_j : S^1 \hookrightarrow \mathbb{T}^n$ defined in (50.7) above by

$$\pi_j \circ i_k = \begin{cases} \mathrm{Id} : S^1 \to S^1 & \text{if } j = k, \\ \text{constant} & \text{if } j \neq k. \end{cases}$$

Lemma 50.16. *We have $\pi_j^* \lambda = \lambda_j$ for each $j = 1,\dots,n$.*

Proof. In light of Lemma 50.15 and the isomorphism $H^1(\mathbb{T}^n) \cong \mathrm{Hom}(H_1(\mathbb{T}^n), R)$, the classes $\pi_j^*\lambda \in H^1(\mathbb{T}^n)$ are characterized by

$$\langle \pi_j^*\lambda, e_k\rangle = \langle \lambda, (\pi_j)_*(i_k)_*[S^1]\rangle = \langle \lambda, (\pi_j \circ i_k)_*[S^1]\rangle = \begin{cases} 1 & \text{if } j = k, \\ 0 & \text{if } j \neq k. \end{cases}$$

At the same time, (50.5) implies

$$\langle \lambda_j, e_k\rangle = \begin{cases} 1 & \text{if } j = k, \\ 0 & \text{if } j \neq k, \end{cases}$$

and the result follows. □

We are now in a position to compute $\lambda_{j_1} \cup \dots \cup \lambda_{j_m} \in H^m(\mathbb{T}^n)$ for any set of integers $1 \leqslant j_1 < \dots < j_m \leqslant n$. Indeed, writing $\alpha_{j_i} = \lambda \in H^1(S^1)$ for each $i = 1,\dots,m$ and $\alpha_j = 1 \in H^0(S^1)$ for all other j, we have $\pi_j^*\alpha_j = \lambda_j \in H^1(\mathbb{T}^n)$ in the first case and $\pi_j^*\alpha_j = \pi_j^* 1 = 1 \in H^0(\mathbb{T}^n)$ in the second case, thus by Theorem 50.14(5),

$$\lambda_{j_1} \cup \dots \cup \lambda_{j_m} = \pi_1^*\alpha_1 \cup \dots \cup \pi_n^*\alpha_n = \alpha_1 \times \dots \times \alpha_n = \lambda_{j_1,\dots,j_m}.$$

This means that all of our basis elements for $H^*(\mathbb{T}^n)$ can be obtained as cup products of the degree 1 elements $\lambda_1,\dots,\lambda_n$, and moreover, this relation fully determines all cup products in $H^*(\mathbb{T}^n)$; indeed, graded commutativity implies

$$\lambda_i \cup \lambda_j = -\lambda_j \cup \lambda_i$$

for all i and j, so in particular $\lambda_i \cup \lambda_i$ always vanishes, and all other products of basis elements $\lambda_{j_i,\dots,j_m}$ can be derived from this via associativity. We've proved:

Theorem 50.17. *For any commutative ring R with unit and any $n \in \mathbb{N}$, the ring $H^*(\mathbb{T}^n; R)$ is isomorphic to the exterior algebra $\Lambda_R[\lambda_1,\dots,\lambda_n]$ on n generators of degree 1, where the generators $\lambda_j \in H^1(\mathbb{T}^n; R)$ can be defined in terms of the projections $\pi_j : \mathbb{T}^n \to S^1$ and the canonical generator $\lambda \in H^1(S^1; R)$ by $\lambda_j = \pi_j^*\lambda$.* □

We will be able to compute more examples of cohomology rings after we discuss Poincaré duality, which provides the most useful geometric interpretation of the cup product in terms of intersections.

51. Relative cross, cup, and cap products (January 26, 2024)

Relative cellular cross products. Our whole discussion of products so far has focused on absolute homology and cohomology, so you may be wondering how it extends to pairs of spaces (X, A) with $A \neq \varnothing$. In singular homology, the answer to this question turns out to be surprisingly subtle, but one gets an important hint about what to do if one starts by asking the same question about cellular homology, where the answer is much easier.

Relative chain complexes are quotient complexes, so let's start with an algebraic observation. If A and B are modules over a fixed ring, and $A_0 \subset A$ and $B_0 \subset B$ are submodules, then the homomorphism

$$A \otimes B \xrightarrow{\pi_A \otimes \pi_B} \frac{A}{A_0} \otimes \frac{B}{B_0}$$

determined by the quotient projections $\pi_A : A \to A/A_0$ and $\pi_B : B \to B/B_0$ descends to a natural isomorphism

$$\frac{A \otimes B}{(A_0 \otimes B) + (A \otimes B_0)} \xrightarrow{\cong} \frac{A}{A_0} \otimes \frac{B}{B_0}.$$

Indeed, its inverse is obtained by noticing that if we compose the canonical bilinear map $A \oplus B \to A \otimes B$ with the quotient projection, it descends to a bilinear map $(A/A_0) \otimes (B/B_0) \to (A \otimes B)/\left((A_0 \otimes B) + (A \otimes B_0)\right)$.

If (X, A) and (Y, B) are CW-pairs, then applying the algebraic observation above to their relative cellular chain complexes with coefficients in a fixed ring R gives a natural isomorphism

$$C_*^{\mathrm{CW}}(X, A) \otimes C_*^{\mathrm{CW}}(Y, B) \cong \frac{C_*^{\mathrm{CW}}(X) \otimes C_*^{\mathrm{CW}}(Y)}{(C_*^{\mathrm{CW}}(A) \otimes C_*^{\mathrm{CW}}(Y)) + (C_*^{\mathrm{CW}}(X) \otimes C_*^{\mathrm{CW}}(B))}.$$

As it happens, the denominator on the right hand side is not only a subcomplex of the chain complex $C_*^{\mathrm{CW}}(X) \otimes C_*^{\mathrm{CW}}(Y)$, but its image under the isomorphism $\times : C_*^{\mathrm{CW}}(X) \otimes C_*^{\mathrm{CW}}(Y) \to C_*^{\mathrm{CW}}(X \times Y)$ is the cellular chain complex of the subset

$$(A \times Y) \cup (X \times B) \subset X \times Y,$$

which is a subcomplex of the CW-complex $X \times Y$. It follows that $\times$ descends to an isomorphism of chain complexes

$$C_*^{\mathrm{CW}}(X, A) \otimes C_*^{\mathrm{CW}}(Y, B) \xrightarrow{\times} C_*^{\mathrm{CW}}((X, A) \times (Y, B))$$

if we define the product of two CW-pairs to be

$$(X, A) \times (Y, B) := (X \times Y, (A \times Y) \cup (X \times B)). \tag{51.1}$$

The most general version of the cross product on relative cellular homology thus takes the form

$$H_k^{\mathrm{CW}}(X, A) \otimes H_\ell^{\mathrm{CW}}(Y, B) \xrightarrow{\times} H_{k+\ell}^{\mathrm{CW}}((X, A) \times (Y, B)),$$

and the Künneth formula (in the case where R is a principal ideal domain) then becomes

$$0 \longrightarrow \bigoplus_{k+\ell=n} H_k^{\mathrm{CW}}(X, A; R) \otimes_R H_\ell^{\mathrm{CW}}(Y, B; R) \xrightarrow{\times} H_{k+\ell}^{\mathrm{CW}}((X, A) \times (Y, B); R)$$
$$\longrightarrow \bigoplus_{k+\ell=n-1} \mathrm{Tor}^R(H_k^{\mathrm{CW}}(X, A; R), H_\ell^{\mathrm{CW}}(Y, B; R)) \longrightarrow 0.$$

Relative singular cross products. Adapting this discussion for singular homology is slightly nontrivial, and it does not completely work for arbitrary pairs (X, A) and (Y, B), but it will work for most pairs that we are actually interested in. We shall adopt (51.1) as a definition of the product of two objects in the category $\mathsf{Top}_{\mathrm{rel}}$.[83] Let $\Phi : C_*(X) \otimes C_*(Y) \to C_*(X \times Y)$ denote a choice of natural chain map associated to any two spaces X, Y, as required for the definition of the cross product on absolute singular homology. Applying naturality to the inclusions $A \hookrightarrow X$ and $B \hookrightarrow Y$, we see that Φ maps $C_*(A) \otimes C_*(Y)$ into $C_*(A \times Y)$ and $C_*(X) \otimes C_*(B)$ into $C_*(X \times B)$, thus it descends to a natural chain map

$$\Phi : C_*(X, A) \otimes C_*(Y, B) \to C_*((X, A) \times (Y, B)),$$

so that the cross product on relative homology is well defined:

$$\times : H_k(X, A) \otimes H_\ell(Y, B) \to H_{k+\ell}((X, A) \times (Y, B)).$$

This part of the story works without any further conditions.

We run into a complication, however, if we want either to define the cross product on relative *cohomology* or to prove a relative Künneth formula. Both require the chain homotopy inverse $\theta : C_*(X \times Y) \to C_*(X) \otimes C_*(Y)$ of Φ, and this does *not* always descend to a map

$$C_*((X, A) \times (Y, B)) \to C_*(X, A) \otimes C_*(Y, B).$$

The problem is that if we are given a chain in the subspace $(A \times Y) \cup (X \times B)$, there is generally no reason to expect that θ will send it into $(C_*(A) \otimes C_*(Y)) + (C_*(X) \otimes C_*(B))$. What we can immediately say instead is that Φ and θ descend to chain homotopy inverses between the two quotient complexes

$$\begin{tikzcd} \dfrac{C_*(X) \otimes C_*(Y)}{(C_*(A) \otimes C_*(Y)) + (C_*(X) \otimes C_*(B))} \arrow[r, bend left, "\Phi"] & \dfrac{C_*(X \times Y)}{C_*(A \times Y) + C_*(X \times B)} \arrow[l, bend left, "\theta"] \end{tikzcd}. \tag{51.2}$$

The complex at the left is just $C_*(X, A) \otimes C_*(Y, B)$, which is what we want, but the one at the right is not the same as $C_*((X, A) \times (Y, B))$. However, the identity map does descend to a natural chain map

$$\frac{C_*(X \times Y)}{C_*(A \times Y) + C_*(X \times B)} \longrightarrow \frac{C_*(X \times Y)}{C_*\big((A \times Y) \cup (X \times B)\big)} = C_*((X, A) \times (Y, B)),$$

and it will happen sometimes that this map is a chain homotopy equivalence, so that Φ descends after all to an isomorphism from the homology of $C_*(X, A) \otimes C_*(Y, B)$ to $H_*((X, A) \times (Y, B))$. The situation should remind you of the technical hurdles we had to overcome in order to prove excision or define the Mayer-Vietoris sequence, and it can be dealt with in a similar way.

Definition 51.1. Given a space X, two subspaces $X_1, X_2 \subset X$ are called an **excisive couple** if the natural chain map defined by the inclusion

$$C_*(X_1; \mathbb{Z}) + C_*(X_2; \mathbb{Z}) \hookrightarrow C_*(X_1 \cup X_2; \mathbb{Z})$$

[83] This definition of $(X, A) \times (Y, B)$ is the right one for talking about the cross product and Künneth's formula, but it is not so good for other purposes: in particular it suffers from the fact that the obvious projection maps from $X \times Y$ to X or Y do not generally define morphisms from $(X, A) \times (Y, B)$ to (X, A) or (Y, B) under this definition. (Think about it.) There is a more obvious alternative definition of the product for an arbitrary collection of pairs of spaces which does not have this problem with projection maps. That is the right definition to use if, say, one wants an explicit description of inverse limits in $\mathsf{Top}_{\mathrm{rel}}$, in the spirit of Exercise 46.7.

induces isomorphisms on the homology groups of these complexes.[84]

LEMMA 51.2. *Two subspaces $X_1, X_2 \subset X$ form an excisive couple if and only if the canonical chain map*

$$\frac{C_*(X;\mathbb{Z})}{C_*(X_1;\mathbb{Z}) + C_*(X_2;\mathbb{Z})} \longrightarrow C_*(X, X_1 \cup X_2;\mathbb{Z})$$

descends to an isomorphism on the homology groups.

PROOF. Assume either that X_1 and X_2 form an excisive couple or that the map of quotient complexes induces isomorphisms on homology. There is a commutative diagram

$$\begin{tikzcd}
0 \arrow[r] & C_*(X_1;\mathbb{Z}) + C_*(X_2;\mathbb{Z}) \arrow[r, hook] \arrow[d, hook] & C_*(X;\mathbb{Z}) \arrow[r] \arrow[d, "\mathbb{1}"] & \frac{C_*(X;\mathbb{Z})}{C_*(X_1;\mathbb{Z})+C_*(X_2;\mathbb{Z})} \arrow[r] \arrow[d] & 0 \\
0 \arrow[r] & C_*(X_1 \cup X_2;\mathbb{Z}) \arrow[r, hook] & C_*(X;\mathbb{Z}) \arrow[r] & C_*(X, X_1 \cup X_2;\mathbb{Z}) \arrow[r] & 0,
\end{tikzcd}$$

where both rows are short exact sequences of chain complexes and all arrows represent chain maps induced by either inclusions or quotient projections. Transforming both rows into long exact sequences of homology groups then produces a diagram in which two out of every three vertical maps are isomorphisms, so the five-lemma implies that the third one is as well. □

EXERCISE 51.3. Show that if $X_1, X_2 \subset X$ are an excisive couple, then the relevant induced maps on homology or cohomology with an arbitrary coefficient group are also isomorphisms.
Hint: Use the naturality of the universal coefficient theorems.

The lemma implies that all important results regarding products $(X, A) \times (Y, B)$ in homology or cohomology will hold as long as the two subsets $A \times Y$ and $X \times B$ in $X \times Y$ form an excisive couple.

LEMMA 51.4. *Given two pairs of spaces (X, A) and (Y, B), the subsets $A \times Y$ and $X \times B$ in $X \times Y$ form an excisive couple whenever any of the following conditions holds:*

(1) $A \subset X$ and $B \subset Y$ are both open subsets;
(2) $A = \varnothing$ or $B = \varnothing$;
(3) (X, A) and (Y, B) are both CW-pairs.

PROOF. The first case follows by barycentric subdivision, as we showed when we proved the excision axiom (see Lemma 24.12 from last semester). The second case is trivial. The third can be proven by replacing singular with cellular chain complexes and appealing to the isomorphism of singular with cellular homology: the necessary condition is obvious on the cellular chain complex since $C_*^{\mathrm{CW}}(X_1)+C_*^{\mathrm{CW}}(X_2)$ and $C_*^{\mathrm{CW}}(X_1 \cup X_2)$ are exactly the same whenever (X, X_1) and (X, X_2) are CW-pairs. □

Now using the quotient complex at the right hand side of (51.2) as a stand-in for $C_*((X, A) \times (Y, B))$, we obtain a relative version of the Eilenberg-Zilber theorem and therefore a relative Künneth formula:

[84] Notice what this definition does not say: it would seem natural to ask for this inclusion to be a chain homotopy equivalence, but we are requiring something slightly weaker. We will in fact want to use the weaker conditions in some situations where the stronger one might not hold; an example is the third case of Lemma 51.4.

THEOREM 51.5. *If R is a principal ideal domain and (X, A) and (Y, B) are pairs such that the subsets $A \times Y$ and $X \times B$ in $X \times Y$ form an excisive couple, then there is a natural short exact sequence*

$$0 \longrightarrow \bigoplus_{k+\ell=n} H_k(X, A; R) \otimes_R H_\ell(Y, B; R) \xrightarrow{\times} H_{k+\ell}((X, A) \times (Y, B); R)$$
$$\longrightarrow \bigoplus_{k+\ell=n-1} \operatorname{Tor}^R(H_k(X, A; R), H_\ell(Y, B; R)) \longrightarrow 0,$$

and the sequence splits. □

EXERCISE 51.6. Using Exercise 51.3 to identify $H^*((X, A) \times (Y, B))$ with the cohomology of the quotient complex on the right hand side of (51.2), write down a chain-level definition of the cross product on relative singular cohomology,

$$\times : H^k(X, A) \otimes H^\ell(Y, B) \to H^{k+\ell}((X, A) \times (Y, B)),$$

under the assumption that $A \times Y, X \times B \subset X \times Y$ form an excisive couple.

Here is an interesting application of the relative Künneth formula. If (X, x_0) and (Y, y_0) are two pointed spaces, their **smash product** $X \wedge Y$ is defined as the quotient space

$$X \wedge Y := (X \times Y) / ((\{x_0\} \times Y) \cup (X \times \{y_0\})).$$

Strictly speaking, this construction depends on the choice of base points, but we shall suppress this in the notation. Notice that the subset being quotiented out is homeomorphic to the wedge sum $X \vee Y$, so it is sensible to write

$$X \wedge Y = (X \times Y)/(X \vee Y).$$

It is now straightforward to check that for any base-point preserving continuous maps $f : (X, x_0) \to (X', x_0')$ and $g : (Y, y_0) \to (Y', y_0')$, the product map $f \times g : X \times Y \to X' \times Y'$ descends to the quotient as a continuous map

$$f \wedge g : X \wedge Y \to X' \wedge Y'.$$

EXAMPLE 51.7. For any integers $k, \ell \geqslant 0$, $S^k \wedge S^\ell \cong S^{k+\ell}$. This is obvious if either k or ℓ is 0, and otherwise, we can identify S^n with $\mathbb{D}^n/\partial\mathbb{D}^n$ for every $n \in \mathbb{N}$ and choose the equivalence class of the boundary to be the base point. The claim then follows easily from the fact that there is a homeomorphism $\mathbb{D}^{k+\ell} \cong \mathbb{D}^k \times \mathbb{D}^\ell$ identifying $\partial\mathbb{D}^{k+\ell}$ with $(\partial\mathbb{D}^k \times \mathbb{D}^\ell) \cup (\mathbb{D}^k \times \partial\mathbb{D}^\ell)$.

Now assume X and Y are both CW-complexes, with base points chosen to be 0-cells in their cell decompositions, so by Lemma 51.4, the Künneth formula is valid for the pairs $(X, \{x_0\})$ and $(Y, \{y_0\})$. Since $(X, \{x_0\}) \times (Y, \{y_0\}) = (X \times Y, X \vee Y)$, the Künneth formula now takes the form

$$0 \to \bigoplus_{k+\ell=n} H_k(X, \{x_0\}) \otimes H_\ell(Y, \{y_0\}) \xrightarrow{\times} H_n(X \times Y, X \vee Y)$$
$$\longrightarrow \bigoplus_{k+\ell=n-1} \operatorname{Tor}(H_k(X, \{x_0\}), H_\ell(Y, \{y_0\})) \to 0,$$

or under the natural isomorphisms $H_*(X, A) = \widetilde{H}_*(X/A)$ for good pairs,

$$0 \to \bigoplus_{k+\ell=n} \widetilde{H}_k(X) \otimes \widetilde{H}_\ell(Y) \xrightarrow{\times} \widetilde{H}_n(X \wedge Y) \longrightarrow \bigoplus_{k+\ell=n-1} \operatorname{Tor}(\widetilde{H}_k(X), \widetilde{H}_\ell(Y)) \to 0.$$

EXERCISE 51.8. Show that for the cross product on reduced homology as described above and the identification of $S^k \wedge S^\ell$ with $S^{k+\ell}$ as indicated in Example 51.7, if $[S^k] \in \widetilde{H}_k(S^k)$ and $[S^\ell] \in \widetilde{H}_\ell(S^\ell)$ are generators, then $[S^k] \times [S^\ell] \in \widetilde{H}_{k+\ell}(S^{k+\ell})$ is also a generator.

EXERCISE 51.9. Suppose $f : S^k \to S^k$ and $g : S^\ell \to S^\ell$ are base-point preserving maps.

(a) Use the naturality of the Künneth formula to prove $\deg(f \wedge g) = \deg(f) \cdot \deg(g)$.
(b) Find an alternative proof of $\deg(f \wedge g) = \deg(f) \cdot \deg(g)$ using the following fact from differential topology: any continuous map $f : S^k \to S^k$ admits a small perturbation to a smooth map such that for almost every point $x \in S^k$, $f^{-1}(x)$ is a finite set of points at which the local degree of f is ± 1. (This follows from Sard's theorem.)
(c) Using the definition of cellular chain maps and the cellular cross product, prove that the cellular cross product is natural, i.e. if $f : X \to X'$ and $g : Y \to Y'$ are cellular maps, then the diagram

$$\begin{tikzcd} C_*^{\mathrm{CW}}(X) \otimes C_*^{\mathrm{CW}}(Y) \arrow[r, "\times"] \arrow[d, "f_* \otimes g_*"] & C_*^{\mathrm{CW}}(X \times Y) \arrow[d, "(f \times g)_*"] \\ C_*^{\mathrm{CW}}(X') \otimes C_*^{\mathrm{CW}}(Y') \arrow[r, "\times"] & C_*^{\mathrm{CW}}(X' \times Y') \end{tikzcd}$$

commutes.

With significantly more effort, one can proceed from Exercise 51.9 to a proof that the cellular cross product matches the cross product on singular homology under the natural isomorphisms $H_*^{\mathrm{CW}}(X; R) \cong H_*(X; R)$ for all CW-complexes X. We will not go into this since we do not intend to use the cellular cross product for anything beyond intuition, but the basic idea (by reducing to the case of wedges of spheres and then computing both explicitly in that case) is outlined in a slightly different context in [**Hat02**, p. 279].

The relative cup product. Recall that the cup product can be defined in terms of any map on cochains $C^k(X) \otimes C^\ell(X) \to C^{k+\ell}(X) : \varphi \otimes \psi \mapsto \varphi \cup \psi$ taking the form

$$\varphi \cup \psi = (\varphi \otimes \psi) \circ \Psi : C_*(X) \to R, \tag{51.3}$$

where Ψ is a choice of diagonal approximation, meaning a natural chain map $C_*(X; \mathbb{Z}) \to C_*(X; \mathbb{Z}) \otimes C_*(X; \mathbb{Z})$ that acts on singular 0-simplices as the diagonal map, or strictly speaking, the R-module homorphism $C_*(X; R) = C_*(X) \to C_*(X) \otimes C_*(X)$ that this uniquely determines on the chain complex with coefficients in the ring R. The formulation of $\cup$ in terms of a diagonal approximation will be convenient in the following, because it avoids any reference to the cross product, whose definition in the relative case we have seen involves some subtleties. We claim that whenever $A, B \subset X$ are two subspaces that form an excisive couple (see Definition 51.1), there is a well-defined relative cup product

$$\cup : H^k(X, A) \otimes H^\ell(X, B) \to H^{k+\ell}(X, A \cup B).$$

Indeed, under this assumption, Exercise 51.3 identifies $H^*(X, A \cup B)$ with the cohomology of the complex $C_*(X)/(C_*(A) + C_*(B))$, and one can then choose any diagonal approximation $\Psi : C_*(X) \to C_*(X) \otimes C_*(X)$ and make sense of

$$\varphi \cup \psi = (\varphi \otimes \psi) \circ \Psi : \frac{C_*(X)}{C_*(A) + C_*(B)} \to R$$

for $\varphi \in C^*(X, A)$ and $\psi \in C^*(X, B)$, the point here being that since Ψ is natural, it sends any chain in either $C_*(A)$ or $C_*(B)$ to something in $C_*(A) \otimes C_*(A)$ or $C_*(B) \otimes C_*(B)$, which is then annihilated by $\varphi \otimes \psi$ since φ vanishes on $C_*(A)$ and ψ vanishes on $C_*(B)$. One can show that this version of $\cup$ satisfies properties analogous to those listed in Theorem 50.14, though we need to be a bit careful about its relation to the cross product. While the diagonal map $d : X \to X \times X$ always gives a well-defined map of pairs

$$(X, A \cup B) \xrightarrow{d} (X, A) \times (Y, B),$$

the formula $\varphi \cup \psi = d^*(\varphi \times \psi)$ might *not* make sense under the assumption above, because $A \times X$ and $X \times B$ might not be an excisive couple in $X \times X$, in which case the cross product in this expression is not well defined. If both $A, B \subset X$ and $A \times X, X \times B \subset X \times X$ are excisive couples, then both definitions of $\varphi \cup \psi \in H^*(X, A \cup B)$ do make sense, and they match.

As a special case, the product

$$\cup : H^k(X, A) \otimes H^\ell(X, A) \to H^{k+\ell}(X, A)$$

is well defined for *every* pair (X, A), as $A, A \subset X$ always trivially forms an excisive couple. This is true even though the cross product $\times : H^k(X, A) \otimes H^\ell(X, A) \to H^{k+\ell}((X, A) \times (X, A))$ might not always make sense.

The cap product. The cap product is another pairing that intertwines cohomology with homology to produce a homology class, and its main property is that it is in some sense *dual* to the cup product.

In this section, it will be important to keep in mind the notational convention that the R-bilinear pairing $\langle\ ,\ \rangle : C^*(X) \otimes C_*(X) \to R$ of cochains and chains with coefficients in the ring R allows cochains φ and chains c of arbitrary degrees, where

$$\langle \varphi, c \rangle := 0 \quad \text{if } |\varphi| \neq |c|.$$

Proposition 51.10. *Given a diagonal approximation $\Psi : C_*(X; \mathbb{Z}) \to C_*(X; \mathbb{Z}) \otimes C_*(X; \mathbb{Z})$ and the associated cup product of cochains defined in (51.3), there exists a unique R-module homomorphism*

$$C^k(X) \otimes C_\ell(X) \to C_{\ell-k}(X) : \varphi \otimes c \mapsto \varphi \cap c \tag{51.4}$$

for each pair of integers k, ℓ such that the relation

$$\langle \psi \cup \varphi, c \rangle = \langle \psi, \varphi \cap c \rangle$$

is satisfied for all $\psi, \varphi \in C^*(X)$ and $c \in C_*(X)$. Moreover, $\cap$ satisfies

$$\partial(\varphi \cap c) = \delta\varphi \cap c + (-1)^k \varphi \cap \partial c \quad \textit{for all} \quad \varphi \in C^k(X),\ c \in C_\ell(X). \tag{51.5}$$

Remark 51.11. The degrees appearing in (51.4) become easy to remember if you regard $C^*(X)$ in the spirit of Remark 47.3 as a *chain* (not cochain) complex by reversing the degrees and writing $C^*(X)_k := C^{-k}(X)$. We can then regard $C^*(X) \otimes C_*(X)$ as a tensor product chain complex, and the Leibniz rule (51.5) becomes the statement that $\cap : C^*(X) \otimes C_*(X) \to C_*(X)$ is a chain map.

Proof of Proposition 51.10. The uniqueness of $\cap$ is easy to see, because if $\cap$ and $\cap'$ are two such maps that both satisfy $\langle \psi \cup \varphi, c \rangle = \langle \psi, \varphi \cap c \rangle = \langle \psi, \varphi \cap' c \rangle$ for all ψ, φ, c, then for every $\varphi \in C^k(X)$ and $c \in C_\ell(X)$ we have

$$\langle \psi, \varphi \cap c - \varphi \cap' c \rangle = 0 \quad \text{for all} \quad \psi \in C^{\ell-k}(X).$$

Since $C_{\ell-k}(X)$ is a free R-module, this cannot happen unless $\varphi \cap c - \varphi \cap' c = 0$; see Exercise 51.12 below.

For existence, we can write down a formula for $\cap$ in terms of the diagonal approximation $\Psi : C_*(X) \to C_*(X) \otimes C_*(X)$, though the formula will take some effort to digest. Fix $\varphi \in C^k(X)$ and regard this as an R-module homomorphism $C_*(X) \to R$ that is trivial on $C_\ell(X)$ for every $\ell \neq k$. We can then form the composed homomorphism

$$\begin{tikzcd} C_*(X) \arrow[r, "\Psi"] \arrow[rr, bend right, "\varphi \cap (\cdot)"'] & C_*(X) \otimes C_*(X) \arrow[r, "\mathbb{1} \otimes \varphi"] & C_*(X) \otimes R = C_*(X). \end{tikzcd}$$

Feeding an ℓ-chain $c \in C_\ell(X)$ into this composition produces at first a finite sum

$$\Psi(c) \in \bigoplus_{p+q=\ell} C_p(X) \otimes C_q(X),$$

and the only term in this sum that does not vanish when fed into $\mathbb{1} \otimes \varphi$ is the one with $q = k$ and $p = \ell - k$, thus $\varphi \cap (\cdot)$ maps $C_\ell(X)$ into $C_{\ell-k}(X)$. Our chain-level formula for the cap product is thus

$$\varphi \cap c = (\mathbb{1} \otimes \varphi) \circ \Psi(c). \tag{51.6}$$

Note that since φ may have either odd or even degree, it is understood in this expression that the usual sign convention must be obeyed when evaluating $\mathbb{1} \otimes \varphi$ on a product chain: in particular, if $\varphi \in C^k(X)$, $a \in C_{\ell-k}(X)$, $b \in C_k(X)$ and $\psi \in C^{\ell-k}(X)$, then

$$\langle \psi, (\mathbb{1} \otimes \varphi)(a \otimes b) \rangle = (-1)^{k(\ell-k)} \langle \psi, a \otimes \varphi(b) \rangle = (-1)^{k(\ell-k)} \psi(a)\varphi(b) = (\psi \otimes \varphi)(a \otimes b),$$

hence the map $C_*(X) \otimes C_*(X) \to R$ defined by $\langle \psi, (\mathbb{1} \otimes \varphi)(\cdot) \rangle$ is the same as $\psi \otimes \varphi$. As a consequence, for every $c \in C_\ell(X)$ we have

$$\langle \psi, \varphi \cap c \rangle = \langle \psi, (\mathbb{1} \otimes \varphi)(\Psi(c)) \rangle = (\psi \otimes \varphi)(\Psi(c)) = \langle \psi \cup \varphi, c \rangle,$$

which proves existence.

The Leibniz rule (51.5) can now be deduced via $\langle \psi \cup \varphi, c \rangle = \langle \psi, \varphi \cap c \rangle$ from the corresponding Leibniz rule for $\cup$, but this is a slightly annoying computation in which getting all the signs right is tricky,[85] so let's instead describe a more "highbrow" version of the same argument. The relation between $\cup$ and $\cap$ can be interpreted as saying that the diagram of R-module homomorphisms

$$\begin{tikzcd}
C^*(X) \otimes C^*(X) \otimes C_*(X) \arrow[r, "\mathbb{1} \otimes \cap"] \arrow[d, "\cup \otimes \mathbb{1}"] & C^*(X) \otimes C_*(X) \arrow[d, "\langle\ ,\ \rangle"] \\
C^*(X) \otimes C_*(X) \arrow[r, "\langle\ ,\ \rangle"] & R
\end{tikzcd}$$

commutes. Notice that if we view $C^*(X)$ as a chain complex with $C^*(X)_n = C^{-n}(X)$ as in Remark 47.3, then all maps in this diagram other than $\mathbb{1} \otimes \cap$ are already known to be chain maps. The composition $\langle\ ,\ \rangle \circ (\mathbb{1} \otimes \cap)$ is therefore also a chain map, implying that if we take any pair of elements $\varphi \in C^*(X)$ and $x \in C^*(X) \otimes C_*(X)$, we will have

$$\left((\mathbb{1} \otimes \cap) \circ \partial - \partial \circ (\mathbb{1} \otimes \cap) \right) (\varphi \otimes x) \in \ker \langle\ ,\ \rangle.$$

Applying the usual graded Leibniz rule for $\partial(\varphi \otimes x)$ gives

$$\left((\mathbb{1} \otimes \cap) \circ \partial - \partial \circ (\mathbb{1} \otimes \cap) \right) (\varphi \otimes x) = (-1)^{|\varphi|} \varphi \otimes \left((\cap \circ \partial - \partial \circ \cap)(x) \right),$$

so after plugging this into $\langle\ ,\ \rangle$, we deduce that

$$\langle \varphi, (\cap \circ \partial - \partial \circ \cap) x \rangle = 0$$

holds for all $\varphi \in C^*(X)$ and $x \in C^*(X) \otimes C_*(X)$. It follows via Exercise 51.12 that $(\cap \circ \partial - \partial \circ \cap)x = 0$ for all x, i.e. $\cap$ is a chain map. $\square$

EXERCISE 51.12. Show that if A is a free R-module and $a \in A$ is a nontrivial element, then $\varphi(a) \neq 0$ for some $\varphi \in \mathrm{Hom}_R(A, R)$.
Hint: This is not true in general without the freeness assumption, e.g. it is clearly false for the $\mathbb{Z}$*-module* $\mathbb{Z}_2$*, since* $\mathrm{Hom}(\mathbb{Z}_2, \mathbb{Z}) = 0$*. Use a basis of* A *in your proof.*

[85] I got it right on the third try.

Exercise 51.13. Prove that $\cap : C^k(X) \otimes C_\ell(X) \to C_{\ell-k}(X)$ has the following naturality property: for any map $f : X \to Y$ with a cochain $\varphi \in C^k(Y)$ and chain $c \in C_\ell(X)$,

$$f_* (f^*\varphi \cap c) = \varphi \cap f_* c.$$

Since $\cap$ is a chain map, it descends to a pairing

$$H^k(X;R) \otimes_R H_\ell(X;R) \to H_{\ell-k}(X;R) : [\varphi] \otimes [c] \mapsto [\varphi] \cap [c] := [\varphi \cap c],$$

which we call the **cap product**. One can use the formula (51.6) and the fact that diagonal approximations are unique up to chain homotopy to prove that different choices of diagonal approximation define the same cap product on the level of homology.

Since the diagonal approximation in the definition of $\cap$ can be chosen freely, we might as well make our "favorite" choice to write down an explicit formula: recall from Lecture 50 the Alexander-Whitney diagonal approximation, defined on singular n-simplices $\sigma : \Delta^n \to X$ by

$$\Psi(\sigma) = \sum_{k+\ell=n} \left(\sigma|_{[0,\ldots,k]}\right) \otimes \left(\sigma|_{[k,\ldots,n]}\right),$$

which leads to the chain-level cup product formula

$$(\varphi \cup \psi)(\sigma) = (-1)^{|\varphi||\psi|} \varphi(\sigma|_{[0,\ldots,k]}) \psi(\sigma|_{[k,\ldots,n]}).$$

This formula has the convenient feature that it satisfies

$$1 \cup \varphi = \varphi = \varphi \cup 1 \tag{51.7}$$

and

$$(\varphi \cup \psi) \cup \eta = \varphi \cup (\psi \cup \eta) \tag{51.8}$$

for all cochains φ, ψ, η, where $1 \in C^0(X)$ is the cochain sending all singular 0-simplices to $1 \in R$,[86] which represents the unit element $1 \in H^0(X;R)$. Under other choices of diagonal approximation, these formulas may not hold except "up to chain homotopy" (so that they do of course hold after descending to cohomology). The corresponding explicit formula for $\cap$ in this setting is

$$\varphi \cap \sigma = (-1)^{k(\ell-k)} \varphi\big(\sigma|_{[\ell-k,\ldots,\ell]}\big) \sigma|_{[0,\ldots,\ell-k]} \in C_{\ell-k}(X) \quad \text{for } \varphi \in C^k(X) \text{ and } \sigma : \Delta^\ell \to X.$$

We will not use this formula for anything, but now you've seen it. The relations (51.7) and (51.8), however, are concretely useful in the following exercise.

Exercise 51.14. Use the relation $\langle \psi \cup \varphi, c \rangle = \langle \psi, \varphi \cap c \rangle$ to prove that if $\cap$ is defined using the Alexander-Whitney diagonal approximation, then it satisfies

$$1 \cap c = c \quad \text{for all} \quad c \in C_*(X)$$

and

$$(\varphi \cup \psi) \cap c = \varphi \cap (\psi \cap c) \quad \text{for all} \quad \varphi, \psi \in C^*(X),\ c \in C_*(X).$$

Letting all these properties descend to the level of homology and cohomology, here is a summary of what we have proved so far about the cap product:

Theorem 51.15. *The R-bilinear cap product $\cap : H^*(X) \otimes H_*(X) \to H_*(X)$ on cohomology and homology with coefficients in a ring R has the following properties.*

1. *$f_*(f^*\varphi \cap A) = \varphi \cap f_* A$ for all continuous maps $f : Y \to X$, $\varphi \in H^*(Y)$ and $A \in H_*(X)$.*
2. *$1 \cap A = A$ for all $A \in H_*(X)$.*
3. *$(\varphi \cup \psi) \cap A = \varphi \cap (\psi \cap A)$ for all $\varphi, \psi \in H^*(X)$ and $A \in H_*(X)$.*
4. *$\langle 1, \varphi \cap A \rangle = \langle \varphi, A \rangle$ for any $\varphi \in H^*(X)$ and $A \in H_*(X)$ with the same degree.*

[86] In different contexts, we have also previously referred to the cochain $1 : C_0(X) \to R$ as the *augmentation* of the singular chain complex.

□

REMARK 51.16. The identity

$$\langle \varphi \cup \psi, A\rangle = \langle \varphi, \psi \cap A\rangle \quad \text{for all} \quad \psi \in H^k(X),\ A \in H_\ell(X), \varphi \in H^{\ell-k}(X) \tag{51.9}$$

was left out of the above theorem only because it is redundant: it follows from the third and fourth identities by writing

$$\langle \varphi \cup \psi, A\rangle = \langle 1, (\varphi \cup \psi) \cap A\rangle = \langle 1, \varphi \cap (\psi \cap A)\rangle = \langle \varphi, \psi \cap A\rangle.$$

Conversely, the fourth identity is a special case of this one.

The relative cap product takes the form

$$\cap : H^*(X, A) \otimes H_*(X, A \cup B) \to H_*(X, B) \tag{51.10}$$

for any two subsets $A, B \subset X$ that form an excisive couple. To see why this works, observe that the chain-level cap product pairing

$$C^*(X) \otimes C_*(X) \to C_*(X) : \varphi \otimes c \mapsto \varphi \cap c = (\mathbb{1} \otimes \varphi) \circ \Psi(c)$$

always descends to a well-defined map on the relative complexes

$$C^*(X, A) \otimes \frac{C_*(X)}{C_*(A) + C_*(B)} \to C_*(X, B),$$

as $\varphi \in C^*(X, A)$ means $\varphi : C_*(X) \to R$ vanishes on $C_*(A) \subset C_*(X)$, so if $c \in C_*(A)$ then $\Psi(c) \in C_*(A) \otimes C_*(A)$ and $\varphi \cap c$ thus vanishes, whereas if $c \in C_*(B)$, then $\Psi(c) \in C_*(B) \otimes C_*(B)$ and $\varphi \cap c \in C_*(B)$. Now if $A, B \subset X$ are an excisive couple, the homology of $C_*(X)/(C_*(A) + C_*(B))$ has a natural identification with $H_*(X, A \cup B)$, thus making sense of (51.10).

52. The orientation bundle (January 30, 2024)

The next few lectures will focus on a new topic: the global topology of finite-dimensional topological manifolds.[87]

There is a basic fact about manifolds that was briefly mentioned in the context of the Lefschetz fixed point theorem (Lecture 41), and now deserves to be repeated: every compact manifold M admits a topological embedding into $\mathbb{R}^N$ for N sufficiently large (see [**Hat02**, Appendix A]), and is therefore a *Euclidean neighborhood retract*. In particular, this means there exists a compact polyhedron P with a retraction $r : P \to M$, and since P necessarily has finitely generated homology, it follows that the same is true for M:

THEOREM 52.1. *For every compact manifold M, $H_*(M; \mathbb{Z})$ is finitely generated.* □

I would now like to discuss what it means for a topological manifold to be *orientable*. We discussed this somewhat in Lecture 31 through the lens of *oriented triangulations*, but that characterization of orientations requires some extra data that might not exist, i.e. not every topological manifold is triangulable. Another natural approach would be to generalize something that we discussed specifically for surfaces in Lecture 20 last semester: one needs to first understand what it means to say that a homeomorphism between two open subsets of $\mathbb{R}^n$ is "orientation preserving," so that an orientation on M can then be defined to mean a covering of M by charts with the property that any two overlaping charts are related by a coordinate transformation that preserves orientations. If we work with smooth manifolds, then it is fairly easy to make this precise, because we can say that a *smooth* coordinate transformation preserves orientations if and only if its derivative at every point is a linear map $\mathbb{R}^n \to \mathbb{R}^n$ with positive determinant. For maps that are continuous

[87] I will typically omit the word "topological" and just say "manifold", as for most of this discussion it will not be at all necessary to mention smooth structures.

but not differentiable, it takes more effort to say precisely what "orientation preserving" means, and the most elegant way to do it uses homology.

Instead of working with coordinate transformations, the standard approach in algebraic topology is via the notion of *local orientations*, which we saw already in our discussion of the mapping degree (Lecture 35). Recall that if $\dim M = n$, then for every choice of coefficient group G and every interior point $x \in M \backslash \partial M$, there is a locally Euclidean neighborhood $\mathbb{R}^n \cong \mathcal{U}_x \subset M$ of x that gives rise (via the usual axioms of homology) to natural isomorphisms

$$\begin{aligned} H_k(M, M\backslash\{x\}; G) &\cong H_k(\mathcal{U}_x, \mathcal{U}_x\backslash\{x\}; G) \cong H_k(\mathbb{R}^n, \mathbb{R}^n\backslash\{0\}; G) \cong H_k(\mathbb{D}^n, \partial\mathbb{D}^n; G) \\ &\cong \widetilde{H}_{k-1}(S^{n-1}; G) \cong \begin{cases} G & \text{if } k = n, \\ 0 & \text{otherwise.} \end{cases} \end{aligned} \tag{52.1}$$

Omitting the coefficient group G from the notation as usual, we call

$$H_n(M \,|\, x) := H_n(M, M\backslash\{x\}) \cong G$$

the **local homology** group of M at x, and a **local orientation** of M at x is defined to be a choice of generator

$$[M]_x \in H_n(M \,|\, x; \mathbb{Z}) \cong \mathbb{Z}.$$

At every interior point x, there clearly are two possible choices of local orientations. The question now is: if we have chosen a local orientation of M at every point $x \in M\backslash\partial M$, what should it mean to say that these orientations *vary continuously* with x?

To answer this question, we can start by viewing the local homology groups

$$H_*(M \,|\, x) := H_*(M, M\backslash\{x\})$$

as an example of "restricting" the homology of M to smaller subsets—in this case, the one-point subset $\{x\} \subset M$. More generally, any subset $A \subset M$ determines relative homology groups

$$H_*(M \,|\, A) := H_*(M, M\backslash A),$$

which we shall call the "homology of M restricted to A". The chain complex underlying $H_*(M \,|\, A)$ does not see any chains that fail to intersect A, and the cycles in this complex are chains whose boundaries are disjoint from A. By subdivision, we can also restrict our attention to arbitrarily "small" singular simplices, which means that $H_*(M \,|\, A)$ really only depends on the topology of arbitrarily small neighborhoods of A in M. (One can of course use the excision property to make this statement more precise.) For any further subset $B \subset A \subset M$, the identity map on M defines a natural inclusion of pairs $(M, M\backslash A) \hookrightarrow (M, M\backslash B)$, which therefore induces natural "restriction" homomorphisms

$$j_{B,A} : H_*(M \,|\, A) \to H_*(M \,|\, B),$$

or in the case where B is a single point $x \in A$,

$$j_{x,A} : H_*(M \,|\, A) \to H_*(M \,|\, x).$$

Note that the absolute homology $H_*(M)$ itself is also an example of a restricted homology group, namely $H_*(M \,|\, M)$.

Since the isomorphism $H_n(M \,|\, x; \mathbb{Z}) \cong \mathbb{Z}$ only holds when $x \in M$ lies in the interior $M\backslash\partial M$, it will be convenient to assume in most of the following discussion that M has empty boundary,

$$\partial M = \varnothing.$$

Now, in order to relate the local homology groups $H_n(M \,|\, x)$ to each other for two distinct but nearby points $x \in M$, suppose $\varphi : \mathcal{U} \xrightarrow{\cong} \mathbb{R}^n$ is a chart defined on some open set $\mathcal{U} \subset M$, and let

$A \subset \mathcal{U}$ denote the subset $\varphi^{-1}(\mathbb{D}^n)$. Then for any point $x \in A$, the maps induced by φ^{-1} and the obvious inclusions of pairs fit together in a commutative diagram

$$\begin{tikzcd} H_n(M \,|\, x) & H_n(\mathcal{U} \,|\, x) \arrow[l, "\cong"'] & H_n(\mathbb{R}^n \,|\, \varphi(x)) \arrow[l, "\cong", "\varphi_*^{-1}"'] \\ H_n(M \,|\, A) \arrow[u, "j_{x,A}"] & H_n(\mathcal{U} \,|\, A) \arrow[l, "\cong"'] & H_n(\mathbb{R}^n \,|\, \mathbb{D}^n), \arrow[l, "\cong", "\varphi_*^{-1}"'] \arrow[u, "\cong"] \end{tikzcd}$$

in which the two horizontal maps at the left are isomorphisms by excision, and the vertical map at the right is an isomorphism due to a combination of homotopy equivalence and the five-lemma, proving that

$$j_{x,A} : H_n(M \,|\, A) \to H_n(M \,|\, x)$$

is an isomorphism. We shall say in this situation that $A \subset M$ is a **disk-like neighborhood** of $x \in M$.

DEFINITION 52.2. An **orientation** of an n-dimensional topological manifold M with empty boundary is a choice of local orientations $[M]_x \in H_n(M \,|\, x; \mathbb{Z})$ for every $x \in M$ satisfying the following consistency condition: for every disk-like neighborhood $A \subset M$ and all $x, y \in A$,

$$j_{x,A}^{-1}[M]_x = j_{y,A}^{-1}[M]_y \in H_n(M \,|\, A; \mathbb{Z}).$$

In the case $\partial M \neq \varnothing$, we define an orientation of M to be an orientation of its interior $M \backslash \partial M$.

A manifold equipped with an orientation will be called an **oriented manifold** *(orientierte Mannigfaltigkeit)*. In light of (52.1), you can imagine an orientation as a choice for every $x \in M \backslash \partial M$ of a favorite generator $[S_x] \in \widetilde{H}_{n-1}(S_x; \mathbb{Z}) \cong \mathbb{Z}$ for some small $(n-1)$-sphere S_x enclosing x, with the property that translating S_x to S_y through a coordinate chart containing x and y produces an isomorphism $\widetilde{H}_{n-1}(S_x; \mathbb{Z}) \to \widetilde{H}_{n-1}(S_y; \mathbb{Z})$ sending $[S_x]$ to $[S_y]$. You should take a moment to contemplate why this description matches Definition 52.2 in the case $M = \mathbb{R}^n$.

The notion of orientation described above admits fruitful generalizations, and in order to express them in the most useful language, let us again assume $\partial M = \varnothing$ and abbreviate

$$\Theta_x^G := H_n(M \,|\, x; G)$$

for each point $x \in M$, or simply

$$\Theta_x := H_n(M \,|\, x)$$

for situations where the choice of coefficient group does not matter. This associates to each point $x \in M$ an abelian group Θ_x that is isomorphic to the chosen coefficient group, and the union of all these groups defines a set that we shall denote by

$$\Theta = \Theta^G := \bigcup_{x \in M} \Theta_x^G.$$

Note that in this definition, we are regarding Θ_x and Θ_y as disjoint sets whenever $x \neq y$, so Θ is set-theoretically their disjoint union; I am avoiding writing it as $\coprod_{x \in M} \Theta_x$ since this notation normally carries connotations about the topology of the union, and those connotations would be inconsistent with the following definition.

DEFINITION 52.3. For an n-manifold M with empty boundary, the **orientation bundle** of M with coefficients in G is the set $\Theta = \Theta^G$ defined above, endowed with the topology generated by the collection of subsets

$$\mathcal{B} := \left\{ \mathcal{U}_c \subset \Theta \;\middle|\; \mathcal{U} \subset M \text{ open and } c \in H_n(M \,|\, \bar{\mathcal{U}}) \right\},$$

where for $\mathcal{U} \subset M$ and $c \in H_n(M \,|\, \bar{\mathcal{U}})$ we define

$$\mathcal{U}_c := \left\{ j_{x,\bar{\mathcal{U}}}(c) \in \Theta_x \;\middle|\; x \in \mathcal{U} \right\}.$$

Proposition 52.4. *The collection of subsets $\mathcal{B} = \{\mathcal{U}_c\}$ appearing in Definition 52.3 is the base of a topology on Θ for which the natural projection map*

$$p : \Theta \to M$$

sending Θ_x to x for each $x \in M$ is continuous and is a covering map.

Proof. To show that $\mathcal{B}$ is the base of a topology, we need to show first that these sets cover all of Θ, and second that any finite intersection of such sets is also a union of such sets. The former is true because for any $x \in M$ and $c \in \Theta_x$, we can pick an open set $\mathcal{U} \subset M$ whose closure is a disk-like neighborhood $\bar{\mathcal{U}} \subset M$ of x, for which we showed above that $j_{x,\bar{\mathcal{U}}} : H_n(M \,|\, \bar{\mathcal{U}}) \to \Theta_x$ is an isomorphism, thus $c \in \mathcal{U}_{c'}$ for $c' := j_{x,\bar{\mathcal{U}}}^{-1}(c)$.

For finite intersections, consider two open sets $\mathcal{U}, \mathcal{V} \subset M$ and classes $a \in H_n(M \,|\, \bar{\mathcal{U}})$ and $b \in H_n(M \,|\, \bar{\mathcal{V}})$. Then $\mathcal{U}_a \cap \mathcal{V}_b \subset \bigcup_{x \in \mathcal{U} \cap \mathcal{V}} \Theta_x$, and we observe that for $x \in \mathcal{U} \cap \mathcal{V}$ and any subset $A \subset \mathcal{U} \cap \mathcal{V}$ containing x, the maps $j_{x,\bar{\mathcal{U}}}$ and $j_{x,\bar{\mathcal{V}}}$ both factor through $H_n(M \,|\, A)$:

$$\begin{tikzcd}
H_n(M \,|\, \bar{\mathcal{U}}) \arrow[dr, "j_{A,\bar{\mathcal{U}}}"] \arrow[drr, bend left, "j_{x,\bar{\mathcal{U}}}"] & & \\
& H_n(M \,|\, A) \arrow[r, "j_{x,A}"] & \Theta_x \\
H_n(M \,|\, \bar{\mathcal{V}}) \arrow[ur, "j_{A,\bar{\mathcal{V}}}"] \arrow[urr, bend right, "j_{x,\bar{\mathcal{V}}}"'] & &
\end{tikzcd}$$

Choose $A \subset \mathcal{U} \cap \mathcal{V}$ to be a disk-like neighborhood, so that $j_{x,A}$ is an isomorphism for every $x \in A$. Now if $x \in A$ and $c \in \Theta_x$ belongs to both $\mathcal{U}_a$ and $\mathcal{V}_b$, it means

$$c = j_{x,\bar{\mathcal{U}}}(a) = j_{x,\bar{\mathcal{V}}}(b) = j_{x,A}(c') \quad \text{where} \quad c' := j_{x,A}^{-1}(c) = j_{A,\bar{\mathcal{U}}}(a) = j_{A,\bar{\mathcal{V}}}(b),$$

hence $c \in \mathring{A}_{c'}$, and conversely, the diagram also demonstrates that $\mathring{A}_{c'} \subset \mathcal{U}_a \cap \mathcal{V}_b$. This proves that $\mathcal{U}_a \cap \mathcal{V}_b$ is a union of sets $\mathring{A}_c \in \mathcal{B}$, where A ranges over disk-like neighborhoods contained in $\mathcal{U} \cap \mathcal{V}$.

To prove that $p : \Theta \to M$ is continuous and is a covering space, the main idea is as follows: for each $x \in M$, choose a disk-like neighborhood $A \subset M$ of x and observe that the isomorphism $j_{x,A} : H_n(M \,|\, A) \to H_n(M \,|\, x)$ factors through $H_n(M \,|\, \bar{\mathcal{U}})$ for any smaller open neighborhood $\mathcal{U} \subset \mathring{A}$ of x, implying that $j_{x,\bar{\mathcal{U}}} : H_n(M \,|\, \bar{\mathcal{U}}) \to H_n(M \,|\, x)$ is also an isomorphism. One can use this to show that for each $x \in \mathring{A}$, assigning the discrete topology to Θ_x makes the map

$$(p, j_{x,A}) : p^{-1}(A) \to A \times \Theta_x$$

a homeomorphism. Using it to identify $p^{-1}(A)$ with $A \times \Theta_x$ turns $p^{-1}(A) \xrightarrow{p} A$ into the trivial covering map $A \times \Theta_x \to A : (a, c) \mapsto a$. ☐

Remark 52.5. The word "bundle" is borrowed from differential geometry, where fiber bundles $p : E \to B$ generalize the notion of a covering space by allowing the fibers $p^{-1}(b) \subset E$ to be more interesting topological spaces (typically manifolds or vector spaces), rather than just discrete sets. In general, a fiber bundle whose fibers are discrete is equivalent to a covering map. The orientation bundle also has a bit more structure than this, since its fibers Θ_x are groups—this makes $p : \Theta \to M$ a *sheaf* of abelian groups, or if we choose G to be a ring R so that each homology group is an R-module, a sheaf of R-modules. For readers who may know what this means and find it interesting: $p : \Theta \to M$ is the *completion* of the *presheaf* that associates to each open subset $\mathcal{U} \subset M$ the abelian group $H_n(M \,|\, \bar{\mathcal{U}})$.

EXERCISE 52.6. Given a point $x \in M$, let I denote the set of all open neighborhoods of x, and write $\mathcal{U} \prec \mathcal{V}$ whenever $\mathcal{V} \subset \mathcal{U}$. This makes $(I, \prec)$ into a directed set, and whenever $\mathcal{U} \prec \mathcal{V}$ there is an associated homomorphism $j_{\bar{\mathcal{V}},\bar{\mathcal{U}}} : H_n(M \,|\, \bar{\mathcal{U}}) \to H_n(M \,|\, \bar{\mathcal{V}})$, so that the collection of abelian groups $\{H_n(M \,|\, \bar{\mathcal{U}})\}_{\mathcal{U} \in I}$ forms a direct system. Find a canonical isomorphism

$$\varinjlim \{H_n(M \,|\, \bar{\mathcal{U}})\} \xrightarrow{\cong} H_n(M \,|\, x).$$

DEFINITION 52.7. For each subset $A \subset M$, we denote

$$\Theta|_A := p^{-1}(A) \subset \Theta$$

and call the covering map $\Theta|_A \xrightarrow{p} A$ the **restriction** of the orientation bundle to A. A **section** *(Schnitt)* of Θ along A is by definition a continuous map $s : A \to \Theta$ such that $p \circ s = \mathrm{Id}_A$, i.e. it continuously associates to each $x \in A$ an element $s(x) \in \Theta_x$. The set of all sections of Θ along A will be denoted by $\Gamma(\Theta|_A)$, with the special case $A = M$ denoted simply by $\Gamma(\Theta)$. We say a section $s \in \Gamma(\Theta|_A)$ has **compact support** if it satisfies $s(x) = 0$ for all x outside some compact subset of A, and denote the set of sections with this property by

$$\Gamma_c(\Theta|_A) \subset \Gamma(\Theta|_A).$$

EXERCISE 52.8. Show that $\Gamma(\Theta|_A)$ and $\Gamma_c(\Theta|_A)$ are both naturally abelian groups, where addition of sections is defined pointwise, i.e.

$$(s_1 + s_2)(x) := s_1(x) + s_2(x) \in \Theta_x.$$

If $G = R$ is a commutative ring with unit, show similarly that $\Gamma(\Theta^R|_A)$ and $\Gamma_c(\Theta^R|_A)$ are modules over R.

Our previous definition of orientations can now be recouched in the following terms.

DEFINITION 52.9. An **orientation** of M along a subset $A \subset M$ is a section $s \in \Gamma(\Theta^{\mathbb{Z}}|_A)$ such that $[M]_x := s(x)$ generates $\Theta^{\mathbb{Z}}_x \cong \mathbb{Z}$ for every $x \in A$.

More generally, if R is a commutative ring with unit, an R**-orientation** of M along $A \subset M$ is a section $s \in \Gamma(\Theta^R|_A)$ such that for every $x \in A$, $s(x)$ generates Θ^R_x as an R-module, i.e. $Rs(x) = \Theta^R_x$. If such a section exists, we say that M is **orientable over** R along A, or simply R**-orientable** along A. A manifold with nonempty boundary can be included in this definition by calling M itself R**-orientable** if it is R-orientable along its interior $M \backslash \partial M$.

REMARK 52.10. One should not be misled by the isomorphism $\Theta^R_x \cong H_n(S^{n-1}; R) \cong R$ into thinking that Θ^R_x is a ring—it is an R-module, and the isomorphism $\Theta^R_x \cong R$ respects that R-module structure, but Θ^R_x does not have any natural ring structure in general.

The geometric meaning of R-orientations when $R \neq \mathbb{Z}$ merits further comment, but let's first look a bit more closely at the case $R = \mathbb{Z}$. There are exactly two possible choices of generators $[M]_x$ in each fiber $\Theta^{\mathbb{Z}}_x \cong \mathbb{Z}$, that is, the two local orientations of M at x. Let us write

$$\widetilde{M} := \left\{ c \in \Theta^{\mathbb{Z}} \;\middle|\; c \text{ is a local orientation} \right\},$$

in other words, $\widetilde{M}$ is the union for all $x \in M$ of the two generators of $\Theta^{\mathbb{Z}}_x \cong \mathbb{Z}$. Assigning to $\widetilde{M} \subset \Theta^{\mathbb{Z}}$ the subspace topology, it is easy to see that the restriction of $p : \Theta^{\mathbb{Z}} \to M$ defines a two-to-one covering map

$$\pi := p|_{\widetilde{M}} : \widetilde{M} \to M : [M]_x \mapsto x.$$

It is called the **orientation double cover** of M. Observe now that if M is orientable over a connected subset $A \subset M$, then there are exactly two choices of orientation, given by some section $s : A \to \widetilde{M}$ and its opposite, $-s : A \to \widetilde{M}$, i.e. the section of $\Theta^{\mathbb{Z}}|_A$ for which $-s + s = 0$. The images of these two sections are disjoint, but by the definition of the topology on $\Theta^{\mathbb{Z}}$, they are both also open subsets of $\pi^{-1}(A) \subset \widetilde{M}$, implying that $\pi^{-1}(A)$ is disconnected. Conversely:

EXERCISE 52.11. If $A \subset M$ is connected and $\pi^{-1}(A) \subset \widetilde{M}$ has more than one connected component, show that each component intersects $\Theta^{\mathbb{Z}}_x$ for every $x \in A$. (Hint: Show that the set of $x \in A$ for which $\Theta^{\mathbb{Z}}_x$ intersects the component is both open and closed.) Conclude that $\pi^{-1}(A) \subset \widetilde{M}$ therefore has exactly two components, each of which is the image of a section of $\Theta^{\mathbb{Z}}$ along A.

Combining the exercise with the previous remarks proves:

PROPOSITION 52.12. *For any connected subset $A \subset M$, $\pi^{-1}(A) \subset \widetilde{M}$ has either one or two connected components, where the latter is the case if and only if M is orientable along A.* □

EXAMPLE 52.13. For $M = \mathbb{RP}^2$, the orientation double cover is equivalent to the standard covering $S^2 \to S^2/\mathbb{Z}_2 = \mathbb{RP}^2$ defined via the antipodal map on S^2. In particular, $\mathbb{RP}^2$ is orientable along a loop $\gamma \subset \mathbb{RP}^2$ if and only if γ has a lift to S^2 that is a loop (instead of a path with distinct end points).

The main advantage of generalizing to other coefficient rings $R \neq \mathbb{Z}$ arises from the following observation about the case $R = \mathbb{Z}_2$:

PROPOSITION 52.14. *Every manifold is orientable over $\mathbb{Z}_2$.*

PROOF. Each fiber $\Theta^{\mathbb{Z}_2}_x$ of the orientation bundle consists only of the trivial element $0 \in \mathbb{Z}_2$ and the nontrivial element $1 \in \mathbb{Z}_2$, so there is a unique nontrivial section $s \in \Gamma(\Theta^{\mathbb{Z}_2})$, defined by $s(x) = 1$ for all x. □

EXERCISE 52.15. Use the universal coefficient theorem to show that for every abelian group G, there is a natural isomorphism $\Phi_x : \Theta^{\mathbb{Z}}_x \otimes G \to \Theta^G_x$ for every $x \in M$ such that if $s \in \Gamma(\Theta^{\mathbb{Z}}|_A)$ is a section and $g \in G$, then $s'(x) := \Phi_x(s(x) \otimes g)$ defines a section $s' \in \Gamma(\Theta^G|_A)$. Deduce that if M is orientable along A, then it is also R-orientable along A for every choice of coefficient ring R.

We would now like to formulate a relationship between the group of sections $\Gamma(\Theta^G|_A)$ and the homology group $H_n(M \,|\, A; G)$.

LEMMA 52.16. *For every closed subset $A \subset M$ and every choice of coefficient group G, there exists a homomorphism*

$$J_A : H_n(M \,|\, A; G) \to \Gamma_c(\Theta^G|_A) : c \mapsto s_c$$

defined by $s_c(x) := j_{x,A}(c)$ for $x \in A$.

PROOF. As usual we shall omit G from the notation. We need to show two things about the map $s_c : A \to \Theta|_A$, first that it is continuous, and second that its support is compact. After this it will be obvious that J_A is a homomorphism. Let's consider first the support of s_c.

Given $[c] \in H_n(M \,|\, A)$ represented by a relative cycle $c \in C_n(M)$ with $\partial c \in C_{n-1}(M \backslash A)$, we can write c as a finite linear combination $\sum_i m_i \sigma_i$ of singular n-simplices $\sigma : \Delta^n \to M$ with coefficients $m_i \in G$. Since Δ^n is compact and the sum is finite, there exists a compact subset $K \subset M$ that contains the images of all the σ_i, so for any $x \in A$ with $x \notin K$, c is an n-chain in $M \backslash \{x\}$, implying that its image under the chain map induced by $(M, M \backslash A) \hookrightarrow (M, M \backslash \{x\})$ is trivial and thus $s_c(x) = j_{x,A}[c] = 0$. The support of s_c is therefore contained in the compact subset $A \cap K \subset A$.

For continuity, we start with the observation that if $A \subset X$ happens to have the property that $j_{x,A}$ is an isomorphism for every $x \in \mathring{A}$, then the same argument as in the proof of Proposition 52.4 identifies $\Theta|_A$ with $A \times \Theta_x$ so that s_c looks like a "constant section" $x \mapsto (x, g)$ for some $g \in \Theta_x$ and is thus obviously continuous. We can reduce the situation to this case as follows. For $[c] \in H_n(M, M \backslash A)$ represented by a relative cycle $c \in C_n(M)$ as above, $\partial c \in C_{n-1}(M \backslash A)$ is a chain in some compact subset $K \subset M \backslash A$, where $M \backslash A$ is open since A is closed. This implies that every $x \in A$ admits a disk-like neighborhood $\bar{\mathcal{U}} \subset M$ disjoint from K, so that ∂c also defines an $(n-1)$-chain in $M \backslash \bar{\mathcal{U}}$ and c can therefore be regarded as a relative n-cycle in $(M, M \backslash \bar{\mathcal{U}})$, representing a class

$$[c] \in H_n(M \,|\, \bar{\mathcal{U}}).$$

There is then a well-defined and necessarily continuous section of $\Theta|_{\mathcal{U}}$ defined by the same formula $x \mapsto j_{x,\bar{\mathcal{U}}}([c])$, and our original $s_c : A \to \Theta|_A$ near $x \in A$ is the restriction of this section to $A \cap \mathcal{U}$, which is therefore continuous. $\square$

Here is the main theorem about the orientation bundle.

THEOREM 52.17. *If M is a topological n-manifold with empty boundary, then for every closed subset $A \subset M$, the map $J_A : H_n(M \,|\, A; G) \to \Gamma_c(\Theta^G|_A)$ is an isomorphism, and $H_k(M \,|\, A; G) = 0$ for all $k > n$.*

The proof of this theorem is a bit long, so we will save it for the next lecture, and focus for now on its corollaries.

COROLLARY 52.18. *Assume M is a topological n-manifold with empty boundary, and G is an abelian group. Then:*

(1) $H_k(M; G) = 0$ for all $k > n$.
(2) If M is noncompact and connected, then additionally $H_n(M; G) = 0$.
(3) If M is compact, connected and orientable, then $H_n(M; G) \cong G$.
(4) If M is compact and connected but not orientable, then $H_n(M; G) \cong \{g \in G \mid 2g = 0\}$.

Moreover, if M is compact and R-orientable for some commutative ring R with unit, then any choice of R-orientation $s \in \Gamma(\Theta^R)$ determines a unique element

$$[M] \in H_n(M; R) \qquad \text{such that} \qquad j_{x,M}[M] = s(x) \text{ for all } x \in M,$$

and if M is additionally connected, then the R-module $H_n(M; R)$ is isomorphic to R and has $[M]$ as a generator.

PROOF. We work through the claims one by one:

(1) Follows from $H_k(M \,|\, A; G) = 0$ with $A = M$.
(2) If M is noncompact and connected then $\Gamma_c(\Theta^G) = 0$, since any section with compact support must equal zero somewhere; indeed, continuity then implies that the subset $\{x \in M \mid s(x) = 0\}$ is both open and closed, so it is all of M.
(3) Taking $A = M$ gives an isomorphism $H_n(M; G) \cong \Gamma(\Theta^G)$, where the compact support condition is irrelevant since M is compact. Then given an orientation $s \in \Gamma(\Theta^{\mathbb{Z}})$ and the natural isomorphisms $\Phi_x : \Theta^{\mathbb{Z}}_x \otimes G \to \Theta^G_x$ from the universal coefficient theorem (cf. Exercise 52.15), we obtain an injective homomorphism

$$G \to \Gamma(\Theta^G) : g \mapsto s_g \qquad \text{where} \qquad s_g(x) := \Phi_x(s(x) \otimes g).$$

This homomorphism is also surjective if M is connected, because in this case, the value of a section $s \in \Gamma(\Theta^{\mathbb{Z}})$ at any one point $x \in M$ uniquely determines the section.

(4) For the subgroup $G_0 = \{g \in G \mid 2g = 0\}$, we can again use the isomorphisms $\Phi_x : \Theta_x^{\mathbb{Z}} \otimes G \to \Theta_x^G$ to define an injective homomorphism

$$G_0 \to \Gamma(\Theta^G) : g \mapsto s_g \quad \text{where} \quad s_g(x) := \Phi_x(\pm[M]_x \otimes g).$$

Here the choice of local orientation $\pm[M]_x \in \Theta_x^{\mathbb{Z}}$ is arbitrary and $s_g(x)$ does not depend on it, since $g = -g$. We leave it as an exercise to show that this map is also surjective whenever M is connected and non-orientable: in particular, since $\widetilde{M}$ is connected, given any $x \in M$, there is no section taking the value $[M]_x \otimes g$ at x for some generator $[M]_x \in \Theta_x^{\mathbb{Z}}$ and $g \in G$ unless $g = -g$.

For the last statement, we observe that any R-orientation $s \in \Gamma(\Theta^R)$ belongs to $\Gamma_c(\Theta^R)$ if M is compact, in which case the distinguished class $[M] \in H_n(M; R)$ can be written as $J_M^{-1}(s)$. If M is also connected, then the map $\Gamma(\Theta^R) \to \Theta_x^R : s \mapsto s(x)$ is an isomorphism for any chosen point $x \in M$, thus s generates $\Gamma(\Theta^R)$ as an R-module, implying that $[M]$ does the same for $H_n(M; R)$. $\square$

Definition 52.19. The generator $[M] \in H_n(M; R) \cong R$ associated to any R-orientation of a closed n-manifold M in the above corollary is called the **fundamental class** of M (over R).

While the construction of fundamental classes $[M] \in H_n(M; R)$ described above works only for manifolds that are closed, just a little bit more effort is required in order to extract from Theorem 52.17 a similar construction for compact manifolds with nonempty boundary. We've already learned from the triangulated manifolds in Lectures 30 and 31 what to expect: when M is compact but $\partial M \neq \varnothing$, $[M]$ should live in the relative homology of the pair $(M, \partial M)$. Recall that when M has nonempty boundary, the orientability of M is defined purely in terms of its interior

$$\mathring{M} := M \backslash \partial M;$$

as the orientation bundle is not defined (or at least it does not have nice properties) along ∂M. We will therefore need a basic observation from point-set topology in order to relate M and $\mathring{M}$: namely, if the boundary ∂M is compact, then it has a so-called **collar neighborhood** *(Kragenumgebung)* in M, meaning a neighborhood $\mathcal{U} \subset M$ of ∂M that is homeomorphic to $(-1, 0] \times \partial M$ via a homeomorphism sending ∂M to $\{0\} \times \partial M$. This is not completely obvious, but the proof is not hard (see e.g. [**Hat02**, Proposition 3.42]). It follows that M is homotopy equivalent to its interior, hence the latter has finitely generated homology if M is compact.

Theorem 52.17 can be applied to $\mathring{M}$, and gives an isomorphism

$$J_A : H_n(\mathring{M}, \mathring{M} \backslash A; G) \to \Gamma_c(\Theta^G|_A)$$

for any closed subset $A \subset \mathring{M}$ and abelian group G. Taking G to be a ring R, any R-orientation $s \in \Gamma(\Theta^R|_{\mathring{M}})$ determines a generator $[M]_x := s(x) \in H_n(M \,|\, x; R) \cong R$ for each interior point $x \in \mathring{M}$. We will refer to a relative homology class

$$[M] \in H_n(M, \partial M; R)$$

as a **relative fundamental class** for M if the natural map $i^x : H_n(M, \partial M; R) \to H_n(M \,|\, x; R)$ defined via the inclusion $(M, \partial M) \hookrightarrow (M, M \backslash \{x\})$ for every $x \in \mathring{M}$ satisfies

$$i_*^x[M] = [M]_x.$$

If $\partial M = \varnothing$, this matches our previous characterization of the fundamental class of a closed manifold, and if M has a triangulation, then it also matches the characterization that was established in Lectures 30 and 31 for relative fundamental classes formed by summing the n-simplices in a triangulation.

THEOREM 52.20. *If M is a compact manifold with boundary carrying an R-orientation $s \in \Gamma(\Theta^R|_{\mathring{M}})$, then s determines a unique relative fundamental class $[M] \in H_n(M, \partial M; R)$, which is a generator of the R-module $H_n(M, \partial M; R) \cong R$ if M is connected.*

PROOF. We shall as usual omit the coefficient ring R from the notation. Identify a neighborhood of ∂M in M with $(-1, 0] \times \partial M$ and for $\epsilon > 0$ small, let $M_\epsilon \subset M$ denote the complement of $(-\epsilon, 0] \times \partial M \subset M$, which is a compact set homotopy equivalent to M. Now if $x \in M_\epsilon$, consider the commuting diagram

$$\begin{tikzcd}
H_n(M, \partial M) \arrow[r, "\cong"] \arrow[dr, "i_x^*"'] & H_n(M, M \backslash M_\epsilon) \arrow[d, "j_{x,M_\epsilon}"] & H_n(\mathring{M}, \mathring{M} \backslash M_\epsilon) \arrow[l, "\cong"'] \arrow[d, "j_{x,M_\epsilon}"] \\
& H_n(M, M \backslash \{x\}) & H_n(\mathring{M}, \mathring{M} \backslash \{x\}), \arrow[l, "\cong"']
\end{tikzcd}$$

where several maps are labeled as isomorphisms due to homotopy invariance. An R-orientation $s \in \Gamma(\Theta|_{\mathring{M}})$ determines a generator $[M]_x := s(x) \in H_n(\mathring{M} \,|\, x) \cong R$ for each $x \in \mathring{M}$, and its restriction to M_ϵ is a compactly supported section of $\Gamma(\Theta|_{M_\epsilon})$, thus it also determines via Theorem 52.17 a unique $[M]_\epsilon \in H_n(\mathring{M} \,|\, M_\epsilon)$ such that $j_{x,M_\epsilon}[M]_\epsilon = [M]_x$ for every $x \in M_\epsilon$. Following the two isomorphisms at the top of the diagram, $[M]_\epsilon$ now determines a class $[M] \in H_n(M, \partial M)$ that satisfies $i_*^x[M] = [M]_x \in H_n(M \,|\, x)$ for all $x \in M_\epsilon$. We leave it as an exercise to check that this definition of $[M] \in H_n(M, \partial M)$ does not depend on the choice of $\epsilon > 0$. (Hint: the isomorphism of Theorem 52.17 can again be used to show that for two $\epsilon, \delta > 0$, $[M]_\epsilon$ and $[M]_\delta$ have the same image under the natural maps to $H_n(M \,|\, M_\epsilon \cap M_\delta)$, which is an isomorphism.) Since any $x \in \mathring{M}$ is in M_ϵ for sufficiently small $\epsilon > 0$, the uniqueness of $[M]$ follows.

Since $H_n(M, \partial M) \cong H_n(\mathring{M} \,|\, M_\epsilon) \cong \Gamma_c(\Theta|_{M_\epsilon}) = \Gamma(\Theta|_{M_\epsilon})$, we deduce $H_n(M, \partial M) \cong R$ with $[M]$ as a generator if M (and therefore also M_ϵ) is connected, because the map $\Gamma(\Theta|_{M_\epsilon}) \to \Theta_x \cong R$ defined by evaluating sections at any point $x \in M_\epsilon$ is then an R-module isomorphism that sends $s|_{M_\epsilon}$ to $[M]_x$. □

EXERCISE 52.21. Prove that if M is a non-orientable connected topological manifold, then $\pi_1(M)$ contains a subgroup of index 2. (In particular, this implies that every simply connected manifold is orientable.)

EXERCISE 52.22. Suppose M is any connected topological manifold of dimension $n \in \mathbb{N}$.

(a) Prove that the torsion subgroup of $H_{n-1}(M; \mathbb{Z})$ is $\mathbb{Z}_2$ if M is compact and non-orientable, and it is otherwise trivial.
Hint: Use the universal coefficient theorem to compute $\mathrm{Tor}(H_{n-1}(M), \mathbb{Z}_p) = 0$ *for various values of $p \geqslant 2$, and see what you can deduce from it (cf. Exercise 43.14). You may want to consider separately the cases where M is noncompact, compact and orientable, or compact and non-orientable.*
(b) Deduce that if $H_*(M; \mathbb{Z})$ is finitely generated and M is orientable, then $H^n(M; \mathbb{Z}) \cong H_n(M; \mathbb{Z})$.

EXERCISE 52.23. Here is an interesting application of Čech cohomology to the question of orientability of manifolds. Fix a space X and abelian group G, and recall that the set $\mathcal{O}(X)$ of all open coverings of X admits an ordering relation $\prec$ that makes it into a directed set: we write $\mathfrak{U} \prec \mathfrak{U}'$ whenever $\mathfrak{U}'$ is a refinement of $\mathfrak{U}$. There is a direct system of $\mathbb{Z}$-graded abelian groups over $\mathcal{O}(X)$ whose direct limit is Čech cohomology, namely

$$\check{H}^*(X; G) := \varinjlim \left\{ H_o^*(\mathcal{N}(\mathfrak{U}); G) \right\}_{\mathfrak{U} \in \mathcal{O}(X)},$$

where $\mathcal{N}(\mathfrak{U})$ is the so-called nerve of the open covering $\mathfrak{U} \in \mathcal{O}(X)$, defining a simplicial complex, and $H^*_o(\mathcal{N}(\mathfrak{U}); G)$ is the cohomology with coefficients in G of its ordered simplicial complex. Concretely, $H^*_o(\mathcal{N}(\mathfrak{U}); G)$ is the homology of a cochain complex $\check{C}^*(\mathfrak{U}; G) := C^*_o(\mathcal{N}(\mathfrak{U}); G)\big)$, where $\check{C}^n(\mathfrak{U}; G) = 0$ for $n < 0$ and, for each $n \geqslant 0$, $\check{C}^n(\mathfrak{U}; G)$ is the additive abelian group of all functions φ that assign an element of G to each ordered $(n+1)$-tuple of sets $\mathcal{U}_0, \dots \mathcal{U}_n \in \mathfrak{U}$ with nonempty intersection:

$$\varphi(\mathcal{U}_0, \dots, \mathcal{U}_n) \in G \quad \text{assuming} \quad \mathcal{U}_0 \cap \dots \cap \mathcal{U}_n \neq \varnothing.$$

The coboundary map $\delta : \check{C}^n(\mathfrak{U}; G) \to \check{C}^{n+1}(\mathfrak{U}; G)$ is defined by

$$(\delta\varphi)(\mathcal{U}_0, \dots, \mathcal{U}_{n+1}) := (-1)^{n+1} \sum_{k=0}^{n+1} (-1)^k \varphi(\mathcal{U}_0, \dots, \widehat{\mathcal{U}}_k, \dots, \mathcal{U}_{n+1}),$$

where the hat over $\widehat{\mathcal{U}}_k$ means that that term is skipped. The homologies of these cochain complexes form a direct system over $(\mathcal{O}(X), \prec)$ because, as mentioned in Lecture 46, refinements $\mathfrak{U}' \succ \mathfrak{U}$ give rise to chain maps $C^o_*(\mathcal{N}(\mathfrak{U}')) \to C^o_*(\mathcal{N}(\mathfrak{U}))$ that are canonical up to chain homotopy, so dualizing these gives chain maps $\check{C}^*(\mathfrak{U}; G) \to \check{C}^*(\mathfrak{U}'; G)$ that are also canonical up to chain homotopy and therefore induce canonical maps on the cohomology groups (see Lecture 48).

Let us call an open covering $\mathfrak{U}$ *admissible* if intersections between two sets in $\mathfrak{U}$ are always connected; this will be a useful technical condition in the following, and one can show that at least if X is a smooth manifold, every open covering of X has an admissible refinement, so assume this from now on.[88] We are going to consider covering[89] maps $f : Y \to X$ of degree 2. Recall that two such covering maps (Y_i, f_i) for $i = 1, 2$ are called **isomorphic** if there exists a homeomorphism $\varphi : Y_1 \to Y_2$ such that the diagram

$$\begin{tikzcd} Y_1 \arrow[rr, "\varphi"] \arrow[dr, "f_1"'] & & Y_2 \arrow[dl, "f_2"] \\ & X & \end{tikzcd}$$

commutes. We will say that a covering map (Y, f) is **trivial** if it is isomorphic to the **trivial double cover**

$$X \times \mathbb{Z}_2 \to X : (x, i) \mapsto x.$$

Given $f : Y \to X$, any open covering $\mathfrak{U} \in \mathcal{O}(X)$ can be replaced with a refinement such that every $\mathcal{U} \in \mathfrak{U}$ is **evenly covered** by $f : Y \to X$, meaning $f^{-1}(\mathcal{U})$ is the union of two disjoint subsets $\mathcal{V}_0, \mathcal{V}_1 \subset Y$ such that $f|_{\mathcal{V}_i} : \mathcal{V}_i \to \mathcal{U}$ is a homeomorphism for $i = 0, 1$. After a further refinement, assume $\mathfrak{U}$ is also admissible. We can now choose for each $\mathcal{U} \in \mathfrak{U}$ a so-called **local trivialization**, meaning a homeomorphism

$$\Phi_{\mathcal{U}} : f^{-1}(\mathcal{U}) \to \mathcal{U} \times \mathbb{Z}_2$$

that sends $f^{-1}(x)$ to $\{x\} \times \mathbb{Z}_2$ for each $x \in \mathcal{U}$. This determines a set of continuous **transition functions** $g_{\mathcal{U},\mathcal{V}} : \mathcal{U} \cap \mathcal{V} \to \mathbb{Z}_2$ for each intersecting pair $\mathcal{U}, \mathcal{V} \in \mathfrak{U}$, defined such that the map

$$(\mathcal{U} \cap \mathcal{V}) \times \mathbb{Z}_2 \xrightarrow{\Phi_{\mathcal{V}} \circ \Phi_{\mathcal{U}}^{-1}} (\mathcal{U} \cap \mathcal{V}) \times \mathbb{Z}_2$$

[88] Alternatively, one could avoid the need for connected intersections by using Čech cohomology with sheaf coefficients, cf. [Spa95, Chapter 6].

[89] Caution! This exercise now contains two distinct meanings of the word "cover": one in the sense of "open covering" *(Überdeckung)* and the other in the sense of "covering map" *(Überlagerung)*. I am trying very hard to ensure that it would be clear in each instance which meaning is intended.

takes the form $(x, i) \mapsto (x, i + g_{\mathcal{U},\mathcal{V}}(x))$. Note that since $\mathcal{U} \cap \mathcal{V}$ is always assumed connected, the transition functions are all constant, i.e. they associate to each ordered pair $(\mathcal{U}, \mathcal{V})$ of sets in $\mathfrak{U}$ with $\mathcal{U} \cap \mathcal{V} \neq \varnothing$ an element $\varphi(\mathcal{U}, \mathcal{V}) := g_{\mathcal{U},\mathcal{V}} \in \mathbb{Z}_2$. See if you can prove the following:

(a) $\varphi \in \check{C}^1(\mathfrak{U}; \mathbb{Z}_2)$ is a cocycle, and choosing different local trivializations changes φ by a coboundary.

(b) Feeding $[\varphi] \in H^1_o(\mathcal{N}(\mathfrak{U}); \mathbb{Z}_2)$ into the canonical map to the direct limit produces a class $w_1(f) \in \check{H}^1(X; \mathbb{Z}_2)$ that is independent of the choice of admissible open covering.

(c) If X is an n-manifold and $f : Y \to X$ is its orientation double cover, then
$$w_1(X) := w_1(f) \in \check{H}^1(X; \mathbb{Z}_2)$$
is zero if and only if X is orientable. (We call $w_1(X)$ the first **Stiefel-Whitney class** of X.)

53. Existence of the fundamental class (February 2, 2024)

The unfinished business of the previous lecture was the proof of Theorem 52.17, that the natural map
$$J_A : H_n(M \,|\, A; G) \to \Gamma_c(\Theta^G|_A) : c \mapsto s_c$$
is an isomorphism and $H_k(M \,|\, A; G) = 0$ for all $k > n$, assuming M is a topological n-manifold without boundary and $A \subset X$ is a closed subset. The choice of coefficient group G plays no role at all in the proof, so we will omit it from the notation from now on.

The proof of Theorem 52.17 follows a certain pattern common to theorems about manifolds: we start by proving by direct means that it holds whenever A is a special type of "small" subset that can be found in some neighborhood of every point in a manifold. One can view this as the first step in a generalized notion of proof by induction, where the "inductive step" involves using a Mayer-Vietoris sequence to extend the validity of the theorem to unions or intersections of sets for which it is already known to hold. As a convenient (but informal) bit of terminology, we shall call a compact subset $A \subset M$ in an n-manifold M **convex** if A is contained in a Euclidean neighborhood $\mathcal{U} \subset M$ with a chart $\varphi : \mathcal{U} \xrightarrow{\cong} \mathbb{R}^n$ such that $\varphi(A) \subset \mathbb{R}^n$ is convex. Taken at face value, this definition just means that A is either empty or is homeomorphic to the disk $\mathbb{D}^n$, but the way in which we will employ this condition in practice uses more information than that. The key detail is that whenever we have two overlapping compact convex subsets $A, B \subset M$ that lie in *the same* Euclidean neighborhood $\mathcal{U} \subset M$ and look convex with respect to *the same* chart $\varphi : \mathcal{U} \xrightarrow{\cong} \mathbb{R}^n$, it follows that $A \cap B \subset M$ is also convex (and therefore homeomorphic to a disk), since intersections of convex subsets in $\mathbb{R}^n$ are always convex. One cannot similarly guarantee for arbitrary overlapping pairs of subsets $A, B \subset M$ homeomorphic to $\mathbb{D}^n$ that $A \cap B$ is also homeomorphic to $\mathbb{D}^n$.

Exercise 53.1. Show that if M is a topological n-manifold, every compact subset $A \subset M$ can be written as $A = \bigcap_{i=1}^{\infty} A_i$ for a nested sequence $M \supset A_1 \supset A_2 \supset A_3 \supset \ldots \supset A$ such that each A_i is a finite union of compact convex subsets.
Hint: Show that for any $\epsilon > 0$, A can be covered by a finite union of compact convex sets that each have diameter less than ϵ with respect to some fixed metric on M.

Proof of Theorem 52.17. The proof is divided into seven steps.

Step 1: We claim that the theorem is true whenever $A \subset M$ is a compact convex subset. Indeed, A is in this case a disk-like neighborhood, so $j_{x,A} : H_n(M, M \backslash A) \to \Theta_x$ is an isomorphism for every $x \in A$, and since two sections of Θ along a connected subset must be identical whenever they match at one point, this makes $J_A : H_n(M \,|\, A) \to \Gamma(\Theta|_A)$ an isomorphism. For $k > n$, we have the usual computation
$$H_k(M, M \backslash A) \cong H_k(\mathbb{R}^n, \mathbb{R}^n \backslash \mathbb{D}^n) \cong \widetilde{H}_{k-1}(S^{n-1}) = 0.$$

Step 2: For the first of three "inductive" steps, we show that if $A, B \subset M$ are two subsets such that the theorem holds for A, B and $A \cap B$, then it also holds for $A \cup B$. The tool required for this is the relative Mayer-Vietoris sequence from the end of Lecture 33. Since A and B are both closed, the complements of these and $A \cap B$ and $A \cup B$ are all open, and the sets $M \backslash A$ and $M \backslash B$ thus form an open covering of $M \backslash (A \cap B)$. The obvious inclusions of pairs

$$\begin{tikzcd} & (M, M\backslash A) \arrow[dr, hook] & \\ (M, M\backslash(A \cup B)) \arrow[ur, hook] \arrow[dr, hook] & & (M, M\backslash(A \cap B)) \\ & (M, M\backslash B) \arrow[ur, hook] & \end{tikzcd}$$

then give rise to a long exact sequence of the form

$$\begin{aligned} \ldots \to H_{k+1}(M \mid A) \oplus H_{k+1}(M \mid B) \to H_{k+1}(M \mid A \cap B) \to H_k(M \mid A \cup B) \\ \to H_k(M \mid A) \oplus H_k(M \mid B) \to H_k(M \mid A \cap B) \to \ldots \end{aligned}$$

If $k > n$, then the sequence places $H_k(M \mid A \cup B)$ in between two vanishing terms and thus proves $H_k(M \mid A \cup B) = 0$. To handle the case $k = n$, observe that the groups of compactly supported sections along these various subsets also fit into a natural exact sequence

$$0 \to \Gamma_c(\Theta|_{A \cup B}) \to \Gamma_c(\Theta|_A) \oplus \Gamma_c(\Theta|_B) \to \Gamma_c(\Theta|_{A \cap B}),$$

where the first map sends $s \in \Gamma_c(\Theta|_{A \cup B})$ to $(s|_A, -s|_B) \in \Gamma_c(\Theta|_A) \oplus \Gamma_c(\Theta|_B)$, and the second sends $(s, t) \in \Gamma_c(\Theta|_A) \oplus \Gamma_c(\Theta|_B)$ to $s|_{A \cap B} + t|_{A \cap B}$. Note that this is not a full "short" exact sequence: we are not claiming that the last map in the sequence is surjective, as it might not be possible to extend a given section along $A \cap B$ to a section along A or B. It should be evident however that the sequence is exact at all other terms. Moreover, the maps in both sequences commute with the natural maps from homology groups to groups of sections; in order to fit the resulting commutative diagram inside the margins, let's abbreviate

$$H_k^X := H_k(M \mid X) \qquad \text{for any } k \geqslant 0 \text{ and subset } X \subset M,$$

so that the diagram in question looks like

$$\begin{tikzcd} H^A_{n+1} \oplus H^B_{n+1} \arrow[r] \arrow[d, "\cong"] & H^{A \cap B}_{n+1} \arrow[r] \arrow[d, "\cong"] & H^{A \cup B}_n \arrow[r] \arrow[d, "J_{A \cup B}"] & H^A_n \oplus H^B_n \arrow[r] \arrow[d, "\cong"', "J_A \oplus J_B"] & H^{A \cap B}_n \arrow[d, "\cong"', "J_{A \cap B}"] \\ 0 \arrow[r] & 0 \arrow[r] & \Gamma_c(\Theta|_{A \cup B}) \arrow[r] & \Gamma_c(\Theta|_A) \oplus \Gamma_c(\Theta|_B) \arrow[r] & \Gamma_c(\Theta|_{A \cap B}) \end{tikzcd}$$

The five-lemma now implies that $J_{A \cup B}$ is also an isomorphism.

Step 3: The second inductive step is to show that if the theorem holds for each set $A_i \subset M$ in a nested sequence of compact subsets $A_1 \supset A_2 \supset A_3 \supset \ldots$, then it also holds for $A_\infty := \bigcap_{i=1}^\infty A_i \subset M$. This requires a direct limit argument. Observe first that the sequence of inclusions

$$(M, M\backslash A_1) \hookrightarrow (M, M\backslash A_2) \hookrightarrow (M, M\backslash A_3) \hookrightarrow \ldots (M, M\backslash A_\infty)$$

induces a sequence of homomorphisms

$$H_*(M \mid A_1) \to H_*(M \mid A_2) \to H_*(M \mid A_3) \to \ldots \to H_*(M \mid A_\infty),$$

so that $\{H_*(M \mid A_i)\}_{i=1}^\infty$ forms a direct system of $\mathbb{Z}$-graded abelian groups, with $H_*(M \mid A_\infty)$ as a target. We claim that the sequence of maps $H_*(M \mid A_i) \to H_*(M \mid A_\infty)$ satisfies the universal property so that $H_*(M \mid A_\infty)$ is in fact the direct limit $\varinjlim \{H_*(M \mid A_i)\}_{i=1}^\infty$. For this, we need to show that if H is another $\mathbb{Z}$-graded abelian group with a sequence of morphisms $\Phi_i : H_*(M \,|\, A_i) \to H$ making the diagram

$$\begin{tikzcd}
H_*(M \,|\, A_1) \arrow[r] \arrow[rrrrd, "\Phi_1"'] & H_*(M \,|\, A_2) \arrow[r] \arrow[rrrd] & H_*(M \,|\, A_3) \arrow[r] \arrow[rrd, "\Phi_3"] & \dots \arrow[r] & H_*(M \,|\, A_\infty) \arrow[d, dashed, "\Phi_\infty"] \\
 & & & & H
\end{tikzcd}$$

commute, then the map Φ_∞ indicated by the dashed arrow exists and is unique. Indeed, we can define $\Phi_\infty[c]$ for any given class $[c] \in H_k(M \,|\, A_\infty)$ as follows: represent $[c]$ by a relative cycle $c \in C_k(M)$, which means $\partial c \in C_{k-1}(M \backslash A_\infty)$, and note that $M \backslash A_\infty$ contains a compact subset K that contains the images of all singular simplices appearing in ∂c. Since A_∞ is compact, we can then find an open neighborhood $\mathcal{U} \subset M$ of A_∞ that is disjoint from K, and we also have $A_N \subset \mathcal{U}$ for $N \in \mathbb{N}$ sufficiently large. It follows that $\partial c \in C_{k-1}(M \backslash A_N)$, so c is also a relative cycle in $(M, M \backslash A_N)$ and thus defines a class $[c] \in H_k(M \,|\, A_N)$. We now define $\Phi_\infty[c]$ as $\Phi_N[c]$ after reinterpreting $[c]$ in this way.

By restricting sections to smaller domains, we also have a sequence of restriction homomorphisms

$$\Gamma(\Theta|_{A_i}) \to \Gamma(\Theta|_{A_2}) \to \Gamma(\Theta|_{A_3}) \to \ldots \to \Gamma(\Theta|_{A_\infty}),$$

and we can use a similar trick to identify $\varinjlim \{\Gamma(\Theta|_{A_i})\}_{i=1}^\infty$ with $\Gamma(\Theta|_{A_\infty})$. Indeed, the problem now is to show that any sequence of homomorphisms $\varphi_i : \Gamma(\Theta|_{A_i}) \to H$ as in the diagram

$$\begin{tikzcd}
\Gamma(\Theta|_{A_1}) \arrow[r] \arrow[rrrrd, "\varphi_1"'] & \Gamma(\Theta|_{A_2}) \arrow[r] \arrow[rrrd] & \Gamma(\Theta|_{A_3}) \arrow[r] \arrow[rrd, "\varphi_3"] & \dots \arrow[r] & \Gamma(\Theta|_{A_\infty}) \arrow[d, dashed, "\varphi_\infty"] \\
 & & & & H
\end{tikzcd}$$

gives rise to a unique map $\varphi_\infty : \Gamma(\Theta|_{A_\infty}) \to H$. The key here is the observation that since $p : \Theta \to M$ is a covering map, A_∞ has an open neighborhood $\mathcal{U} \subset M$ such that every section $A \to \Theta$ has a unique extension over $\mathcal{U}$, which is therefore defined on A_N for $N \in \mathbb{N}$ sufficiently large. The desired map φ_∞ is thus defined on any $s \in \Gamma(\Theta|_{A_\infty})$ by extending s to A_N and then applying φ_N.

With these preliminaries in place, we can combine both direct systems into a commuting diagram

$$\begin{tikzcd}
H_n(M \,|\, A_1) \arrow[r] \arrow[d, "J_{A_1}", "\cong"'] & H_n(M \,|\, A_2) \arrow[r] \arrow[d, "J_{A_2}", "\cong"'] & H_n(M \,|\, A_3) \arrow[r] \arrow[d, "J_{A_3}", "\cong"'] & \dots \arrow[r] & H_n(M \,|\, A_\infty) \arrow[d, "J_{A_\infty}"] \\
\Gamma(\Theta|_{A_1}) \arrow[r] & \Gamma(\Theta|_{A_2}) \arrow[r] & \Gamma(\Theta|_{A_3}) \arrow[r] & \dots \arrow[r] & \Gamma(\Theta|_{A_\infty}),
\end{tikzcd}$$

so that the sequence of isomorphisms $J_{A_i} : H_n(M \,|\, A_i) \to \Gamma(\Theta|_{A_i})$ defines an isomorphism between the two direct systems, and its limit is therefore an isomorphism between the direct limits. One can make this precise by composing maps in this diagram so as to understand $\Gamma(\Theta|_{A_\infty})$ as a target of the system $\{H_n(M \,|\, A_i)\}_{i=1}^\infty$, whose limit map is necessarily J_{A_∞}, but since the J_{A_i} are all invertible, one can similarly understand $H_n(M \,|\, A_\infty)$ as a target of $\{\Gamma(\Theta|_{A_i})\}_{i=1}^\infty$ and obtain from this a limit map $\Gamma(\Theta|_{A_\infty}) \to H_n(M \,|\, A_\infty)$ that is the inverse of J_{A_∞}.

Step 4: Steps 1 and 3 applied only to *compact* subsets $A \subset M$, but the next inductive step introduces noncompact subsets by allowing infinite disjoint unions. Let us call a collection of compact subsets $\{A_\alpha \subset M\}_{\alpha \in I}$ **separated** if they admit a collection of open neighborhoods $\{A_\alpha \subset \mathcal{U}_\alpha \subset M\}_{\alpha \in I}$ such that $\mathcal{U}_\alpha \cap \mathcal{U}_\beta = \varnothing$ for all $\alpha \neq \beta$. The claim now is that if the theorem holds for every A_α in a separated collection of compact subsets, then it also holds for their union

$A := \bigcup_{\alpha \in I} A_\alpha$. The point of the separation condition is that if we write $\mathcal{U} := \bigcup_{\alpha \in I} \mathcal{U}_\alpha$, then $(\mathcal{U}, \mathcal{U} \backslash A) \cong \coprod_\alpha (\mathcal{U}_\alpha, \mathcal{U}_\alpha \backslash A_\alpha)$, so the excision and additivity axioms give natural isomorphisms

$$H_*(M \,|\, A) \cong H_*(\mathcal{U} \,|\, A) \cong \bigoplus_\alpha H_*(\mathcal{U}_\alpha \,|\, A_\alpha) \cong \bigoplus_\alpha H_*(M \,|\, A_\alpha).$$

This already implies $H_k(M \,|\, A) = 0$ for all $k > n$. For degree n, these isomorphisms fit together into a commutative diagram

$$\begin{tikzcd}
H_n(M \,|\, A) \arrow[d, "J_A"] & H_n(\mathcal{U} \,|\, A) \arrow[l, "\cong"'] \arrow[d, "J_A"] & \bigoplus_\alpha H_n(\mathcal{U}_\alpha \,|\, A_\alpha) \arrow[l, "\cong"'] \arrow[r, "\cong"] \arrow[d, "\bigoplus_\alpha J_{A_\alpha}"] & \bigoplus_\alpha H_n(M \,|\, A_\alpha) \arrow[d, "\bigoplus_\alpha J_{A_\alpha}", "\cong"'] \\
\Gamma_c(\Theta|_A) \arrow[r, equal] & \Gamma_c(\Theta|_A) & \bigoplus_\alpha \Gamma_c(\Theta|_{A_\alpha}) \arrow[l, "\cong"'] \arrow[r, equal] & \bigoplus_\alpha \Gamma_c(\Theta|_{A_\alpha}),
\end{tikzcd}$$

where the isomorphism $\bigoplus_\alpha \Gamma_c(\Theta|_{A_\alpha}) \to \Gamma_c(\Theta|_A)$ sends each $\sum_\alpha s_\alpha \in \bigoplus_\alpha \Gamma_c(\Theta|_{A_\alpha})$ to the unique section $s \in \Gamma(\Theta|_A)$ such that $s|_{A_\alpha} = s_\alpha$ for every $\alpha \in I$. (Note that s necessarily has compact support since only finitely many of the summands in $\sum_\alpha s_\alpha$ can be nonzero. Conversely, a section $s \in \Gamma(\Theta|_A)$ with compact support can be nonzero only on finitely many of the A_α, and is therefore in the image of the map from the direct sum.) This proves that $J_A : H_n(M \,|\, A) \to \Gamma_c(\Theta|_A)$ is an isomorphism.

Step 5: We claim that the theorem holds for every compact set $A \subset M$ that is contained in a Euclidean neighborhood. According to Exercise 53.1, any such set is the intersection of a nested sequence of sets that are each finite unions of compact convex sets, where we can assume all the convex sets are contained in the *same* Euclidean neighborhood. In this case, all intersections of these sets are also compact and convex, so combining steps 1 and 2 proves that the theorem holds for all the finite unions of convex sets, and step 3 then establishes it for A.

Step 6: We extend the theorem to arbitrary compact subsets $A \subset M$. In light of Exercise 53.1, this now follows directly from steps 5, 2 and 3, as A is the intersection of a nested sequence of compact sets that are each finite unions of sets contained in Euclidean neighborhoods. (The fact that those sets can be assumed convex is no longer relevant, but since any intersection between them is contained in a Euclidean neighborhood, step 5 now replaces step 1.)

Step 7: The extension of the theorem to an arbitrary closed $A \subset M$ can now be achieved as follows. I need to appeal to a slightly nontrivial point-set topological fact about manifolds: every finite-dimensional topological manifold M has a one-point compactification M^* that is *metrizable*. Recall that the one-point compactification of any space X is defined as the union of X with one extra point $X^* := X \cup \{\infty\}$, where a subset of X^* is considered open if it is either an open set in X or takes the form $(X \backslash K) \cup \{\infty\}$ for some closed and compact set $K \subset X$. While X^* is always compact, it can easily have horrible topological properties unless X is an especially nice space, e.g. X^* is Hausdorff if and only if X is both Hausdorff and locally compact (cf. Exercise 7.27 from last semester's *Topologie I* class). The one-point compactification M^* of a manifold M is not usually a manifold (the major exception being $(\mathbb{R}^n)^* \cong S^n$), but it is always a metrizable space. This is easy to see if you believe the (also nontrivial) theorem that every n-manifold admits a proper topological embedding into a Euclidean space $\mathbb{R}^N$ of sufficiently high dimension N. A proof of this is sketched in [**Lee11**, p. 116], with several details either left as exercises or outsourced to other references. Since the embedding $M \hookrightarrow \mathbb{R}^N$ is proper, it extends to an embedding $M^* \hookrightarrow (\mathbb{R}^N)^* \cong S^N$, so a metric on M^* can be defined as the restriction of a metric on S^N.

With this detail in place, let dist(,) denote a metric on M^* and exhaust A by the countable sequence of subsets

$$\begin{aligned} A_1 &:= \left\{ x \in A \;\middle|\; 1 \leqslant \operatorname{dist}(x,\infty) < \infty \right\}, \\ A_2 &:= \left\{ x \in A \;\middle|\; 1/2 \leqslant \operatorname{dist}(x,\infty) \leqslant 1 \right\}, \\ A_3 &:= \left\{ x \in A \;\middle|\; 1/3 \leqslant \operatorname{dist}(x,\infty) \leqslant 1/2 \right\}, \\ &\dots \end{aligned}$$

all of which are intersections of A with closed (and therefore compact) subsets of M^*, so they are compact, and the theorem holds for each of them by step 6. We can now apply step 4 to conclude that the theorem also holds for the noncompact subsets

$$B := \bigcup_{j=1}^{\infty} A_{2j-1}, \qquad C := \bigcup_{j=1}^{\infty} A_{2j}, \qquad B \cap C = \bigcup_{j=1}^{\infty} \left\{ x \in A \;\middle|\; \operatorname{dist}(x,\infty) = 1/j \right\},$$

all of which are unions of separated collections of compact sets. We then conclude from step 2 that the theorem also holds for $A = B \cup C$. □

Exercise 53.2. Assume M satisfies the hypotheses of Theorem 52.20 and thus has a relative fundamental class $[M] \in H_n(M, \partial M; R)$.

(a) Show that if M and ∂M are both connected and ∂M is nonempty, then ∂M is also R-orientable, and the connecting homomorphism $\partial_* : H_n(M, \partial M; R) \to H_{n-1}(\partial M; R)$ in the long exact sequence of $(M, \partial M)$ is an isomorphism sending $[M]$ to the fundamental class $[\partial M]$ of ∂M (for a suitable choice of orientation of ∂M).
Hint: Focus on the case $R = \mathbb{Z}$. It is easy to prove that ∂_* is injective; show that if it were not surjective, then $H_{n-1}(M; \mathbb{Z})$ would have torsion, contradicting the result of Exercise 52.22(a).

(b) Generalize the result of part (a) to prove $\partial_*[M] = [\partial M]$ without assuming ∂M is connected.
Hint: For any connected component $N \subset \partial M$, consider the exact sequence of the triple $(M, \partial M, \partial M \backslash N)$ and notice that $H_{n-1}(\partial M, \partial M \backslash N) \cong H_{n-1}(N)$ by excision.

(c) Conclude that for any compact manifold M with boundary and an R-orientation, the map $H_{n-1}(\partial M; R) \to H_{n-1}(M; R)$ induced by the inclusion $\partial M \hookrightarrow M$ sends $[\partial M]$ to 0. In other words, "the boundary of a compact oriented n-manifold M represents the trivial homology class in $H_{n-1}(M)$."
Remark: We discussed a similar result in the setting of triangulable manifolds in Lecture 31 (see Exercise 31.12), but here we are not assuming that any of our manifolds admit triangulations.

54. Poincaré duality (February 6, 2024)

The statement. The classical perspective on Poincaré duality is demonstrated by Figure 31. The picture shows a portion of a closed triangulated manifold M of dimension $n = 2$, with the 1-simplices and vertices of the triangulation depicted in black. We've then added a red dot at the barycenter of each n-simplex and drawn a red line segment connecting the barycenters of any two n-simplices that share a boundary face. Note that since M is assumed to be a manifold without boundary, every $(n-1)$-simplex in the triangulation is a boundary face of exactly two n-simplices. As a consequence, there is a one-to-one correspondence between the $(n-1)$-simplices in the triangulation and the red line segments joining the red dots. Moreover, every vertex of the triangulation is contained in a unique polygon bounded by the red segments. If we think of the red dots as 0-cells, the red line segments as 1-cells and the polygons bounded by them as 2-cells, they

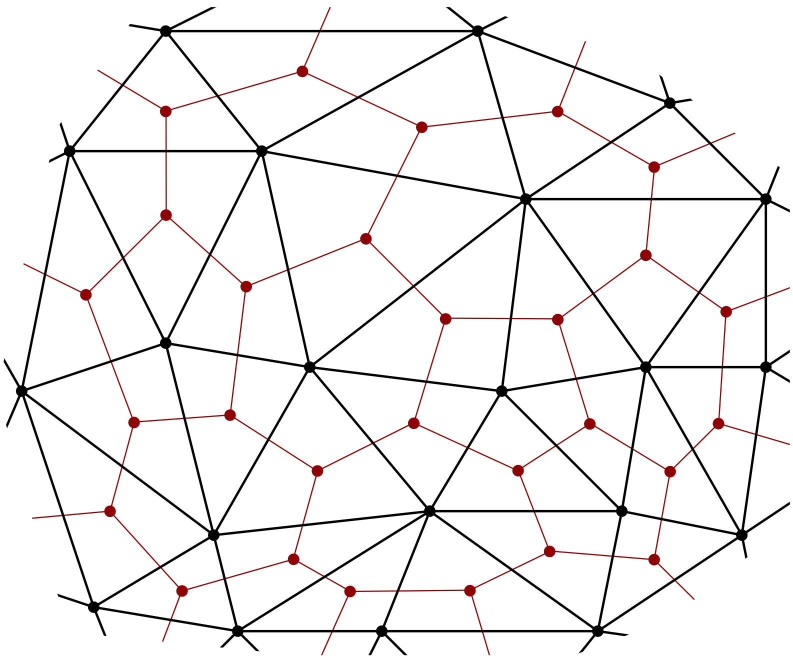

Figure 31. A triangulation of a surface and its dual cell decomposition.

form what is called the **dual cell decomposition** of M determined by the original triangulation. We could now write down two quite different chain complexes to compute the homology of M: let us denote by $C_*^\Delta(M)$ the simplicial chain complex of the original triangulation, and by $C_*^{\mathrm{CW}}(M)$ the cellular chain complex for its dual cell complex. Evidently, there is a natural bijection

$$C_k^\Delta(M) \to C_{n-k}^{\mathrm{CW}}(M),$$

defined by sending each k-simplex of the triangulation to its dual $(n-k)$-cell. You will notice an interesting thing, however, if you try to understand what happens to the boundary map under this bijection: it transforms the boundary map of $C_*^\Delta(M)$ into the *coboundary* map of $C^*_{\mathrm{CW}}(M)$. Thus it can be more properly interpreted as a bijective chain map

$$C_*^\Delta(M) \to C_{\mathrm{CW}}^{n-*}(M),$$

therefore giving rise to an isomorphism $H_k(M) \xrightarrow{\cong} H^{n-k}(M)$ for each $k = 0, \dots, n$.

Remark 54.1. Did you notice where we used the assumption that M is compact in the above discussion? The notion of the dual cell decomposition makes sense on any triangulated manifold, compact or not, so there is still a bijection $C_k^\Delta(M) \to C_{n-k}^{\mathrm{CW}}(M)$, and simplicial and cellular homology also still make sense in the noncompact case. A problem emerges, however, if the triangulation is infinite and we try to pay attention to the boundary map by defining a chain isomorphism $C_*^{\Delta}(M) \to C_{\mathrm{CW}}^{n-*}(M)$. If you don't immediately see why, then keep this question in mind as you read the rest of this lecture, and we'll come back to it at the end.

It would be a bit of an effort to make the idea of the dual cell decomposition precise and general enough to prove an actual theorem, and it would then be a theorem that applies only to *triangulated* manifolds, which is more restrictive than we would like. The key feature that makes Poincaré duality possible is not the triangulation—there are many examples of compact n-dimensional polyhedra X for which $H^k(X) \not\cong H_{n-k}(X)$. The important detail is rather that we are talking specifically about *manifolds*, e.g. it is the locally Euclidean structure of M in the above example that enables us to identify the regions surrounded by dual 1-cells as 2-cells in bijective correspondence with the original vertices. Now that we know there is good reason to expect an isomorphism $H^k(M) \to H_{n-k}(M)$, we observe that a candidate for this isomorphism arises naturally from the previous two topics we discussed in this course: the fundamental class, and the cap product, neither of which had anything directly to do with triangulations. Here's the main theorem in its standard form.

THEOREM 54.2 (Poincaré duality). *For any closed n-manifold M with an R-orientation and corresponding fundamental class $[M] \in H_n(M; R)$ for some commutative ring R with unit, the map*

$$H^k(M; R) \xrightarrow{\mathrm{PD}} H_{n-k}(M; R) : \varphi \mapsto \varphi \cap [M]$$

is an isomorphism for every $k \in \mathbb{Z}$.

Applications. Before getting into the proof, let's pick some low-hanging fruit and state a few corollaries. Recall that by the universal coefficient theorem, the Betti numbers of a space can be expressed as ranks of either the homology or the cohomology groups, which are the same in corresponding degrees. Poincaré duality thus gives a nontrivial relation between them:

COROLLARY 54.3. *For every closed orientable n-manifold M,*

$$b_k(M) = b_{n-k}(M)$$

for all $k \in \mathbb{Z}$. Moreover, without any orientability assumption, the same relation also holds for the so-called "$\mathbb{Z}_2$ Betti numbers," i.e.

$$\dim_{\mathbb{Z}_2} H_k(M; \mathbb{Z}_2) = \dim_{\mathbb{Z}_2} H_{n-k}(M; \mathbb{Z}_2)$$

for all $k \in \mathbb{Z}$. □

COROLLARY 54.4. *Every closed odd-dimensional manifold M satisfies $\chi(M) = 0$.*

PROOF. In the oriented case, this follows because $b_k(M)$ and $b_{n-k}(M)$ cancel each other in the alternating sum that defines $\chi(M)$.

If M is not orientable, then as mentioned in Corollary 54.3, the $\mathbb{Z}_2$ Betti numbers $\dim_{\mathbb{Z}_2} H_k(M; \mathbb{Z}_2)$ still satisfy the corresponding relation, so the result follows if one can show that their alternating sum is also the Euler characteristic, i.e.

$$\chi(M) = \sum_{k=0}^{n} (-1)^k \dim_{\mathbb{Z}_2} H_k(M; \mathbb{Z}_2). \tag{54.1}$$

One way to see this is from the universal coefficient theorem, as the isomorphism

$$H_k(M; \mathbb{Z}_2) \cong (H_k(M; \mathbb{Z}) \otimes \mathbb{Z}_2) \oplus \mathrm{Tor}^{\mathbb{Z}}(H_{k-1}(M; \mathbb{Z}), \mathbb{Z}_2)$$

reveals on close examination that the differences between $b_k(M)$ and $\dim_{\mathbb{Z}_2} H_k(M; \mathbb{Z}_2)$ cancel out in the alternating sum. If you are content to assume M has a cell decomposition, then two somewhat quicker arguments are available: first, Proposition 40.9 identifies $\chi(M)$ with the alternating sum of the numbers of cells, which does not change if we switch to $\mathbb{Z}_2$-coefficients, and then matches

(54.1) by another application of Proposition 40.9. Alternatively, one could argue in terms of the orientation double cover $\pi : \widetilde{M} \to M$, which satisfies $\chi(\widetilde{M}) = 2\chi(M)$ by Theorem 40.15, but also $\chi(\widetilde{M}) = 0$ since $\widetilde{M}$ is in this case also a closed n-manifold and is *always* orientable. (In fact, $\widetilde{M}$ is the space of local orientations of M and thus has a "tautological" orientation.) By an important (but difficult) theorem of [**KS69**], the assumption that M has a cell decomposition is not actually a loss of generality; every compact topological manifold is homotopy equivalent to a finite CW-complex. □

Here is an application of Corollary 54.4 that plays a fundamental role in bordism theory, which studies the equivalence relation $M \sim N$ among closed n-manifolds defined via the condition that $M \amalg N$ is the boundary of a compact $(n+1)$-manifold. For context, observe that every closed oriented surface is the boundary of a compact oriented 3-manifold; we know this because we have a complete classification of such surfaces (see Lecture 19 from last semester), and we can realize all of them as boundaries of compact regions in $\mathbb{R}^3$. It is harder to see, but nonetheless true, that every closed oriented 3-manifold bounds a compact 4-manifold. This fact is originally due to Rokhlin, and it can also be deduced easily from a slightly later result known as the Lickorish-Wallace theorem, which is fundamental in the study of 3-manifolds: it states that every closed oriented 3-manifold M can be obtained by performing surgery along a link in S^3, and from this it is a short step to presenting M as the boundary of a 4-manifold constructed by attaching 4-dimensional 2-handles to $\mathbb{D}^4$. The following result shows that the question of which oriented manifolds are boundaries becomes more interesting from dimension four upwards.

Corollary 54.5. *There is no compact 5-manifold with boundary homeomorphic to $\mathbb{CP}^2$.*

Proof. Suppose M is a compact manifold with $\partial M \cong \mathbb{CP}^2$. We can then construct a closed 5-manifold $\widehat{M}$ by gluing M to a copy of itself along the matching boundary,

$$\widehat{M} := M \cup_{\mathbb{CP}^2} M,$$

and by Corollary 54.4, $\chi(\widehat{M}) = 0$. But according to Exercise 54.6 below, we also have

$$\chi(\widehat{M}) = 2\chi(M) - \chi(\mathbb{CP}^2),$$

a formula that admits an easy interpretation of we assume M has a cell decomposition with ∂M as a subcomplex: counting the cells in $M \amalg M$ with appropriate signs gives $2\chi(M)$, but this overcounts each cell in ∂M by a factor of 2, leading to $-\chi(\mathbb{CP}^2)$ as a correction term. Thanks to its cell decomposition with one cell in each even dimension, we know the homology of $\mathbb{CP}^2$ and therefore know its Euler characteristic: it is 3, implying that $2\chi(M) - \chi(\mathbb{CP}^2)$ can never be 0. □

Exercise 54.6. Suppose X and Y are two compact n-manifolds with homeomorphic boundaries $\partial X \cong \partial Y \cong M$, and $Z := X \cup_M Y$ is a closed n-manifold constructed by gluing them together along their boundaries. Prove the formula

$$\chi(Z) = \chi(X) + \chi(Y) - \chi(M).$$

Hint: This is easy if you assume X and Y have cell decompositions that restrict to their boundaries as matching cell decompositions of M. Without that assumption, you can consider the Mayer-Vietoris sequence for $Z = A \cup B$, where $B \cong (-1,1) \times M$ is the union of collar neighborhoods of ∂X and ∂Y, and A is a disjoint union of open subsets homotopy equivalent to X and Y. Don't try to compute $H_*(Z)$ with this, just view the Mayer-Vietoris sequence itself as a chain complex whose homology is trivial. What does Proposition 40.9 then tell you?

Exercise 54.7. Show that the Klein bottle is homeomorphic to the boundary of a compact (and necessarily non-orientable) 3-manifold, but $\mathbb{RP}^2$ is not.

Poincaré duality also provides considerable information about the ring structure of $H^*(M;R)$ as a consequence of the relation $\langle \psi \cup \varphi, [M] \rangle = \langle \psi, \varphi \cap [M] \rangle$. For each $k = 0, \ldots, n$, consider the bilinear form $Q : H^k(M;R) \times H^{n-k}(M;R) \to R$ defined via the R-module homomorphism

$$\begin{aligned} H^k(M;R) \otimes_R H^{n-k}(M;R) &\xrightarrow{Q} R \\ \varphi \otimes \psi &\mapsto Q(\varphi,\psi) := \langle \varphi \cup \psi, [M] \rangle. \end{aligned}$$

For reasons that we will discuss in the next lecture, this is called the **intersection form** on M. In the case $R = \mathbb{Z}$, $Q(\varphi,\psi)$ vanishes whenever either φ or ψ is torsion, thus it descends to a bilinear map on the free part $H^*_{\text{free}}(M) := H^*(M)/\text{torsion}$,

$$Q : H^k_{\text{free}}(M) \otimes H^{n-k}_{\text{free}}(M) \to \mathbb{Z}.$$

For a general pair of abelian groups A and B, a bilinear map $Q : A \times B \to G$ (or equivalently a group homomorphism $Q : A \otimes B \to G$) is called **nonsingular** if the maps $A \to \text{Hom}(B,G) : a \mapsto Q(a,\cdot)$ and $B \to \text{Hom}(A,G) : b \mapsto Q(\cdot,b)$ are both isomorphisms. There is an obvious analogue of this definition for R-modules and R-module homomorphisms.

Corollary 54.8. *For any closed n-manifold M with a $\mathbb{K}$-orientation and corresponding fundamental class $[M] \in H_n(M;\mathbb{K})$ for some field $\mathbb{K}$, the intersection form*

$$Q : H^k(M;\mathbb{K}) \otimes_{\mathbb{K}} H^{n-k}(M;\mathbb{K}) \to \mathbb{K}$$

is nonsingular for every $k = 0, \ldots, n$, and if M is oriented, Q descends to the free part of $H^*(M;\mathbb{Z})$ as a nonsingular bilinear form $H^k_{\text{free}}(M;\mathbb{Z}) \otimes H^{n-k}_{\text{free}}(M;\mathbb{Z}) \to \mathbb{Z}$.

Proof. With integer coefficients, we saw in Lecture 49 that the canonical map $h : H^{n-k}_{\text{free}}(M) \to \text{Hom}(H^{\text{free}}_{n-k}(M),\mathbb{Z}) : \varphi \mapsto \langle \varphi, \cdot \rangle$ is an isomorphism. Since the duality map $\text{PD} : H^k(M) \to H_{n-k}(M)$ is also an isomorphism, it and its inverse each map torsion to torsion and thus descend to the free parts as isomorphisms $H^k_{\text{free}}(M) \cong H^{\text{free}}_{n-k}(M)$. We can then compose h with the dualization of PD to form an isomorphism

$$\begin{tikzcd} H^{n-k}_{\text{free}}(M) \arrow[r, "h", "\cong"'] \arrow[rr, bend right, "\Phi"', "\cong"] & \text{Hom}\left(H^{\text{free}}_{n-k}(M),\mathbb{Z}\right) \arrow[r, "\text{PD}^*", "\cong"'] & \text{Hom}\left(H^k_{\text{free}}(M),\mathbb{Z}\right) \end{tikzcd}$$

To see what this map actually is, we choose $\psi \in H^{n-k}_{\text{free}}(M)$ and $\varphi \in H^k_{\text{free}}(M)$ and compute:

$$\Phi(\psi)(\varphi) = (\text{PD}^* \circ h(\psi))(\varphi) = h(\psi) \circ \text{PD}(\varphi) = \langle \psi, \varphi \cap [M] \rangle = \langle \psi \cup \varphi, [M] \rangle = Q(\psi,\varphi),$$

so this proves the first of two statements required for showing that Q is nonsingular on the free parts with integer coefficients. But the second required statement is equivalent to this, since $Q(\psi,\varphi) = (-1)^{k(n-k)} Q(\varphi,\psi)$. The argument with field coefficients is completely analogous since, in that case as well, the canonical map $h : H^{n-k}(M;\mathbb{K}) \to \text{Hom}_{\mathbb{K}}(H_{n-k}(M;\mathbb{K}),\mathbb{K})$ is a vector space isomorphism. ☐

Corollary 54.9. *If M is a closed oriented n-manifold and $\varphi \in H^k(M;\mathbb{Z})$ is a primitive*[90] *element for some $k \in \{0, \ldots, n\}$, then there exists some $\psi \in H^{n-k}(M;\mathbb{Z})$ with $Q(\varphi,\psi) = 1$. The same result holds with coefficients in a field $\mathbb{K}$ for every $\varphi \neq 0 \in H^k(M;\mathbb{K})$ if M is $\mathbb{K}$-oriented.*

[90] Recall that $\varphi \in H^k(M;\mathbb{Z})$ is **primitive** if φ is not $m\psi$ for any $\psi \in H^k(M;\mathbb{Z})$ and an integer $m \geqslant 2$. In particular, this rules out that φ is torsion, since $m\varphi = 0$ would imply $(m+1)\varphi = \varphi$.

PROOF. The primitivity hypothesis means that the projection of φ to $H^k_{\text{free}}(M)$ is nontrivial and generates a subgroup $H \subset H^k_{\text{free}}(M)$ such that $H^k_{\text{free}}(M)/H$ has no torsion, implying that it is free (see e.g. [**Lan02**, Chapter I, Theorem 8.4]). It follows that φ can be taken as the first element in a basis of $H^k_{\text{free}}(M)$, so that there exists a homomorphism $\Phi : H^k_{\text{free}}(M) \to \mathbb{Z}$ satisfying $\Phi(\varphi) = 1$. The result then follows from the nonsingularity of Q. In the field case, one can instead appeal to the fact that every nonzero element in a vector space can be an element of a basis. $\square$

EXERCISE 54.10. We can now compute the ring structure of $H^*(\mathbb{CP}^n;\mathbb{Z})$. Take the usual cell decomposition $\mathbb{CP}^n = e^0 \cup e^2 \cup \ldots \cup e^{2n}$, and for $k = 1,\ldots,n$, let $\alpha_k \in H^{2k}(\mathbb{CP}^n;\mathbb{Z}) \cong \mathbb{Z}$ denote the generator that evaluates to 1 on the generator of $H_{2k}(\mathbb{CP}^n;\mathbb{Z})$ represented by the $2k$-cell.

(a) Use Corollary 54.9 to prove $\alpha_k \cup \alpha_{n-k} = \pm\alpha_n$ for every k.
(b) Generalize part (a) to show that $\alpha_k \cup \alpha_\ell = \pm\alpha_{k+\ell}$ for every $k,\ell \in \mathbb{N}$ with $k+\ell \leqslant n$. *Hint: There is a natural inclusion $\mathbb{CP}^{k+\ell} \hookrightarrow \mathbb{CP}^n$ that is a cellular map. How does it act on cohomology?*

This proves that the ring $H^*(\mathbb{CP}^n;\mathbb{Z})$ is generated by the single element $\alpha := \alpha_1 \in H^2(\mathbb{CP}^n;\mathbb{Z})$, subject only to the relation $\alpha^{n+1} = 0$ since $H^k(\mathbb{CP}^n;\mathbb{Z}) = 0$ for all $k > 2n$. We conclude that there is an isomorphism of $\mathbb{Z}$-graded rings[91]

$$H^*(\mathbb{CP}^n;\mathbb{Z}) \cong \mathbb{Z}[\alpha]/(\alpha^{n+1}), \qquad |\alpha| = 2,$$

where $\mathbb{Z}[\alpha]$ denotes the ring of integer-valued polynomials in one variable α, $(\alpha^{n+1}) \subset \mathbb{Z}[\alpha]$ is the ideal generated by α^{n+1}, and the grading is determined by the condition that the variable α has degree 2 while all coefficients have degree 0.

(c) Use inclusions $\mathbb{CP}^n \hookrightarrow \mathbb{CP}^\infty$ to find a graded ring isomorphism $H^*(\mathbb{CP}^\infty;\mathbb{Z}) \cong \mathbb{Z}[\alpha]$, where again $|\alpha| = 2$.

REMARK 54.11. The computation in Exercise 54.10 fills in the last remaining gap in our proof from Lecture 47 (see Theorem 47.1) that all maps $f : \mathbb{CP}^n \to \mathbb{CP}^n$ have fixed points when n is even.

EXERCISE 54.12. Compute each of the following cohomology rings:

(a) $H^*(\mathbb{RP}^n;\mathbb{Z}_2) \cong \mathbb{Z}_2[\alpha]/(\alpha^{n+1})$ with $|\alpha| = 1$.
(b) $H^*(\mathbb{RP}^\infty;\mathbb{Z}_2) \cong \mathbb{Z}_2[\alpha]$ with $|\alpha| = 1$.

The noncompact version. Like the construction of the fundamental class, the proof of Poincaré duality starts by showing that the result is in some sense true "locally," and then uses a form of induction based on Mayer-Vietoris sequences and direct limits to piece together local results into a global result. We therefore need to formulate a more general version of the theorem that can make sense for small neighborhoods in manifolds, rather than just for an entire *closed* manifold.

Fix a coefficient ring R, and suppose M is an n-manifold without boundary that is not necessarily compact, but is endowed with an R-orientation $s \in \Gamma(\Theta) := \Gamma(\Theta^R)$. This section does not have compact support if M is noncompact, but if we choose a compact subset $K \subset M$, then $s|_K \in \Gamma(\Theta|_K)$ trivially does have compact support, and therefore corresponds under Theorem 52.17 to a distinguished homology class

$$[M]_K := J_K^{-1}(s) \in H_n(M \,|\, K).$$

[91] A $\mathbb{Z}$-**graded ring** is a ring R that is split into a direct sum $R = \bigoplus_{n\in\mathbb{Z}} R_n$ such that any $a \in R_k$ and $b \in R_\ell$ have product $ab \in R_{k+\ell}$.

Recall from Lecture 51 that the relative cohomology with coefficients in the ring R admits a cap product pairing

$$\cap : H^k(M, M\backslash K) \otimes H_n(M, M\backslash K) \to H_n(M),$$

which is well defined in this case because the subsets $M\backslash K$ and $\varnothing$ in M trivially form an excisive couple. We can therefore define a "restricted" duality map by

$$\mathrm{PD}_K : H^k(M \,|\, K) \to H_n(M) : \varphi \mapsto \varphi \cap [M]_K.$$

Now consider what happens to this map if we replace K by a larger compact subset $K' \subset M$ that contains K: first, since $[M]_K \in H_n(M \,|\, K)$ and $[M]_{K'} \in H_n(M \,|\, K')$ are determined by the same globally-defined section $s \in \Gamma(\Theta)$, the map induced by the inclusion

$$i : (M, M\backslash K') \hookrightarrow (M, M\backslash K)$$

satisfies

$$i_*[M]_{K'} = [M]_K.$$

The naturality property of the cap product (i.e. Theorem 51.15(1)) then implies that for all $\varphi \in H^k(M \,|\, K)$,

$$i_* \left(i^*\varphi \cap [M]_{K'}\right) = \mathrm{PD}_{K'}(i^*\varphi) = \varphi \cap [M]_K = \mathrm{PD}_K(\varphi),$$

where "i_*" has disappeared in the second expression since $\mathrm{PD}_{K'}(i^*\varphi)$ is an *absolute* homology class and $i : M \to M$ is just the identity map. The result is a commutative diagram

$$\begin{tikzcd} H^k(M \,|\, K) \arrow[r, "i^*"] \arrow[dr, "\mathrm{PD}_K"'] & H^k(M \,|\, K') \arrow[d, "\mathrm{PD}_{K'}"] \\ & H_{n-k}(M), \end{tikzcd} \tag{54.2}$$

which means that we can view the maps $\mathrm{PD}_K : H^k(M \,|\, K) \to H_{n-k}(M)$ as defining a target of a direct system of R-modules $\{H^k(M \,|\, K)\}_K$ over the directed set of compact subsets $K \subset M$, with the partial order defined by inclusion. By the universal property of the direct limit, there is then a uniquely determined homomorphism

$$\mathrm{PD} : \varinjlim \left\{H^k(M \,|\, K)\right\}_K \to H_{n-k}(M).$$

Definition 54.13. For any space X, we define the **compactly supported cohomology** of X with coefficients in an abelian group G as the direct limit

$$H_c^*(X) = H_c^*(X; G) := \varinjlim \left\{H^*(X \,|\, K; G)\right\}_K,$$

where K ranges over the set of all compact subsets of X, ordered by inclusion and forming a direct system via the maps $H^*(X \,|\, K; G) \to H^*(X \,|\, K'; G)$ induced by inclusions $(X, X\backslash K') \hookrightarrow (X, X\backslash K)$ whenever $K \subset K'$.

With this definition in place, the previous discussion produces natural homomorphisms

$$\mathrm{PD} : H_c^k(M; R) \to H_{n-k}(M; R)$$

for every $k \in \mathbb{Z}$ whenever M is a (possibly noncompact) manifold of dimension n with a fixed R-orientation.

Exercise 54.14. Show that if M is compact, there is a natural isomorphism $H_c^*(M) \cong H^*(M)$ which identifies the map $\mathrm{PD} : H_c^k(M) \to H_{n-k}(M)$ defined above with the usual map $\varphi \mapsto \varphi \cap [M]$.

Exercise 54.15. In the following, suppose G is any abelian group.

(a) Prove that $H_c^n(\mathbb{R}^n; G) \cong G$ and $H_c^k(\mathbb{R}^n; G) = 0$ for all $k \neq n$.

(b) Construct a canonical isomorphism between $H^*_c(X;G)$ and the homology of the subcomplex $C^*_c(X;G) \subset C^*(X;G)$ consisting of every cochain $\varphi : C_k(X) \to G$ that vanishes on all simplices with images outside some compact subset $K \subset X$. (Note that K may depend on φ).

(c) Recall that a continuous map $f : X \to Y$ is called **proper**[92] if for every compact set $K \subset Y$, $f^{-1}(K) \subset X$ is also compact. Show that proper maps $f : X \to Y$ induce homomorphisms $f^* : H^*_c(Y;G) \to H^*_c(X;G)$, making $H^*_c(\cdot;G)$ into a contravariant functor on the category of topological spaces with morphisms defined as proper maps.

(d) Deduce from part (c) that $H^*_c(\cdot;G)$ is a topological invariant, i.e. $H^*_c(X;G)$ and $H^*_c(Y;G)$ are isomorphic whenever X and Y are homeomorphic. Give an example showing that this need not be true if X and Y are only homotopy equivalent.

(e) In contrast to part (c), show that $H^*_c(\cdot;G)$ does *not* define a functor on the usual category of topological spaces with morphisms defined to be continuous (but not necessarily proper) maps.
Hint: Think about maps between $\mathbb{R}^n$ and the one-point space.

Here is the noncompact version of Poincaré duality, which has the compact version as a corollary in light of Exercise 54.14.

THEOREM 54.16. *For every R-oriented topological n-manifold M and every $k \in \mathbb{Z}$, the map*
$$\mathrm{PD} : H^k_c(M;R) \to H_{n-k}(M;R),$$
defined as the direct limit of the maps $\mathrm{PD}_K : H^k(M\,|\,K;R) \to H_{n-k}(M;R) : \varphi \mapsto \varphi \cap [M]_K$ for all compact subsets $K \subset M$, is an isomorphism.

The proof. The argument that follows will bear a resemblance to the inductive construction of the fundamental class in the previous lecture. The choice of coefficient ring R plays no significant role, so we shall omit it from the notation from now on. We start with a purely local result to begin the induction.

LEMMA 54.17. *For either choice of orientation of $\mathbb{R}^n$, the map $\mathrm{PD} : H^k_c(\mathbb{R}^n) \to H_{n-k}(\mathbb{R}^n)$ is an isomorphism for every $k \in \mathbb{Z}$.*

PROOF. There is an obvious cofinal set of compact subsets to use in computing $H^k_c(\mathbb{R}^n) = \varinjlim\{H^k(\mathbb{R}^n\,|\,K)\}_K$: every compact subset $K \subset \mathbb{R}^n$ is contained in the disk $\mathbb{D}^n_r$ of sufficiently large radius $r > 0$, and the natural maps $H^k(\mathbb{R}^n\,|\,\mathbb{D}^n_r) \to H^k(\mathbb{R}^n\,|\,\mathbb{D}^n_{r'})$ are isomorphisms for all $r' > r$, thus
$$H^k_c(\mathbb{R}^n) \cong H^k(\mathbb{R}^n\,|\,\mathbb{D}^n) \cong \begin{cases} R & \text{if } k = n, \\ 0 & \text{if } k \ne n. \end{cases}$$
Similarly, $H_{n-k}(\mathbb{R}^n)$ is R if $k = n$ and vanishes otherwise, so it suffices to prove that for any chosen pair of generators $\varphi \in H^n(\mathbb{R}^n\,|\,\mathbb{D}^n) \cong R$ and $[\mathbb{R}^n]_{\mathbb{D}^n} \in H_n(\mathbb{R}^n\,|\,\mathbb{D}^n) \cong R$, $\varphi \cap [\mathbb{R}^n]_{\mathbb{D}^n}$ is also a generator of $H_0(\mathbb{R}^n) \cong R$. This is true since the universal coefficient theorem gives an isomorphism $H^n(\mathbb{R}^n\,|\,\mathbb{D}^n) \cong \mathrm{Hom}(H_n(\mathbb{R}^n\,|\,\mathbb{D}^n), R)$ by evaluation of cohomology classes on homology classes, so that $\langle \varphi, \cdot \rangle$ generates $\mathrm{Hom}_(H_n(\mathbb{R}^n\,|\,\mathbb{D}^n), R)$ and thus
$$\langle 1, \varphi \cap [\mathbb{R}^n]_{\mathbb{D}^n} \rangle = \langle \varphi, [\mathbb{R}^n]_{\mathbb{D}^n} \rangle \in R$$
is a generator of R. □

[92] This definition of properness is standard in the study of manifolds, though for certain purposes, it is sometimes considered an inadequate definition if considering spaces that are not assumed second countable and Hausdorff (the general definition of properness is then a slightly stronger condition). As far as I can tell, it's still an adequate definition for the purposes of Exercise 54.15.

The inductive step unsurprisingly requires Mayer-Vietoris sequences. To prepare for this, we first need to understand the functoriality of H_c^* slightly better. Exercise 54.15 reveals that continuous maps $f : X \to Y$ do not always induce homomorphisms $f^* : H_c^*(Y) \to H_c^*(X)$ unless an additional condition is imposed, i.e. $f : X \to Y$ needs to be proper. We will be especially interested in inclusion maps $A \hookrightarrow X$ for subspaces $A \subset X$, and these are typically *not* proper, e.g. if A is open but not closed, which will be the main case of interest. In this situation, however, there is a natural map going the other direction, from $H_c^*(A)$ to $H_c^*(X)$. This follows from excision: if X is a Hausdorff space with subsets $K \subset A \subset X$ such that A is open and K is compact, then $X\backslash A$ is a closed subset contained in the open set $X\backslash K$, hence the inclusion $(A, A\backslash K) \hookrightarrow (X, X\backslash K)$ is an excision map and induces an isomorphism

$$H^*(X \,|\, K) \xrightarrow{\cong} H^*(A \,|\, K).$$

Now for any compact set $L \subset X$ that contains K, composing the inverse of this isomorphism with the natural map $H^*(X \,|\, K) \to H^*(X \,|\, L)$ induced by the inclusion $(X, X\backslash L) \hookrightarrow (X, X\backslash K)$ produces a map $H^*(A \,|\, K) \to H^*(X \,|\, L)$:

$$\begin{tikzcd}
 & H^*(X \,|\, K) \arrow[dl, "\cong"'] \arrow[dr] & \\
H^*(A \,|\, K) \arrow[rr] & & H^*(X \,|\, L).
\end{tikzcd}$$

If we then compose this with the natural map of $H^*(X \,|\, L)$ to the direct limit $H_c^*(X)$, it produces a map $H^*(A \,|\, K) \to H_c^*(X)$ for every compact $K \subset A$, and one can easily check that this map is independent of the choice of compact subset $L \subset X$ containing K; moreover, if $K' \subset A$ is another compact set containing K, then the diagram

$$\begin{tikzcd}
H^*(A \,|\, K) \arrow[r] \arrow[dr] & H^*(A \,|\, K') \arrow[d] \\
 & H_c^*(X)
\end{tikzcd}$$

commutes. This makes $H_c^*(X)$ a target of the direct system $\{H^*(A \,|\, K)\}_K$, so that there is a uniquely determined limit map

$$H_c^*(A) \to H_c^*(X).$$

We will refer to this always as the *natural map induced by the inclusion* $A \hookrightarrow X$, and it is important to understand that it is only well defined when $A \subset X$ is open.

LEMMA 54.18. *If M is an R-oriented n-manifold and $A \subset M$ is an open subset, then for every $k \in \mathbb{Z}$, the natural maps on H_c^* and H_* induced by the inclusion $A \hookrightarrow M$ fit into a commutative diagram of the form*

$$\begin{tikzcd}
H_c^k(A) \arrow[r] \arrow[d, "\mathrm{PD}"] & H_c^k(M) \arrow[d, "\mathrm{PD}"] \\
H_{n-k}(A) \arrow[r] & H_{n-k}(M)
\end{tikzcd}$$

PROOF. Given a compact set $K \subset A$, pick any compact set $L \subset M$ that contains K, and denote the obvious inclusions

$$A \overset{i}{\hookrightarrow} M, \qquad (A, A\backslash K) \overset{i}{\hookrightarrow} (M, M\backslash K), \qquad (M, M\backslash L) \overset{j}{\hookrightarrow} (M, M\backslash K).$$

We then claim that the diagram

$$\begin{tikzcd}[column sep=small]
& H^k(M\,|\,K) \arrow[dl, "\cong"', "i^*"] \arrow[dr, "j^*"] & \\
H^k(A\,|\,K) \arrow[d, "\mathrm{PD}_K"] & & H^k(M\,|\,L) \arrow[d, "\mathrm{PD}_L"] \\
H_{n-k}(A) \arrow[rr, "i_*"] & & H_{n-k}(M)
\end{tikzcd}$$

commutes. To see this, observe that there is another map we could add to this diagram and sensibly denote by PD_K, namely $H^k(M\,|\,K) \to H_{n-k}(M) : \varphi \mapsto \varphi \cap [M]_K$; let's call this one PD'_K to avoid confusion, and note that by (54.2), it satisfies

$$\mathrm{PD}_L \circ j^* = \mathrm{PD}'_K\,.$$

Viewing i as a map of pairs, we also have $i_*[A]_K = [M]_K$, and naturality of the cap product then implies that for all $\varphi \in H^k(M\,|\,K)$,

$$i_* \circ \mathrm{PD}_K \circ i^*\varphi = i_*\left(i^*\varphi \cap [A]_K\right) = \varphi \cap i_*[A]_K = \varphi \cap [M]_K = \mathrm{PD}'_K(\varphi),$$

thus proving the claim. This implies in particular that for the natural maps $H^*(A\,|\,K) \to H^*(M\,|\,L)$ that determine $H^*_c(A) \to H^*_c(M)$ via the direct limit, the diagram

$$\begin{tikzcd}
H^k(A\,|\,K) \arrow[r] \arrow[d, "\mathrm{PD}_K"] & H^k(M\,|\,L) \arrow[d, "\mathrm{PD}_L"] \\
H_{n-k}(A) \arrow[r] & H_{n-k}(M)
\end{tikzcd}$$

always commutes. The rest is essentially abstract nonsense: if we let $\Psi : H^k_c(A) \to H_{n-k}(M)$ denote the difference between the maps defined via the two possible paths in the diagram of the lemma, we can now view Ψ as the limiting map for a family of maps $H^k(A\,|\,K) \to H_{n-k}(M)$ over the directed set of compact subsets $K \subset A$, and the diagram above forces all these maps to vanish, hence so does Ψ. $\square$

Now suppose $M = A \cup B$, where $A, B \subset M$ are open subsets (and therefore also n-manifolds without boundary). The Mayer-Vietoris sequence we need for H^*_c arises from the natural maps induced by the inclusions of $A \cap B$ into A and B and of each of these into M. Concretely, given any compact subsets $K \subset A$ and $L \subset B$, there are natural inclusions of pairs

$$\begin{tikzcd}[column sep=small]
& (M, M\backslash K) \arrow[dr, hook] & \\
(M, M\backslash(K \cup L)) \arrow[ur, hook] \arrow[dr, hook] & & (M, M\backslash(K \cap L)) \\
& (M, M\backslash L) \arrow[ur, hook] &
\end{tikzcd}$$

which give rise to a relative Mayer-Vietoris sequence in cohomology. The following diagram combines this sequence with the natural excision isomorphisms and localized duality maps:
(54.3)

$$\begin{tikzcd}
\ldots \arrow[r] & H^k(M\,|\,K\cap L) \arrow[r] \arrow[d, "\cong"] & H^k(M\,|\,K)\oplus H^k(M\,|\,L) \arrow[r] \arrow[d, "\cong"] & H^k(M\,|\,K\cup L) \arrow[r] \arrow[d, equal] & H^{k+1}(M\,|\,K\cap L) \arrow[r] \arrow[d, "\cong"] & \ldots \\
\ldots \arrow[r] & H^k(A\cap B\,|\,K\cap L) \arrow[r] \arrow[d, "\mathrm{PD}_{K\cap L}"] & H^k(A\,|\,K)\oplus H^k(B\,|\,L) \arrow[r] \arrow[d, "\mathrm{PD}_K\oplus\mathrm{PD}_L"] & H^k(M\,|\,K\cup L) \arrow[r] \arrow[d, "\mathrm{PD}_{K\cup L}"] & H^{k+1}(A\cap B\,|\,K\cap L) \arrow[r] \arrow[d, "\mathrm{PD}_{K\cap L}"] & \ldots \\
\ldots \arrow[r] & H_{n-k}(A\cap B) \arrow[r] & H_{n-k}(A)\oplus H_{n-k}(B) \arrow[r] & H_{n-k}(M) \arrow[r] & H_{n-k-1}(A\cap B) \arrow[r] & \ldots
\end{tikzcd}$$

We take the horizontal maps in the bottom row to be the usual maps in the Mayer-Vietoris sequence for $H_*(A\cup B)$, and if the signs are chosen appropriately,[93] then the same arguments as in the proof of Lemma 54.18 imply that this diagram commutes, with the possible exception of the bottom right square involving connecting homomorphisms. It turns out that this square also commutes, and the proof is not especially deep, but it is a tedious chain-level calculation involving barycentric subdivision, so we will skip it and simply refer to [**Hat02**, pp. 246–247]. The result is:

LEMMA 54.19. *The diagram in* (54.3) *commutes, and passing to the direct limit over all choices of compact subsets $K\subset A$ and $L\subset B$ then produces a commutative diagram*

$$\begin{tikzcd}
\ldots \arrow[r] & H^k_c(A\cap B) \arrow[r] \arrow[d, "\mathrm{PD}"] & H^k_c(A)\oplus H^k_c(B) \arrow[r] \arrow[d, "\mathrm{PD}\oplus\mathrm{PD}"] & H^k_c(M) \arrow[r] \arrow[d, "\mathrm{PD}"] & H^{k+1}_c(A\cap B) \arrow[r] \arrow[d, "\mathrm{PD}"] & \ldots \\
\ldots \arrow[r] & H_{n-k}(A\cap B) \arrow[r] & H_{n-k}(A)\oplus H_{n-k}(B) \arrow[r] & H_{n-k}(M) \arrow[r] & H_{n-k-1}(A\cap B) \arrow[r] & \ldots
\end{tikzcd} \tag{54.4}$$

in which both rows are exact.

SKETCH OF THE PROOF. Aside from the tedious verification that (54.3) commutes, the claim that the top row of (54.4) is exact is slightly nontrivial: this follows from the general fact that direct limits of exact sequences in the category of R-modules are always exact. Indeed, we proved in Proposition 39.18 that the functor $H_* : \mathsf{Chain} \to \mathsf{Ab}_{\mathbb{Z}}$ is continuous under direct limits, and an exact sequence is nothing other than a chain complex with trivial homology. (Recall from Lecture 46 however that the analogous statement for inverse limits is false, so this detail should not be taken for granted.) $\square$

Applying the five-lemma now gives:

COROLLARY 54.20. *If the duality map is an isomorphism on A, B and $A\cap B$, then it is also an isomorphism on $M = A\cup B$.* $\square$

Open convex sets in Euclidean neighborhoods are homeomorphic to $\mathbb{R}^n$, and so is the intersection of any two such sets in the same Euclidean neighborhood, so Lemmas 54.17 and 54.19 are enough to prove that PD is an isomorphism on any finite union of open convex sets in a single Euclidean neighborhood. Now observe that *any* open set in a Euclidean neighborhood is the union of a countable collection of convex open sets: indeed, just take any covering collection of open balls and reduce it to a countable subcover. Something similar is true in fact for any manifold M: since manifolds are second countable, every open cover of M has a countable subcover (see

[93] Recall that in the Mayer-Vietoris sequence for $H_*(A\cup B)$, there needs to be a minus sign in the definition of either of the maps $H_k(A\cap B)\to H_k(A)\oplus H_k(B)$ or $H_k(A)\oplus H_k(B)\to H_k(A\cup B)$. For most purposes it does not matter which term gets the minus sign, but since we are now relating *two* Mayer-Vietoris sequences to each other, the signs in both need to be consistent.

Lemma 5.25), so one can start with any covering by convex sets in Euclidean neighborhoods and reduce to a countable subcover. Since these coverings consist of countable collections $\mathcal{V}_1, \mathcal{V}_2, \mathcal{V}_3, \ldots$, one can also arrange them into nested sequences of open subsets

$$\mathcal{U}_1 := \mathcal{V}_1 \subset \mathcal{U}_2 := \mathcal{V}_1 \cup \mathcal{V}_2 \subset \mathcal{U}_3 := \mathcal{V}_1 \cup \mathcal{V}_2 \cup \mathcal{V}_3 \subset \ldots$$

whose unions cover everything. In other words, every manifold is the union of a nested sequence of open subsets that are each finite unions of convex sets. We therefore need a lemma for passing from a nested sequence of open subsets to its union.

Lemma 54.21. *Suppose $\mathcal{U}_1 \subset \mathcal{U}_2 \subset \mathcal{U}_3 \subset \ldots \subset M$ is a nested sequence of open subsets of an R-oriented n-manifold M such that $\bigcup_{i=1}^{\infty} \mathcal{U}_i = M$. If the duality map is an isomorphism on $\mathcal{U}_i$ for every $i \in \mathbb{N}$, then it is also an isomorphism on M.*

Proof. The idea is to present $H_{n-k}(M)$ and $H_c^k(M)$ as direct limits of the sequences of R-modules $H_{n-k}(\mathcal{U}_i)$ and $H_c^k(\mathcal{U}_i)$ respectively. In the former case, we already know how to do this: it is easy to check that the direct limit of the spaces $\{\mathcal{U}_i\}_{i=1}^{\infty}$ with respect to inclusion is M, and since every compact subset of M must be contained in $\mathcal{U}_i$ for i sufficiently large, Theorem 39.20 provides a natural isomorphism

$$\varinjlim \{H_*(\mathcal{U}_i)\}_{i=1}^{\infty} \xrightarrow{\cong} H_*\big(\varinjlim \{\mathcal{U}_i\}_{i=1}^{\infty}\big) = H_*(M).$$

For the cohomology, the fact that every $\mathcal{U}_i$ is open in $\mathcal{U}_j$ for $j > i$ and also open in M gives rise to natural maps

$$H_c^*(\mathcal{U}_1) \to H_c^*(\mathcal{U}_2) \to H_c^*(\mathcal{U}_3) \to \ldots \to H_c^*(M),$$

making $\{H_c^*(\mathcal{U}_i)\}_{i=1}^{\infty}$ a direct system, and we claim that $H_c^*(M)$ is its direct limit. This can proved by establishing the universal property: if we have a sequence of morphisms $f_i : H_c^*(\mathcal{U}_i) \to A$ to some other $\mathbb{Z}$-graded abelian R-module A such that the diagram

$$\begin{tikzcd}
H_c^*(\mathcal{U}_1) \arrow[r] \arrow[drrrr, "f_1"'] & H_c^*(\mathcal{U}_2) \arrow[r] \arrow[drrr] & H_c^*(\mathcal{U}_3) \arrow[r] \arrow[drr, "f_3"] & \ldots \arrow[r] & H_c^*(M) \arrow[d, dashed, "f_\infty"] \\
& & & & A
\end{tikzcd}$$

commutes, then we need to show that the map f_∞ in this diagram exists and is unique. To define $f_\infty(\varphi)$ for some $\varphi \in H_c^k(M)$, observe that φ is necessarily in the image of the natural map $H^k(M \,|\, K) \to H_c^k(M)$ for some compact set $K \subset M$, and since K is compact, it must be contained in $\mathcal{U}_N$ for $N \in \mathbb{N}$ sufficiently large. Excision then allows us to regard φ as an element of $H^k(\mathcal{U}_N \,|\, K)$, which therefore represents some element of $H_c^k(\mathcal{U}_N)$, so we define $f_\infty(\varphi)$ by applying f_N to this element. Proving that this is independent of choices is now a routine matter of writing down diagrams to check that they commute, so we shall leave it as an exercise.

By Lemma 54.18, we now obtain a commutative diagram

$$\begin{tikzcd}
H_c^k(\mathcal{U}_1) \arrow[r] \arrow[d, "\mathrm{PD}"] & H_c^k(\mathcal{U}_2) \arrow[r] \arrow[d, "\mathrm{PD}"] & H_c^k(\mathcal{U}_3) \arrow[r] \arrow[d, "\mathrm{PD}"] & \ldots \\
H_{n-k}(\mathcal{U}_1) \arrow[r] & H_{n-k}(\mathcal{U}_2) \arrow[r] & H_{n-k}(\mathcal{U}_3) \arrow[r] & \ldots
\end{tikzcd} \tag{54.5}$$

in which the vertical maps are all isomorphisms, thus it defines an isomorphism between the two direct systems. These therefore have a limiting map which is also an isomorphism, and one can check that the limiting map is PD:

$$\varinjlim \left\{H_c^k(\mathcal{U}_i)\right\}_{i=1}^{\infty} = H_c^k(M) \underset{\cong}{\xrightarrow{\mathrm{PD}}} H_{n-k}(M) = \varinjlim \left\{H_{n-k}(\mathcal{U}_i)\right\}_{i=1}^{\infty}.$$

$\square$

Proof of Theorem 54.16. Lemmas 54.17 and 54.19 prove the theorem for all finite unions of convex open sets $\mathbb{R}^n$, and feeding this into Lemma 54.21 then establishes it for all open subsets of $\mathbb{R}^n$. In a manifold M, the intersection of two open sets contained in Euclidean neighborhoods is also contained in a Euclidean neighborhood, so another application of Lemma 54.19 now proves the theorem for all finite unions of open subsets in Euclidean neighborhoods, and we can then present M is a nested union of such subsets and establish the theorem for M via a second application of Lemma 54.21. □

Exercise 54.22. Fix a coefficient ring R and assume M is a compact R-oriented n-manifold with boundary, with $[M] \in H_n(M, \partial M) = H_n(M, \partial M; R)$ as its relative fundamental class. The relative cap product with $[M]$ then gives rise to two natural maps

$$\mathrm{PD} : H^k(M, \partial M; R) \to H_{n-k}(M; R), \tag{54.6}$$

$$\mathrm{PD} : H^k(M; R) \to H_{n-k}(M, \partial M; R), \tag{54.7}$$

both defined by $\mathrm{PD}(\varphi) = \varphi \cap [M]$. The theorem that both are isomorphisms is sometimes called **Lefschetz duality**.

(a) Find a cofinal family of compact subsets $A \subset \mathring{M} := M \backslash \partial M$ such that the natural maps in the diagram

$$H^*(\mathring{M} \,|\, A) \longleftarrow H^*(M \,|\, A) \longrightarrow H^*(M, \partial M)$$

are isomorphisms. Use this to find a natural isomorphism (cf. Exercise 39.9)

$$H^*_c(M) \cong H^*(M, \partial M),$$

and deduce via Theorem 54.16 that (54.6) is an isomorphism.

(b) Show that the long exact sequenes of the pair $(M, \partial M)$ in homology and cohomology fit together into a commutative diagram of the form

$$\begin{tikzcd}
\ldots \arrow[r] & H^k(M,\partial M) \arrow[r, "j^*"] \arrow[d, "\cdot\cap[M]"] & H^k(M) \arrow[r, "i^*"] \arrow[d, "\cdot\cap[M]"] & H^k(\partial M) \arrow[r, "\delta^*"] \arrow[d, "\cdot\cap[\partial M]"] & H^{k+1}(M,\partial M) \arrow[r] \arrow[d, "\cdot\cap[M]"] & \ldots \\
\ldots \arrow[r] & H_{n-k}(M) \arrow[r, "j_*"] & H_{n-k}(M,\partial M) \arrow[r, "\partial_*"] & H_{n-k-1}(\partial M) \arrow[r, "i_*"] & H_{n-k-1}(M) \arrow[r] & \ldots,
\end{tikzcd}$$

where $i : \partial M \hookrightarrow M$ and $j : (M, \varnothing) \hookrightarrow (M, \partial M)$ denote the usual inclusions. *Hint: Work directly with chains and cochains. It helps to know that if $c \in C_n(M)$ is a relative n-cycle representing $[M] \in H_n(M, \partial M)$, then the $(n-1)$-cycle $\partial c \in C_{n-1}(\partial M)$ represents $[\partial M] \in H_{n-1}(\partial M)$; see Exercise 53.2.*

(c) Deduce from the diagram in part (b) that the map in (54.7) is also an isomorphism.

(d) If M has a triangulation, interpret the isomorphisms (54.6) and (54.7) in terms of the dual cell decomposition.

Remark 54.23. Here is the promised addendum to Remark 54.1. When M is compact and has an oriented triangulation, $C^k_\Delta(M; \mathbb{Z})$ has an obvious identification with the free abelian group generated by all the k-simplices in the triangulation: indeed, if we fix an orientation on each k-simplex and call $\mathcal{K}_k(M)$ the resulting set of oriented k-simplices so that $C_k^\Delta(M; \mathbb{Z}) = \bigoplus_{\sigma \in \mathcal{K}_k(M)} \mathbb{Z}$, then the dual elements $\varphi_\sigma : C_k^\Delta(M; \mathbb{Z}) \to \mathbb{Z}$ defined on generators $\tau \in \mathcal{K}_k(M)$ by

$$\varphi_\sigma(\tau) := \begin{cases} 1 & \text{if } \tau = \sigma, \\ 0 & \text{if } \tau \neq \sigma \end{cases}$$

form a basis for $C^k_\Delta(M;\mathbb{Z})$. In this case, we obtain a chain isomorphism

$$C^k_\Delta(M;\mathbb{Z}) \to C^{\mathrm{CW}}_{n-k}(M;\mathbb{Z})$$

by sending each of the k-cochains φ_σ to the $(n-k)$-cell dual to σ, and the isomorphism $H^k(M;\mathbb{Z}) \cong H_{n-k}(M;\mathbb{Z})$ follows. The trouble if M is not compact is that $C_k^\Delta(M;\mathbb{Z})$ is now an infinitely-generated free abelian group, so its dual $C^k_\Delta(M;\mathbb{Z})$ is not isomorphic to it, but is actually much larger: the cochains φ_σ do not form a basis for $C^k_\Delta(M;\mathbb{Z})$ since they only span the subgroup of homomorphisms $C_k^\Delta(M;\mathbb{Z}) \to \mathbb{Z}$ that are nonzero on finitely many simplices. As a consequence, $C^k_\Delta(M;\mathbb{Z})$ and $C^{\mathrm{CW}}_{n-k}(M;\mathbb{Z})$ are not isomorphic, but now that you've seen how Poincaré duality works for singular homology on noncompact manifolds, you may be able to guess how to fix this: the cochains φ_σ do span a *subcomplex* of $C^*_\Delta(M;\mathbb{Z})$, whose homology is the simplicial version of $H^*_c(M;\mathbb{Z})$.

55. The intersection product (February 9, 2024)

Introduction. Today's topic is an addendum to Poincaré duality: I want to describe the natural product structure on *homology* that arises from the combination of Poincaré duality with the cup product. Unlike the cup product, the product on $H_*(X)$ will not be defined for arbitrary spaces X, but makes sense only when Poincaré duality holds, i.e. when X is a closed manifold with an orientation over the chosen coefficient ring. This is a bit restrictive, but the restriction pays off: in fact, if we restrict further and assume X is a *smooth* manifold, then the intersection product provides the nicest possible geometric interpretation of the cup product, namely as something that measures (in homological terms) the intersection between submanifolds.

Our standing assumptions throughout this lecture are as follows: M is a closed, connected and R-oriented smooth manifold of dimension n, where the coefficient ring R will always be either $\mathbb{Z}$, $\mathbb{Z}_2$, $\mathbb{Q}$ or $\mathbb{R}$. One can allow more general choices for R, but these are the main ones of interest. The assumption of an R-orientation actually just means that M is oriented if R is $\mathbb{Z}$, $\mathbb{Q}$ or $\mathbb{R}$, and in the case $R = \mathbb{Z}_2$ it is a vacuous assumption. The smoothness assumption can also be relaxed somewhat, but it is the quickest way to achieve the conditions that we will actually need at various points: namely that suitable intersections $A \cap B$ between two submanifolds $A, B \subset M$ are also submanifolds (of the "correct" dimension), and that neighborhoods of submanifolds can be identified with certain vector bundles (the tubular neighborhood theorem). For this reason, the present lecture will assume some knowledge of the basic theory of smooth manifolds and their tangent spaces.

Definition 55.1. The **intersection product** on M associates to each pair of integers $k, \ell = 0, \dots, n$ a bilinear map

$$H_{n-k}(M;R) \otimes_R H_{n-\ell}(M;R) \to H_{n-(k+\ell)}(M;R) : A \otimes B \mapsto A \cdot B$$

uniquely defined by the condition

$$\mathrm{PD}(\varphi) \cdot \mathrm{PD}(\psi) = \mathrm{PD}(\psi \cup \varphi),$$

for $\varphi \in H^k(M;R)$ and $\psi \in H^\ell(M;R)$, where $\mathrm{PD} : H^m(M;R) \to H_{n-m}(M;R)$ denotes the Poincaré duality isomorphism. In the case $k + \ell = n$, it is useful to recall that M is assumed connected, so there is a canonical isomorphism $H_0(M;R) \xrightarrow{\cong} R : c \mapsto \langle 1, c\rangle$ which can be used to regard $A \cdot B$ as a number in R, the so-called **intersection number** between A and B,

$$H_{n-k}(M;R) \otimes_R H_k(M;R) \to R : A \otimes B \mapsto A \cdot B.$$

The intersection number is equivalent to what we called the *intersection form* $Q : H^k(M;R) \otimes_R H^{n-k}(M;R) \to R$ in the previous lecture: if $\varphi \in H^k(M;R)$ and $\psi \in H^{n-k}(M;R)$ have Poincaré dual classes $A := \mathrm{PD}(\varphi) \in H_{n-k}(M;R)$ and $B := \mathrm{PD}(\psi) \in H_k(M;R)$, then a precise relation is given by

$$\begin{aligned} A \cdot B &= \langle 1, A \cdot B \rangle = \langle 1, \mathrm{PD}(\psi \cup \varphi) \rangle \\ &= \langle 1, (\psi \cup \varphi) \cap [M] \rangle = \langle 1, \psi \cap (\varphi \cap [M]) \rangle = \langle \psi, \varphi \cap [M] \rangle \\ &= \langle \psi \cup \varphi, [M] \rangle = Q(\psi, \varphi). \end{aligned} \tag{55.1}$$

Corollary 54.8 therefore implies:

COROLLARY 55.2. *If M is a closed, connected and oriented n-manifold, the intersection number defines a nonsingular bilinear form*

$$H_{n-k}^{\mathrm{free}}(M) \otimes H_k^{\mathrm{free}}(M) \to \mathbb{Z} : A \otimes B \mapsto A \cdot B,$$

for every $k = 0, \dots, n$, so in particular, for every primitive element $A \in H_{n-k}(M)$ there exists a class $B \in H_k(M)$ with $A \cdot B = 1$. If M is instead assumed to be orientable over a field $\mathbb{K}$, then the intersection number similarly defines a nonsingular $\mathbb{K}$-bilinear form

$$H_{n-k}(M;\mathbb{K}) \otimes_{\mathbb{K}} H_k(M;\mathbb{K}) \to \mathbb{K}.$$

□

We can also extract from the end of the second line of (55.1) the following useful formula: for every $\varphi \in H^k(M;R)$ and $A \in H_k(M;R)$,

$$\langle \varphi, A \rangle = A \cdot \mathrm{PD}(\varphi). \tag{55.2}$$

Once we have understood how to interpret the intersection product geometrically, this formula will yield some intuitive insight into the isomorphism $\mathrm{PD} : H^k(M;R) \to H_{n-k}(M;R)$, i.e. it transforms the natural evaluation of cohomology classes on homology classes into the operation of counting intersections between homology classes. In many situations, this can be used to compute Poincaré dual classes explicitly.

To understand what $A \cdot B$ means in terms of "counting intersections," we consider the following scenario: suppose M has a smooth structure and A and B are closed *smooth submanifolds* of M with dimensions $n - k$ and $n - \ell$ respectively, i.e. their *codimensions* are k and ℓ. We also assume that A and B carry R-orientations, in which case they have well-defined fundamental classes $[A] \in H_{n-k}(A;R)$ and $[B] \in H_{n-\ell}(B;R)$, and we shall use the same notation for the classes in $H_*(M;R)$ obtained by feeding these into the maps induced by the inclusions $A, B \hookrightarrow M$, that is,

$$[A] \in H_{n-k}(M;R), \qquad [B] \in H_{n-\ell}(M;R).$$

Note that in practice, the orientation assumption just means that A and B are both oriented if the coefficient ring R is chosen from among $\mathbb{Z}$, $\mathbb{Q}$ and $\mathbb{R}$, whereas there is no orientation assumption at all if we use $R = \mathbb{Z}_2$.

We now need a few basic notions from differential topology. The first is the *smooth* version of an orientation, which is simpler than what we have defined for topological manifolds. For our purposes, it is best to express this in terms of tangent spaces: since M has a smooth structure, there is a tangent space T_xM associated to every point $x \in M$, which is a real vector space of dimension n. In general, an **orientation** of an n-dimensional vector space V is defined to be an equivalence class of bases of V, where two bases are equivalent if and only if one can be deformed to the other through a continuous family of bases. There are always two equivalence classes, due to the fact that the group $\mathrm{GL}(n,\mathbb{R})$ has two connected components, distinguished by the sign of the determinant: we call the bases in the preferred equivalence class **positively oriented** and all others **negatively oriented**. An orientation of the smooth manifold M can then be defined as a choice of orientation for every tangent space T_xM that varies continuously with respect to x.

To make this precise, we would need to define the appropriate topology on the *tangent bundle* $TM = \bigcup_{x \in M} T_x M$ and discuss what it means for a vector field to be continuous, but it should at least be intuitively clear what is meant, and since this is only meant as a survey, we'll leave it at that. One can check that this notion of orientation is equivalent to the various other notions of orientation that we've seen before, i.e. in terms of local homology groups or orientation-preserving coordinate transformations.

An important fact about tangent spaces is that if $A \subset M$ is a smooth submanifold, then each of its tangent spaces $T_x A$ is naturally a linear subspace of $T_x M$. This is enough background information to define the important notion of *transversality.*

DEFINITION 55.3. We say that two smooth submanifolds $A, B \subset M$ are **transverse** and write "$A \pitchfork B$" if for every $x \in A \cap B$,

$$T_x A + T_x B = T_x M.$$

It is an easy exercise in linear algebra to show that if the condition $T_x A + T_x B = T_x M$ holds where $T_x A \subset T_x M$ and $T_x B \subset T_x M$ have codimensions k and ℓ respectively, then $T_x A \cap T_x B \subset T_x M$ is a subspace with codimension $k + \ell$. The following nonlinear version of this observation is a standard application of the implicit function theorem.

PROPOSITION 55.4. *If $A, B \subset M$ are closed smooth submanifolds of codimensions k and ℓ respectively and are transverse to each other, then $A \cap B \subset M$ is also a closed smooth submanifold, with codimension $k + \ell$.* □

We can also add orientations to this picture. Suppose $V, W \subset \mathbb{R}^n$ are two oriented linear subspaces of codimensions k and ℓ such that $V + W = \mathbb{R}^n$, so $V \cap W$ has codimension $k + \ell$. Define the orientation of $\mathbb{R}^n$ so that the standard basis $(e_1, \dots, e_n)$ is considered positively oriented. We can then define an orientation of $V \cap W$ as follows. If either of V or W is contained in the other, then their intersection is the smaller subspace and thus already has an orientation, so assume this is not the case. Then given a basis $(X_1, \dots, X_{n-(k+\ell)})$ of $V \cap W$, it is always possible to choose additional vectors $Y_1, \dots, Y_\ell \in V$ and $Z_1, \dots, Z_k \in W$ such that the ordered tuples

$$(X_1, \dots, X_{n-(k+\ell)}, Y_1, \dots, Y_\ell) \text{ in } V, \qquad (X_1, \dots, X_{n-(k+\ell)}, Z_1, \dots, Z_k) \text{ in } W$$

both form positively oriented bases, and

$$(X_1, \dots, X_{n-(k+\ell)}, Y_1, \dots, Y_\ell, Z_1, \dots, Z_k)$$

is then a basis of $\mathbb{R}^n$. We define the orientation of $V \cap W$ such that $(X_1, \dots, X_{n-(k+\ell)})$ is positively oriented if and only if this basis of $\mathbb{R}^n$ is positively oriented. One can check that the only choices this definition depends on are the orientations of V, W and $\mathbb{R}^n$, plus the choice to write V in front of W instead of vice versa. Doing the same thing with the tangent spaces at all points $x \in A \cap B$ of a transverse intersection gives:

PROPOSITION 55.5. *In the setting of Proposition 55.4, if A, B and M are all oriented, then $A \cap B$ inherits from this data a natural orientation.* □

REMARK 55.6. It should be emphasized that according to the definition above, the orientation of $A \cap B$ may be different from that of $B \cap A$, i.e. they are the same submanifold, but the choice to write A in front of B or vice versa may determine different orientations. It is an easy exercise to show that the orientations differ if and only if the codimensions of A and B are both odd. (This should remind you of a graded commutativity relation—there is good reason for that!)

We should add a word about the case $k + \ell = n$, which does not quite fit into the above discussion since we have not properly defined what an orientation of a 0-dimensional vector space should mean. The definition is consistent with the notion that there should always be exactly two choices of orientation: if $\dim V = 0$, we define an **orientation** of V to mean a choice of sign ± 1. To see that this is a sensible definition, consider the situation where $V, W \subset \mathbb{R}^n$ are transverse linear subspaces of dimensions k and $n-k$, which means $V \oplus W = \mathbb{R}^n$, and $V \cap W$ is thus a 0-dimensional subspace. There are exactly two possibilities: choosing positively oriented bases $(X_1, \dots, X_k)$ of V and $(Y_1, \dots, Y_{n-k})$ of W, the basis

$$(X_1, \dots, X_k, Y_1, \dots, Y_{n-k})$$

of $\mathbb{R}^n$ is either positively or negatively oriented, and we define the orientation of $V \cap W$ to be $+1$ or -1 accordingly. Applying this idea to tangent spaces, the intersection $A \cap B$ between two transverse oriented closed submanifolds $A, B \subset M$ of complementary dimensions $\dim A + \dim B = \dim M$ is simply a finite set of points $x \in A \cap B$ with attached signs

$$\epsilon(x) = \pm 1,$$

determined as described above from the orientations of the complementary tangent spaces $T_x A \oplus T_x B = T_x M$. You should take a moment to convince yourself that this notion of the orientation of a 0-manifold is consistent with the definition we already had for *topological* 0-manifolds: indeed, if M is a discrete set, then each of the local homology groups $H_0(M, M \backslash \{x\}; \mathbb{Z})$ is *canonically* isomorphic to $H_0(\{\text{pt}\}; \mathbb{Z}) = \mathbb{Z}$, so a local orientation is a choice of generator of the group $\mathbb{Z}$, i.e. either $+1$ or -1.

We can now state the main theorem of this lecture.

THEOREM 55.7. *Assume M is a closed, connected, smooth and R-oriented manifold of dimension n, and $A, B \subset M$ are closed, smooth, R-oriented submanifolds of codimensions k and ℓ respectively, such that $A \pitchfork B$. Then for the induced R-orientation on $A \cap B$ from Proposition 55.5,*

$$[A] \cdot [B] = [A \cap B].$$

As usual, the case $k + \ell = n$ deserves special comment. If A and B are oriented in this case, $A \cap B$ is a finite set of points x with attached signs $\epsilon(x) = \pm 1$, and since M is connected, the canonical isomorphism $H_0(M; \mathbb{Z}) = \mathbb{Z}$ identifies $[A \cap B]$ with the integer

$$[A] \cdot [B] = \sum_{x \in A \cap B} \epsilon(x) \in \mathbb{Z}.$$

The right hand side of this expression is sometimes called the *algebraic* (or *signed*) count of transverse intersections between A and B.

Applications. In many situations of geometric interest, the intersection product provides an easy criterion for recognizing when a homology class is nontrivial.

EXERCISE 55.8. Draw some examples of pairs of transversely intersecting closed oriented 1-dimensional submanifolds $A, B \subset S^2$, and convince yourself that the signed count of intersections between them will always be 0. (Indeed, this must be true for at least two reasons: first, since $H_1(S^2) = 0$, both submanifolds represent the trivial homology class and it follows that $[A] \cdot [B] = 0$. Alternatively, if you believe the implication of Theorem 55.7 that the signed count of intersections between A and B only depends on their homology classes, then you can easily adjust A or B by a homotopy (i.e. contracting A so that it lies in an arbitrarily small neighborhood) so that $A \cap B = \varnothing$.)

EXAMPLE 55.9. Figure 32 shows a closed, connected and orientable surface Σ with four oriented 1-dimensional submanifolds $\alpha, \beta, \gamma, \delta \subset \Sigma$, or equivalently, loops $S^1 \hookrightarrow \Sigma$. Since α bounds a disk, it is clearly nullhomotopic, and therefore also nullhomologous, i.e. $[\alpha] = 0 \in H_1(\Sigma; \mathbb{Z})$. One can show by computations of $\pi_1(\Sigma)$ that β is not nullhomotopic, but it clearly is nullhomologous: this follows from the observation that β splits Σ into two connected components, a pair of compact

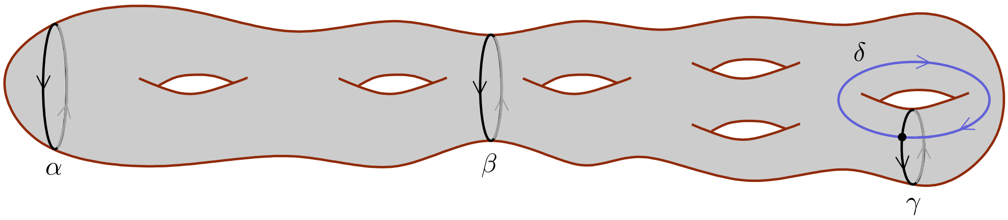

FIGURE 32. The surface and 1-dimensional submanifolds discussed in Example 55.9.

oriented surfaces $\Sigma_\pm$ with boundary $\partial\Sigma_\pm = \beta$ such that $\Sigma = \Sigma_+ \cup_\beta \Sigma_-$. If we factor the inclusion $i : \beta \hookrightarrow \Sigma$ through the inclusions $\beta \hookrightarrow \Sigma_+$ and $\Sigma_+ \hookrightarrow \Sigma$, we notice that the induced map $H_1(\beta;\mathbb{Z}) \to H_1(\Sigma;\mathbb{Z})$ is zero because the map $H_1(\beta;\mathbb{Z}) \to H_1(\Sigma_+;\mathbb{Z})$ is zero (see Exercise 53.2(c)),

$$\begin{tikzcd} H_1(\beta;\mathbb{Z}) \arrow[r, "0"] \arrow[rr, bend right, "i_*"'] & H_1(\Sigma_+;\mathbb{Z}) \arrow[r] & H_1(\Sigma), \end{tikzcd}$$

hence $i_*[\beta] = 0$. The case of $\gamma \subset \Sigma$ is less obvious: it does not split Σ in two pieces, as $\Sigma\backslash\gamma$ is connected, thus it is hard to imagine a 2-chain in Σ that would have γ as its boundary, but this on its own is not a proof that no such chain exists. The intersection product, however, provides a clear criterion showing that $[\gamma] \in H_1(\Sigma;\mathbb{Z})$ cannot be zero: the reason is that there is another loop, $\delta \subset \Sigma$, which intersects γ exactly once transversely, hence their intersection product must satisfy

$$[\gamma] \cdot [\delta] = \pm 1.$$

This proves that both of the classes $[\gamma], [\delta] \in H_1(\Sigma;\mathbb{Z})$ are not only nontrivial but also primitive.

The nonseparating loops in Example 55.9 admit the following interesting generalization. If M is an n-manifold, a submanifold $\Sigma \subset M$ is called a **hypersurface** if $\dim \Sigma = n - 1$. Assuming M is connected, we say that $\Sigma \subset M$ **separates** M if $M\backslash\Sigma$ is disconnected.

THEOREM 55.10. *Suppose M is a closed, connected and R-oriented smooth n-manifold containing a closed, connected and R-oriented smooth hypersurface $\Sigma \subset M$. Then the homology class $[\Sigma] \in H_{n-1}(M;R)$ is trivial if and only if Σ separates M.*

PROOF. If Σ separates M then we can write $M = M_+ \cup_\Sigma M_-$ where $M_\pm$ are two compact R-oriented n-manifolds with boundary $\partial M_\pm = \Sigma$, so the same argument as in Example 55.9 implies that $[\Sigma] = 0$. On the other hand, if Σ does not separate M, then $M\backslash\Sigma$ is connected, so we can fix a point $z \in \Sigma$ and two nearby points $z_\pm \in M\backslash\Sigma$ that lie in a common Euclidean neighborhood with z identifying Σ with $\mathbb{R}^{n-1} \times \{0\} \subset \mathbb{R}^n$,[94] but on opposite sides of Σ, i.e. the nth coordinates of z_+ and z_- have opposite signs. We can then find a smooth path γ joining z_+ to z_- in $M\backslash\Sigma$, and then complete it with a path in the Euclidean neighborhood that passes through Σ once, producing (as in Figure 32) a smooth loop $\gamma \subset M$ that intersects Σ exactly once and transversely. It follows that

$$[\Sigma] \cdot [\gamma] = \pm 1,$$

hence $[\Sigma] \in H_{n-1}(M;R)$ and $[\gamma] \in H_1(M;R)$ are both nontrivial. $\square$

[94] One of the standard ways of characterizing a *smooth submanifold* $\Sigma \subset M$ is through the existence of **slice charts**: for every $x \in \Sigma$, some neighborhood $\mathcal{U} \subset M$ of x admits a smooth chart $\varphi : \mathcal{U} \xrightarrow{\cong} \varphi(\mathcal{U}) \overset{\text{open}}{\subset} \mathbb{R}^n$ that identifies a neighborhood of x in Σ with an open subset of the linear subspace $\mathbb{R}^k \times \{0\} \subset \mathbb{R}^n$ for $k = \dim \Sigma$.

Remark 55.11. Did you notice where we cheated a bit in the proof of Theorem 55.10? We constructed a smooth map $\gamma : S^1 \to M$ and argued that the intersection number $[\Sigma] \cdot [\gamma]$ is ± 1 for obvious geometric reasons, but strictly speaking, a smooth map $\gamma : S^1 \to M$ is not the same thing as a smooth closed 1-dimensional submanifold $S \subset M$, which is what we'd actually need in order to apply Theorem 55.7 as it was stated. If $\dim M \geqslant 3$, then one can appeal to general facts about transversality in order to make $\gamma : S^1 \to M$ an embedding without loss of generality, which is just as good as a submanifold; however, this trick does not work if M is a surface, as there is not enough room in two dimensions to perturb every map $S^1 \to M$ to an embedding. There is nonetheless a different cheap trick that works without any dimensional assumptions: instead of counting intersections between a submanifold $\Sigma \subset M$ and a map $\gamma : S^1 \to M$, one can count the intersections between two submanifolds of $S^1 \times M$, namely

$$S^1 \times \Sigma \subset S^1 \times M, \qquad \text{and} \qquad \Gamma_\gamma := \left\{ (t, \gamma(t)) \in S^1 \times M \;\middle|\; t \in S^1 \right\}.$$

This trick can be used more generally to extend the validity of Theorem 55.7 to homology classes in $H_*(M)$ represented by transversely intersecting smooth maps $f : A \to M$ and $g : B \to M$ defined on closed manifolds A and B; the manifolds A and B need not be given as submanifolds of M, and the maps f and g need not be embeddings.

Remark 55.12. In the case of Theorem 55.10 with coefficients in $\mathbb{Z}$, the argument shows that the class $[\Sigma] \in H_{n-1}(M; \mathbb{Z})$ represented by a non-separating oriented hypersurface is not just nontrivial, but also primitive.

If you've been wondering why non-orientable surfaces like $\mathbb{RP}^2$ and the Klein bottle cannot be embedded in $\mathbb{R}^3$, we can now answer this question. If you can embed them in $\mathbb{R}^3$, then you can also embed them in its one-point compactification, S^3, which is prevented by the following corollary:

Corollary 55.13. *For every $n \geqslant 2$, closed smooth hypersurfaces in S^n are always orientable.*

Proof. Suppose to the contrary that $\Sigma \subset S^n$ is a closed non-orientable smooth hypersurface, and without loss of generality assume Σ is connected. Then one can find (as in the proof of Theorem 55.10) a path in $S^n \backslash \Sigma$ that stays within a small neighborhood of Σ but starts and ends on opposite sides of it, thus giving rise to a loop $\gamma : S^1 \to S^n$ that intersects Σ once transversely. Using $\mathbb{Z}_2$ coefficients (since Σ is orientable over $\mathbb{Z}_2$), the intersection number of Σ with γ is then

$$[\Sigma] \cdot [\gamma] = 1 \in \mathbb{Z}_2,$$

implying $[\Sigma] \neq 0 \in H_{n-1}(S^n; \mathbb{Z}_2)$ and $[\gamma] \neq 0 \in H_1(S^n; \mathbb{Z}_2)$. This contradicts are computation of $H_*(S^n; \mathbb{Z}_2)$. □

Remark 55.14. The fact that S^n is orientable is not the decisive factor in Corollary 55.13, as there is no obstruction in general to embedding closed non-orientable hypersurfaces into closed orientable manifolds. An easy example is $\mathbb{RP}^2 \hookrightarrow \mathbb{RP}^3$.

You may have noticed a resemblance between the signed count of intersections $[A] \cdot [B]$ and the sum of local mapping degrees discussed in Lecture 35. Indeed, for any smooth map $f : M \to N$ between a pair of closed, oriented and connected n-manifolds, counting points in $f^{-1}(y)$ for $y \in N$ is equivalent to counting intersections of the oriented n-dimensional submanifolds

$$\Gamma_f := \left\{ (x, f(x)) \;\middle|\; x \in M \right\} \subset M \times N, \qquad \text{and} \qquad M_y := M \times \{y\} \subset M \times N, \tag{55.3}$$

which are transverse if and only if y is a regular value of f. Theorem 35.14 on the relation between global and local mapping degrees thus follows in this setting from Theorem 55.7 and the following exercise; it should go without saying that a version without orientation assumptions also holds if one uses $\mathbb{Z}_2$ coefficients.

Exercise 55.15. For a smooth map $f : M \to N$ between two closed, oriented and connected n-manifolds and $k \in \mathbb{Z}$ the unique integer such that $f_*[M] = k[N] \in H_n(N;\mathbb{Z})$, show that the submanifolds defined in (55.3) satisfy $[\Gamma_f] \cdot [M_y] = \pm k$. (Life is too short to worry about the sign.)

As a fancier variation on the same theme, we should mention briefly that the intersection product provides a beautiful reinterpretation of the Lefschetz fixed point theorem in the smooth category. Suppose M is a closed, connected and oriented smooth n-manifold and $f : M \to M$ is a smooth map. The fixed point set $\mathrm{Fix}(f) = \{x \in M \mid f(x) = x\}$ is then in one-to-one correspondence with the intersection of the two smooth submanifolds

$$\Delta := \{(x,x) \mid x \in M\} \subset M \times M,$$
$$\Gamma_f := \{(x,f(x)) \mid x \in M\} \subset M \times M.$$

Both are closed and inherit from M obvious orientations.

Theorem 55.16. *The homological intersection number $[\Delta] \cdot [\Gamma_f] \in \mathbb{Z}$ is the Lefschetz number $L(f)$.*

The smooth oriented case of the Lefschetz fixed point theorem follows from this and Theorem 55.7 immediately, since $L(f) \neq 0$ now implies that Δ and Γ_f cannot be disjoint. The orientation condition can also be dropped by using coefficients in $\mathbb{Z}_2$, in which case the analogous theorem identifies $[\Delta]\cdot[\Gamma_f] \in \mathbb{Z}_2$ with $L_{\mathbb{Z}_2}(f)$. In either case, the Lefschetz number is no longer just a criterion for the *existence* of a fixed point, but is actually a quantitative *count* of fixed points, so long as the concept of "counting" is understood in the proper generalized sense (e.g. counting with signs). For a proof of Theorem 55.16, see [**Hut**].

Tubular neighborhoods. The coefficient ring R will be omitted from the notation wherever possible in the remainder of this lecture.

To prove Theorem 55.7, we need another basic result from differential topology, which gives a model for the neighborhood of any smooth submanifold $A \subset M$. Assume $\dim M = n$ and $\dim A = n - k$, so A has codimension k. The **normal bundle** of A is defined as the union of k-dimensional quotient vector spaces

$$N^M A := \bigcup_{x \in A} N^M_x A, \qquad \text{where} \qquad N^M_x A := T_x M / T_x A.$$

When there is no ambiguity about the ambient manifold, we will usually abbreviate

$$NA := N^M A, \qquad N_x A := N^M_x A.$$

This is an example of a *vector bundle*; its **fibers** are the individual vector spaces $N_x A$, which are also the preimages of points under the natural projection map

$$\pi : NA \to A : N_x A \mapsto x \text{ for all } x \in A.$$

One can define a natural topology and smooth structure on NA so that it becomes a smooth n-dimensional manifold and $\pi : NA \to A$ is a smooth map. The determining feature of this topology and smooth structure is that the map $\pi : NA \to A$ is **locally trivializable**, meaning every point $x \in A$ has a neighborhood $\mathcal{U} \subset A$ admitting a diffeomorphism

$$\Phi_{\mathcal{U}} : \pi^{-1}(\mathcal{U}) \to \mathcal{U} \times \mathbb{R}^k$$

that restricts to a linear isomorphism $N_y A \to \{y\} \times \mathbb{R}^k$ for each $y \in \mathcal{U}$. A local trivialization thus identifies the map $\pi^{-1}(\mathcal{U}) \xrightarrow{\pi} \mathcal{U}$ with the obvious projection map $\mathcal{U} \times \mathbb{R}^k \to \mathcal{U}$. The existence of local trivializations is a straightforward consequence of the existence of slice charts near every point of A. If NA and A are both endowed with orientations, then this local product structure also determines an orientation of each of the fibers $N_x A \cong \mathbb{R}^k$, which can be viewed as a family of orientations varying continuously with $x \in A$. The subset of NA consisting of all 0-vectors in the spaces $N_x A$ is a smooth submanifold that is canonically diffeomorphic to A, called the **zero-section**. We shall often regard A itself as a submanifold of NA and write

$$A \subset NA$$

by identifying A with the zero-section.

If $A \subset M$ and $B \subset M$ are submanifolds that intersect transversely, then for each $x \in A \cap B$, the diagonal map $T_x M \to T_x M \oplus T_x M$ descends to a canonical isomorphism

$$N_x^M(A \cap B) \xrightarrow{\cong} N_x^M A \oplus N_x^M B,$$

producing a vector bundle isomorphism of $N^M(A \cap B)$ with the direct sum of the restrictions of $N^M A$ and $N^M B$ along the submanifold $A \cap B$. Since $A \cap B$ is also a submanifold of A and B, we can similarly consider the normal bundles $N^A(A \cap B)$ and $N^B(A \cap B)$, and notice that for each $x \in A \cap B$, the inclusions $T_x A \hookrightarrow T_x M$ and $T_x B \hookrightarrow T_x M$ descend to canonical isomorphisms

$$N_x^A(A \cap B) \xrightarrow{\cong} N_x^M B, \qquad N_x^B(A \cap B) \xrightarrow{\cong} N_x^M A,$$

giving vector bundle isomorphisms $N^A(A \cap B) \cong N^M B|_{A \cap B}$ and $N^B(A \cap B) \cong N^M A|_{A \cap B}$. In this way we can regard $N^B(A \cap B)$ as a subset of $N^M A$ and $N^A(A \cap B)$ as a subset of $N^M B$; in fact, both are smooth submanifolds whose codimensions match the codimensions of B and A respectively.

This is enough background to state the **tubular neighborhood theorem**. Its proof is a fairly straightforward matter of defining a smooth map $N^M A \to M$ whose derivative along the zero-section is the identity map, and then citing the inverse function theorem.

THEOREM 55.17. *There exists a smooth embedding $N^M A \hookrightarrow M$ that is a diffeomorphism onto a neighborhood of A and matches the inclusion $A \hookrightarrow M$ along the zero-section. Moreover, if $B \subset M$ is another smooth submanifold that intersects A transversely, then the embedding $N^M A \hookrightarrow M$ can be arranged so that it maps $N^B(A \cap B) \subset N^M A$ onto a neighborhood of $A \cap B$ in B.* □

The Thom isomorphism. We shall use the tubular neighborhood theorem in the following to identify an open neighborhood of $A \subset M$ with the normal bundle NA. It will be useful also to shrink this to a smaller *compact* neighborhood that is a smooth manifold with boundary. To define this, choose a family of inner products $\langle\ ,\ \rangle$ on the tangent spaces $T_x M$ that vary smoothly with x; this is what is called a *Riemannian metric* on M. The inner product on each $T_x M$ determines an isomorphism of $N_x^M A$ with the orthogonal complement of $T_x A \subset T_x M$ and thus (by restriction) also determines an inner product on $N_x^M A$. We can then define

$$\mathbb{D}N_x A := \{X \in N_x A \mid \langle X, X\rangle \leqslant 1\}, \qquad \mathbb{D}NA := \bigcup_{x \in A} \mathbb{D}N_x A,$$

$$\mathbb{S}N_x A := \{X \in N_x A \mid \langle X, X\rangle = 1\}, \qquad \mathbb{S}NA := \bigcup_{x \in A} \mathbb{S}N_x A.$$

Each $\mathbb{D}N_x A$ is now a compact k-disk with boundary $\mathbb{S}N_x A \cong S^{k-1}$, and their union $\mathbb{D}NA$ is a compact n-manifold with boundary $\mathbb{S}NA$, called the **unit disk bundle** in NA. The tubular neighborhood theorem then identifies $\mathbb{D}NA$ with a compact neighborhood of A in M, while simultaneously identifying the subset $\mathbb{D}N^B(A \cap B) \subset \mathbb{D}NA$ with a compact neighborhood of $A \cap B$ in B.

Since $\mathbb{D}NA$ is an n-dimensional submanifold of M and is also compact with boundary, it inherits an R-orientation from M, and therefore has a relative fundamental class $[\mathbb{D}NA] \in H_n(\mathbb{D}NA, \mathbb{S}NA)$ and (thanks to Exercise 54.22) a Poincaré-Lefschetz duality isomorphism

$$\mathrm{PD} : H^k(\mathbb{D}NA, \mathbb{S}NA) \xrightarrow{\cong} H_{n-k}(\mathbb{D}NA) : \varphi \mapsto \varphi \cap [\mathbb{D}NA].$$

In the following, it will be important to specify umanbiguously which homology classes belong to which spaces or pairs, so let us reserve the notation $[A]$ for the fundamental class in $H_{n-k}(A)$, and denote the class in $H_{n-k}(M)$ that appears in the statement of Theorem 55.7 by

$$(i_A^M)_*[A] \in H_{n-k}(M).$$

We shall use the notation $i_X^Y : X \hookrightarrow Y$ for the inclusion of any subspace $X \subset Y$, so for instance the homology class represented by the zero-section in NA is

$$(i_A^{\mathbb{D}NA})_*[A] \in H_{n-k}(\mathbb{D}NA).$$

Note that since orientations on NA and A determine orientations on the fibers, the individual disks $\mathbb{D}N_xA \cong \mathbb{D}^k$ for each $x \in A$ also have well-defined fundamental classes $[\mathbb{D}N_xA] \in H_k(\mathbb{D}N_xA, \mathbb{S}N_xA)$. For $x \in A$, let us abbreviate the inclusion of the fiber by

$$f_x^A : (\mathbb{D}N_xA, \mathbb{S}N_xA) \hookrightarrow (\mathbb{D}NA, \mathbb{S}NA),$$

so $(f_x^A)_*[\mathbb{D}N_xA] \in H_k(\mathbb{D}NA, \mathbb{S}NA)$.

Definition 55.18. The **Thom class** of the normal bundle NA is

$$\tau(NA) := \mathrm{PD}^{-1}\left((i_A^{\mathbb{D}NA})_*[A]\right) \in H^k(\mathbb{D}NA, \mathbb{S}NA).$$

The Thom class is determined by the vector bundle NA, i.e. it depends on its properties as a vector bundle, but not on the fact that it is the normal bundle of a submanifold $A \subset M$. More generally, for any vector bundle E with R-oriented k-dimensional fibers over a closed R-oriented manifold M, one can choose inner products on the fibers to define a unit disk bundle $\mathbb{D}E \subset E$ and define the Thom class $\tau(E) \in H^k(\mathbb{D}E, \mathbb{S}E)$ as the Poincaré dual of the homology class of the zero-section $M \subset E$. By excision, this can be identified with a class in $H^k(E, E \backslash M)$, thus it does not depend on the choice of inner products used to define $\mathbb{D}E$ and $\mathbb{S}E$. We have chosen to formulate the definition above in a less general way, since it can then be understood without any concrete knowledge of the theory of vector bundles, i.e. you can simply think of $\mathbb{D}NA$ as a compact neighborhood of A in M, and $\mathbb{S}NA$ as the boundary of that neighborhood.

The next theorem is similarly true for all vector bundles that have a well-defined Thom class, but we are stating it for the specific case that we need for applications.

Theorem 55.19 (Thom isomorphism theorem). *The Thom class $\tau(NA) \in H^k(\mathbb{D}NA, \mathbb{S}NA)$ is uniquely determined by the condition*

$$\langle \tau(NA), (f_x^A)_*[\mathbb{D}N_xA] \rangle = 1 \quad \textit{for every } x \in A.$$

Moreover, the map

$$H^m(A) \to H^{m+k}(\mathbb{D}NA, \mathbb{S}NA) : \varphi \mapsto \pi^*\varphi \cup \tau(NA)$$

is an isomorphism.

Proof. One can contract all fibers $\mathbb{D}N_xA$ to their origins to produce a deformation retraction of $\mathbb{D}NA$ to its zero-section $A \subset \mathbb{D}NA$, thus the inclusion $i_A^{\mathbb{D}NA} : A \hookrightarrow \mathbb{D}NA$ and projection $\pi : \mathbb{D}NA \to A$ are homotopy inverses, implying that

$$\pi^* : H^*(A) \to H^*(\mathbb{D}NA) \quad \text{and} \quad (i_A^{\mathbb{D}NA})_* : H_*(A) \to H_*(\mathbb{D}NA)$$

are inverses. Given $\varphi \in H^m(A)$, we now set $\psi = \pi^*\varphi \in H^m(\mathbb{D}NA)$ and can feed $\psi \cup \tau(NA) \in H^{m+k}(\mathbb{D}NA, \mathbb{S}NA)$ into the duality isomorphism $\mathrm{PD} : H^{m+k}(\mathbb{D}NA, \mathbb{S}NA) \to H_{n-(m+k)}(\mathbb{D}NA)$. Using the naturality of the cap product, this gives

$$\begin{aligned}\mathrm{PD}\left(\pi^*\varphi \cup \tau(NA)\right) &= (\psi \cup \tau(NA)) \cap [\mathbb{D}NA] = \psi \cap (\tau(NA) \cap [\mathbb{D}NA]) = \psi \cap (i_A^{\mathbb{D}NA})_*[A] \\ &= (i_A^{\mathbb{D}NA})_*\left((i_A^{\mathbb{D}NA})^*\psi \cap [A]\right) = (i_A^{\mathbb{D}NA})^* \,\mathrm{PD}(\varphi),\end{aligned}$$

which presents the map $\pi^*(\cdot) \cup \tau(NA) : H^m(A) \to H^{m+k}(\mathbb{D}NA, \mathbb{S}NA)$ as a composition

$$\begin{tikzcd} H^m(A) \arrow[r, "\mathrm{PD}"] \arrow[rrr, bend right, "\pi^*(\cdot)\cup\tau(NA)"'] & H_{n-k-m}(A) \arrow[r, "(i_A^{\mathbb{D}NA})_*"] & H_{n-(k+m)}(\mathbb{D}NA) \arrow[r, "\mathrm{PD}^{-1}"] & H^{k+m}(\mathbb{D}NA, \mathbb{S}NA). \end{tikzcd}$$

All three maps in this composition are isomorphisms, thus so is $\varphi \mapsto \pi^*\varphi \cup \tau(NA)$.

To check that $\langle \tau(NA), (f_x^A)_*[\mathbb{D}N_xA] \rangle = 1$ for all $x \in A$, it suffices to prove this for the case when $A \subset M$ is connected, as it is then an easy exercise to generalize to a finite disjoint union of connected submanifolds. The advantage of assuming A connected is that $H^0(A)$ is then canonically isomorphic to the coefficient ring R, with the unit $1 \in H^0(A)$ as a generator, so the isomorphism above then implies that $\tau(NA)$ generates $H^k(\mathbb{D}NA, \mathbb{S}NA) \cong R$. It is clear that at most one element $\tau \in H^k(\mathbb{D}NA, \mathbb{S}NA)$ can satisfy $\langle \tau, (f_x^A)_*[\mathbb{D}N_xA] \rangle = 1$ for any given $x \in A$; moreover, if this is satisfied for one $x \in A$ then it is also satisfied for every $y \in A$, as the path-connectedness of A produces a homotopy of inclusions $(\mathbb{D}^k, S^{k-1}) \hookrightarrow (\mathbb{D}NA, \mathbb{S}NA)$ relating f_x^A and some reparametrization of f_y^A, so that $(f_y^A)_*[\mathbb{D}N_yA] = (f_x^A)_*[\mathbb{D}N_xA]$. We shall now show that there exists a generator τ satisfying this relation for some $x \in A$. This will prove $\tau = \tau(NA)$ in the case $R = \mathbb{Z}_2$; if $R = \mathbb{Z}$, then there is a sign ambiguity $\tau = \pm\tau(NA)$ that can be resolved by paying more careful attention to orientation conventions, and the remaining cases follow from this via the universal coefficient theorem.

The idea is to realize the isomorphism $H^*(A) \to H^{*+k}(\mathbb{D}NA, \mathbb{S}NA)$ via a cell decomposition. Notice first that $(\mathbb{D}NA, \mathbb{S}NA)$ is a good pair, so its relative cohomology is naturally isomorphic to the reduced cohomology of the quotient space

$$\mathrm{Th}(NA) := \mathbb{D}NA/\mathbb{S}NA.$$

This is known as the **Thom space** of the vector bundle NA. Since the interior of $\mathbb{D}NA$ is homeomorphic to NA, one can imagine $\mathrm{Th}(NA)$ as the one-point compactification of NA, and we shall label the point represented by $\mathbb{S}NA = \partial(\mathbb{D}NA)$ accordingly as $\infty \in \mathrm{Th}(NA)$. The isomorphism $\widetilde{H}^*(\mathrm{Th}(NA)) \cong H^*(\mathbb{D}NA, \mathbb{S}NA)$ is then equivalent to the isomorphism

$$H^*(\mathrm{Th}(NA), \{\infty\}) \to H^*(\mathbb{D}NA, \mathbb{S}NA)$$

induced by the quotient map $(\mathbb{D}NA, \mathbb{S}NA) \to (\mathrm{Th}(NA), \{\infty\})$; cf. Lecture 32.

Now since A is a smooth manifold, it has a triangulation, and after barycentric subdivision[95] we can assume without loss of generality that every simplex in the triangulation is small enough to be contained in a region $\mathcal{U} \subset A$ where there exists a local trivialization identifying $\pi^{-1}(\mathcal{U}) \subset \mathbb{D}NA$ with $\mathcal{U} \times \mathbb{D}^k$. Regarding $\{\infty\}$ as the 0-skeleton of $\mathrm{Th}(NA)$, we can then associate to each m-simplex σ in the triangulation of A a product $(m+k)$-cell in $\mathrm{Th}(NA)$, defined by identifying $\pi^{-1}(\sigma) \subset \mathbb{D}NA$ with $\sigma \times \mathbb{D}^k$ and attaching $\sigma \times \partial\mathbb{D}^k$ along the constant map to ∞. This gives $(\mathrm{Th}(NA), \{\infty\})$ the structure of a CW-pair, and while the characteristic maps of the cells in this decomposition depend on choices of local trivializations, their images do not, thus in a meaningful sense, our cell decomposition of $(\mathrm{Th}(NA), \{\infty\})$ depends only on the chosen triangulation of A. The most important observation is that since $C_*^{\mathrm{CW}}(\mathrm{Th}(NA), \{\infty\})$ is generated by the cells in the interior $\mathbb{D}NA \backslash \mathbb{S}NA$ while ignoring the 0-cell at ∞, the association of each simplex to its product with $\mathbb{D}^k$ defines a chain isomorphism

$$C_*^{\Delta}(A) \to C_{*+k}^{\mathrm{CW}}(\mathrm{Th}(NA), \{\infty\}),$$

[95]The subdivision is not actually necessary at all if you know a little bit more about the properties of vector bundles: in particular, they are always trivializable over a contractible space. For similar reasons, one could just as well work with an arbitrary cell decomposition of A instead of a triangulation.

which can be dualized to define an isomorphism $H^{m+k}(\mathrm{Th}(NA),\{\infty\}) \cong H^m(A)$ for every m. Take the generator $\tau \in H^k(\mathrm{Th}(NA),\{\infty\})$ that corresponds to $1 \in H^0(A)$ under this isomorphism, then pull it back through the quotient map $q : (\mathbb{D}NA, \mathbb{S}NA) \to (\mathrm{Th}(NA),\{\infty\})$ to define a generator $q^*\tau \in H^k(\mathbb{D}NA, \mathbb{S}NA)$. We now have an explicit cellular cochain representative of τ and can thus check that for any point $x \in A$ in the 0-skeleton of A, the associated product k-cell in $\mathrm{Th}(NA)$ represents $q_*(f_x^A)_*[\mathbb{D}N_xA] \in H_k(\mathrm{Th}(NA),\{\infty\})$ and satisfies

$$\langle q^*\tau, (f_x^A)_*[\mathbb{D}N_xA]\rangle = \langle \tau, q_*(f_x^A)_*[\mathbb{D}N_xA]\rangle = 1.$$

$\square$

Remark 55.20. The Thom space $\mathrm{Th}(E)$ can also be defined for more general vector bundles $\pi : E \to M$ without any reference to a submanifold or normal bundle. Together with the cell decomposition constructed in the proof above, $\mathrm{Th}(E)$ is known as the *Thom complex.*

Localizing Poincaré dual classes. Let us now associate to the closed R-oriented submanifold $A \subset M$ of codimension k its Poincaré dual class in M,

$$\tau_A^M := \mathrm{PD}^{-1}\left((i_A^M)_*[A]\right) \in H^k(M), \qquad \text{i.e.} \quad \tau_A^M \cap [M] = (i_A^M)_*[A].$$

The Thom class provides a way of "localizing" τ_A^M, in the following sense. Consider the inclusions

$$(M,\varnothing) \xhookrightarrow{j_A^M} (M, M\backslash A) \xhookleftarrow{i_{\mathbb{D}NA}^M} (\mathbb{D}NA, \mathbb{S}NA),$$

where $i_{\mathbb{D}NA}^M$ is provided by the tubular neighborhood theorem. This is an excision map, so it induces an isomorphism on relative cohomology and thus identifies $\tau(NA)$ with a class

$$\hat{\tau}_A^M \in H^k(M\,|\,A) \quad \text{ such that } \quad (i_{\mathbb{D}NA}^M)^*\hat{\tau}_A^M = \tau(NA).$$

Lemma 55.21. $\tau_A^M = (j_A^M)^*\hat{\tau}_A^M$.

Proof. The three fundamental classes $[M] \in H_n(M)$, $[M]_A \in H_n(M\,|\,A)$ and $[\mathbb{D}NA] \in H_n(\mathbb{D}NA, \mathbb{S}NA)$ in this picture are related by

$$(j_A^M)_*[M] = [M]_A = (i_{\mathbb{D}NA}^M)_*[\mathbb{D}NA].$$

Since $\mathrm{PD}((j_A^M)^*\hat{\tau}_A^M) \in H_{n-k}(M)$ is an *absolute* homology class and $j_A^M : M \to M$ is the identity map, the naturality of the cap product with respect to $j_A^M : (M,\varnothing) \hookrightarrow (M, M\backslash A)$ and $i_{\mathbb{D}NA}^M : (\mathbb{D}NA, \mathbb{S}NA) \hookrightarrow (M, M\backslash A)$ gives

$$\begin{aligned}
\mathrm{PD}((j_A^M)^*\hat{\tau}_A^M) &= (j_A^M)_*\,\mathrm{PD}((j_A^M)^*\hat{\tau}_A^M) = (j_A^M)_*\big((j_A^M)^*\hat{\tau}_A^M \cap [M]\big) = \hat{\tau}_A^M \cap (j_A^M)_*[M] \\
&= \hat{\tau}_A^M \cap (i_{\mathbb{D}NA}^M)_*[\mathbb{D}NA] = (i_{\mathbb{D}NA}^M)_*\big((i_{\mathbb{D}NA}^M)^*\hat{\tau}_A^M \cap [\mathbb{D}NA]\big) \\
&= (i_{\mathbb{D}NA}^M)_*\,\mathrm{PD}(\tau(NA)) = (i_{\mathbb{D}NA}^M)_*(i_A^{\mathbb{D}NA})_*[A] = (i_A^M)_*[A],
\end{aligned}$$

so $(j_A^M)^*\hat{\tau}_A^M$ satisfies the defining property of τ_A^M. $\square$

The message of this lemma is that the cohomology class Poincaré dual to A is determined by a class in $H^k(M\,|\,A)$, or equivalently $H^k(\mathbb{D}NA, \mathbb{S}NA)$; in either case it depends only on a neighborhood of A. It is now easy to see why Theorem 55.7 holds in the case $A \cap B = \varnothing$. It is equivalent in that case to $\tau_A^M \cup \tau_B^M = 0$, and by naturality of the cup product we have

$$\tau_A^M \cup \tau_B^M = (j_A^M)^*\hat{\tau}_A^M \cup (j_B^M)^*\hat{\tau}_B^M = (j_{A\cap B}^M)^*(\hat{\tau}_A^M \cup \hat{\tau}_B^M) = 0,$$

where the relative cup product

$$\hat{\tau}_A^M \cup \hat{\tau}_B^M \in H^{k+\ell}(M, (M\backslash A) \cup (M\backslash B)) = H^{k+\ell}(M, M\backslash(A \cap B)) = H^{k+\ell}(M, M) = 0$$

is well defined since $M\backslash A$ and $M\backslash B$ are both open in M and thus form an excisive couple. To handle the general case where $A \pitchfork B$ so that $A \cap B$ is a submanifold of A with normal bundle $N^A(A \cap B) = N^M B|_{A\cap B}$, we need the following application of the Thom isomorphism theorem:

LEMMA 55.22. $\tau^A_{A\cap B} = (i^M_A)^* \tau^M_B$.

PROOF. Let $f^B_{A\cap B} : (\mathbb{D}N^A A \cap B, \mathbb{S}N^A(A \cap B)) \hookrightarrow (\mathbb{D}N^M B, \mathbb{S}N^M B)$ denote the inclusion defined by identifying the fiber of $N^A(A \cap B)$ over each point in $A \cap B$ with the corresponding fiber of $N^M B$. For each $x \in A \cap B$ we then have

$$\begin{aligned}\langle (f^B_{A\cap B})^* \tau(N^M B), (f^{A\cap B}_x)_* [\mathbb{D}N^A_x (A \cap B)] \rangle &= \langle \tau(N^M B), (f^B_{A\cap B} \circ f^{A\cap B}_x)_* [\mathbb{D}N^A_x (A \cap B)] \rangle \\ &= \langle \tau(N^M B), (f^B_x)_* [\mathbb{D}N^M_x B] \rangle = 1,\end{aligned}$$

so $(f^B_{A\cap B})^* \tau(N^M B) = \tau(N^A(A \cap B))$ according to Theorem 55.19.[96] The result then follows from the commutative diagram

$$\begin{tikzcd}
H^k(M) \arrow[d, "(i^M_A)^*"] & H^k(M \,|\, B) \arrow[l, "(j^M_B)^*"'] \arrow[r, "(i^M_{\mathbb{D}NB})^*", "\cong"'] \arrow[d, "(i^M_A)^*"] & H^k(\mathbb{D}NB, \mathbb{S}NB) \arrow[d, "(f^B_{A\cap B})^*"] \\
H^k(A) & H^k(A \,|\, A \cap B) \arrow[l, "(j^A_{A\cap B})^*"] \arrow[r, "(i^A_{\mathbb{D}N^A(A\cap B)})^*", "\cong"'] & H^k(\mathbb{D}N^A(A \cap B), \mathbb{S}N^A(A \cap B))
\end{tikzcd}$$

□

This is enough preparation to prove the main theorem.

PROOF OF THEOREM 55.7. Combining Lemma 55.22 with the usual naturality and associativity properties of the cap and cup products, we compute:

$$\begin{aligned}(i^M_{A\cap B})_*[A \cap B] &= (i^M_A)_* (i^A_{A\cap B})_* [A \cap B] = (i^M_A)_* \operatorname{PD}(\tau^A_{A\cap B}) \\ &= (i^M_A)_* (\tau^A_{A\cap B} \cap [A]) = (i^M_A)_* ((i^M_A)^* \tau^M_B \cap [A]) = \tau^M_B \cap (i^M_A)_*[A] \\ &= \tau^M_B \cap \operatorname{PD}(\tau^M_A) = \tau^M_B \cap (\tau^M_A \cap [M]) \\ &= (\tau^M_B \cup \tau^M_A) \cap [M] = \operatorname{PD}(\tau^M_B \cup \tau^M_A) = (i^M_A)_*[A] \cdot (i^M_B)_*[B].\end{aligned}$$

□

56. Higher homotopy groups (February 13, 2024)

The last two lectures in this course will have more the character of a survey, as I want to mention several important things but will not have time to prove many of them.

The higher homotopy groups $\pi_n(X)$ were mentioned informally last semester in Lecture 21. Let's give a more formal definition. It will help to have the following popular notation at our disposal: given spaces X and Y, we define the set

$$[X, Y] := \{\text{continuous maps } X \to Y\} \big/ \sim,$$

where the equivalence relation is homotopy. Similarly, for pairs of spaces (X, A) and (Y, B),

$$[(X, A), (Y, B)]$$

will denote the set of homotopy classes of maps of pairs. Here one can also specialize to the case where A and B are each a single point (homotopy classes of base-point preserving maps), or extend

[96] What this argument actually shows is that for any vector bundle $\pi : E \to M$ and a map $f : N \to M$ for which the Thom classes of E and the *pullback bundle* $f^*E \to N$ are well defined, $\tau(f^*E) = \widehat{f}^* \tau(f)$, where $\widehat{f} : (\mathbb{D}(f^*E), \mathbb{S}(f^*E)) \to (\mathbb{D}(E), \mathbb{S}(E))$ is the natural map of disk bundles covering the map $f : N \to M$.

the definition in an obvious way to allow triples (X, A, B) where $B \subset A \subset X$. In this notation, the fundamental group of a pointed space (X, x_0) can be expressed in two equivalent ways as

$$\pi_1(X, x_0) = [(S^1, \mathrm{pt}), (X, x_0)] = [(I, \partial I), (X, x_0)],$$

where pt denotes an arbitrary choice of base point in S^1, and I is the unit interval $[0, 1]$. Since the latter is homeomorphic to the 1-dimensional unit disk $\mathbb{D}^1$, we could also equivalently write

$$\pi_1(X, x_0) = [(\mathbb{D}^1, \partial\mathbb{D}^1), (X, x_0)].$$

These definitions are equivalent to the definition in terms of S^1 because $S^1 \cong \mathbb{D}^1/\partial\mathbb{D}^1$. Note that there are also higher-dimensional analogues of this statement: S^n is homeomorphic to $\mathbb{D}^n/\partial\mathbb{D}^n$ and $I^n/\partial I^n$ for all $n \in \mathbb{N}$, where I^n here denotes the n-fold product of I, i.e. an n-dimensional unit cube.

Definition 56.1. For each integer $n \geqslant 0$, we define the set

$$\pi_n(X, x_0) := [(S^n, \mathrm{pt}), (X, x_0)].$$

When $n \geqslant 1$, this can be expressed equivalently as

$$\pi_n(X, x_0) = [(\mathbb{D}^n, \partial\mathbb{D}^n), (X, x_0)] = [(I^n, \partial I^n), (X, x_0)].$$

As yet this is only a set; we have not given it a group structure. The case $n = 0$ has occasionally been mentioned before: since $S^0 = \{1, -1\}$ and one of these two points must be chosen as a base point and thus mapped to x_0, $\pi_0(X, x_0)$ is just the set of homotopy classes of maps of the other point to X, so it has a natural bijective correspondence with the set of path-components of X. This is indeed *only* a set, and not a group. The group structure of $\pi_1(X, x_0)$ as we learned it in *Topologie I* is based on the notion of concatenation of paths, which makes sense due to the fact that if I_1 and I_2 denote two copies of the unit interval $I = [0, 1]$, then the space obtained by gluing them together end-to-end,

$$(I_1 \amalg I_2)\big/(I_1 \ni 1 \sim 0 \in I_2)$$

is homeomorphic to I. One can do the same thing with the cube I^n by singling out one of the coordinates as the one to be concatenated, e.g. if I_1^n and I_2^n denote two copies of I^n, we have

$$(I_1^n \amalg I_2^n)\big/(I_1^n \ni (1, t_2, \dots, t_n) \sim (0, t_2, \dots, t_n) \in I_2^n) \cong I^n,$$

where the equivalence relation now applies for all values of $(t_2, \dots, t_n) \in I^{n-1}$. This observation leads to the natural group structure on $\pi_n(X, x_0)$. We shall state it here only for $n \geqslant 2$, since the fundmental group is already familiar, and the standard notation for its group structure is slightly different for reasons that we'll get into in a moment.

Definition 56.2. For $n \geqslant 2$ and two elements $[f], [g] \in \pi_n(X, x_0)$ represented by maps $f, g : (I^n, \partial I^n) \to (X, x_0)$, we define $[f] + [g] \in \pi_n(X, x_0)$ to be the homotopy class of the map

$$(I^n, \partial I^n) \to (X, x_0) : (t_1, \dots, t_n) \mapsto \begin{cases} f(2t_1, t_2, \dots, t_n) & \text{if } 0 \leqslant t_1 \leqslant 1/2, \\ g(2t_1 - 1, t_2, \dots, t_n) & \text{if } 1/2 \leqslant t_1 \leqslant 1. \end{cases}$$

This definition seems a bit arbitrary at first, e.g. one might wonder why the coordinate t_1 is singled out for special treatment when any of the other coordinates would work just as well. The answer is that one could indeed formulate the definition in various alternative ways, but one would always obtain the same result *up to homotopy.* This is easy to see once you've absorbed the proof of the following related fact, which justifies our use of additive notation:

Proposition 56.3. *For all $n \geqslant 2$, the operation in Definition 56.2 makes $\pi_n(X, x_0)$ an abelian group.*

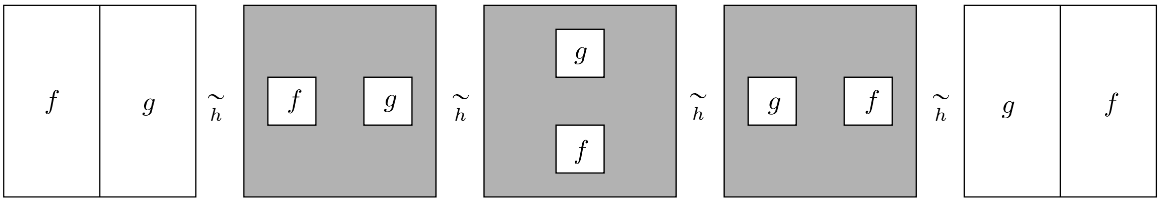

FIGURE 33. The homotopy in the proof of Proposition 56.3.

PROOF. The proof that $\pi_n(X, x_0)$ is a group can be carried out by ignoring $n - 1$ of the coordinates and repeating the same arguments with which we proved last semester that $\pi_1(X, x_0)$ is a group. The identity element is exactly what you think it should be: it is represented by the constant map of S^n to x_0.

The novel feature is that $\pi_n(X, x_0)$ is abelian for $n \geqslant 2$; as we've seen, the fundamental group does not generally have this property. The proof is a homotopy depicted in Figure 33. The shaded region in each picture represents a subset of I^n on which the map takes a constant value, namely the base point x_0. The leftmost picture shows the map representing $[f] + [g]$ as specified in Definition 56.2, with the cube I^n divided into two halves on which the map restricts to f or g. We then homotop this map by shrinking the two halves to smaller cubes and mapping everything outside the smaller cubes to the base point—this is possible because $f|_{\partial I^n}$ and $g|_{\partial I^n}$ are also constant maps to the base point. After shrinking both cubes far enough, there is enough room to move them past each other so that the roles of f and g are reversed. It should be clear why this trick does not work when $n = 1$. □

With this group structure, $\pi_n(X, x_0)$ is called the n**th homotopy group** of X.

There are also **relative homotopy groups** $\pi_n(X, A, x_0)$ associated to any pair of spaces (X, A) with a base point $x_0 \in A$. One can define this as a mild generalization of $\pi_n(X, x_0) = [(\mathbb{D}^n, \partial\mathbb{D}^n), (X, x_0)]$ by choosing a base point $\mathrm{pt} \in \partial\mathbb{D}^n$ and setting

$$\pi_n(X, A, x_0) := [(\mathbb{D}^n, \partial\mathbb{D}^n, \mathrm{pt}), (X, A, x_0)].$$

This reduces to $\pi_n(X, x_0)$ if $A = \{x_0\}$, but in all other cases, we need to be aware that it only makes sense for $n \geqslant 1$; there is no definition of $\pi_0(X, A, x_0)$ for $A \neq \{x_0\}$ since $n = 0$ is the one case where the relation $S^n \cong \mathbb{D}^n/\partial\mathbb{D}^n$ fails to hold. For $n = 1$, we can identify $\mathbb{D}^1$ with I and choose $0 \in I$ as the base point so that $\pi_1(X, A, x_0)$ becomes the set of all homotopy classes of paths in X from x_0 to arbitrary points in A. Since these paths do not need to be loops, there is no obvious notion of concatenation here, so that $\pi_1(X, A, x_0)$ does not have a natural group structure—it is only a set. A group structure can be defined for $\pi_n(X, A, x_0)$ if $n \geqslant 2$. To explain this, we reformulate the definition as a generalization of $\pi_n(X, x_0) = [(I^n, \partial I^n), (X, x_0)]$ by singling out a particular boundary face of I^n to play the role of $\partial\mathbb{D}^n = S^{n-1} \cong I^{n-1}/\partial I^{n-1}$ and regarding the rest of ∂I^n as the base point: let

$$J_n := I^{n-1} \times \{0\} \subset \partial I^n$$

and redefine $\pi_n(X, A, x_0)$ as

$$\pi_n(X, A, x_0) := [I^n, \partial I^n, \overline{\partial I^n \backslash J_n}, (X, A, x_0)].$$

By this definition, the formula in Definition 56.2 still makes sense for $n \geqslant 2$ and defines a group structure on $\pi_n(X, A, x_0)$, though Proposition 56.3 no longer works in the $n = 2$ case. You can see why not if you look again at Figure 33 and imagine that the maps on the bottom edge of each square are not required to be constant, but only to have their images in A: there is now

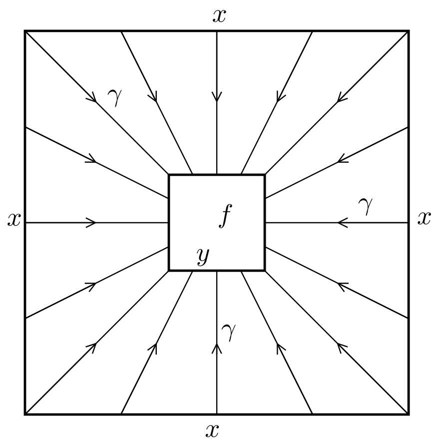

FIGURE 34. The isomorphism $\pi_n(X, y) \to \pi_n(X, x)$ determined by a path $x \overset{\gamma}{\rightsquigarrow} y$.

no obvious way to define the map on the shaded areas so that it gives a well-defined homotopy. The argument can be rescued, however, if $n \geqslant 3$, as we can then assume the two small cubes are "rooted" to the bottom face J_n, but there are still enough dimensions to move them past each other. To summarize:

PROPOSITION 56.4. *For general pairs of spaces (X, A) with a base point $x_0 \in A$, $\pi_n(X, A, x_0)$ has a natural group structure for every $n \geqslant 2$, and it is abelian for $n \geqslant 3$.* $\square$

Like the fundamental group, the higher homotopy groups depend on a choice of base point, but there is an isomorphism

$$\Phi_\gamma : \pi_n(X, y) \xrightarrow{\cong} \pi_n(X, x)$$

determined by any path γ from x to y in X. The definition is best explained with a picture: Figure 34 shows a recipe for transforming any map $f : (I^n, \partial I^n) \to (X, y)$ into a map $(I^n, \partial I^n) \to (X, x)$ by shrinking the domain of the original map f to a smaller cube within I^n, and then filling the region between this and ∂I^n with copies of the path $x \overset{\gamma}{\rightsquigarrow} y$. The picture shows the $n = 2$ case, but if you draw the analogous picture for $n = 1$, you will find that it reproduces exactly the isomorphism $\Phi_\gamma : \pi_1(X, y) \to \pi_1(X, x)$ described in last semester's Lecture 9. We leave it as an exercise to verify that this really is a well-defined isomorphism, and that it only depends on the (end-point preserving) homotopy class of the path γ. With this in mind, we will sometimes abbreviate

$$\pi_n(X) := \pi_n(X, x_0)$$

when the space X is path-connected and the base point does not play a major role.

There is a fairly obvious way to view π_n as a functor from the category Top_* of pointed spaces to the category Grp of groups (or Ab for $n \geqslant 2$). Namely, every base-point preserving map $f : (X, x_0) \to (Y, y_0)$ induces a homomorphism

$$f_* : \pi_n(X, x_0) \to \pi_n(Y, y_0) : [\varphi] \mapsto [f \circ \varphi].$$

It is similarly easy to see that this homomorphism only depends on the (base-point preserving!) homotopy class of f. The following property is less obvious, but important to know:

THEOREM 56.5. *If $f : X \to Y$ is a homotopy equivalence, then $f_* : \pi_n(X, x_0) \to \pi_n(Y, f(x_0))$ is an isomorphism for all $n \geqslant 0$ and $x_0 \in X$.* $\square$

Since we've been talking about homology for the rest of this course, you may have forgotten why Theorem 56.5 is already a nontrivial statement in the $n = 1$ case, which took some effort to prove in *Topologie I*. The annoying detail is the base point: if $g : Y \to X$ is a homotopy inverse for f, then it does not automatically induce an inverse for f_* since g need not take $f(x_0)$ back to the base point x_0; in general, g_* sends $\pi_n(Y, f(x_0))$ to a different group, $\pi_n(X, g(f(x_0)))$. But this headache can be dealt with in the same way as in the $n = 1$ case, using the isomorphism $\Phi_\gamma : \pi_n(X, g(f(x_0))) \to \pi_n(X, x_0)$ induced by a path $x_0 \rightsquigarrow g(f(x_0))$, which necessarily exists due to the homotopy inverse condition. The proof is then a direct adaptation of what we already did for the $n = 1$ case in Lecture 9, so we'll leave it as an exercise. The reason this detail was easier in homology theory is that homology does not care about base points, so the homotopy invariance of induced maps $f_* : H_*(X) \to H_*(Y)$ immediately implied that $H_*(X)$ depends only on the homotopy type of X.

Let's look at some examples now. It should be said that, in general, higher homotopy groups are not easy to compute—there is nothing quite analogous to cellular homology to produce a practical algorithm for computing $\pi_n(X)$. But to start with, there are some easy cases where theorems that we've proved for other purposes imply computations of $\pi_n(X)$.

Example 56.6. For every $k \geqslant 2$ and $n \in \mathbb{N}$, $\pi_k(\mathbb{T}^n) = 0$. This is a consequence of the fact that $\mathbb{T}^n$ has a contractible universal cover, namely $p : \mathbb{R}^n \to \mathbb{T}^n$. Since S^k is simply connected for $k \geqslant 2$, every map $f : S^k \to \mathbb{T}^n$ has a lift $\widetilde{f} : S^k \to \mathbb{R}^n$, which is homotopic to a constant map since $\mathbb{R}^n$ is contractible. Composing this homotopy with $p : \mathbb{R}^n \to \mathbb{T}^n$ then gives a homotopy of f to a constant map $S^k \to \mathbb{T}^n$. (Strictly speaking, one should pay a bit more attention to the base point in this discussion, but that is easy to do.) Note that the circle $S^1 = \mathbb{T}^1$ is a special case of this computation, so we now know *all* the homotopy groups of S^1.

Example 56.7. For $n \in \mathbb{N}$ and $k < n$, $\pi_k(S^n) = 0$. One can see this by proving that every map $f : S^k \to S^n$ with $n > k$ is homotopic to a map $g : S^k \to S^n$ that is not surjective: then if $p \in S^n \backslash g(S^k)$, it follows that the image of g is in $S^n \backslash \{p\} \cong \mathbb{R}^n$, and is then homotopic to a constant since $\mathbb{R}^n$ is contractible. Here are two possible ways to prove the claim that f is homotopic to something non-surjective: (1) The simplicial approximation theorem (see Lecture 41) implies that for suitable choices of triangulations of S^k and S^n, f is homotopic to a simplicial map $g : S^k \to S^n$, which is therefore also a cellular map and thus has image in the k-skeleton of S^n. When $n > k$, the k-skeleton cannot cover all of S^n, thus g is not surjective. (2) There is a very easy proof using basic results of differential topology as in [**Mil97**]: $f : S^k \to S^n$ is homotopic to a smooth map $g : S^k \to S^n$ that is C^0-close to f, and Sard's theorem then implies that almost every point $y \in S^n$ is a regular value of g. This means the derivative $dg(x) : T_x S^k \to T_y S^n$ is surjective for every $x \in g^{-1}(y)$, but since that condition can never be satisfied for $n > k$, it follows that $g^{-1}(y) = \varnothing$.

Example 56.8. Viewing elements of $\pi_n(S^n)$ as represented by maps $f : S^n \to S^n$, the mapping degree determines an isomorphism

$$\pi_n(S^n) \xrightarrow{\cong} \mathbb{Z} : [f] \mapsto \deg(f)$$

for every $n \in \mathbb{N}$. This does not immediately follow from anything we've covered in this course, but here are two ways to see it: (1) Using differential topology as in [**Mil97**], the so-called *Pontryagin-Thom* construction elegantly defines a bijection for any closed, connected and oriented k-manifold M between the set of homotopy classes $[M, S^n]$ and the set of "framed bordism classes" in M, where the latter have a natural correspondence with the integers when $k = n$. In particular, when $\dim M = n$ this proves that the map $\deg : [M, S^n] \to \mathbb{Z}$ is a bijection. (One must transform arbitrary homotopies into base-point preserving homotopies before this becomes a statement about $\pi_n(S^n)$, but the gap is not hard to fill.) (2) In the next lecture we will state the *Hurewicz* theorem, which defines a natural homomorphism $\pi_n(X) \to H_n(X)$ and gives conditions for it to be an isomorphism, which hold in the case $X = S^n$ due to the computation of $H_*(S^n)$.

In Example 56.16 at the end of this lecture, we will discuss the interesting case of $\pi_3(S^2)$, which is fairly easy to compute, but the answer may contradict the intuition you've developed from homology, i.e. it is not trivial. Unlike $H_k(M)$, there is no reason in general why $\pi_k(M)$ should vanish when $k > \dim M$.

EXAMPLE 56.9. The following is way beyond the scope of this course, but just to give you a taste of what is studied in modern homotopy theory: it turns out that there are natural isomorphisms

$$\pi_{n+k}(S^n) \to \pi_{n+k+1}(S^{n+1})$$

for all $k \geqslant 0$ as soon as n is sufficiently large. The resulting groups that depend only on k are known as the *stable homotopy groups* of the spheres. They have been computed in many cases, but they are not known in general for $k > 64$. The computation of higher homotopy groups of spheres is considered one of the most important open problems in algebraic topology.

The following definition makes the notions of path-connectedness ($n = 0$) and simple connectedness ($n = 1$) into the first two items on an infinite hierarchy of conditions.

DEFINITION 56.10. For integers $n \geqslant 0$, a space X is called **n-connected** if $\pi_k(X) = 0$ for all $k \leqslant n$.

We can now give an example of the kind of problem for which computing higher homotopy groups is useful.

THEOREM 56.11. *If X is a CW-complex of dimension at most n and Y is an n-connected space, then all maps $X \to Y$ are homotopic.*

PROOF. We need to show that any two given maps $f, g : X \to Y$ are homotopic. The method of the proof is known as "induction over the skeleta".[97] As preparation, one needs to think through the following exercise: if $f|_{X^k} : X^k \to Y$ is homotopic to $g|_{X^k} : X^k \to Y$ for some $k \geqslant 0$, then f is also homotopic on X to a map $f' : X \to Y$ such that $f'|_{X^k} = g|_{X^k}$. This can be done by using cutoff functions to extend the homotopy from the k-skeleton to all higher-dimensional cells.

Now to start the induction, note that since Y is path-connected, $f|_{X^0}$ and $g|_{X^0}$ are clearly homotopic, as one can just pick a path from $f(x)$ to $g(x)$ for every $x \in X^0$. Now for a given $k \in \{1, \dots, n\}$, we need to show that if f has already been adjusted by a homotopy so that $f|_{X^{k-1}} = g|_{X^{k-1}}$, then $f|_{X^k}$ is also homotopic to $g|_{X^k}$. It suffices to show that the restrictions of f and g to each k-cell $e^k_\alpha \subset X$ are homotopic via a homotopy that is fixed at the boundary of the cell, i.e. on the $(k-1)$-skeleton. Let $\Phi_\alpha : (\mathbb{D}^k, S^{k-1}) \to (X^k, X^{k-1})$ denote the characteristic map of e^k_α. Then $f \circ \Phi_\alpha$ and $g \circ \Phi_\alpha$ are two maps $\mathbb{D}^k \to Y$ that match at the boundary S^{k-1}, hence we can glue their domains together to form a sphere $S^k \cong \mathbb{D}^k_+ \cup_{S^{k-1}} \mathbb{D}^k_-$ and define on this sphere a continuous map

$$F : S^k \to Y : x \mapsto \begin{cases} f \circ \Phi_\alpha(x) & \text{if } x \in \mathbb{D}^k_+, \\ g \circ \Phi_\alpha(x) & \text{if } x \in \mathbb{D}^k_-. \end{cases}$$

Since $\pi_k(Y) = 0$, the map $F : S^k \to Y$ is homotopic to a constant, which is equivalent to saying that it extends to a map $\mathbb{D}^{k+1} \to Y$, and this extension can be used to define a homotopy between $f \circ \Phi_\alpha$ and $g \circ \Phi_\alpha$ that is fixed along the boundary. This completes the induction. □

[97] It seems that the plural of the English word "skeleton" is different in topology than it is in the rest of the English language. Dictionaries list both "skeletons" and "skeleta," but I have never heard the latter outside of mathematical contexts, e.g. one would not say that a politician with potentially damaging secrets has "skeleta in the closet".

You may notice that Theorem 56.11 has an obvious converse: if Y is not n-connected, then there clearly also exists a CW-complex X of dimension at most n (in particuar a sphere) such that not all maps $X \to Y$ are homotopic. This example is the beginning of the subject known as *obstruction theory*, which finds necessary and sufficient conditions for the existence and/or uniqueness (up to homotopy) of various geometric structures, particularly on manifolds. An example of such a geometric structure is an orientation, whose existence on a manifold M is equivalent to the vanishing of a particular element of $H^1(M;\mathbb{Z}_2)$, called the first Stiefel-Whitney class (see Exercise 52.23). The standard procedure is to express the geometric structure of interest in terms of sections of some fiber bundle associated to the manifold, so that the important question to answer is whether a section of this bundle exists and under what conditions two such sections must be homotopic. By induction over the skeleta, these questions are typically equivalent to the vanishing of certain higher homotopy groups. For a detailed exposition of this subject, I recommend [**Ste51**].

We have not yet talked much about the relative homotopy groups, and we won't, but I should mention that they appear in a fairly obvious exact sequence. Given a pair of spaces (X, A) and a base point $x_0 \in A$, denote by

$$(A, x_0) \overset{i}{\hookrightarrow} (X, x_0) \quad \text{and} \quad (X, x_0, x_0) \overset{j}{\hookrightarrow} (X, A, x_0)$$

the obvious inclusions. For each $n \geqslant 1$ there is also a natural homomorphism

$$\partial : \pi_n(X, A, x_0) \to \pi_{n-1}(A, x_0) : [f] \mapsto [f|_{S^{n-1}}],$$

where we regard elements of $\pi_n(X, A, x_0)$ as represented by maps $f : (\mathbb{D}^n, S^{n-1}, \mathrm{pt}) \to (X, A, x_0)$. You can easily check by translating this into the corresponding formula with $f : (I^n, \partial I^n, \overline{\partial I^n \backslash J_n}) \to (X, A, x_0)$ that it really is a homomorphism.

Theorem 56.12. *For $x_0 \in A \subset X$, the sequence*

$$\ldots \to \pi_{n+1}(X, A, x_0) \xrightarrow{\partial} \pi_n(A, x_0) \xrightarrow{i_*} \pi_n(X, x_0) \xrightarrow{j_*} \pi_n(X, A, x_0) \xrightarrow{\partial} \pi_{n-1}(A, x_0) \to \ldots$$
$$\ldots \to \pi_1(X, x_0) \xrightarrow{j_*} \pi_1(X, A, x_0) \xrightarrow{\partial} \pi_0(A, x_0) \xrightarrow{i_*} \pi_0(X, x_0).$$

is exact. □

Some comments on interpretation are required since the last three terms in this sequence are not groups, but only sets. They do have a bit more structure than this, as the constant map to x_0 defines in each case a distinguished element: if one interprets the kernel of each map in this part of the sequence to mean the preimage of the distinguished element, then it makes sense to say that the sequence is exact. The proof of exactness is more straightforward than for most exact sequences that arise in homology theory: instead of constructing chain complexes with a short exact sequence and chasing diagrams, one can just check directly that the image of each map equals the kernel of the next. For details, see [**Hat02**, Theorem 4.3].

A particular application of this exact sequence leads to one of the most popular tools for computing homotopy groups, called the *homotopy exact sequence of a fibration*. I will express the theorem in the form that arises most often in geometric applications, though it is somewhat less general than what is actually true. In the previous lecture we saw some examples of vector bundles, which one can imagine as families of vector spaces parametrized by an underlying space, carrying a topology determined by the notion of *local trivialization*. If one replaces vector spaces with arbitrary topological spaces in this picture, one arrives at the following notion.

Definition 56.13. A **fiber bundle** consists of the following data: topological spaces E, B and F known as the **total space**, **base** and **standard fiber** respectively, and a continuous map

$p : E \to B$, such that B can be covered by open sets $\mathcal{U}$ that admit **local trivializations**, meaning homeomorphisms

$$\Phi : p^{-1}(\mathcal{U}) \to \mathcal{U} \times F$$

that send $p^{-1}(b)$ homeomorphically to $\{b\} \times F$ for each $b \in \mathcal{U}$. The **fibers** of the bundle are the subspaces $p^{-1}(b) \cong F$ for $b \in B$.

Fiber bundles are often abbreviated with the notation

$$F \hookrightarrow E \xrightarrow{p} B,$$

where the inclusion $F \hookrightarrow E$ is not canonical but is defined by choosing any $b \in B$ and a local trivialization near b to identify $p^{-1}(b)$ with F. Note that while every fiber of a fiber bundle is homeomorphic to the standard fiber, there is typically no *canonical* homeomorphism, since there may be many choices of local trivializations covering each $b \in B$. If we choose base points $b_0 \in B$ and $x_0 \in p^{-1}(b_0) \subset E$, then it is natural to identify F with $p^{-1}(b_0)$ so that we obtain base-point preserving maps

$$(F, x_0) \hookrightarrow (E, x_0) \xrightarrow{p} (B, b_0).$$

A **trivial fiber bundle** is one that admits a single trivialization covering all of B, so that E can be identified globally with $B \times F$ and the map $p : E \to B$ becomes the obvious projection map $B \times F \to B$.

There's at least one general class of fiber bundles that you've definitely seen plenty of before: a covering map $p : E \to B$ is simply a fiber bundle whose standard fiber F is a discrete space, and this is for instance why our terminology for the "orientation bundle" of a manifold makes sense. A fiber bundle of this type admits a local trivialization over a subset $\mathcal{U} \subset B$ if and only if that subset is evenly covered, and it is a trivial fiber bundle if and only if $p : E \to B$ can be identified with the trivial covering map, i.e. the one for which E is a disjoint union of copies of B and $p : E \to B$ is the identity map on each copy.

Here is a popular example of a fiber bundle that is not a covering map and is also not trivial—we know it is not trivial since we know several ways of proving that S^3 is not homeomorphic to $S^2 \times S^1$.

EXAMPLE 56.14. The **Hopf fibration** $p : S^3 \to S^2$ is defined by identifying S^3 with the unit sphere in $\mathbb{C}^2$ and S^2 with the extended complex plane $\mathbb{C} \cup \{\infty\}$, and then writing

$$p : S^3 \to S^2 : (z_1, z_2) \mapsto \frac{z_1}{z_2}.$$

Equivalently, one can identify S^2 with $\mathbb{CP}^1$ so that this becomes the map

$$p : S^3 \to \mathbb{CP}^1 : (z_1, z_2) \mapsto [z_1 : z_2].$$

The fiber containing any given point $(z_1, z_2) \in S^3$ is the set

$$\left\{ (e^{i\theta} z_1, e^{i\theta} z_2) \in S^3 \;\middle|\; \theta \in \mathbb{R} \right\} \cong S^1.$$

We leave it as an exercise to check that local trivializations exist near every point.

THEOREM 56.15. *Given a fiber bundle $(F, x_0) \xhookrightarrow{i} (E, x_0) \xrightarrow{p} (B, b_0)$ with base points, the map $p : (E, F, x_0) \to (B, b_0, b_0)$ induces an isomorphism*

$$p_* : \pi_n(E, F, x_0) \xrightarrow{\cong} \pi_n(B, b_0)$$

for every $n \in \mathbb{N}$. Plugging this into the exact sequence of (E, F, x_0) thus produces an exact sequence

$$\ldots \to \pi_{n+1}(B, b_0) \xrightarrow{\partial} \pi_n(F, x_0) \xrightarrow{i_*} \pi_n(E, x_0) \xrightarrow{p_*} \pi_n(B, b_0) \xrightarrow{\partial} \pi_{n-1}(F, x_0) \to \ldots$$
$$\ldots \to \pi_1(E, x_0) \xrightarrow{p_*} \pi_1(B, b_0) \xrightarrow{\partial} \pi_0(F, x_0) \xrightarrow{i_*} \pi_0(E, x_0),$$

where the maps $\partial : \pi_n(B, b_0) \to \pi_{n-1}(F, x_0)$ *send each* $[f]$ *to* $[\widetilde{f}|_{S^{n-1}}]$ *for* $f : (\mathbb{D}^n, S^{n-1}) \to (B, b_0)$ *and* $\widetilde{f} : (\mathbb{D}^n, S^{n-1}, \mathrm{pt}) \to (E, F, x_0)$ *solving the lifting problem*

$$\begin{tikzcd} & E \arrow[d, "p"] \\ \mathbb{D}^n \arrow[ur, "\widetilde{f}"] \arrow[r, "f"] & B \end{tikzcd} \tag{56.1}$$

□

I will not say anything about the proof of this theorem except that the most important topological property of fiber bundles is the solvability of the lifting problem indicated in (56.1). The proper formulation of this condition is something called the *homotopy lifting property*, and Theorem 56.15 is true in fact for any map $p : E \to B$ that has the homotopy lifting property for maps of disks into B. Maps with this property are called **Serre fibrations**, and they are somewhat more general than fiber bundles. We saw in *Topologie I* that the lifting problem (56.1) is solvable in the special case of covering maps since $\mathbb{D}^n$ is simply connected; in fact there exists a *unique* lift that sends a given base point on $\partial\mathbb{D}^n$ to the base point $x_0 \in E$. For more general Serre fibrations, the lift is not always unique, but it is unique up to homotopy, which is why the map $\partial : \pi_n(B, b_0) \to \pi_{n-1}(E, x_0)$ described in the theorem is well defined.

Example 56.16. Returning to the Hopf fibration of Example 56.14, the homotopy exact sequence has a segment of the form

$$0 = \pi_3(S^1) \to \pi_3(S^3) \xrightarrow{p_*} \pi_3(S^2) \to \pi_2(S^1) = 0,$$

proving that $p_* : \pi_3(S^3) \to \pi_3(S^2)$ is an isomorphism. Since $\pi_3(S^3) \cong \mathbb{Z}$ is generated by the identity map $S^3 \to S^3$, this implies that $\pi_3(S^2) \cong \mathbb{Z}$, with the Hopf fibration itself representing a generator.

Example 56.17. In obstruction theory, one often needs to know the homotopy groups of certain topological groups that arise as "structure groups" of fiber bundles. For example, the structure group of any oriented vector bundle with n-dimensional fibers is

$$\mathrm{GL}_+(n, \mathbb{R}) := \left\{ \mathbf{A} \in \mathrm{GL}(n, \mathbb{R}) \;\middle|\; \det \mathbf{A} > 0 \right\}.$$

Here is a trick for computing $\pi_1(\mathrm{GL}_+(n, \mathbb{R}))$. Polar decomposition provides a deformation retraction of $\mathrm{GL}_+(n, \mathbb{R})$ to $\mathrm{SO}(n)$, the special orthogonal group, thus it suffices to compute $\pi_1(\mathrm{SO}(n))$. For $n = 1$ and $n = 2$, this is easy because $\mathrm{SO}(1) \cong \{\mathrm{pt}\}$ and $\mathrm{SO}(2) \cong S^1$. For $n = 3$, it is not hard to find a homeomorphism of $\mathrm{SO}(3)$ to $\mathbb{RP}^3$: this arises from the fact that every element of $\mathrm{SO}(3)$ defines a rotation about some axis in $\mathbb{R}^3$, so there is a natural map

$$\mathbb{D}^3 \to \mathrm{SO}(3)$$

that sends the origin to $\mathbb{1}$ and sends the point $r\mathbf{x}$ for $0 < r \leqslant 1$ and $\mathbf{x} \in S^2$ to the rotation by angle πr about the axis spanned by $\mathbf{x}$. By this definition, a rotation of angle πr about $\mathbf{x}$ is the same as a rotation of angle $-\pi r$ about $-\mathbf{x}$, so the map is injective on the interior of $\mathbb{D}^3$ but it sends antipodal points on $\partial\mathbb{D}^3$ to the same point, thus descending to a homeomorphism

$$\mathbb{D}^3/\!\sim\; \cong \mathrm{SO}(3)$$

where $\mathbf{x} \sim -\mathbf{x}$ for all $\mathbf{x} \in \partial\mathbb{D}^3$. This quotient space is homeomorphic to $\mathbb{RP}^3$, thus $\pi_1(\mathrm{SO}(3)) \cong \pi_1(\mathbb{RP}^3) \cong \mathbb{Z}_2$.

The remaining cases of $\pi_1(\mathrm{SO}(n))$ can now be deduced from the case $n = 3$ via a homotopy exact sequence. The fiber bundle we need for this purpose has the form

$$\mathrm{SO}(n) \overset{i}{\hookrightarrow} \mathrm{SO}(n+1) \xrightarrow{p} S^n,$$

where
$$i(\mathbf{A}) := \begin{pmatrix} 1 & 0 \\ 0 & \mathbf{A} \end{pmatrix} \quad \text{and} \quad p(\mathbf{A}) = \mathbf{A}e_1,$$
for $e_1 = (1, 0, \ldots, 0) \in S^n \subset \mathbb{R}^{n+1}$. The homotopy exact sequence then has segments of the form
$$\ldots \to \pi_{k+1}(S^n) \to \pi_k(\mathrm{SO}(n)) \xrightarrow{i_*} \pi_k(\mathrm{SO}(n+1)) \to \pi_k(S^n) \to \ldots,$$
and taking $k = 1$, both $\pi_2(S^n)$ and $\pi_1(S^n)$ vanish if $n \geqslant 3$. This produces an infinite sequence of isomorphisms
$$\mathbb{Z}_2 \cong \pi_1(\mathrm{SO}(3)) \cong \pi_1(\mathrm{SO}(4)) \cong \pi_1(\mathrm{SO}(5)) \cong \ldots,$$
proving that $\pi_1(\mathrm{GL}_+(n, \mathbb{R})) \cong \mathbb{Z}_2$ for all $n \geqslant 3$.

57. The theorems of Hurewicz and Whitehead (February 16, 2024)

I have more to say about higher homotopy groups, but I want to focus the discussion around a particular application:

THEOREM 57.1. *Every closed simply connected 3-manifold is homotopy equivalent to S^3.*

You may have heard of the Poincaré conjecture, which was open for most of the 20th century and proved by Perelman early in the 21st: it strengthens the theorem above to the statement that every closed simply connected 3-manifold is *homeomorphic* to S^3. Actually, Poincaré himself was originally more ambitious and suggested that every closed 3-manifold M with $H_*(M) \cong H_*(S^3)$ should be homeomorphic to S^3, but he found a counterexample to this conjecture a few years later, now known as the *Poincaré homology sphere*. It was not simply connected and therefore, obviously, not homotopy equivalent to S^3. Theorem 57.1 thus made Poincaré's strengthened conjecture seem plausible, but in general there is a very wide gap between homotopy equivalence and homeomorphism, i.e. even in dimension three, there are many known examples of pairs of closed manifolds that are homotopy equivalent but not homeomorphic. The proper statement of Poincaré's conjecture is thus that there is something special about spheres which makes homotopy equivalence imply homeomorphism, and in fact, that is also the right way to state the *higher-dimensional* Poincaré conjecture, proved by Smale around 1960 for dimensions $n \geqslant 5$ and Freedman around 1980 for dimension 4. From dimension four upwards, it is easy to see that simple connectedness would not be enough, e.g. $\mathbb{CP}^2$ is an easy example of a closed simply connected 4-manifold that is not a sphere, and there are many more. But we can easily distinguish $\mathbb{CP}^2$ from S^4 via its homology, of course. Part of the interest in Theorem 57.1, for our purposes, is the way that the condition $\pi_1(M) = 0$ in dimension three produces just enough constraints on $H_*(M)$ to make all the familiar obstructions to a homotopy equivalence between M and S^3 vanish, starting with the homology and cohomology groups, and then continuing with the higher homotopy groups. Several of the important theorems we've proved in this course have some role to play in the proof, thus it will serve both as a review of the course and as motivation to introduce two new and powerful theorems involving the higher homotopy groups.

Part 1: From simply connected to homology sphere. This part of the argument will be a review of techniques developed in the course. Our first main objective is to prove the following lemma:

LEMMA 57.2. *If M is a closed and connected 3-manifold with $\pi_1(M) = 0$, then $H_*(M; \mathbb{Z}) \cong H_*(S^3; \mathbb{Z})$.*

Any manifold for which this conclusion holds is called a **homology 3-sphere**. We shall prove this as an amalgamation of several smaller lemmas. Assume henceforth that M is a closed and simply connected 3-manifold.

LEMMA 57.3. *$H_n(M;\mathbb{Z})$ is finitely generated for all n and vanishes for $n > 3$.*

PROOF. The homology of every compact n-manifold is finitely generated since all such manifolds are Eulidean neighborhood retracts; see Theorem 52.1. The groups $H_k(M;\mathbb{Z})$ for $k > n$ vanish by Corollary 52.18. Alternatively, one could in the present case appeal to the (much harder) fact that all topological 3-manifolds are triangulable (see e.g. [Moi77]), thus M is a 3-dimensional finite cell complex and the lemma therefore follows from cellular homology. $\square$

LEMMA 57.4. *$H_1(M;\mathbb{Z}) = 0$.*

PROOF. This is immediate from the isomorphism of $H_1(M;\mathbb{Z})$ with the abelianization of $\pi_1(M)$. $\square$

LEMMA 57.5. *M is orientable.*

PROOF. If it is not orientable, then its orientation double cover $\pi : \widetilde{M} \to M$ is a connected 3-manifold. But the Galois correspondence identifies the set of connected covers of M up to isomorphism with the set of all subgroups of $\pi_1(M)$, and the latter has only one element, hence the only connected cover of M is the identity map (which is the universal cover). $\square$

LEMMA 57.6. *For every choice of coefficient group G, $H_3(M;G) \cong G$.*

PROOF. This is true in the top dimension for every closed, connected and oriented manifold, by Corollary 52.18. $\square$

LEMMA 57.7. *$H_2(M;\mathbb{Z})$ is torsion free.*

PROOF. This is true for $H_{n-1}(M;\mathbb{Z})$ whenever M is a closed oriented n-manifold; see Exercise 52.22(a). Since every closed manifold is the disjoint union of its finitely many connected components, it suffices to consider the case where M is connected. The idea is then to apply the universal coefficient theorem for homology with coefficients $\mathbb{Z}_p$ for any prime number p: it gives an isomorphism

$$H_n(M;\mathbb{Z}_p) \cong (H_n(M;\mathbb{Z}) \otimes \mathbb{Z}_p) \oplus \operatorname{Tor}(H_{n-1}(M;\mathbb{Z}), \mathbb{Z}_p),$$

where we are working in the setting of $\mathbb{Z}$-modules and thus abbreviating $\operatorname{Tor} := \operatorname{Tor}^{\mathbb{Z}}$. Since $H_n(M;\mathbb{Z}_p) \cong \mathbb{Z}_p$ and $H_n(M;\mathbb{Z}) \cong \mathbb{Z}$ by Corollary 52.18, this isomorphism implies the vanishing of $\operatorname{Tor}(H_{n-1}(M;\mathbb{Z}), \mathbb{Z}_p)$. Since $H_{n-1}(M;\mathbb{Z})$ is finitely generated, we can then use the classification of finitely-generated abelian groups to write

$$H_{n-1}(M;\mathbb{Z}) \cong F \oplus \left(\bigoplus_{i=1}^N \mathbb{Z}_{k_i} \right)$$

for some free abelian group F and integers $N \geqslant 0$, $k_1, \dots, k_N \geqslant 2$, where $N > 0$ if and only if $H_{n-1}(M;\mathbb{Z})$ has torsion. According to the properties of Tor proved in Lecture 43, we then have

$$0 = \operatorname{Tor}(H_{n-1}(M;\mathbb{Z}), \mathbb{Z}_p) \cong \bigoplus_{i=1}^N \operatorname{Tor}(\mathbb{Z}_{k_i}, \mathbb{Z}_p),$$

implying $\operatorname{Tor}(\mathbb{Z}_{k_i}, \mathbb{Z}_p) = 0$ for every $i = 1, \dots, N$ and every prime p. But if p is chosen to be any prime factor of k_1, then Exercise 43.5 also gives

$$\operatorname{Tor}(\mathbb{Z}_{k_1}, \mathbb{Z}_p) = \ker \left(\mathbb{Z}_p \xrightarrow{\cdot k_1} \mathbb{Z}_p \right) = \ker \left(\mathbb{Z}_p \xrightarrow{0} \mathbb{Z}_p \right) = \mathbb{Z}_p \neq 0,$$

which is a contradiction unless $N = 0$. $\square$

The last step is to apply Poincaré duality and the universal coefficient theorem for cohomology: the former gives

$$H^2(M;\mathbb{Z}) \cong H_1(M;\mathbb{Z}) = 0,$$

and the latter then implies

$$0 = H^2(M;\mathbb{Z}) \cong \operatorname{Hom}(H_2(M;\mathbb{Z}),\mathbb{Z}) \oplus \operatorname{Ext}(H_1(M;\mathbb{Z}),\mathbb{Z}),$$

hence $\operatorname{Hom}(H_2(M;\mathbb{Z}),\mathbb{Z}) = 0$. Since $H_2(M;\mathbb{Z})$ is torsion free, it follows that $H_2(M;\mathbb{Z}) = 0$. We already have isomorphisms $H_n(M;\mathbb{Z}) \cong H_n(S^3;\mathbb{Z})$ for $n \geqslant 3$ by Lemmas 57.3 and 57.6, and $H_0(M;\mathbb{Z}) \cong H_0(S^3;\mathbb{Z}) \cong \mathbb{Z}$ is immediate since M is connected, so this completes the proof of Lemma 57.2.

Part 2: From homology sphere to homotopy sphere. The step from $\pi_1(M) = 0$ and $H_*(M) \cong H_*(S^3)$ to $M \underset{h.e.}{\simeq} S^3$ requires two theorems about homotopy groups that we will need to quote without proof, though the proofs (explained e.g. in [**Hat02**, Chapter 4]) do not require substantial machinery beyond what we have discussed in this course.

Definition 57.8. A map $f : X \to Y$ is called a **weak homotopy equivalence** if for all choices of base points $x_0 \in X$ and $y_0 = f(x_0) \in Y$, $f_* : \pi_n(X,x_0) \to \pi_n(Y,y_0)$ is an isomorphism for all $n \geqslant 0$.

Theorem 56.5 in the previous lecture implies that every homotopy equivalence is also a weak homotopy equivalence. We also know of course that if $f : X \to Y$ is a homotopy equivalence, then the induced maps on homology and cohomology groups are isomorphisms, but we are not giving any name to the latter condition because it is not sufficiently useful on its own. By contrast, the notion of a weak homotopy equivalence justifies itself through the following result:

Theorem 57.9 (Whitehead's theorem). *If X and Y are both homotopy equivalent to CW-complexes, then every weak homotopy equivalence $f : X \to Y$ is a homotopy equivalence.* □

While I do not intend to discuss the proof of this theorem, you will hopefully gain some intuition about it from Theorem 56.11 in the previous lecture; in particular, it should be clear why having cell decompositions of X and Y might be useful in the proof.

With Whitehead's theorem added to our toolbox, it would suffice to find a map $f : M \to S^3$ that induces isomorphisms $\pi_n(M) \to \pi_n(S^3)$ for all n. This project seems hopeless if we don't yet even know how to compute $\pi_n(M)$ for $n \geqslant 2$, so we first need another tool for transforming our computation of $H_*(M)$ into information about the higher homotopy groups. The obvious tool to consider is the so-called **Hurewicz map**,

$$h : \pi_n(X,x_0) \to H_n(X;\mathbb{Z}) : [f] \mapsto f_*[S^n],$$

defined in terms of the fundamental class $[S^n] \in H_n(S^n;\mathbb{Z})$ and maps $f : (S^n,\mathrm{pt}) \to (X,x_0)$ representing elements of $\pi_n(X,x_0)$. We've seen that for $n = 1$, this map cannot generally be an isomorphism since $H_1(X;\mathbb{Z})$ is always abelian while $\pi_1(X)$ is not, but the next best thing is true: when $\pi_0(X) = 0$, $h : \pi_1(X) \to H_1(X;\mathbb{Z})$ descends to an isomorphism on the abelianization of $\pi_1(X)$. For $n \geqslant 2$, both groups are abelian, so there is some hope of $h : \pi_n(X) \to H_n(X;\mathbb{Z})$ actually being an isomorphism, though we've also seen cases where this is not true: e.g. $\pi_2(\mathbb{T}^2) = 0$ but $H_2(\mathbb{T}^2;\mathbb{Z}) \cong \mathbb{Z}$. The Hurewicz theorem gives sufficient conditions for h to be an isomorphism, or to put it another way, for every n-dimensional homology class in X to correspond to a unique *spherical* homology class.

Theorem 57.10 (Hurewicz's theorem). *Suppose (X,x_0) is a pointed space that is $(n-1)$-connected for some $n \geqslant 2$. Then $\widetilde{H}_k(X) = 0$ for all $k \leqslant n-1$, and the Hurewicz map $h : \pi_n(X,x_0) \to H_n(X;\mathbb{Z})$ is an isomorphism.* □

Here are a couple of applications before we get back to discussing 3-manifolds homotopy equivalent to S^3.

COROLLARY 57.11. *If X is path-connected and has universal cover $\widetilde{X} \to X$, then $\pi_2(X) \cong H_2(\widetilde{X};\mathbb{Z})$.*

PROOF. Since S^2 is simply connected, any map $S^2 \to X$ or homotopy of such maps can be lifted to $\widetilde{X}$, implying $\pi_2(X) \cong \pi_2(\widetilde{X})$. Since $\widetilde{X}$ is simply connected, the Hurewicz theorem then identifies $\pi_2(\widetilde{X})$ with $H_2(\widetilde{X};\mathbb{Z})$. □

COROLLARY 57.12. *If X is a simply connected CW-complex with $\widetilde{H}_*(X;\mathbb{Z}) = 0$, then X is contractible.*

PROOF. The Hurewicz theorem gives an isomorphism $\pi_2(X) \cong H_2(X;\mathbb{Z}) = 0$, proving X is 2-connected, so one can then apply the theorem again and conclude $\pi_3(X) \cong H_3(X;\mathbb{Z}) = 0$, and then again…by induction, we deduce $\pi_n(X) = 0$ for all $n \geqslant 0$. It follows that the unique map $\epsilon : X \to \{\mathrm{pt}\}$ induces isomorphisms $\epsilon_* : \pi_n(X) \to \pi_n(\{\mathrm{pt}\}) = 0$ for all $n \geqslant 0$ and is therefore a weak homotopy equivalence. Whitehead's theorem then implies that it is also a homotopy equivalence. □

You can now imagine at least part of a strategy to complete the proof of Theorem 57.1: instead of the map $\epsilon : X \to \{\mathrm{pt}\}$ in the proof of Corollary 57.12, one could take any map $f : M \to S^3$ of degree 1 and try to prove that $f_* : \pi_n(M) \to \pi_n(S^3)$ is an isomorphism for all $n \geqslant 0$. This idea can be carried out for all $n \leqslant 3$, as Hurewicz now transforms the computation $H_*(M;\mathbb{Z}) \cong H_*(S^3;\mathbb{Z})$ into $\pi_1(M) = \pi_2(M) = 0$ and $\pi_3(M) \cong \mathbb{Z}$. For $n \geqslant 4$, however, we get stuck, among other reasons because it is not so clear what $\pi_n(S^3)$ is, and the Hurewicz theorem provides no information about this above the lowest dimension where $\widetilde{H}_n(S^3;\mathbb{Z}) \neq 0$.

To make further progress, we need a relative version of the Hurewicz theorem. Given $x_0 \in A \subset X$, there is a relative Hurewicz map defined for each $n \in \mathbb{N}$ by

$$h : \pi_n(X, A, x_0) \to H_n(X, A;\mathbb{Z}) : [f] \mapsto f_*[\mathbb{D}^n],$$

where $[f] \in \pi_n(X, A, x_0)$ is represented by a map $f : (\mathbb{D}^n, \partial\mathbb{D}^n, \mathrm{pt}) \to (X, A, x_0)$ and $[\mathbb{D}^n] \in H_n(\mathbb{D}^n, \partial\mathbb{D}^n;\mathbb{Z})$ denotes the relative fundamental class of $\mathbb{D}^n$. One can check that this map is a homomorphism for each $n \geqslant 2$. Let us say that the pair (X, A) is n-**connected** if $\pi_k(X, A) = 0$ for all $k \leqslant n$. Since $\pi_2(X, A, x_0)$ is not always abelian, we cannot generally expect $h : \pi_2(X, A, x_0) \to H_2(X, A;\mathbb{Z})$ to be an isomorphism, even if (X, A) is 1-connected. Observe however that if A is additionally assumed to be simply connected, then the long exact sequence of homotopy groups for (X, A) has a segment of the form

$$\ldots \to \pi_2(X) \to \pi_2(X, A) \to \pi_1(A) = 0,$$

implying that $\pi_2(X, A)$ is the surjective image of a homomorphism defined on the *abelian* group $\pi_2(X)$, and is therefore also abelian. This serves as a sanity check for the following generalization of Theorem 57.10:

THEOREM 57.13. *Suppose (X, A) is an $(n-1)$-connected pair of spaces for some $n \geqslant 2$, where $A \subset X$ is also simply connected and $x_0 \in A$ is a base point. Then $H_k(X, A) = 0$ for all $k \leqslant n-1$, and the relative Hurewicz map $h : \pi_n(X, A, x_0) \to H_n(X, A;\mathbb{Z})$ is an isomorphism.* □

COROLLARY 57.14. *Suppose X and Y are two simply connected spaces that are both homotopy equivalent to CW-complexes, and $f : X \to Y$ is a map that induces isomorphisms $f_* : H_n(X;\mathbb{Z}) \to H_n(Y;\mathbb{Z})$ for every $n \geqslant 0$. Then f is a homotopy equivalence.*

Proof. We first prove it under the simplifying assumption that $X \subset Y$ is a subspace with $f : X \hookrightarrow Y$ as the inclusion map. The long exact sequence of the pair (Y, X) in homology converts the assumption $f_* : H_n(X;\mathbb{Z}) \xrightarrow{\cong} H_n(Y;\mathbb{Z})$ into

$$H_n(Y, X;\mathbb{Z}) = 0 \quad \text{for all } n \geqslant 0.$$

Similarly, the long exact sequence of relative homotopy groups includes a segment of the form

$$0 = \pi_1(Y) \to \pi_1(Y, X) \to \pi_0(X) = 0,$$

implying $\pi_1(Y, X) = 0$, so that the relative Hurewicz theorem can be applied to the pair (Y, X) with $n = 2$, producing an isomorphism $\pi_2(Y, X) \cong H_2(Y, X;\mathbb{Z}) = 0$ and thus proving that (Y, X) is 2-connected. One can then apply the relative Hurewicz theorem again with $n = 3$, and continue this process inductively to prove $\pi_n(Y, X) = 0$ for all $n \geqslant 0$. In light of the exact sequence

$$0 = \pi_{n+1}(Y, X) \to \pi_n(X) \xrightarrow{f_*} \pi_n(Y) \to \pi_n(Y, X) = 0,$$

this proves that $f : X \hookrightarrow Y$ is a weak homotopy equivalence, so Whitehead's theorem implies that it is a homotopy equivalence.

To generalize beyond the case where $f : X \to Y$ is an inclusion, we consider the **mapping cylinder** of f, defined as the space

$$M_f := \left((X \times I) \amalg Y\right)/\!\sim \quad \text{where} \quad (x, 1) \sim f(x) \text{ for all } x \in X.$$

This space contains disjoint homeomorphic copies of X and Y, namely the images of the inclusion maps

$$i_X : X \hookrightarrow M_f : x \mapsto [(x, 0)], \qquad \text{and} \qquad i_Y : Y \hookrightarrow M_f : y \mapsto [y],$$

and the latter is a homotopy equivalence due to the obvious deformation retraction of M_f to $i_Y(Y) \subset M_f$ defined by pushing the t-coordinate of each $(x, t) \in X \times I$ upward toward 1. This is one of two crucial properties that the mapping cylinder has; the other is that the diagram

$$\begin{tikzcd} Y \arrow[r, "i_Y"] & M_f \\ X \arrow[u, "f"] \arrow[ur, "i_X"'] & \end{tikzcd}$$

commutes *up to homotopy*, meaning that while the two maps $i_X : X \to M_f$ and $i_Y \circ f : X \to M_f$ have disjoint images and are thus obviously not equal, they are homotopic. As a consequence, $f : X \to Y$ is a homotopy equivalence if and only if $i_X : X \to M_f$ is a homotopy equivalence, and the inclusion map $i_X : X \hookrightarrow M_f$ can thus be used as a substitute for $f : X \to Y$ in arguments that depend only on homotopy type. In particular, if X and Y are simply connected and $f_* : H_*(X;\mathbb{Z}) \to H_*(Y;\mathbb{Z})$ is an isomorphism, then M_f is also simply connected and $(i_X)_* : H_*(X;\mathbb{Z}) \to H_*(M_f;\mathbb{Z})$ is also an isomorphism, so the argument of the previous paragraph makes i_X a homotopy equivalence, and so therefore is f. □

Conclusion of the proof of Theorem 57.1. We have shown thus far that if M is a closed simply connected 3-manifold, then $H_*(M;\mathbb{Z}) \cong H_*(S^3;\mathbb{Z})$. Now pick any map $f : M \to S^3$ that has degree 1. Such maps are easily found by identifying S^3 with the one-point compactification $\mathbb{R}^3 \cup \{\infty\}$, then choosing a Euclidean neighborhood $\mathcal{U} \subset M$ and defining $f : M \to S^3$ to be a homeomorphism $\mathcal{U} \xrightarrow{\cong} \mathbb{R}^3$ on this neighborhood while sending every other point to ∞. The characterization of the mapping degree via local degrees in Lecture 35 implies $\deg(f) = 1$.

It is trivial that $f_* : H_0(M;\mathbb{Z}) \to H_0(S^3;\mathbb{Z})$ is an isomorphism, and so is $f_* : H_3(M;\mathbb{Z}) \to H_3(S^3;\mathbb{Z})$ due to the degree assumption. In all other dimensions, both homology groups vanish, so we conclude that $f_* : H_*(M;\mathbb{Z}) \to H_*(S^3;\mathbb{Z})$ is an isomorphism. Since M and S^3 are both simply connected, Corollary 57.14 now implies that f is a homotopy equivalence. $\square$

Bibliography

[Are46] R. Arens, *Topologies for homeomorphism groups*, Amer. J. Math. **68** (1946), 593–610. MR19916

[Art91] M. Artin, *Algebra*, Prentice Hall, Inc., Englewood Cliffs, NJ, 1991.

[AD14] M. Audin and M. Damian, *Morse theory and Floer homology*, Universitext, Springer, London; EDP Sciences, Les Ulis, 2014. Translated from the 2010 French original by Reinie Erné.

[Bae40] R. Baer, *Abelian groups that are direct summands of every containing abelian group*, Bull. Amer. Math. Soc. **46** (1940), 800–806.

[Bla79] A. Blass, *Injectivity, projectivity, and the axiom of choice*, Trans. Amer. Math. Soc. **255** (1979), 31–59.

[Boo77] B. Booss, *Topologie und Analysis*, Springer-Verlag, Berlin-New York, 1977 (German). Einführung in die Atiyah-Singer-Indexformel; Hochschultext.

[BB85] B. Booss and D. D. Bleecker, *Topology and analysis*, Universitext, Springer-Verlag, New York, 1985. The Atiyah-Singer index formula and gauge-theoretic physics; Translated from the German by Bleecker and A. Mader.

[Bre93] G. E. Bredon, *Topology and geometry*, Springer-Verlag, New York, 1993.

[BT] A. D. Brooke-Taylor, *Products of CW complexes*. Preprint arXiv:1710.05296.

[Che92] P. R. Chernoff, *A simple proof of Tychonoff's theorem via nets*, Amer. Math. Monthly **99** (1992), no. 10, 932–934.

[DK90] S. K. Donaldson and P. B. Kronheimer, *The geometry of four-manifolds*, Oxford Mathematical Monographs, The Clarendon Press Oxford University Press, New York, 1990. Oxford Science Publications.

[ES53] B. Eckmann and A. Schopf, *Über injektive Moduln*, Arch. Math. (Basel) **4** (1953), 75–78 (German).

[ES52] S. Eilenberg and N. Steenrod, *Foundations of algebraic topology*, Princeton University Press, Princeton, New Jersey, 1952.

[FW99] G. K. Francis and J. R. Weeks, *Conway's ZIP proof*, Amer. Math. Monthly **106** (1999), no. 5, 393–399.

[Gal87] D. Gale, *The Teaching of Mathematics: The Classification of 1-Manifolds: A Take-Home Exam*, Amer. Math. Monthly **94** (1987), no. 2, 170–175.

[Hat02] A. Hatcher, *Algebraic topology*, Cambridge University Press, Cambridge, 2002.

[Hir94] M. W. Hirsch, *Differential topology*, Springer-Verlag, New York, 1994.

[Hut] M. Hutchings, *Cup product and intersections*. lecture notes, available at `https://math.berkeley.edu/~hutching/teach/215b-2011/cup.pdf`.

[Jän05] K. Jänich, *Topologie*, 8th ed., Springer-Verlag, Berlin, 2005 (German).

[JW58] I. M. James and J. H. C. Whitehead, *Homology with zero coefficients*, Quart. J. Math. Oxford Ser. (2) **9** (1958), 317–320.

[Kel50] J. L. Kelley, *The Tychonoff product theorem implies the axiom of choice*, Fund. Math. **37** (1950), 75–76.

[Kel75] J. L. Kelley, *General topology*, Graduate Texts in Mathematics, vol. 27, Springer-Verlag, New York, 1975. Reprint of the 1955 edition [Van Nostrand, Toronto, Ont.]

[KS69] R. C. Kirby and L. C. Siebenmann, *On the triangulation of manifolds and the Hauptvermutung*, Bull. Amer. Math. Soc. **75** (1969), 742–749.

[Lan02] S. Lang, *Algebra*, 3rd ed., Graduate Texts in Mathematics, vol. 211, Springer-Verlag, New York, 2002.

[Lee11] J. M. Lee, *Introduction to topological manifolds*, 2nd ed., Graduate Texts in Mathematics, vol. 202, Springer, New York, 2011.

[LL01] E. H. Lieb and M. Loss, *Analysis*, 2nd ed., Graduate Studies in Mathematics, vol. 14, American Mathematical Society, Providence, RI, 2001.

[Man14] C. Manolescu, *Triangulations of manifolds*, ICCM Not. **2** (2014), no. 2, 21–23.

[Mas91] W. S. Massey, *A basic course in algebraic topology*, Graduate Texts in Mathematics, vol. 127, Springer-Verlag, New York, 1991.

[MS12] D. McDuff and D. Salamon, *J-holomorphic curves and symplectic topology*, 2nd ed., American Mathematical Society Colloquium Publications, vol. 52, American Mathematical Society, Providence, RI, 2012.

[Mil62] J. W. Milnor, *On axiomatic homology theory*, Pacific J. Math. **12** (1962), 337–341.

[Mil97] ______, *Topology from the differentiable viewpoint*, Princeton Landmarks in Mathematics, Princeton University Press, Princeton, NJ, 1997. Based on notes by David W. Weaver; Revised reprint of the 1965 original.

[MS74] J. W. Milnor and J. D. Stasheff, *Characteristic classes*, Princeton University Press, Princeton, N. J., 1974. Annals of Mathematics Studies, No. 76.

[Moi77] E. E. Moise, *Geometric topology in dimensions* 2 *and* 3, Springer-Verlag, New York-Heidelberg, 1977. Graduate Texts in Mathematics, Vol. 47.

[Pal66] R. S. Palais, *Homotopy theory of infinite dimensional manifolds*, Topology **5** (1966), 1–16.

[Rad25] T. Radó, *Über den Begriff der Riemannschen Fläche*, Acta Szeged **2** (1925), 101–121.

[Rud87] W. Rudin, *Real and complex analysis*, 3rd ed., McGraw-Hill Book Co., New York, 1987.

[Sch93] M. Schwarz, *Morse homology*, Progress in Mathematics, vol. 111, Birkhäuser Verlag, Basel, 1993.

[Spa95] E. H. Spanier, *Algebraic topology*, Springer-Verlag, New York, 1995. Corrected reprint of the 1966 original.

[Spa48] ______, *Cohomology theory for general spaces*, Ann. of Math. (2) **49** (1948), 407–427.

[Ste51] N. Steenrod, *The Topology of Fibre Bundles*, Princeton University Press, Princeton, N. J., 1951.

[Ste67] N. E. Steenrod, *A convenient category of topological spaces*, Michigan Math. J. **14** (1967), 133–152.

[Tho54] R. Thom, *Quelques propriétés globales des variétés différentiables*, Comment. Math. Helv. **28** (1954), 17–86 (French).

[Tho92] C. Thomassen, *The Jordan-Schönflies theorem and the classification of surfaces*, Amer. Math. Monthly **99** (1992), no. 2, 116–130.

[Vic94] J. W. Vick, *Homology theory*, 2nd ed., Graduate Texts in Mathematics, vol. 145, Springer-Verlag, New York, 1994. An introduction to algebraic topology.

[Wen18] C. Wendl, *Holomorphic curves in low dimensions: from symplectic ruled surfaces to planar contact manifolds*, Lecture Notes in Mathematics, vol. 2216, Springer-Verlag, 2018.