

Topology I—III, HU Berlin

Chris Wendl

Contents

Preface	v
About the current version	v
Disclaimer and acknowledgements	vi
First semester (Topologie I)	1
1. Introduction and motivation	1
2. Metric spaces	5
3. Topological spaces	12
4. Products, sequential continuity and nets	17
5. Compactness	24
6. Tychonoff's theorem and the separation axioms	29
7. Connectedness and local compactness	37
8. Paths, homotopy and the fundamental group	46
9. Some properties of the fundamental group	52
10. Retractions and homotopy equivalence	57
11. The easy part of van Kampen's theorem	63
12. Normal subgroups, generators and relations	68
13. Proof of the Seifert-van Kampen theorem	75
14. Surfaces and torus knots	80
15. Covering spaces and the lifting theorem	91
16. Classification of covers	96
17. The universal cover and group actions	101
18. Manifolds	106
19. Surfaces and triangulations	116
20. Orientations	123
21. Higher homotopy, bordism, and simplicial homology	130
22. Singular homology	140
23. Relative homology and long exact sequences	146
24. Homotopy invariance and excision	153
25. The homology of the spheres, and applications	163
26. Axioms, cells, and the Euler characteristic	166
Second semester (Topologie II)	177
27. Categories and functors	177
28. Axioms for homology theories	190
29. Simplicial homology	203
30. Triangulated manifolds and subdivision	215
31. Chain homotopy and simplicial approximation	222
32. Acyclic models and relative homology	231
33. Singular homology	239

34. Pairs, triples, and the Mayer-Vietoris sequence	250
35. Mapping tori and maps between spheres	263
36. Local and global mapping degree	269
37. CW-complexes	279
38. Invariance of cellular homology	288
39. Direct limits and infinite-dimensional cell complexes	297
40. Euler characteristic and fixed points	311
41. Singular cohomology	319
42. Axioms for cohomology	327
43. Inverse limits, CW-complexes and Čech theory	336
44. Universal coefficient theorems and exact functors	352
45. The derived functors Tor and Ext	364
46. The cross product	390
47. Products in singular cohomology	405
48. Relative cross, cup and cap products	416
49. The orientation bundle	424
50. Existence of the fundamental class	435
51. Poincaré duality	445
52. The Thom class	458
53. The intersection product	466
54. Higher homotopy groups	478
55. The theorems of Hurewicz and Whitehead	487
Third semester (Topologie III)	493
56. Motivation: Detecting exotic 7-spheres	493
Bibliography	495

Preface

The original version of these notes was created in 2018–19 for a two-semester sequence of topology courses at Humboldt University, Berlin. It has undergone substantial revisions during a few repetitions of the two-semester course since then, plus the addition of a third semester in 2025. The topics are divided among the three semesters roughly as follows:

- First semester: basic point-set topology, fundamental group and covering spaces, manifolds of dimension one and two, introduction to homology
- Second semester: homology and cohomology
- Third semester: homotopy theory, higher homotopy groups, fiber bundles and characteristic classes

A few topics appear in multiple semesters, e.g. while the end of the first semester contains material on singular homology, the second semester does not assume previous knowledge of homology, and thus starts that subject from the beginning, though at a slightly higher level of sophistication. This reflects the fact that at our university, Topology I is technically a Bachelor-level course and Topology II is technically Master-level, though in practice, the audience for both courses is typically a mixture.

There is a nearly exact one-to-one correspondence between the chapters in these notes and the actual 90-minute lectures given in the course, though for some chapters that are a bit fatter, some portions had to be skipped or mentioned only briefly in class.

Since the notes were designed for use at a German university, I have made an effort to include the German translations (*geschrieben in dieser Schriftart*) of important terms wherever they are introduced. The reader may notice that this effort subsides later in the course, as the deeper one gets into algebraic topology, the harder it becomes to find authoritative German sources for clarifying the terminology (and I am not linguistically qualified to invent terms in German myself).

About the current version

The version you are looking at right now is being updated regularly in order to serve as lecture notes for the HU’s course “Topics in Topology” (unofficially known as *Topology III*) in the Summer 2025 semester.

One innovation of the current version—implemented in the notes for the second and third semesters but not yet for the first semester—is that all exercises now appear in their own subsection at the end of each lecture, and some of them are marked with an asterisk (like this (*)). The asterisk means that the exercise is *essential*, e.g. because it contains a proof of some important result that will be used again in the course, perhaps multiple times. Exercises without an asterisk are intended to be helpful and/or informative, but not essential for the logical continuity of the notes.

Most recent update: **March 24, 2025**

Disclaimer and acknowledgements

These lecture notes were written quickly, and while many typos have in the mean time been eliminated due to careful reading by a few motivated students, some probably remain. If you notice any, please send me an e-mail and I will correct. Thanks for corrections already received are due to Lennard Henze, Jens Lücke, Mateusz Majchrzak, Marie Christin Schmidlein, Rens Breur, Maxim Nevkrytyh, Laurenz Upmeier zu Belzen, Florian Kaufmann, Ben Eltschig and Daniel Acker. (Apologies if I forgot anyone!)


First semester (Topologie I)

1. Introduction and motivation

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*) $U \subset X$ and *continuous maps/mappings* (*stetige Abbildungen*) $f : X \rightarrow Y$ (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 $f : X \rightarrow Y$ whose inverse $f^{-1} : Y \rightarrow X$ 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 $g \geq 0$.

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 $\mathbb{D}^2 := \{\mathbf{x} \in \mathbb{R}^2 \mid |\mathbf{x}| \leq 1\}$ 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*)

$$\mathbb{M} := \{(\theta, t \cos(\pi\theta), t \sin(\pi\theta)) \in \mathbb{R}/\mathbb{Z} \times \mathbb{R}^2 \mid \theta \in \mathbb{R}, t \in [-1, 1]\}.$$

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 $n \geq 0$ one can define the **n -sphere**

$$S^n = \{\mathbf{x} \in \mathbb{R}^{n+1} \mid |\mathbf{x}| = 1\}$$

and the **real projective n -space**

$$\mathbb{RP}^n = S^n / \{\mathbf{x} \sim -\mathbf{x}\} = \{\text{lines through 0 in } \mathbb{R}^{n+1}\}.$$

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 $f : A \rightarrow X$ 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 $f : \Sigma_g \rightarrow \Sigma_h$ 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 $f : X \rightarrow Y$ induces a homomorphism $f_* : H(X) \rightarrow H(Y)$, and such that compositions of continuous maps satisfy

$$(f \circ g)_* = f_* \circ g_*$$

and the identity map $\text{Id} : X \rightarrow X$ gives rise to the identity map $H(X) \rightarrow H(X)$. These properties imply that whenever $f : X \rightarrow Y$ is a homeomorphism, $f_* : H(X) \rightarrow H(Y)$ 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 $g \geq 0$.*

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 $\mathbb{R}P^2$, one can understand $\mathbb{R}P^2$ 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 $n \geq 0$; obvious examples include \mathbb{R}^n , S^n and $\mathbb{R}P^n$. The classification problem becomes much harder when $n \geq 3$, 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 $f : S^1 \rightarrow X$ can be extended continuously to $\mathbb{D}^2 \rightarrow X$). 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 equivalence” (*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 $n \geq 5$, 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 $f : X \rightarrow Y$ 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 $X \rightarrow Y$ 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 $n \leq 3$, 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 $f : X \rightarrow X$ 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 $\mathbb{D}^n := \{\mathbf{x} \in \mathbb{R}^n \mid |\mathbf{x}| \leq 1\}$.

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

The case $n = 1$ is an easy consequence of the intermediate value theorem, but for $n \geq 2$, 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 $f : \mathbb{D}^n \rightarrow \mathbb{D}^n$ 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

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

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

$$(1.1) \quad g \circ i = \text{Id}_{S^{n-1}}.$$

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 $f : X \rightarrow Y$ induce group homomorphisms $f_* : H_{n-1}(X) \rightarrow H_{n-1}(Y)$ satisfying $(f \circ g)_* = f_* \circ g_*$ and $\text{Id}_* = \mathbb{1}$. Crucially, one can also compute this invariant for both \mathbb{D}^n and S^{n-1} , and the answers are

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

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

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 $d : X \times X \rightarrow \mathbb{R}$ that satisfies the following conditions for all $x, y, z \in X$:

- (i) $d(x, y) \geq 0$;

- (ii) $d(x, x) = 0$;
- (iii) $d(x, y) = d(y, x)$, i.e. “symmetry”;
- (iv) $d(x, z) \leq d(x, y) + d(y, z)$, 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

$$B_r(x) := \{y \in X \mid d(x, y) < r\}.$$

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 $X \setminus A$ 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

$$d(x, z) \leq d(x, y) + d(y, z) < d(x, y) + \epsilon < d(x, y) + r - d(x, y) = r.$$

- The function $d : X \times X \rightarrow [0, \infty)$ is *continuous* (see below for a review of the definition of continuity), since one can use the triangle inequality to show that for every $x, y, x', y' \in X$,

$$|d(x, y) - d(x', y')| \leq d(x, x') + d(y, y').$$

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, $X \setminus \{x\}$ is an open subset of X because for every $y \in X \setminus \{x\}$, the ball $B_\epsilon(y)$ is contained in $X \setminus \{x\}$ 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.

¹Yes, the empty set $\emptyset \subset X$ 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 $d : X \times X \rightarrow \mathbb{R}$ 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

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

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*) $f : X \rightarrow Y$ 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 $x_n \in X$, $x_n \rightarrow x$ implies $f(x_n) \rightarrow f(x)$.

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) \Leftrightarrow (C). To show that (b) \Rightarrow (c), suppose $x_n \rightarrow x$ 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 $f(x_n) \rightarrow f(x)$ 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 $f(x)$. \square

I want to point out two things about the above proof. First, the proof that (b) \Rightarrow (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) \Rightarrow (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 $f : X \rightarrow Y$ is a **homeomorphism** (*Homöomorphismus*) if it is continuous and bijective and its inverse $f^{-1} : Y \rightarrow X$ 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 $\mathbf{x} = (x_1, \dots, x_n)$ and $\mathbf{y} = (y_1, \dots, y_n)$ 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 $f : [0, r) \rightarrow [0, \infty)$ and define $F : B_r(\mathbf{x}) \rightarrow \mathbb{R}^n$ by

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

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 $f : X \rightarrow Y$ 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 $f(K)$. \square

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

$$S^n = \{\mathbf{x} \in \mathbb{R}^{n+1} \mid |\mathbf{x}| = 1\},$$

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

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

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 $(X, d_1) \rightarrow (X, d_2)$ if and only if the following condition is satisfied: for every sequence $x_n \in X$ and $x \in X$,

$$x_n \rightarrow x \text{ in } (X, d_1) \iff x_n \rightarrow x \text{ in } (X, d_2).$$

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

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

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 $[-1, 1] \rightarrow \mathbb{R}$. 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 $f_n \rightarrow f$ in (X, d_{∞}) , then

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

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

$$(2.1) \quad d([x], [y]) := \inf \{d_X(x, y) \mid x \in [x], y \in [y]\}.$$

- (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 $\pi : S^n \rightarrow \mathbb{RP}^n$ 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

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

with the metric d_X induced by the Euclidean metric on \mathbb{R}^2 . Now fix the function $f : X \rightarrow \mathbb{R} : (x, y) \mapsto xy$ and define the relation $p_0 \sim p_1$ for $p_0, p_1 \in X$ to mean that there exists a continuous curve $\gamma : [0, 1] \rightarrow X$ 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 $\pi : X \rightarrow X/\sim$ 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

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

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

$$d_X(x, x') \leq 1 \text{ for all } x, x' \in X, \quad d_Y(y, y') \leq 1 \text{ for all } y, y' \in Y.$$

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 $U \subset Z$ 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 $\gamma : [0, 1] \rightarrow Z$ 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

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

$$\overset{\circ}{A} = \{x \in A \mid \text{some neighborhood of } x \text{ in } X \text{ is contained in } A\}.$$

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** (*Berührungspunkte*) 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 $\overset{\circ}{A}$ 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.

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

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 \emptyset 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 collection² \mathcal{T} of subsets of X satisfying the following axioms:

- (i) $\emptyset \in \mathcal{T}$ 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 $X \setminus A \in \mathcal{T}_X$.

²I 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 $f : X \rightarrow Y$ 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 “ $x_n \rightarrow x$ ”) 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 $\mathcal{T}' := \{X \setminus \mathcal{U} \mid \mathcal{U} \in \mathcal{T}\}$. 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 $\alpha \in I$,

$$\mathcal{U}_\alpha = \mathcal{U}_\alpha^1 \cap \dots \cap \mathcal{U}_\alpha^{N_\alpha}$$

for some $N_\alpha \in \mathbb{N}$ and $\mathcal{U}_\alpha^1, \dots, \mathcal{U}_\alpha^{N_\alpha} \in \mathcal{B}$.

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 $\{(a, b) \subset \mathbb{R} \mid -\infty \leq a < b \leq \infty\}$ 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?)

³Things 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 $(X, d) \rightarrow (X, d')$ 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 $f : X \rightarrow \mathbb{R}$ is continuous.
- A map $f : \mathbb{R} \rightarrow X$ 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 $\mathcal{T} = \{\emptyset, X\}$. 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 $f : X \rightarrow \mathbb{R}$ 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 \emptyset and is therefore X . This proves

$$f(X) \subset \bigcap_{\epsilon > 0} (t - \epsilon, t + \epsilon) = \{t\}.$$

- All maps $f : \mathbb{R} \rightarrow X$ are continuous.
- $x_n \rightarrow x$ 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 $X \setminus \{x_1, \dots, x_N\}$ 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 $\mathbb{R} \setminus \{x_1, \dots, x_N\}$ 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 $X_\alpha \cap X_\beta = \emptyset$ for all $\alpha \neq \beta$, the **disjoint union** (*disjunkte Vereinigung*) is the set $X := \bigcup_{\alpha \in I} X_\alpha$ with the topology

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

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

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

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

EXERCISE 3.16. Show that the disjoint union topology \mathcal{T} on $X = \bigsqcup_{\alpha} X_\alpha$ 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 $\gamma : [0, 1] \rightarrow X$ 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 $\bigsqcup_{\alpha} X_\alpha$ 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

$$X := \{(\alpha, x) \mid \alpha \in I, x \in X_{\alpha}\},$$

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 $X_{\alpha} \rightarrow X : x \mapsto (\alpha, x)$ 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

$$\begin{aligned} \prod_{\alpha \in \mathbb{R}} \{0\} &:= \{(\alpha, 0) \in \mathbb{R} \times \{0\}\} \rightarrow \mathbb{R}, \\ &(\alpha, 0) \mapsto \alpha. \end{aligned}$$

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

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 $f : X \rightarrow Y$ 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 $f : I \rightarrow X$ 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

$$\mathcal{B} := \{\mathcal{U}_1 \times \mathcal{U}_2 \subset X_1 \times X_2 \mid \mathcal{U}_1 \in \mathcal{T}_1, \mathcal{U}_2 \in \mathcal{T}_2\}.$$

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

$$\begin{aligned} \pi_1 : X_1 \times X_2 &\rightarrow X_1 : (x_1, x_2) \mapsto x_1 \\ \pi_2 : X_1 \times X_2 &\rightarrow X_2 : (x_1, x_2) \mapsto x_2 \end{aligned}$$

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

$$\{\pi_1^{-1}(\mathcal{U}) \mid \mathcal{U} \in \mathcal{T}_1\} \cup \{\pi_2^{-1}(\mathcal{U}) \mid \mathcal{U} \in \mathcal{T}_2\}$$

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

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

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

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

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

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

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 $\alpha \in I$. \square*

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 $n \rightarrow \infty$ 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 $f : Y \rightarrow \prod_{\alpha \in I} X_\alpha$ is continuous if and only if $\pi_\alpha \circ f : Y \rightarrow X_\alpha$ 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 $I \rightarrow X$, and we write

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

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 $\{0, 1\}^I = \{f : I \rightarrow \{0, 1\}\}$ 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.

$$f_n(\alpha) \rightarrow f(\alpha) \quad \text{for every } \alpha \in I.$$

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 $\prod_{\alpha \in I} X_\alpha$, one can define the so-called *box topology*, which has a base of the form

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

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

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 $f : X \rightarrow Y$ between topological spaces equivalent?*

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

The converse is trickier. Look again at the proof in Lecture 2 that (c) \Rightarrow (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 $x_n \in X \setminus A$ such that $x_n \rightarrow x$.*

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 $f : [0, 1] \rightarrow [-1, 1]$, and we assign to it the subspace topology as a subset of the space $[-1, 1]^{[0, 1]}$ of all functions $f : [0, 1] \rightarrow [-1, 1]$. 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 :

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

If $f_n \rightarrow f$ 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 $\int_0^1 |f_n(t) - f(t)|^2 dt \rightarrow 0$. 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 $|g_n| \leq G$ 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 $d_2(f_n, f) \rightarrow 0$, 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}_α 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 $f : [0, 1] \rightarrow [0, 1]$ that vanishes on the finite subset I but equals 1 on a “large” subset of $[0, 1] \setminus I$. 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 $x \in A$. \square

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 $x_n \in X \setminus A$ such that $x_n \rightarrow x$.*

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 $x_n \in X \setminus A$ 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 $n \geq N$,

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

□

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 $X \rightarrow Y$ 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 $<$ 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) $x < x$ for all x , (ii) $x < y$ and $y < x$ implies $x = y$, and (iii) $x < y$ and $y < z$ implies $x < z$. We write “ $x > y$ ” as a synonym for “ $y < x$ ”, and the set I together with its partial order $<$ is called a **partially ordered set** (*partiell geordnete Menge*). One obvious example is (\mathbb{N}, \leq) , though unlike this example (which is *totally* ordered), it is not generally required in a partially ordered set $(I, <)$ that every pair of elements $x, y \in I$ satisfy either $x < y$ or $y < x$. We will see more exotic examples below.

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

The natural numbers (\mathbb{N}, \leq) 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 $(I, <)$ where I is the set of all neighborhoods of x in X , and $\mathcal{U} < \mathcal{V}$ 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 $\mathcal{U} < \mathcal{V}$ or $\mathcal{V} < \mathcal{U}$.

DEFINITION 4.20. Given a space X , a **net** (*Netz*) $\{x_\alpha\}_{\alpha \in I}$ in X is a function $I \rightarrow X : \alpha \mapsto x_\alpha$, where $(I, <)$ 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 $\alpha > \alpha_0$.

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 $(I, <) = (\mathbb{N}, \leq)$ 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 $(I, <)$ 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 ,

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

□

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 $f : X \rightarrow Y$ 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 $\alpha > \alpha_0$ 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 $f : X \rightarrow Y$ 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 $\Phi : X \rightarrow Y$ 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 $A < B$ for $A, B \subset [0, 1]$ defined to mean $A \subset B$. Note that $(I, <)$ 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

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

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 $A > A_0 \in I$,

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

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

5. Compactness

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 $\{\alpha_1, \dots, \alpha_N\} \subset I$ such that $A \subset \mathcal{U}_{\alpha_1} \cup \dots \cup \mathcal{U}_{\alpha_N}$. 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 $A = \{a_1, \dots, a_N\}$ 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 $i = 1, \dots, N$, then the sets $\mathcal{U}_{\alpha_1}, \dots, \mathcal{U}_{\alpha_N}$ 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 $\langle \cdot, \cdot \rangle$ 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 $\bar{B}_1(0) = \{x \in \mathcal{H} \mid \langle x, x \rangle \leq 1\}$ 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, \dots , i.e. satisfying

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

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 closure⁴ of A is X .

- 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 .
- Deduce that every separable and metrizable space is second countable.
- Show that \mathbb{R}^n with its standard topology is separable.
- 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, $X \setminus A$ is open, so that supplementing the collection $\{\mathcal{U}_\alpha\}_{\alpha \in I}$ with $X \setminus A$ defines an open cover of X , and therefore also an open cover of K . Since K is compact, there is then a finite subset $\{\alpha_1, \dots, \alpha_N\} \subset I$ such that

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

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

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 $f : X \rightarrow Y$ is continuous and $K \subset X$ is compact, then so is $f(K) \subset Y$.* \square

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

⁴We 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 $\pi : X \rightarrow X/\sim$ 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 $\pi : X \rightarrow X/\sim$, 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. \square

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

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

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 $(I, <)$ 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 $\alpha > \alpha_0$ 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 $(I, <) = (\mathbb{N}, \leq)$, where the net becomes a sequence, so the notion should be already familiar.

DEFINITION 5.19. Given two directed sets $(I, <)$ and $(J, <)$, 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 $\phi : J \rightarrow I$ with the property that for every $\alpha_0 \in I$, there exists $\beta_0 \in J$ for which $\beta > \beta_0$ implies $\phi(\beta) > \alpha_0$.

If $(I, <)$ and $(J, <)$ in the above definition are both (\mathbb{N}, \leq) 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 $\lim_{k \rightarrow \infty} n_k = \infty$. 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 $(I, <) = (\mathbb{N}, \leq)$, Definition 5.19 allows much more general choices for the directed set $(J, <)$ and the function $\phi : J \rightarrow \mathbb{N}$ 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 $<$ defined by

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

This makes $(J, <)$ a directed set since $(I, <)$ 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 $\phi : J \rightarrow I$ such that for all $(\beta, \mathcal{U}) \in J$, $\phi(\beta, \mathcal{U}) =: \alpha$ satisfies $\alpha > \beta$ 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. \square

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 $\alpha > \alpha_x$, $x_\alpha \notin \mathcal{U}_x$. 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 $(I, <)$ is a directed set, we can find an element $\beta \in I$ satisfying

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

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

$$(5.1) \quad x_{\mathcal{A}} \notin \bigcup_{\mathcal{U} \in \mathcal{A}} \mathcal{U}.$$

Define a partial order $<$ on I by

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

and notice that $(I, <)$ 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 $\mathcal{A} > \{\mathcal{V}\}$ 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. \square

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

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

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

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

\square

COROLLARY 5.23. *If X is compact and first countable, then it is also sequentially compact.* \square

EXAMPLE 5.24. Though it is not so easy to see this, the space $[0, 1]^{\mathbb{R}}$ of (not necessarily continuous) functions $\mathbb{R} \rightarrow [0, 1]$ 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}_{α} is a union of sets in \mathcal{B} , and the collection of all sets in \mathcal{B} that are contained in some \mathcal{U}_{α} is a countable subcollection $\mathcal{B}' \subset \mathcal{B}$ that also covers X . Let us denote $\mathcal{B}' = \{\mathcal{V}_1, \mathcal{V}_2, \mathcal{V}_3, \dots\}$. 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 $\{\mathcal{U}_{\alpha}\}_{\alpha \in I}$. \square

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 $\mathcal{U}_1, \mathcal{U}_2, \mathcal{U}_3, \dots \subset X$. Now arguing by contradiction, suppose that X is sequentially compact but the sets $\mathcal{U}_1, \dots, \mathcal{U}_n$ do not cover X for any $n \in \mathbb{N}$, hence there exists a sequence $x_n \in X$ such that

$$(5.2) \quad x_n \notin \mathcal{U}_1 \cup \dots \cup \mathcal{U}_n$$

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 $k_n \geq N$. \square

EXERCISE 5.27. Consider the space

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

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 $\{\mathcal{U}_x \subset X \mid x \in \mathbb{R}\}$ 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.* \square

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

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 $\phi : J \rightarrow I$. 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) . \square

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

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

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 $f_n : \mathbb{R} \rightarrow [0, 1]$ 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 $n \rightarrow \infty$ to a sequence $\{x_k\}_{k \in \mathbb{N}}$ if $\lim_{n \rightarrow \infty} x_k^n = x_k$ for every $k \in \mathbb{N}$. Now observe that $[-1, 1]^{\mathbb{N}}$ also contains the unit ball in the infinite-dimensional Hilbert space

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

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 sequences

$$(1, 0, 0, \dots), (0, 1, 0, \dots), (0, 0, 1, 0, \dots), \dots,$$

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

$$\{f \in [-1, 1]^{\mathbb{N}} \mid f(k) \in \mathcal{U}_k \text{ for all } k \in \mathbb{N}\}$$

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,

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

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

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

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

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, \dots 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 $f : \mathbb{N} \rightarrow \bigcup_{i=1}^{\infty} X_i$ 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 $k \geq 2$, $\{f_n^k\}_{n=1}^{\infty}$ is a subsequence of $\{f_n^{k-1}\}_{n=1}^{\infty}$ and the sequence $f_n^k(k)$ in X_k converges as $n \rightarrow \infty$. It follows that for every fixed $k \in \mathbb{N}$, the sequence $\{f_n^k(k)\}_{n=1}^{\infty}$ in X_k converges, thus $\{f_n^k\}_{n=1}^{\infty}$ is a convergent subsequence of the original sequence f_n in $\prod_{i=1}^{\infty} X_i$. \square

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 [Kel150], 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 $(\mathcal{P}, <)$ 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 $x < y$ or $y < x$, there exists an element $p \in \mathcal{P}$ such that $p > a$ 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 $q > p$. \square*

PROOF OF THEOREM 6.1, GENERAL CASE. We shall continue to denote elements of $\prod_{\alpha \in I} X_\alpha$ by functions $f : I \rightarrow \bigcup_{\alpha \in I} X_\alpha$ 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 $f_\beta : I \rightarrow \bigcup_{\alpha \in I} X_\alpha$ 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 $(J, g) \leq (J', g')$ 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

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

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

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

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 $J_0 \subset J_{\mathcal{J}}$. 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 $\beta > \beta_0$ for which

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

It follows that $f_{\beta}|_{J_{\mathcal{J}}} \in \prod_{\alpha \in J_{\mathcal{J}}} \mathcal{U}_{\alpha}$ as well, hence $(J_{\mathcal{J}}, g_{\mathcal{J}})$ 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 $\alpha_0 \in I \setminus J_M$, 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 $(J_M \cup \{\alpha_0\}, g'_M) > (J_M, g_M)$, 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 $f : \mathbb{R} \rightarrow [0, 1]$, 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 $x_{(n)} \in \{0, \dots, 9\}$ denote the n th 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 $x \in \mathbb{R} \setminus \{0\}$ 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 $X \setminus \{0_1\}$ 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 \emptyset 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. \square

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 $\mathcal{U}_x \cap \mathcal{U}_y = \emptyset$.

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 $x_n \rightarrow x$ and $x_n \rightarrow y$, 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 $\pi : \mathbb{R} \times \{0, 1\} \rightarrow X : (x, i) \mapsto [(x, i)]$ 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 $X \supset \mathcal{U}_1 \supset \mathcal{U}_2 \supset \dots \ni x$ and $X \supset \mathcal{V}_1 \supset \mathcal{V}_2 \dots \ni y$. 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 $x_n \rightarrow x$ and $x_n \rightarrow y$. \square

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

$$(6.1) \quad T_4 \Rightarrow T_3 \Rightarrow T_2 \Rightarrow T_1 \Rightarrow T_0.$$

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 $X \setminus K$ is open, or equivalently, that every $x \in X \setminus K$ 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 $y_1, \dots, y_N \in K$ such that

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

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

EXERCISE 6.19. Prove:

- A finite topological space satisfies the axiom T_1 if and only if it carries the discrete topology.
- 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 $X \times X$.
- Every compact Hausdorff space is regular, i.e. $\text{compact} + T_2 \Rightarrow T_3$.
Hint: The argument needed for this was already used in the proof of Theorem 6.18.
- 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 $\pi : X \rightarrow X/\sim$ the canonical projection. Given a subset $A \subset X$, we will sometimes also use the notation X/A explained in Exercise 5.16.

- A map $s : X/\sim \rightarrow X$ 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.
- Show that if X is also regular and $A \subset X$ is a closed subset, then X/A is Hausdorff.
- Consider $X = \mathbb{R}$ with the non-closed subset $A = (0, 1]$. Which of the separation axioms T_0, \dots, 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 $X/\sim \rightarrow X$?

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 century 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

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 $\gamma : [0, 1] \rightarrow X$ 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 $f : X \rightarrow \{0, 1\}$ 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 $X \rightarrow \{0, 1\}$ 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 \emptyset 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 $f : X \rightarrow \{0, 1\}$ 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 $X_1 := X \setminus X_0$ 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 $f : X \rightarrow \{0, 1\}$ 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 $f : X \rightarrow Y$ 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 $f : A \cup B \rightarrow \{0, 1\}$ 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 continuous function $A \cup B \rightarrow \{0, 1\}$ 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. \square

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 $X \setminus X_\alpha = \bigcup_{\beta \neq \alpha} X_\beta$ 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

$$(7.1) \quad \coprod_{\alpha \in I} i_\alpha : \coprod_{\alpha \in I} X_\alpha \rightarrow X,$$

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 $r \in \mathbb{R} \setminus \mathbb{Q}$ 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

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

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 disjoint 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 $f : \bar{A} \rightarrow \{0, 1\}$ 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. \square

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.* \square

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 $f : X \rightarrow \{0, 1\}$ such that $f(x) = 0$ and $f(y) = 1$. But if X is path-connected, then there also exists a continuous map $\gamma : [0, 1] \rightarrow X$ with $\gamma(0) = x$ and $\gamma(1) = y$. The composition $g := f \circ \gamma$ is then a continuous function $g : [0, 1] \rightarrow \{0, 1\}$ satisfying $g(0) = 0$ and $g(1) = 1$, and this violates the intermediate value theorem. \square

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 $x \rightarrow 0$. 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 $\gamma : [0, 1] \rightarrow X$ 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 $\gamma : [0, \tau] \rightarrow X$ and assume that it lies on the sine curve for all $0 \leq t < \tau$ but ends on the vertical line at $t = \tau$. Now observe that due to the rapid oscillation as $x \rightarrow 0$, we can find for any $y \in [-1, 1]$ a sequence $t_n \in [0, \tau)$ with $t_n \rightarrow \tau$ such that $\gamma(t_n) \rightarrow (0, y)$. The point y here is arbitrary, yet continuity of γ requires $\gamma(t_n) \rightarrow \gamma(\tau)$, 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 \bar{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 $f : \mathcal{U} \rightarrow \{0, 1\}$ 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, $X \setminus A$ 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.

- 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 $X \times Y \rightarrow \{0, 1\}$.
- 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.
- 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 $f : \mathbb{N} \rightarrow \mathbb{R}$ that satisfy $\lim_{n \rightarrow \infty} f(n) = 0$ 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

$$X_J := \left\{ \{x_\alpha\}_{\alpha \in I} \in \prod_{\alpha \in I} X_\alpha \mid x_\beta = c_\beta \text{ for all } \beta \in I \setminus J \right\},$$

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

- 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).
- 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 neighborhood.

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 neighborhood 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

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

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 $X \setminus \mathcal{U}$, x is not in the closure $\overline{X \setminus \mathcal{U}}$ and thus

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

Since X is Hausdorff, for every $y \in \partial\bar{U}$ 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 $y \in \partial\bar{U}$ form an open cover of the compact set $\partial\bar{U}$, hence there exists a finite subset $\{y_1, \dots, y_N\} \subset \partial\bar{U}$ such that

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

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 $\partial\bar{U}$. The latter implies that for any $y \in \partial\bar{U}$, 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 $\bar{X} \setminus \bar{U}$, so we conclude $\bar{\mathcal{V}} \subset \mathcal{U}$. The compactness of $\bar{\mathcal{V}}$ follows because it is a closed subset of \bar{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 $(X \setminus K) \cup \{\infty\} \subset X^*$ where $K \subset X$ is closed and compact.

- Verify that \mathcal{T}^* is a topology and that X^* is always compact.
- 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.
- 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 .)
- Show that if X is already compact, then X^* is homeomorphic to the disjoint union $X \amalg \{\infty\}$.
- 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.

- 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*

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

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

$$\mathcal{U}_{K, V} := \{f \in C(X, Y) \mid f(K) \subset V\},$$

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 $(f_0, x_0) \in \text{ev}^{-1}(V)$ where $V \subset Y$ is open, then $(f_0, x_0) \in \mathcal{O} \times \mathcal{U} \subset \text{ev}^{-1}(V)$ 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 $C(\mathbb{Q}, \mathbb{R})$, $\text{ev} : C(\mathbb{Q}, \mathbb{R}) \times \mathbb{Q} \rightarrow \mathbb{R}$ is not continuous.

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

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

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

$$(7.2) \quad \Phi(x)(y) = F(x, y).$$

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 $X \rightarrow Y$, with the compact-open topology, which has a subbase consisting of all sets of the form

$$\mathcal{U}_{K, V} := \{f \in C(X, Y) \mid f(K) \subset V\}$$

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 $F : X \times Y \rightarrow Z$ is continuous, then the correspondence (7.2) defines a continuous map $\Phi : X \rightarrow C(Y, Z)$.
- (b) Prove that if Y is locally compact and Hausdorff, then the converse also holds: any continuous map $\Phi : X \rightarrow C(Y, Z)$ defines a continuous map $F : X \times Y \rightarrow Z$ 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 $f_0, f_1 : Y \rightarrow Z$, a continuous map

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

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 $Y \rightarrow Z$ 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 $Y \rightarrow Z$ 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* $\text{ev} : C(X, Y) \times X \rightarrow Y : (f, x) \mapsto f(x)$ 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

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

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

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

are both continuous. Common examples include the standard matrix groups $\text{GL}(n, \mathbb{R})$, $\text{GL}(n, \mathbb{C})$ 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

$$\text{Homeo}(X) = \{f \in C(X, X) \mid f \text{ is bijective and } f^{-1} \in C(X, X)\}$$

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 $\text{Homeo}(X) \subset C(X, X)$, 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.

⁵Since $C(X \times Y, Z)$ 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 $\text{Homeo}(X)$ with the topology \mathcal{T}_H is a topological group.

Hint: Notice that $f(K) \subset V$ if and only if $f^{-1}(X \setminus V) \subset X \setminus K$. 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 $\text{Homeo}(X)$ 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 $f_k \in \text{Homeo}(X)$ defined for $k \in \mathbb{N}$ by $f_k(0) = 0$, $f_k(e^n) = e^{n-1}$ for $n \leq -k$ 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 $C(X, X)$, $f_k \rightarrow \text{Id}$ but $f_k^{-1} \not\rightarrow \text{Id}$ as $k \rightarrow \infty$, hence the map $\text{Homeo}(X) \rightarrow \text{Homeo}(X) : f \mapsto f^{-1}$ is not continuous.

8. Paths, homotopy and the fundamental group


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 $\mathbb{R} \rightarrow \mathbb{R}^3$. 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

⁶The 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 exist a continuous map

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

such that $\varphi(t, \cdot) : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ is a homeomorphism for every $t \in [0, 1]$, $\varphi(0, \cdot)$ is the identity map on \mathbb{R}^3 and $\varphi(1, \cdot) : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ 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 $\mathbb{R}^3 \setminus K$ and $\mathbb{R}^3 \setminus K_0$, and we will see that the latter is impossible. The reason is because we can associate to these spaces groups $\pi_1(\mathbb{R}^3 \setminus K)$ and $\pi_1(\mathbb{R}^3 \setminus K_0)$, which would need to be isomorphic if $\mathbb{R}^3 \setminus K$ and $\mathbb{R}^3 \setminus K_0$ 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 $F \hookrightarrow \mathbb{R}^3 \setminus A$ 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” $\gamma : S^1 \rightarrow \mathbb{R}^3 \setminus A$, meaning γ represents a nontrivial element in the fundamental group $\pi_1(\mathbb{R}^3 \setminus A)$.

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 $\gamma : I \rightarrow X$ satisfying $\gamma(0) = x$ and $\gamma(1) = y$.⁷ We will sometimes use the notation

$$x \rightsquigarrow y$$

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

The **inverse** of a path $x \rightsquigarrow y$ is the path

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

⁷This 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 $x \xrightarrow{\alpha} y$ and $y \xrightarrow{\beta} z$, we can define the product path $x \xrightarrow{\alpha \cdot \beta} z$ by

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

This operation is also called a **concatenation** of paths. The **trivial path** at a point $x \in X$ is defined as the constant path $x \xrightarrow{e_x} x$, 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 $x \xrightarrow{\alpha} y$, $y \xrightarrow{\beta} z$ and $z \xrightarrow{\gamma} a$, 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 $f, g : X \rightarrow Y$ are **homotopic** (*homotop*) if there exists a map

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

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 $H : I \times I \rightarrow X$ 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 $H : I \times X \rightarrow Y$ is a homotopy from $f_0 := H(0, \cdot) : X \rightarrow Y$ to $f_1 := H(1, \cdot) : X \rightarrow Y$, then we obtain a **continuous family of maps** $f_s := H(s, \cdot) : X \rightarrow Y$ 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 $x \underset{\alpha_s}{\rightsquigarrow} y$ and $y \underset{\beta_s}{\rightsquigarrow} z$ 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

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

□

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

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

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

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

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

□

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 $x \underset{\gamma}{\rightsquigarrow} y$, 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:

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

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 :

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

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:

$$\pi_1(X, p) := \left\{ \text{paths } p \rightsquigarrow p \right\} / \sim_{h_+}.$$

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 $f : X \rightarrow Y$ satisfying $f(p) = q$ is called a **pointed map** or **map of pointed spaces**, and can be denoted by

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

We also sometimes refer to such objects as **base-point preserving** maps. Finally, given two pointed maps $f, g : (X, p) \rightarrow (Y, q)$, a homotopy $H : I \times X \rightarrow Y$ 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 $f_s := H(s, \cdot) : (X, p) \rightarrow (Y, q)$ 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 $f : (X, p) \rightarrow (Y, q)$ a group homomorphism*

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

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

$h : I \times I \rightarrow Y : (s, t) \mapsto H(s, \gamma(t))$ defines a homotopy with fixed end points between $f \circ \gamma$ and $g \circ \gamma$. This shows that $f_* : \pi_1(X, p) \rightarrow \pi_1(Y, q)$ 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 $p \xrightarrow{\alpha, \beta} p$, 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 $f : X \rightarrow Y$, 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 $f : (X, p) \rightarrow (Y, q)$ is a pointed map, and since its inverse is continuous, $f^{-1} : (Y, q) \rightarrow (X, p)$ is also a pointed map. Using Theorem 8.11, the commutative diagram (see Remark 8.14 below) of continuous maps

$$(8.2) \quad \begin{array}{ccc} & (Y, q) & \\ f \nearrow & & \searrow f^{-1} \\ (X, p) & \xrightarrow{\text{Id}} & (X, p) \end{array}$$

then gives rise to a similar commutative diagram of group homomorphisms

$$(8.3) \quad \begin{array}{ccc} & \pi_1(Y, q) & \\ f_* \nearrow & & \searrow f_*^{-1} \\ \pi_1(X, p) & \xrightarrow{\mathbb{1}} & \pi_1(X, p) \end{array}$$

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 $f^{-1} \circ f = \text{Id}$, and (8.3) means $f_*^{-1} \circ f_* = \mathbb{1}$. These were especially simple examples, but later we will also encounter larger diagrams like

$$\begin{array}{ccccc} A & \xrightarrow{f} & B & \xrightarrow{g} & C_* \\ \downarrow \alpha & & \downarrow \beta & & \downarrow \gamma \\ A & \xrightarrow{f'} & B' & \xrightarrow{g'} & C' \end{array}$$

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 $\pi_1(\mathbb{R}^2 \setminus \{0\})$ 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

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 $p \xrightarrow{\sim} q$ determines a group isomorphism*

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

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. \square

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

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 $\pi_1(X)$.

⁸Note 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 $I/\partial I \rightarrow S^1$ can be written most easily by thinking of S^1 as the unit circle in \mathbb{C} :

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

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

$$\begin{array}{ccc} X & \xrightarrow{\pi} & X/A \\ & \searrow f & \swarrow g \\ & & Y \end{array}$$

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

PROOF. The diagram determines the correspondence: given $g : X/A \rightarrow Y$, we can define $f := g \circ \pi$ to obtain a map $X \rightarrow Y$ that is automatically constant on A , and conversely, if $f : X \rightarrow Y$ is given and is constant on A , then there is a well-defined map $g : X/A \rightarrow Y : [x] \mapsto f(x)$. 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 continuous. \square

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

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

It is easy to check moreover that two paths $p \xrightarrow{\sim} p$ are homotopic with fixed end points if and only if they correspond to maps $(I/\partial I, [0]) \rightarrow (X, p)$ 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

$$\pi_1(X, p) = \{ \text{pointed maps } \gamma : (S^1, 1) \rightarrow (X, p) \} / \sim_{h_+},$$




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

where \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* $\gamma : S^1 \rightarrow X$.

THEOREM 9.4. A loop $\gamma : (S^1, 1) \rightarrow (X, p)$ represents the identity element in $\pi_1(X, p)$ if and only if there exists a continuous map $u : \mathbb{D}^2 \rightarrow X$ 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 $f : I^2 \rightarrow \mathbb{D}^2 \subset \mathbb{C}$ 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

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

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 $\gamma : I \rightarrow X$ with $\gamma(0) = \gamma(1) = p$ represents the identity in $\pi_1(X, p)$ if and only if there exists a homotopy $H : I^2 \rightarrow X$ with $H(0, \cdot) = \gamma$ and $H|_A \equiv p$. Applying Lemma 9.3 again, such a map is equivalent to a map $h : I^2/A \rightarrow X$ 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 $\mathbb{D}^2 \rightarrow X$ whose restriction to $\partial\mathbb{D}^2$ is the loop $S^1 \cong I/\partial I \rightarrow X$ determined by $\gamma : I \rightarrow X$. \square

REMARK 9.5. Maps $\gamma : S^1 \rightarrow X$ 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

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

By Theorem 9.4, this is equivalent to the condition that every map $\gamma : S^1 \rightarrow X$ admits a continuous extension $u : \mathbb{D}^2 \rightarrow X$ 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 $\gamma : S^1 \rightarrow X$ 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, $\mathbb{R}^2 \setminus \{0\}$ is not simply connected: we will see that the natural inclusion map $\gamma : S^1 \hookrightarrow \mathbb{R}^2 \setminus \{0\}$ is an example of a loop that cannot be extended to a map $u : \mathbb{D}^2 \rightarrow \mathbb{R}^2 \setminus \{0\}$. Of course, it *can* be extended to a map $\mathbb{D}^2 \rightarrow \mathbb{R}^2$, 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 $\mathbb{R}^2 \setminus \{0\}$. 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 $x \xrightarrow{\alpha, \beta} y$ in X are homotopic with fixed end points if and only if the concatenated path $x \xrightarrow{\alpha \cdot \beta^{-1}} x$ represents the identity element in $\pi_1(X, x)$.*

PROOF. The condition $\alpha \underset{h+}{\sim} \beta$ means the existence of a homotopy $H : I^2 \rightarrow X$ 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 $\mathbb{D}^2 \rightarrow X$. The loop $\gamma := H|_{\partial \mathbb{D}^2} : S^1 \rightarrow X$ 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 $n \geq 0$, 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 $0 \rightsquigarrow 0$ is homotopic to the constant loop via the continuous family of loops

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

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

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

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

Now for any $[\gamma] \in \pi_1(S^1, 1)$ represented by a path $1 \rightsquigarrow 1$, there is a unique lift to a path $0 \rightsquigarrow \tilde{\gamma}(1)$ 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

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

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 $\gamma_k(t) = e^{2\pi ikt}$ 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 $f : \mathbb{R} \rightarrow S^1$ 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 $p^X : X \times Y \rightarrow X$ and $p^Y : X \times Y \rightarrow Y$ 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 $(S^1, \text{pt}) \rightarrow (X, p)$, and let $[S^1, X]$ denote the set of homotopy classes of maps $S^1 \rightarrow X$, 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

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

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 $H : [0, 1] \times S^1 \rightarrow X$ 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 $\gamma := H(\cdot, t_0) : [0, 1] \rightarrow X$ 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

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 $f_* : \pi_1(X, p) \rightarrow \pi_1(Y, q)$ injective or surjective.

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

$$(10.1) \quad \begin{array}{ccc} A & \xrightarrow{\text{Id}} & A \\ & \searrow i & \nearrow f \\ & & X \end{array}$$

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 $f : \mathbb{R}^2 \rightarrow A : (x, y) \mapsto (x, 0)$ 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

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

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 $\{f_\alpha : (X_\alpha, p_\alpha) \rightarrow (Y, q)\}_{\alpha \in J}$, the unique map $f : \bigvee_{\alpha \in J} X_\alpha \rightarrow Y$ 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

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

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 $X \vee Y \rightarrow Y$, and for a wedge sum of arbitrarily many spaces, a natural retraction $\bigvee_{\beta \in J} X_\beta \rightarrow X_\alpha$ for each $\alpha \in J$.

EXERCISE 10.7. Convince yourself that the map $f : X \vee Y \rightarrow X$ 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 $S^1 \vee S^1 \rightarrow S^1$ 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 $f : \mathbb{D}^n \rightarrow \mathbb{D}^n$ with no fixed point, then one can use it to define a map $g : \mathbb{D}^n \rightarrow \partial\mathbb{D}^n = S^{n-1}$ 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 $f : X \rightarrow A$ is a retraction and $i : A \hookrightarrow X$ denotes the inclusion, then for any choice of base point $a \in A$, the induced homomorphism $i_* : \pi_1(A, a) \rightarrow \pi_1(X, a)$ is injective, while $f_* : \pi_1(X, a) \rightarrow \pi_1(A, a)$ 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:

$$\begin{array}{ccc} \pi_1(A, a) & \xrightarrow{\mathbb{1}} & \pi_1(A, a) \\ & \searrow i_* & \nearrow f_* \\ & & \pi_1(X, a) \end{array}$$

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

PROOF OF THE BROUWER FIXED POINT THEOREM FOR $n = 2$. If there is a map $f : \mathbb{D}^2 \rightarrow \mathbb{D}^2$ with no fixed point, then there is also a retraction $g : \mathbb{D}^2 \rightarrow \partial\mathbb{D}^2 = S^1$ as explained in Example 10.9, so Theorem 10.10 implies that the induced homomorphism $g_* : \pi_1(\mathbb{D}^2) \rightarrow \pi_1(S^1)$ 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. \square

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 $H : I \times X \rightarrow X$ such that $H(s, \cdot)|_A = \text{Id}_A$ for every $s \in I$, $H(1, \cdot) = \text{Id}_X$ and $H(0, \cdot) = i \circ f$ for some retraction $f : X \rightarrow A$. If a deformation retraction exists, we say that A is 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 $f : S^1 \vee S^1 \rightarrow S^1$ 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 $S^1 \vee S^1 \rightarrow S^1 \vee S^1$ 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 $S^{n-1} \subset \mathbb{R}^n \setminus \{0\}$ is a deformation retract of the punctured Euclidean space. A suitable homotopy $H : I \times (\mathbb{R}^n \setminus \{0\}) \rightarrow \mathbb{R}^n \setminus \{0\}$ 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 $\mathbb{R}^n \setminus \{0\}$ 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

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

If we identify $S^n \setminus \{p_+, p_-\}$ 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 $S^n \setminus \{p_+, p_-\}$. This follows from the fact that $\{0\}$ is a deformation retract of $(-1, 1)$.

DEFINITION 10.16. A map $f : X \rightarrow Y$ is a **homotopy equivalence** (*Homotopieäquivalenz*) if there exists a map $g : Y \rightarrow X$ 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 $f : X \rightarrow Y$ is obviously also a homotopy equivalence, with homotopy inverse f^{-1} .

EXAMPLE 10.20. If $H : I \times X \rightarrow X$ is a deformation retraction with $H(0, \cdot) = f \circ i$ for a retraction $f : X \rightarrow A$, 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 $i \circ f \underset{h}{\sim} \text{Id}_X$.

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 $f : X \rightarrow \{x\}$ is a homotopy equivalence with homotopy inverse $g : \{x\} \rightarrow X$ and a homotopy $H : I \times X \rightarrow X$ 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 $\{x\} \rightarrow \{x\}$.) Then if $p := g(x) \in X$, $F : X \rightarrow \{p\}$ denotes the constant map at p and $i : \{p\} \hookrightarrow X$ is the inclusion, we have $F \circ i = \text{Id}_{\{p\}}$, 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 $f : X \rightarrow Y$ is a homotopy equivalence with $f(p) = q$, then the induced homomorphism $f_* : \pi_1(X, p) \rightarrow \pi_1(Y, q)$ 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 , $\pi_1(X) = 0$.* \square

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 $f_* : \pi_1(X, p) \rightarrow \pi_1(Y, q)$ 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 $g : Y \rightarrow X$ is a homotopy inverse of f , then there is no reason to expect $g(q) = p$, in which case $g_* : \pi_1(Y, q) \rightarrow \pi_1(X, g(q))$ cannot be an inverse of $f_* : \pi_1(X, p) \rightarrow \pi_1(Y, q)$, 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 $f : X \rightarrow Y$ is a map with homotopy inverse $g : Y \rightarrow X$, 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

$$(10.2) \quad \pi_1(X, p) \xrightarrow{f_*} \pi_1(Y, q) \xrightarrow{g_*} \pi_1(X, r).$$

By assumption there exists a homotopy $H : I \times X \rightarrow X$, which we shall write as a 1-parameter family of maps

$$h_s := H(s, \cdot) : X \rightarrow X \quad \text{for } s \in I,$$

satisfying $h_0 = \text{Id}_X$ and $h_1 = g \circ f$. We can therefore define a path $p \rightsquigarrow r$ by

$$\gamma(t) := h_t(p),$$

and by Theorem 9.1, this gives rise to an isomorphism

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

We claim that the diagram

$$\begin{array}{ccc} \pi_1(X, p) & \xrightarrow{f_*} & \pi_1(Y, q) \\ & \searrow \Phi_\gamma^{-1} & \downarrow g_* \\ & & \pi_1(X, r) \end{array}$$

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 $p \rightsquigarrow \gamma(s)$ 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 $g \circ f : X \rightarrow X$ is homotopic to the identity. We have not yet used the assumption that $f \circ g : Y \rightarrow Y$ 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 $\gamma_z : S^1 \rightarrow \mathbb{C} \setminus \{z\} : e^{i\theta} \mapsto z + e^{i\theta}$. Since these are pointed maps $(S^1, 1) \rightarrow (\mathbb{C} \setminus \{z\}, z + 1)$, they represent elements $[\gamma_z] \in \pi_1(\mathbb{C} \setminus \{z\}, z + 1)$. 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 $\mathbb{C} \setminus \{z\}$, by a construction analogous to Example 10.14, hence γ_z is a homotopy equivalence and therefore induces an isomorphism $\pi_1(S^1, 1) \rightarrow \pi_1(\mathbb{C} \setminus \{z\}, z + 1)$. Since the identity map $(S^1, 1) \rightarrow (S^1, 1)$ represents a generator of $\pi_1(S^1, 1)$, composing this with γ_z now represents a generator of $\pi_1(\mathbb{C} \setminus \{z\}, z + 1)$ as claimed.

EXERCISE 10.27. For a point $z \in \mathbb{C}$ and a continuous map $\gamma : [0, 1] \rightarrow \mathbb{C} \setminus \{z\}$ 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 $\theta : [0, 1] \rightarrow \mathbb{R}$ is any choice of continuous function such that

$$\gamma(t) = z + r(t)e^{2\pi i\theta(t)}$$

for some function $r : [0, 1] \rightarrow (0, \infty)$. 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

$$\pi_1(S^1, 1) \rightarrow \mathbb{Z} : [\gamma] \mapsto \text{wind}(\gamma; 0)$$

is an isomorphism to the abelian group $(\mathbb{Z}, +)$. Assume in the following that $\Omega \subset \mathbb{C}$ is an open set and $f : \Omega \rightarrow \mathbb{C}$ is a continuous function.

- (a) Suppose $f(z) = w$ and $w \notin f(\mathcal{U} \setminus \{z\})$ for some neighborhood $\mathcal{U} \subset \Omega$ of z . This implies that the loop $f \circ \gamma_\epsilon$ for $\gamma_\epsilon : [0, 1] \rightarrow \Omega : t \mapsto z + \epsilon e^{2\pi i t}$ has image in $\mathbb{C} \setminus \{w\}$ 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 $0 < \epsilon \leq \epsilon_0$.

- (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 $\text{wind}(f \circ \gamma; w) \neq 0$ 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 $S^1 \rightarrow X$, then such a map represents the identity in $\pi_1(X, p)$ if and only if it admits a continuous extension to a map $\mathbb{D}^2 \rightarrow X$. Define X in the present case to be $\mathbb{C} \setminus \{w\}$.

- (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 $f : \Omega \rightarrow \mathbb{C}$ if $f(z_0) = 0$ but $0 \notin f(\mathcal{U} \setminus \{z_0\})$ for some neighborhood $\mathcal{U} \subset \Omega$ of z_0 . Let us say that such a zero has **order** $k \in \mathbb{Z}$ if $\text{wind}(f \circ \gamma_\epsilon; 0) = k$ 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 $g : \Omega \rightarrow \mathbb{C}$ 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 $\mathbb{C} \setminus \{w\}$ is homotopy equivalent to S^1 for every $w \in \mathbb{C}$. Thus all questions about $\pi_1(\mathbb{C} \setminus \{w\})$ can be reduced to questions about $\pi_1(S^1)$.

11. The easy part of van Kampen's theorem

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) $\bigcap_{\alpha \in J} A_\alpha \neq \emptyset$.

Let $A_\alpha \xrightarrow{i_\alpha} X$ 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

$$A := \{z > -\epsilon\} \subset S^n, \quad B := \{z < \epsilon\} \subset S^n$$

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 $n \geq 2$. (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 $n \geq 2$, S^n is simply connected. \square

Here is an easy application:

THEOREM 11.5. For every $n \geq 3$, \mathbb{R}^2 is not homeomorphic to \mathbb{R}^n .

PROOF. The complement of one point in \mathbb{R}^n is homotopy equivalent to S^{n-1} , thus $\pi_1(\mathbb{R}^n \setminus \{\text{pt}\}) \cong \pi_1(S^{n-1}) = 0$ if $n \geq 3$, while $\pi_1(\mathbb{R}^2 \setminus \{\text{pt}\}) \cong \pi_1(S^1) \cong \mathbb{Z}$. It follows that $\mathbb{R}^2 \setminus \{\text{pt}\}$ and $\mathbb{R}^n \setminus \{\text{pt}\}$ for $n \geq 3$ are not homeomorphic, hence neither are \mathbb{R}^2 and \mathbb{R}^n . \square

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 inclusions⁹

$$i_X : A \hookrightarrow X, \quad i_Y : A \hookrightarrow Y.$$

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

$$\mathbb{D}^n \cup_{S^{n-1}} \mathbb{D}^n \cong S^n.$$

⁹The 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

$$\mathbb{M} := \{(e^{i\theta}, t \cos(\theta/2), t \sin(\theta/2)) \in S^1 \times \mathbb{R}^2 \mid e^{i\theta} \in S^1, t \in [-1, 1]\}$$

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

$$H : I \times \mathbb{M} \rightarrow \mathbb{M} : (s, (e^{i\theta}, t \cos(\theta/2), t \sin(\theta/2))) \mapsto (e^{i\theta}, st \cos(\theta/2), st \sin(\theta/2)),$$

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 $\gamma : (S^1, 1) \rightarrow (B, p)$ 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:

$$(11.1) \quad \begin{aligned} \gamma : S^1 &\xrightarrow{\cong} S^1 \times \{0\} \subset \mathbb{M} \subset \mathbb{RP}^2, \\ e^{i\theta} &\mapsto (e^{i\theta}, 0). \end{aligned}$$

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 $\alpha : S^1 \rightarrow S^2$ parametrizes this loop around the equator, then there is obviously an extension of α to a map $u : \mathbb{D}^2 \rightarrow S^2$ satisfying $u|_{\partial\mathbb{D}^2} = \alpha$, namely the inclusion of either the northern or southern hemisphere of S^2 .

The map $\pi \circ u : \mathbb{D}^2 \rightarrow \mathbb{R}\mathbb{P}^2$ 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 $\pi_1(\mathbb{R}\mathbb{P}^2)$ 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. \square

DEFINITION 11.11. Given a space X , the **suspension** (*Einhangung*) 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 $n \geq 2$ 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.

$$A := (X \times (-\epsilon, 1])/(X \times \{1\}), \quad B := (X \times [-1, \epsilon))/(X \times \{-1\})$$

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. \square

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 $p \xrightarrow{\gamma} p$ 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 $\gamma : I \rightarrow X$ 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

¹⁰Remember 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

$$I = I_1 \cup \dots \cup I_N$$

such that for each $j = 1, \dots, N$, $\gamma(I_j) \subset A_{\alpha_j}$ for some $\alpha_j \in J$. After relabeling the α_j 's if necessary, we can then find a finite increasing sequence

$$0 =: t_0 < t_1 < \dots < t_{N-1} < t_N := 1$$

such that $\gamma([t_{j-1}, t_j]) \subset A_{\alpha_j}$ for each $j = 1, \dots, N$. In particular, for $j = 1, \dots, N-1$, 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

$$\gamma = \gamma_1 \cdot \dots \cdot \gamma_N \underset{h_+}{\sim} \gamma_1 \cdot \beta_1 \cdot \beta_1^{-1} \cdot \gamma_2 \cdot \beta_2 \cdot \beta_2^{-1} \cdot \dots \cdot \beta_{N-2} \cdot \beta_{N-2}^{-1} \cdot \gamma_{N-1} \cdot \beta_{N-1} \cdot \beta_{N-1}^{-1} \cdot \gamma_N.$$

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} . \square

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 $N \geq 0$, an ordered set $b_1 b_2 \dots b_N$ together with a corresponding ordered set $\alpha_1, \alpha_2, \dots, \alpha_N \in J$ is called a **word** in $\{G_\alpha\}_{\alpha \in J}$ if $b_i \in G_{\alpha_i}$ for each $i = 1, \dots, N$. Informally, we call the elements of the sequence *letters*, and denote the word by $b_1 \dots b_N$ even though, strictly speaking, the set of indices $\alpha_1, \dots, \alpha_N \in J$ 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 $a_1 \dots a_N$ 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 $b_1 \dots b_N$ 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*) $\ast_{\alpha \in J} G_\alpha$ 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 $w = b_1 \dots b_N$ and $w' = b'_1 \dots b'_{N'}$ in this group is defined to be the reduction of the concatenated word $ww' = b_1 \dots b_N b'_1 \dots b'_{N'}$. The identity element is the empty word, and will be denoted by

$$e \in \ast_{\alpha \in J} G_\alpha.$$

We will typically deal with collections of only finitely many groups G_1, \dots, G_N , in which case the free product is usually denoted by

$$G_1 \ast \dots \ast G_N.$$

¹¹This 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 $b_1 \dots b_N$ 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

$$\mathbb{Z} * \mathbb{Z}_2 \cong G * H = \{e, a^p, b, a^p b, ba^p, a^p ba^q, ba^p ba^q, a^p ba^q ba^r, \dots \mid p, q, r, \dots \in \mathbb{Z}\}.$$

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*

$$*_{\alpha \in J} \pi_1(A_\alpha, p) \xrightarrow{\Phi} \pi_1(X, p)$$

*sending each reduced word $[\gamma_1] \dots [\gamma_N] \in *_{\alpha \in J} \pi_1(A_\alpha, p)$ 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. \square*

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

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 $*_{\alpha \in J} G_\alpha$ 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 $*_{\alpha \in J} G_\alpha$ are formed by concatenating words and then

¹²This 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 $*_{\alpha \in J} G_\alpha$ 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 $*_{\beta \in J} G_\beta$ 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 $*_{\gamma \in J} G_\gamma$ as described above, then for every $\alpha, \beta \in J$ with $\alpha \neq \beta$, the intersection $G_\alpha \cap G_\beta$ in $*_{\gamma \in J} G_\gamma$ 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 $*_{\gamma \in J} G_\gamma$.*
- (3) *For any group H with a collection of homomorphisms $\{\Phi_\alpha : G_\alpha \rightarrow H\}_{\alpha \in J}$, there exists a unique homomorphism*

$$\Phi : *_{\alpha \in J} G_\alpha \rightarrow H$$

*whose restriction to each of the subgroups $G_\alpha \subset *_{\beta \in J} G_\beta$ is Φ_α .*

The third item in this list deserves brief comment: the homomorphism $\Phi : *_{\alpha \in J} G_\alpha \rightarrow H$ exists and is unique because every element of $*_{\alpha \in J} G_\alpha$ 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 straightforward 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

$$(12.1) \quad *_{\alpha \in J} \pi_1(A_\alpha, p) \xrightarrow{\Phi} \pi_1(X, p)$$

is determined as in the proposition above by the homomorphisms $(i_\alpha)_* : \pi_1(A_\alpha, p) \rightarrow \pi_1(X, p)$ 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 $\Psi : G \rightarrow H$, $\ker \Psi$ 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 $\Psi : G \rightarrow H$ 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 well-defined map

$$G/\ker \Psi \rightarrow H : gK \mapsto \Psi(g),$$

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 $G/K \rightarrow H$, 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 $G/\ker \Psi \rightarrow H$ 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,

$$\begin{array}{ccc} & A_\alpha & \\ j_{\alpha\beta} \nearrow & & \searrow i_\alpha \\ A_\alpha \cap A_\beta & & X \\ j_{\beta\alpha} \searrow & & \nearrow i_\beta \\ & A_\beta & \end{array}$$

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 $p \xrightarrow{\sim} p$ 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*

$$\Phi : \ast_{\alpha \in J} \pi_1(A_\alpha, p) \rightarrow \pi_1(X, p)$$

induced by the homomorphisms $(i_\alpha)_ : \pi_1(A_\alpha, p) \rightarrow \pi_1(X, p)$ 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*

$$\ker \Phi = \left\langle \left\{ (j_{\alpha\beta})_* [\gamma] (j_{\beta\alpha})_* [\gamma]^{-1} \mid \alpha, \beta \in J, [\gamma] \in \pi_1(A_\alpha \cap A_\beta, p) \right\} \right\rangle_N.$$

In particular, Φ then descends to an isomorphism

$$\ast_{\alpha \in J} \pi_1(A_\alpha, p) / \ker \Phi \xrightarrow{\cong} \pi_1(X, p).$$

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

$$\pi_1(A, p) * \pi_1(B, p) \rightarrow \pi_1(S^1 \vee S^1, p)$$

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

$$\pi_1(S^1 \vee S^1, p) \cong \mathbb{Z} * \mathbb{Z} = \{e, \alpha^p, \beta^q, \alpha^p \beta^q, \beta^p \alpha^q, \alpha^p \beta^q \alpha^r, \dots \mid p, q, r, \dots \in \mathbb{Z}\}.$$

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 $\mathbb{R}^n \xrightarrow{i} S^n$ with image $S^n \setminus \{\infty\}$. We claim that for any compact subset $K \subset \mathbb{R}^3$ such that $\mathbb{R}^3 \setminus K$ is path-connected, and any choice of base point $p \in \mathbb{R}^3 \setminus K$,

$$i_* : \pi_1(\mathbb{R}^3 \setminus K, p) \rightarrow \pi_1(S^3 \setminus K, p)$$

is an isomorphism. To see this, define the open subset $A := \mathbb{R}^3 \setminus K \subset S^3 \setminus K$, and choose $B_0 \subset S^3 \setminus K$ to be an open ball about ∞ , i.e. a set of the form $(\mathbb{R}^3 \setminus \bar{B}_R(0)) \cup \{\infty\}$ where $\bar{B}_R(0) \subset \mathbb{R}^3$ is any closed ball large enough to contain K . Since p might not be contained in B_0 but $\mathbb{R}^3 \setminus K$ is path-connected, we can then define a larger set B by adjoining to B_0 the neighborhood in $\mathbb{R}^3 \setminus K$ 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 $B \setminus \{\infty\}$ and is thus homoeomorphic to $\mathbb{R}^3 \setminus \{0\}$ and homotopy equivalent to S^2 , implying $\pi_1(A \cap B) = 0$. The Seifert-van Kampen theorem therefore gives an isomorphism $\pi_1(\mathbb{R}^3 \setminus K, p) * \pi_1(B, p) \rightarrow \pi_1(S^3 \setminus K, p)$, but $\pi_1(B, p)$ is the trivial group, so this proves the claim.

A frequently occurring 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 $\pi_1(\mathbb{R}^3 \setminus K)$ 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 $\pi_1(S^3 \setminus K)$. 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

$$F_S := \ast_{\alpha \in S} \mathbb{Z},$$

or in other words, the set of all reduced words $a_1^{p_1} a_2^{p_2} \dots a_N^{p_N}$ for $N \geq 0$, $p_i \in \mathbb{Z}$ 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 $\phi : S \rightarrow G$, there is a unique group homomorphism $\Phi : F_S \rightarrow G$ 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 $a_1^{p_1} a_2^{p_2} \dots a_N^{p_N}$, there is clearly only one formula for $\Phi : F_S \rightarrow G$ that will match ϕ on single-letter words and also be a homomorphism, namely

$$\Phi(a_1^{p_1} \dots a_N^{p_N}) = \phi(a_1)^{p_1} \dots \phi(a_N)^{p_N}.$$

It is straightforward to check that this defines a homomorphism. □

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 $\Phi : F_S \rightarrow G$ sending each $g \in S \subset F_S$ to $g \in G$ is surjective, thus Φ descends to an isomorphism $F_S / \ker \Phi \rightarrow G$. □

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

$$\{S \mid R\} := F_S / \langle R' \rangle_N$$

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 necessary 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 $\Phi : F_S \rightarrow G$, 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 $F_S \rightarrow G$ 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 $\{a\} := \{a \mid \emptyset\}$ 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 $\Phi : F_{\{a, b\}} \rightarrow \mathbb{Z}^2$ with $\Phi(a) = (1, 0)$ and $\Phi(b) = (0, 1)$, and Φ is clearly surjective since it necessarily 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 $\{a, b \mid ab = ba\} \rightarrow \mathbb{Z}^2$. 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 $F_{\{a\}} \rightarrow \mathbb{Z}_p$ 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), $\pi_1(S^3 \setminus K) \cong \{a, b \mid a^2 = b^3\}$, 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 $\pi_1(S^3 \setminus K_0) \cong \mathbb{Z}$, which is abelian. This implies via Example 12.7 that $\pi_1(\mathbb{R}^3 \setminus K) \not\cong \pi_1(\mathbb{R}^3 \setminus K_0)$, so $\mathbb{R}^3 \setminus K$ and $\mathbb{R}^3 \setminus K_0$ 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*

$$\pi_1(A) \cong \{\{a_i\} \mid \{R_j\}\}, \quad \pi_1(B) \cong \{\{b_k\} \mid \{S_\ell\}\}, \quad \pi_1(A \cap B) \cong \{\{c_p\} \mid \{T_q\}\},$$

with the indices i, j, k, ℓ, p, q each ranging over finite sets. Then

$$\pi_1(X) \cong \{\{a_i\} \cup \{b_k\} \mid \{R_j\} \cup \{S_\ell\} \cup \{(j_A)_*c_p = (j_B)_*c_p\}\}.$$

□

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 $H \rightarrow H'$, which is therefore an isomorphism since $|H| \leq 6$.

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

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*

$$\Phi : \ast_{\alpha \in J} \pi_1(A_\alpha, p) \rightarrow \pi_1(X, p)$$

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

$$S := \left\{ (j_{\alpha\beta})_* [\gamma] (j_{\beta\alpha})_* [\gamma]^{-1} \mid \alpha, \beta \in J, [\gamma] \in \pi_1(A_\alpha \cap A_\beta, p) \right\}.$$

In particular, if we abbreviate $F := \ast_{\alpha \in J} \pi_1(A_\alpha, p)$, then Φ descends to an isomorphism

$$F / \langle S \rangle_N \rightarrow \pi_1(X, p).$$

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 $F / \langle S \rangle_N \rightarrow \pi_1(X, p)$. 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.

¹³Note 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 $p \xrightarrow{\sim} p$ 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 $p \xrightarrow{\sim} p$ is a loop in A_{α_i} for each $i = 1, \dots, N$, and

$$\gamma \underset{h+}{\sim} \gamma_1 \cdot \dots \cdot \gamma_N.$$

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 $p \xrightarrow{\sim} p$. (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 $\gamma_1 \cdot \dots \cdot \gamma_N \underset{h+}{\sim} \gamma'_1 \cdot \dots \cdot \gamma'_{N'}$, so after choosing suitable parametrizations of both of these concatenations on the unit interval I ,¹⁴ there exists a homotopy

$$H : I^2 \rightarrow X$$

with $H(0, \cdot) = \gamma_1 \cdot \dots \cdot \gamma_N$, $H(1, \cdot) = \gamma'_1 \cdot \dots \cdot \gamma'_{N'}$ 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

¹⁴Recall that concatenation of paths is associative up to homotopy, so the N -fold concatenation $\gamma_1 \cdot \dots \cdot \gamma_N$ is not a uniquely determined path $I \rightarrow X$ if $N > 2$, but it is unique up to homotopy with fixed end points.

¹⁵I 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 $H : I^2 \rightarrow X$ between two factorizations $\gamma_1 \cdot \dots \cdot \gamma_N$ and $\gamma'_1 \cdot \dots \cdot \gamma'_{N'}$, of a loop $p \xrightarrow{\sim} p$ 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.

$\gamma_1 \cdot \dots \cdot \gamma_N$ and $\gamma'_1 \cdot \dots \cdot \gamma'_{N'}$. 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 $H(0, \cdot) = \gamma_1 \cdot \dots \cdot \gamma_N$ contains $(0, t)$. For vertices $(1, t)$ with $t \notin \{0, 1\}$, choose $A_\gamma := A_{\alpha'_i}$ similarly in terms of the concatenation $\gamma'_1 \cdot \dots \cdot \gamma'_{N'}$. 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 path¹⁶

$$H(s, t) \xrightarrow{\delta_{(s,t)}} p \quad \text{in} \quad A_\alpha \cap A_\beta \cap A_\gamma.$$

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 $H|_E : E \rightarrow X$ as a path, and thus associate to E a loop

$$p \xrightarrow{\gamma_E} p \quad \text{in} \quad A_\alpha \cap A_\beta, \quad \gamma_E := \delta_{(s_0, t_0)}^{-1} \cdot H|_E \cdot \delta_{(s_1, t_1)},$$

where $\alpha, \beta \in J$ are the two (not necessarily distinct) elements associated to the boxes bordered by E .

¹⁶This 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, \dots, 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 $\gamma_1 \cdot \dots \cdot \gamma_N$, 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 $(0, t_1), \dots, (0, t_{m-1})$ the particular grid vertices in the domain of γ_i , thus splitting up γ_i into a concatenation of paths $\gamma_i = \gamma_i^1 \cdot \dots \cdot \gamma_i^m$ which have these vertices as starting and/or end points. Then

$$\gamma_i \underset{h+}{\sim} (\gamma_i^1 \cdot \delta_{(0, t_1)}) \cdot (\delta_{(0, t_1)}^{-1} \cdot \gamma_i^2 \cdot \delta_{(0, t_2)}) \cdot \dots \cdot (\delta_{(0, t_{m-1})}^{-1} \cdot \gamma_i^m) \quad \text{in } A_{\alpha_i}.$$

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 $\gamma_1 \cdot \dots \cdot \gamma_N$, 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.

□




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

$$(13.1) \quad \pi_1(X \vee Y) \cong \pi_1(X) * \pi_1(Y)$$

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 $X \setminus \{x\}$ and $Y \setminus \{y\}$ are both open subsets. Consider loops $\gamma : [0, 1] \rightarrow X \vee Y$ 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


$$\pi_1 \left(\bigvee_{\alpha \in J} X_\alpha \right) \cong \bigast_{\alpha \in J} \pi_1(X_\alpha)$$

holds for any (possibly infinite) collection of nice pointed spaces $\{(X_\alpha, x_\alpha)\}_{\alpha \in J}$.

¹⁷Not 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 $S_n^1 \subset \mathbb{R}^2$ 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*

$$H := \bigcup_{n \in \mathbb{N}} S_n^1 \subset \mathbb{R}^2.$$



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

$$f : X \rightarrow H$$

that is a bijection sending the natural base point of $\bigvee_n S^1$ to $(0, 0) \in \bigcap_n S_n^1$, but that X (unlike H) is a “nice” space, hence $f : X \rightarrow H$ 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 $S^1 \rightarrow H$, but continuous maps $S^1 \rightarrow X$ 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

$$r_J : H \rightarrow \bigcup_{n \in J} S_n^1 \subset H,$$

and deduce from this that the map $f_* : \pi_1(X) \rightarrow \pi_1(H)$ is injective.

Hint: Unlike H , $\bigcup_{n \in J} S_n^1$ really is homeomorphic to a wedge sum of circles, the crucial detail in this case being that there are only finitely many.

- (g) Writing $r_n := r_{\{n\}} : H \rightarrow S_n^1$ for each individual value of $n \in \mathbb{N}$, show that the homomorphism

$$\pi_1(H) \rightarrow \prod_{n \in \mathbb{N}} \pi_1(S_n^1) \cong \prod_{n \in \mathbb{N}} \mathbb{Z}$$

determined by the maps $(r_n)_* : \pi_1(H) \rightarrow \pi_1(S_n^1)$ is surjective, and deduce from this that $f_* : \pi_1(X) \rightarrow \pi_1(H)$ 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

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, \dots, 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 $i = 1, \dots, N$, 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 $a_1^{p_1} \dots a_N^{p_N} = e$:*

$$\pi_1(X) \cong \{G \mid a_1^{p_1} \dots a_N^{p_N} = e\}.$$

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)),

$$\pi_1(P^1) \cong \pi_1(S^1) * \dots * \pi_1(S^1) \cong \mathbb{Z} * \dots * \mathbb{Z} = F_G,$$

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 identifies $\pi_1(X, p)$ 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 $p \xrightarrow{\gamma} p$ 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 $a_1^{p_1} \dots a_N^{p_N}$ defined by composing a traversal of ∂P with the quotient projection $\partial P \rightarrow P^1$, thus producing the relation $a_1^{p_1} \dots a_N^{p_N} = e$. \square

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 $\gamma : S^1 \rightarrow \mathbb{RP}^2$, 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 $\mathbb{D}^2 \rightarrow S^2$ since $\pi_1(S^2) = 0$. Composing the map $\mathbb{D}^2 \rightarrow S^2$ with the projection $S^2 \rightarrow \mathbb{RP}^2$ 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 $g \geq 0$, 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

$$a_1, b_1, a_1, b_1, a_2, b_2, a_2, b_2, \dots, a_g, b_g, a_g, b_g,$$

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 $\Sigma_0 = S^2$, $\Sigma_1 = \mathbb{T}^2$, Σ_2 , Σ_3, \dots 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

$$\Sigma \# \Sigma' := \left(\Sigma \setminus \mathring{\mathbb{D}}^2 \right) \cup_{S^1} \left(\Sigma' \setminus \mathring{\mathbb{D}}^2 \right).$$

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 $M \setminus \mathring{\mathbb{D}}^n$ and $M' \setminus \mathring{\mathbb{D}}^n$ 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

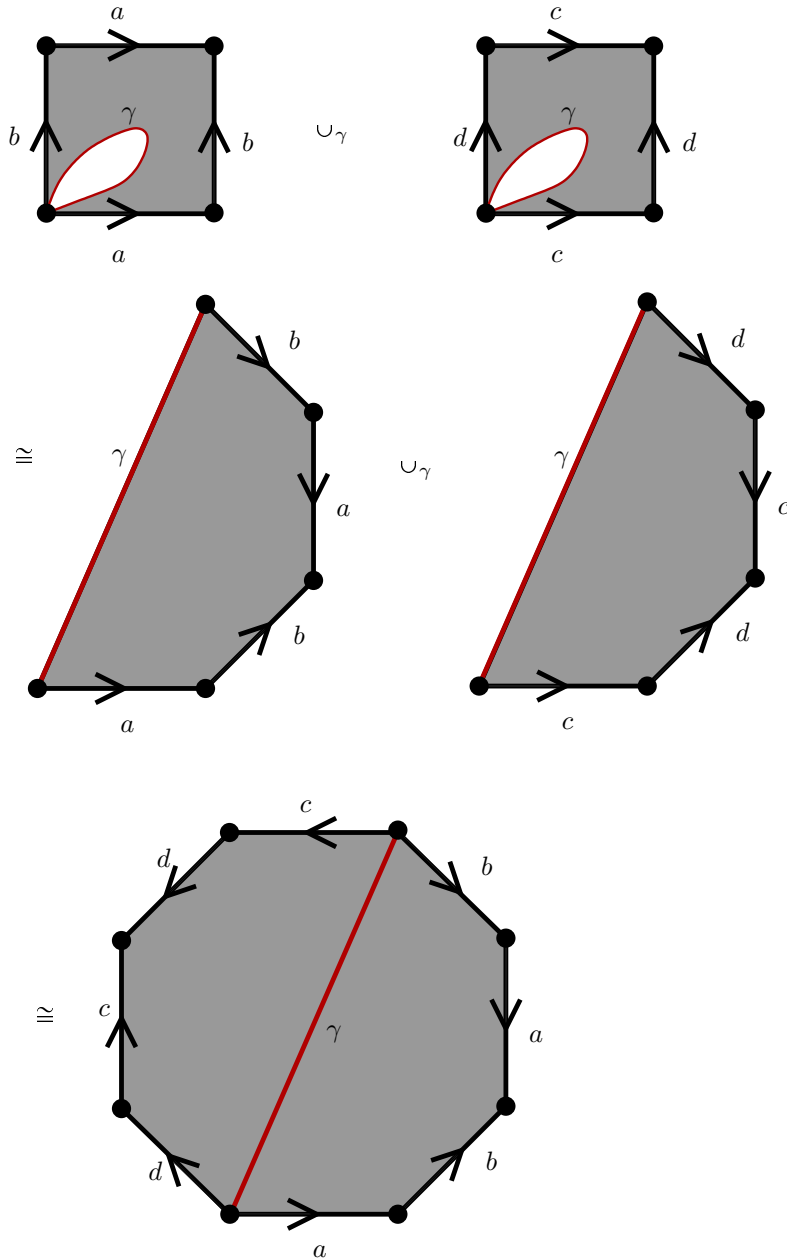


FIGURE 5. The connected sum $\mathbb{T}^2 \# \mathbb{T}^2$ is formed by cutting holes \mathbb{D}^2 out of two copies of \mathbb{T}^2 along some loop γ , and then gluing together the two copies of $\mathbb{T}^2 \setminus \mathbb{D}^2$. The result is Σ_2 , the closed orientable surface of genus 2.

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}_r^2 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 $\mathbb{R}^2 \rightarrow \mathbb{R}^2$, 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

$$\Sigma \#_t \Sigma' := \left(\Sigma \setminus i_t(\mathring{\mathbb{D}}^2) \right) \cup_{S^1} \left(\Sigma' \setminus j_t(\mathring{\mathbb{D}}^2) \right)$$

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 $\Sigma \#_t \Sigma' \rightarrow \Sigma \#_s \Sigma'$ for $t \neq s$ is easy to define if we can first find a homeomorphism $\Sigma \rightarrow \Sigma$ 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:*

$$\Sigma_g \cong \underbrace{\mathbb{T}^2 \# \dots \# \mathbb{T}^2}_g$$

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 $g \geq 0$ has a fundamental group with $2g$ generators and one relation, namely*

$$\pi_1(\Sigma_g) \cong \{a_1, b_1, \dots, a_g, b_g \mid a_1 b_1 a_1^{-1} b_1^{-1} a_2 b_2 a_2^{-1} b_2^{-1} \dots a_g b_g a_g^{-1} b_g^{-1} = e\}.$$

\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 .


- Use the Seifert-Van Kampen theorem to show that if $n \geq 3$, then $\pi_1(X \# Y) \cong \pi_1(X) * \pi_1(Y)$. Where does your proof fail in the cases $n = 1$ and $n = 2$?
- Show that the formula of part (a) is false in general for $n = 1, 2$.

EXERCISE 14.13. For integers $g, m \geq 0$, 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.


- Show that $\pi_1(\Sigma_{g,1})$ is a free group with $2g$ generators, and if $g \geq 1$, 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.

¹⁸Terminology: 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 $h, k \geq 0$. (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 $g \geq 2$ and G denotes the group $\{a_1, b_1, \dots, a_g, b_g \mid \prod_{i=1}^g [a_i, b_i] = e\}$, then for any proper subset $J \subset \{1, \dots, g\}$, $\prod_{i \in J} [a_i, b_i]$ is a nontrivial element of G .

Hint: Given $j \in J$ and $\ell \in \{1, \dots, g\} \setminus J$, there is a homomorphism $\Phi : F_{\{a_1, b_1, \dots, a_g, b_g\}} \rightarrow F_{\{x, y\}}$ 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 $m \geq 1$, $\pi_1(\Sigma_{g,m})$ 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

$$K_{p,q} := \{f(e^{pi\theta}, e^{qi\theta}) \mid \theta \in \mathbb{R}\} \subset \mathbb{R}^3.$$

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).

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*, $\pi_1(\mathbb{R}^3 \setminus K)$. 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 $\pi_1(\mathbb{R}^3 \setminus K)$ to $\pi_1(S^3 \setminus K)$, 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 $\pi_1(S^3 \setminus K)$. We shall now answer the question: given relatively prime integers p and q , what is $\pi_1(S^3 \setminus K_{p,q})$?

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 $S^3 \setminus K_{p,q}$ into two open subsets. They will be defined as open neighborhoods of the two subsets

$$A_0 := (S^1 \times \mathring{\mathbb{D}}^2) \setminus K_{p,q}, \quad B_0 := (\mathring{\mathbb{D}}^2 \times S^1) \setminus K_{p,q}.$$

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

$$A := \left(S^1 \times \mathring{\mathbb{D}}^2 \right) \cup \left((-1, 1) \times (\mathbb{T}^2 \setminus f^{-1}(K_{p,q})) \right),$$

and

$$B := \left(\mathring{\mathbb{D}}^2 \times S^1 \right) \cup \left((-1, 1) \times (\mathbb{T}^2 \setminus f^{-1}(K_{p,q})) \right).$$




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 contracting 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,

$$\pi_1(A) \cong \{a \mid \emptyset\}, \quad \pi_1(B) \cong \{b \mid \emptyset\}.$$

The intersection is

$$A \cap B = (-1, 1) \times (\mathbb{T}^2 \setminus f^{-1}(K_{p,q})) \underset{h.e.}{\simeq} \mathbb{T}^2 \setminus f^{-1}(K_{p,q}) \underset{h.e.}{\simeq} S^1.$$

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 ,

$$\pi_1(A \cap B) \cong \{c \mid \emptyset\}.$$

According to the Seifert-van Kampen theorem (in particular Corollary 12.19, the version for finitely-presented groups), we can now write

$$\pi_1(S^3 \setminus K_{p,q}) \cong \{a, b \mid (j_A)_*c = (j_B)_*c\},$$

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. $\pi_1(S^3 \setminus K_{p,q}) \cong \{a, b \mid a^p = b^q\}$. □

EXAMPLE 14.17. For $(p, q) = (1, 0)$, we obtain the knot group of the unknot: $\pi_1(S^3 \setminus K_{1,0}) \cong \{a, b \mid a = e\} = \{b \mid \emptyset\} = \mathbb{Z}$. In particular, this is an abelian group.

EXAMPLE 14.18. The knot group of the trefoil is $\pi_1(S^3 \setminus K_{2,3}) \cong \{a, b \mid a^2 = b^3\}$. We proved in Exercise 12.20 that this group is not abelian, in contrast to Example 14.17, hence $\pi_1(S^3 \setminus K_{2,3})$ and $\pi_1(S^3 \setminus K_{1,0})$ are not isomorphic.

COROLLARY 14.19. *The knot complements $\mathbb{R}^3 \setminus K_{1,0}$ and $\mathbb{R}^3 \setminus K_{2,3}$ are not homeomorphic.* □

Before moving on¹⁹ 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

$$\pi_1(X_0, p) \cong \{\{a_i\} \mid \{R_j\}\}.$$

Then for any loop $\gamma : (S^1, 1) \rightarrow (X_0, p)$, we claim that the space

$$X := \mathbb{D}^2 \cup_\gamma X_0 := (\mathbb{D}^2 \amalg X_0) / z \sim \gamma(z) \in X_0 \text{ for all } z \in \partial\mathbb{D}^2$$

is compact and Hausdorff with

$$\pi_1(X, p) \cong \{\{a_i\} \mid \{R_j\}, [\gamma] = e\},$$

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 easily²⁰ from the Seifert-van Kampen

¹⁹We 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.

²⁰I 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 $A = \mathbb{D}^2$ 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. \square

15. Covering spaces and the lifting theorem

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 $f : Y \rightarrow X$ 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 $f|_{\mathcal{V}_\alpha} : \mathcal{V}_\alpha \rightarrow \mathcal{U}$ 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 $f : \mathbb{R} \rightarrow S^1 : \theta \mapsto e^{i\theta}$ is a covering map of S^1 .

EXAMPLE 15.3. The map $S^1 \rightarrow S^1$ 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 $\mathbb{T}^n := \underbrace{S^1 \times \dots \times S^1}_n$ admits a covering map

$$\mathbb{R}^n \rightarrow \mathbb{T}^n : (\theta_1, \dots, \theta_n) \mapsto (e^{i\theta_1}, \dots, e^{i\theta_n}).$$

More generally, it is straightforward to show that given any two covering maps $f_i : Y_i \rightarrow X_i$ 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 $X \rightarrow X$ 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 $f : \coprod_{\alpha \in J} X_\alpha \rightarrow X$ 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 $S^n \rightarrow \mathbb{R}P^n = S^n / \sim$ is a covering map.

THEOREM 15.8. *If X is connected and $f : Y \rightarrow X$ 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 $n \in \{0, 1, 2, 3, \dots, \infty\}$, 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 $S^1 \rightarrow S^1 : z \mapsto z^k$ from Example 15.3 has degree $|k|$, while the quotient projection $S^n \rightarrow \mathbb{R}P^n$ has degree 2 and the cover $\mathbb{R} \rightarrow S^1$ from Example 15.2 has infinite degree.

REMARK 15.10. Some authors strengthen the definition of a covering map $f : Y \rightarrow X$ 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.

$$f^{-1}(\mathcal{U}) \cong \bigsqcup_{\alpha \in J} \mathcal{U}.$$

This is true specifically because each of the sets \mathcal{V}_α 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 $f : Y \rightarrow X$ 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 $f : A \rightarrow X$ and a covering map $p : Y \rightarrow X$, can f be “lifted” to a map $\tilde{f} : A \rightarrow Y$ satisfying $p \circ \tilde{f} = f$? This problem can be summarized with the diagram

$$(15.1) \quad \begin{array}{ccc} & & Y \\ & \nearrow \tilde{f} & \downarrow p \\ A & \xrightarrow{f} & X \end{array}$$

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 $p : \mathbb{R} \rightarrow S^1 : \theta \mapsto e^{i\theta}$ and let $f : S^1 \rightarrow S^1$ be the identity map. A lift $\tilde{f} : S^1 \rightarrow \mathbb{R}$ 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 $f : (A, a) \rightarrow (X, x)$ into a pointed map. Then if a lift $\tilde{f} : A \rightarrow Y$ 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

$$(15.2) \quad \begin{array}{ccc} & & \pi_1(Y, y) \\ & \nearrow \tilde{f}_* & \downarrow p_* \\ \pi_1(A, a) & \xrightarrow{f_*} & \pi_1(X, x). \end{array}$$

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 $\text{im } f_* \subset \text{im } p_*$, 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, $p : Y \rightarrow X$ is a covering map and $f : (A, a_0) \rightarrow (X, x_0)$ 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 $\tilde{f} : (A, a_0) \rightarrow (Y, y_0)$ if and only if*

$$f_*(\pi_1(A, a_0)) \subset p_*(\pi_1(Y, y_0)),$$

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 $p : Y \rightarrow X$ between path-connected spaces and any space A that is simply connected and locally path-connected, every map $f : A \rightarrow X$ can be lifted to Y .* \square

COROLLARY 15.13. *For every base-point preserving covering map $p : (Y, y_0) \rightarrow (X, x_0)$ between path-connected spaces, the homomorphism $p_* : \pi_1(Y, y_0) \rightarrow \pi_1(X, x_0)$ is injective.*

PROOF. Suppose $\tilde{\gamma} : (S^1, 1) \rightarrow (Y, y_0)$ is a loop such that $p_*[\tilde{\gamma}] = e \in \pi_1(X, x_0)$. Then $\gamma := p \circ \tilde{\gamma} : (S^1, 1) \rightarrow (X, x_0)$ admits an extension $u : (\mathbb{D}^2, 1) \rightarrow (X, x_0)$ with $u|_{\partial\mathbb{D}^2} = \gamma$. But \mathbb{D}^2 is simply connected, so u admits a lift $\tilde{u} : (\mathbb{D}^2, 1) \rightarrow (Y, y_0)$ satisfying $p \circ \tilde{u} = u$, thus $p \circ \tilde{u}|_{\partial\mathbb{D}^2} = \gamma$ implies that $\tilde{u}|_{\partial\mathbb{D}^2} : (S^1, 1) \rightarrow (Y, y_0)$ is a lift of γ . Uniqueness of lifts then implies $\tilde{u}|_{\partial\mathbb{D}^2} = \tilde{\gamma}$ and thus $[\tilde{\gamma}] = e \in \pi_1(Y, y_0)$. \square

COROLLARY 15.14. *If X is simply connected, then every path-connected covering space of X is also simply connected.* \square

EXAMPLE 15.15. Corollary 15.14 implies that there does not exist any covering map $S^1 \rightarrow \mathbb{R}$.

Here is an application important in complex analysis. Observe that

$$p : \mathbb{C} \rightarrow \mathbb{C}^* := \mathbb{C} \setminus \{0\} : z \mapsto e^z$$

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 $\{\text{Im } z = \text{const}\}$ to a ray in \mathbb{C}^* emanating from the origin, and every vertical line $\{\text{Re } z = \text{const}\}$ 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

$$\log : \mathcal{U} \rightarrow \mathbb{C} \quad \text{such that} \quad e^{\log z} = z \text{ for all } z \in \mathcal{U}.$$

For simplicity, we will restrict our attention to path-connected²¹ 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 $\log : (\mathcal{U}, 1) \rightarrow (\mathbb{C}, 0)$ will then be the unique solution to the lifting problem

$$\begin{array}{ccc} & & (\mathbb{C}, 0) \\ & \nearrow \log & \downarrow p \\ (\mathcal{U}, 1) & \xrightarrow{f} & (\mathbb{C}^*, 1) \end{array}$$

Theorem 15.11 now gives the answer: $\log : \mathcal{U} \rightarrow \mathbb{C}$ 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 $\mathbb{D}^2 \rightarrow \mathbb{C}^*$. Using the notion of the *winding number* from Exercise 10.27, this is the same as saying every loop $\gamma : S^1 \rightarrow \mathcal{U}$ satisfies $\text{wind}(\gamma; 0) = 0$. For example, $\log : \mathcal{U} \rightarrow \mathbb{C}$ 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 $\log : \mathcal{U} \rightarrow \mathbb{C}$ 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 $\gamma : (I, 0) \rightarrow (X, x_0)$ has a unique lift $\tilde{\gamma} : (I, 0) \rightarrow (Y, y_0)$.*

PROOF. Since I is compact, we can find a finite partition $0 =: t_0 < t_1 < \dots < t_{N-1} < t_N := 1$ such that for each $j = 1, \dots, N$, 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 $\tilde{\gamma}_j : [t_{j-1}, t_j] \rightarrow Y$ 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. \square

LEMMA 15.17 (the homotopy lifting property). *Suppose $H : I \times A \rightarrow X$ is a homotopy with $H(0, \cdot) = f : A \rightarrow X$, and $\tilde{f} : A \rightarrow Y$ is a lift of f . Then there exists a unique lift $\tilde{H} : I \times A \rightarrow Y$ of H satisfying $\tilde{H}(0, \cdot) = \tilde{f}$.*

PROOF. The previous lemma implies that each of the paths $s \mapsto H(s, a) \in X$ for $a \in A$ have unique lifts $s \mapsto \tilde{H}(s, a) \in Y$ with $\tilde{H}(0, a) = \tilde{f}(a)$. One should then check that the map $\tilde{H} : I \times A \rightarrow Y$ defined in this way is continuous; I leave this as an exercise. \square

PROOF OF THEOREM 15.11. We shall first define an appropriate map $\tilde{f} : A \rightarrow Y$ 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.

²¹Since $\mathcal{U} \subset \mathbb{C}^*$ 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 $a_0 \xrightarrow{\alpha} a$, giving a path $x_0 \xrightarrow{f \circ \alpha} f(a)$, 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 $\tilde{f} : A \rightarrow Y$. Indeed, suppose $a_0 \xrightarrow{\beta} a$ 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 $\text{im } f_* \subset \text{im } p_*$ then implies the existence of a loop $y_0 \xrightarrow{\tilde{\gamma}} y_0$ in Y such that

$$[(f \circ \alpha) \cdot (f \circ \beta^{-1})] = p_*[\tilde{\gamma}] = [p \circ \tilde{\gamma}],$$

so there is a homotopy $H : I^2 \rightarrow X$ 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 $\gamma : (I, 0) \rightarrow (X, x_0)$. Now Lemma 15.17 provides a lift $\tilde{H} : I^2 \rightarrow Y$ of H with $\tilde{H}(0, \cdot) = \tilde{\gamma}$. In this homotopy, the paths $s \mapsto \tilde{H}(s, 0)$ and $s \mapsto \tilde{H}(s, 1)$ 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 $\tilde{H}(s, 0) = \tilde{H}(s, 1) = y_0$ 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 $\tilde{H}(1, \cdot)$. 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 $\tilde{f} : A \rightarrow Y$ 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 $p|_{\mathcal{V}} : \mathcal{V} \rightarrow \mathcal{U}$ 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 $a_0 \xrightarrow{\gamma} a$ in A and, for any $a' \in \mathcal{O}$, choose a path $a \xrightarrow{\beta} a'$ in \mathcal{O} . Now $\gamma \cdot \beta$ is a path from a_0 to a' , so

$$\tilde{f}(a) = \widetilde{f \circ \gamma}(1) = y \in \mathcal{V} \quad \text{and} \quad \tilde{f}(a') = \widetilde{f \circ \gamma \cdot f \circ \beta}(1),$$

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 $\tilde{f} : A \rightarrow Y$, 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 $\mathbb{R} \rightarrow S^1 : \theta \mapsto e^{i\theta}$ and a continuous map $f : A \rightarrow S^1$ 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 $\tilde{f} : A \rightarrow \mathbb{R}$ 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

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 $p_i : Y_i \rightarrow X$ for $i = 1, 2$, a **map of covers** from p_1 to p_2 is a map $f : Y_1 \rightarrow Y_2$ such that $p_2 \circ f = p_1$, i.e. the following diagram commutes:


$$(16.1) \quad \begin{array}{ccc} Y_1 & \xrightarrow{f} & Y_2 \\ & \searrow p_1 & \swarrow p_2 \\ & & X \end{array}$$

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 $f : Y_1 \rightarrow Y_2$ is a homeomorphism, since its inverse $f^{-1} : Y_2 \rightarrow Y_1$ 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 $p_i : (Y_i, y_i) \rightarrow (X, x)$ and $f : (Y_1, y_1) \rightarrow (Y_2, y_2)$ 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 $p : Y \rightarrow X$, an isomorphism of covers from p to itself is called a **deck transformation**²³ (*Decktransformation*) of $p : Y \rightarrow X$.

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 $p : Y \rightarrow X$ forms a

²²This is not a universally standard term.

²³This 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**

$$\text{Aut}(p) := \{f : Y \rightarrow Y \mid f \text{ is a homeomorphism such that } p \circ f = p\},$$

where the group operation is defined by composition of maps.

EXAMPLE 16.3. For the cover $p : \mathbb{R} \rightarrow S^1 : \theta \mapsto e^{i\theta}$, $\text{Aut}(p)$ consists of all maps $f_k : \mathbb{R} \rightarrow \mathbb{R}$ of the form $f_k(\theta) = \theta + 2\pi k$ for $k \in \mathbb{Z}$, so in particular, $\text{Aut}(p)$ is isomorphic to \mathbb{Z} .

EXAMPLE 16.4. Figure 10 illustrates a covering map $p : Y \rightarrow S^1 \vee S^1$ 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 $f : Y \rightarrow Y$ 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, $\text{Aut}(p)$ 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 $f : Y_1 \rightarrow Y_2$ is a *lift* of the map $p_1 : Y_1 \rightarrow X$ 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 $f : Y_1 \rightarrow Y_2$ 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 $p_i : Y_i \rightarrow X$ 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.* \square

Next we use the same perspective to study deck transformations of a single cover $p : Y \rightarrow X$. 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 $f : Y \rightarrow Y$ 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 $p : Y \rightarrow X$ 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 $p : Y \rightarrow X$ 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 $y_1, y_2 \in p^{-1}(x) \subset Y$ satisfy $y_2 = f(y_1)$ for some deck transformation $f \in \text{Aut}(p)$.

If $\deg(p) < \infty$, the previous remarks about uniqueness of deck transformations imply $|\text{Aut}(p)| \leq \deg(p)$, 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

$$(16.2) \quad p_*\pi_1(Y, y_1) = p_*\pi_1(Y, y_2).$$

Let us try to translate this into a condition for recognizing when p is regular. Recall that any path $y_1 \xrightarrow{\tilde{\gamma}} y_2$ in Y determines an isomorphism

$$\Phi_{\tilde{\gamma}} : \pi_1(Y, y_2) \rightarrow \pi_1(Y, y_1) : [\alpha] \mapsto [\tilde{\gamma} \cdot \alpha \cdot \tilde{\gamma}^{-1}].$$

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 $y_2 \xrightarrow{\alpha} y_2$, 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 $y_i \xrightarrow{\tilde{\gamma}} y_2$. Now observe: *any* loop $x \rightsquigarrow x$ 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 $p : (Y, y_0) \rightarrow (X, x_0)$ is regular if and only if the subgroup $p_*\pi_1(Y, y_0) \subset \pi_1(X, x_0)$ is normal.* \square

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 $p : (Y, y_0) \rightarrow (X, x_0)$ to the set of all subgroups of $\pi_1(X, x_0)$: it is defined by*

$$[p : (Y, y_0) \rightarrow (X, x_0)] \mapsto p_*\pi_1(Y, y_0).$$

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 $p_i : (Y_i, y_i) \rightarrow (X, x_0)$ 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 \tilde{X} .

EXAMPLES 16.11. The universal cover \tilde{S}^1 of S^1 is \mathbb{R} , due to the covering map $\mathbb{R} \rightarrow S^1 : \theta \mapsto e^{i\theta}$. Similarly, $\widetilde{\mathbb{R}P^n} \cong S^n$ for $n \geq 2$, 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 $g \geq 2$: these have universal cover $\tilde{\Sigma}_g \cong \mathbb{R}^2$. 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 $\mathring{\mathbb{D}}^2 \rightarrow \Sigma_g$ as the quotient projection.

For the remainder of this lecture, fix a base-point preserving covering map $p : (Y, y_0) \rightarrow (X, x_0)$ where X and Y are assumed reasonable, and denote

$$G := \pi_1(X, x_0), \quad H := p_*\pi_1(Y, y_0) \subset G.$$

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

$$G/H = \{gH \subset G \mid g \in G\},$$

where gH denotes the subset $\{gh \mid h \in H\} \subset G$. One can similarly consider the set of right cosets

$$H \backslash G = \{Hg \subset G \mid g \in G\}.$$

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/H and $H \backslash G$ have the same number (finite or infinite) of elements, which is called the **index** of H in G and denoted by

$$[G : H] := |G/H| = |H \backslash G|.$$

In the following we will make repeated use of the fact that for any $y \in p^{-1}(x_0)$, any path $y_0 \xrightarrow{\tilde{\gamma}} y$ 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*

$$\Phi : p^{-1}(x_0) \rightarrow H \backslash G : y \mapsto H[\gamma],$$

where $x_0 \xrightarrow{\tilde{\gamma}} x_0$ is any loop that lifts to a path $y_0 \xrightarrow{\tilde{\gamma}} y$.

COROLLARY 16.13. $\deg(p) = [G : H]$. □

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 $\Phi(y) = H[\tilde{\gamma}]$.

To see that Φ is injective, suppose $\Phi(y) = \Phi(y')$, choose paths $y_0 \xrightarrow{\tilde{\alpha}} y$ and $y_0 \xrightarrow{\tilde{\beta}} y'$, 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 $y_0 \xrightarrow{\tilde{\gamma}} y_0$ 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 $y = \tilde{\alpha}(1) = \tilde{\beta}(1) = y'$. \square

If the cover is regular so $H \subset G$ is normal, then $\deg(p) = |\text{Aut}(p)|$, and Corollary 16.13 therefore implies that $\text{Aut}(p)$ 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 $p : (Y, y_0) \rightarrow (X, x_0)$ 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*

$$\Psi : \text{Aut}(p) \rightarrow G/H : f \mapsto [\gamma]H,$$

where $x_0 \xrightarrow{\tilde{\gamma}} x_0$ is any loop that has a lift to a path from y_0 to $f(y_0)$.

Notice that the universal cover $p : (\tilde{X}, \tilde{x}_0) \rightarrow (X, x_0)$ 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 $p : (\tilde{X}, \tilde{x}_0) \rightarrow (X, x_0)$, there is an isomorphism $\text{Aut}(p) \rightarrow \pi_1(X, x_0)$ sending each deck transformation f to the homotopy class of any loop $x_0 \rightsquigarrow x_0$ that lifts to a path $\tilde{x}_0 \rightsquigarrow f(\tilde{x}_0)$.* \square

PROOF OF THEOREM 16.14. Regularity implies that the map $\text{Aut}(p) \rightarrow p^{-1}(x_0) : f \mapsto f(y_0)$ is bijective, so Ψ is then well defined and bijective due to Lemma 16.12. For the identity element $\text{Id} \in \text{Aut}(p)$, we have $\Psi(\text{Id}) = [\gamma]H$ for any loop γ that lifts to a loop from y_0 to $\text{Id}(y_0) = y_0$, 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 $f, g \in \text{Aut}(p)$. Choose loops α, β based at x_0 which lift to paths $y_0 \xrightarrow{\tilde{\alpha}} f(y_0)$ and $y_0 \xrightarrow{\tilde{\beta}} g(y_0)$. 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

$$y_0 \xrightarrow{\tilde{\alpha} \cdot (f \circ \tilde{\beta})} f \circ g(y_0).$$

Now since $f \in \text{Aut}(p)$, $p \circ f = p$ implies $p \circ (f \circ \tilde{\beta}) = p \circ \tilde{\beta} = \beta$, thus

$$\Psi(f \circ g) = [p \circ (\tilde{\alpha} \cdot (f \circ \tilde{\beta}))] = [\alpha][\beta] = \Psi(f)\Psi(g).$$

\square

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 $\text{Aut}(p) \rightarrow p^{-1}(x_0)$ 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 \tilde{X} have the same end point, and a loop is trivial in $\pi_1(X, x_0)$ if and only if its lift to \tilde{X} is also a loop.

EXAMPLE 16.16. Applying Corollary 16.15 to the cover $p : \mathbb{R} \rightarrow S^1 : \theta \mapsto e^{i\theta}$ reproduces the isomorphism $\pi_1(S^1, 1) \cong \mathbb{Z}$ we discussed at the end of Lecture 9. The loop $\gamma_k(t) := e^{2\pi ikt}$ in S^1 for each $k \in \mathbb{Z}$ lifts to \mathbb{R} with base point 0 as the path $\tilde{\gamma}_k(t) = 2\pi kt$.

EXAMPLE 16.17. For each $n \geq 2$, Corollary 16.15 implies $\pi_1(\mathbb{R}P^n) \cong \mathbb{Z}_2$, as this is the automorphism group of the universal cover $p : S^n \rightarrow \mathbb{R}P^n$, defined as the natural quotient projection. Concretely, after fixing base points $x_0 \in \mathbb{R}P^n$ and $y_0 \in p^{-1}(x_0) \subset S^n$, each loop in $\mathbb{R}P^n$ 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 $\pi_1(\mathbb{R}P^n, x_0)$ is thus represented by any loop whose lift to S^n starts and ends at antipodal points.

17. The universal cover and group actions

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 $p : (Y, y_0) \rightarrow (X, x_0)$ of reasonable spaces can be fit into a diagram

$$(17.1) \quad \begin{array}{ccccc} (Z, z_0) & \xrightarrow{q} & (Y, y_0) & \xrightarrow{p} & (X, x_0), \\ & & \searrow & \nearrow & \\ & & & P & \end{array}$$

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 $q : Z \rightarrow Y$ to be the universal cover of Y , which makes $P : Z \rightarrow X$ 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 $\text{Aut}(P) \cong \pi_1(X, x_0)$ and $\text{Aut}(q) \cong \pi_1(Y, y_0)$. 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 $\text{Aut}(p)$ 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

$$G := \text{Aut}(P), \quad \text{and} \quad H := \text{Aut}(q).$$

- (a) Use the path-lifting property to prove the following lemma: If $\Psi \in G$ and $\psi \in \text{Aut}(p)$ 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 $\Psi : Z \rightarrow Z$ for P that satisfy $\Psi(z_0) \in q^{-1}(y_0)$.
- (c) Show that if $P : Z \rightarrow X$ is regular then so is $q : Z \rightarrow Y$. 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.

$$N(H) := \{g \in G \mid gHg^{-1} = H\}.$$

Show that if the cover $q : Z \rightarrow Y$ is regular, then for any $\Psi \in N(H)$, there exists a deck transformation $\psi : Y \rightarrow Y$ 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 $N(H) \rightarrow \text{Aut}(p)$ whose kernel is H .

- (e) Show that if the cover $P : Z \rightarrow X$ is also regular, then the homomorphism $N(H) \rightarrow \text{Aut}(p)$ in part (d) is also surjective, and thus descends to an isomorphism

$$N(H)/H \xrightarrow{\cong} \text{Aut}(p).$$

Applying Exercise 17.1 with Z simply connected now gives:

COROLLARY 17.2. *For any covering map $p : (Y, y_0) \rightarrow (X, x_0)$ 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 $\text{Aut}(p) \cong N(H)/H$. \square*

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 $\text{Aut}(p)$ 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 $p_* : \pi_1(Y, y_0) \rightarrow \pi_1(X, x_0)$ induced by a covering map $p : (Y, y_0) \rightarrow (X, x_0)$ 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 $p : (\tilde{X}, \tilde{x}_0) \rightarrow (X, x_0)$, 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 $p^{-1}(x_0) \subset \tilde{X}$ 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

$$\tilde{X} := \{\text{paths } \gamma : (I, 0) \rightarrow (X, x_0)\} / \sim_{h+},$$

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

$$p : \tilde{X} \rightarrow X : [\gamma] \mapsto \gamma(1),$$

which is obviously surjective since X is path-connected. Notice that $p^{-1}(x_0) = \pi_1(X, x_0)$.

We claim that \tilde{X} can be assigned a topology that makes $p : \tilde{X} \rightarrow X$ 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 $i_*^{\mathcal{U}} : \pi_1(\mathcal{U}, x) \rightarrow \pi_1(X, x)$ is trivial if and only if every loop $S^1 \rightarrow \mathcal{U}$ based at x can be extended to a map $\mathbb{D}^2 \rightarrow X$. Notice that this is weaker in general than demanding an extension $\mathbb{D}^2 \rightarrow \mathcal{U}$; 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 $i_*^{\mathcal{U}} : \pi_1(\mathcal{U}) \rightarrow \pi_1(X)$ is trivial. With this understood, consider the collection of sets

$$\mathcal{B} := \{ \mathcal{U} \subset X \mid \mathcal{U} \text{ is open and path-connected and } i_*^{\mathcal{U}} : \pi_1(\mathcal{U}) \rightarrow \pi_1(X) \text{ is trivial} \}.$$

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

$$\mathcal{U}_{[\gamma]} := \left\{ [\gamma \cdot \alpha] \in \tilde{X} \mid \alpha \text{ is a path in } \mathcal{U} \text{ starting at } x \right\}.$$

Notice that $\mathcal{U}_{[\gamma]}$ depends only on the homotopy class $[\gamma] \in \tilde{X}$; 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 $p : \tilde{X} \rightarrow X$ restricts to a bijection

$$\mathcal{U}_{[\gamma]} \xrightarrow{p} \mathcal{U}.$$

With all this in mind, one can now show that

$$\tilde{\mathcal{B}} := \left\{ \mathcal{U}_{[\gamma]} \subset \tilde{X} \mid \mathcal{U} \in \mathcal{B} \text{ and } [\gamma] \in \tilde{X} \text{ with } \gamma(1) \in \mathcal{U} \right\}$$

is a base for a topology on \tilde{X} such that each $\mathcal{U} \in \mathcal{B}$ is evenly covered by $p : \tilde{X} \rightarrow X$. We leave the details of this as an exercise.

There is an obvious choice of base point in \tilde{X} : define $\tilde{x}_0 \in \tilde{X}$ as the homotopy class of the constant path at x_0 . It remains to prove that $\pi_1(\tilde{X}, \tilde{x}_0) = 0$. Since we now know that $p : (\tilde{X}, \tilde{x}_0) \rightarrow (X, x_0)$ is a covering map, Corollary 15.13 implies that $p_* : \pi_1(\tilde{X}, \tilde{x}_0) \rightarrow \pi_1(X, x_0)$ is injective, thus it will suffice to show that the subgroup $p_*\pi_1(\tilde{X}, \tilde{x}_0)$ 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 \tilde{X} can be written as

$$\tilde{\gamma}(t) = [\gamma_t] \in \tilde{X},$$

where for each $t \in I$ we define

$$\gamma_t(s) := \begin{cases} \gamma(s) & \text{for } 0 \leq s \leq t, \\ \gamma(t) & \text{for } t \leq s \leq 1. \end{cases}$$

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. \square

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 \tilde{X} . One can thus construct \tilde{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 \tilde{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

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

are both continuous.

Popular examples of topological groups include the various subgroups of the real or complex general linear groups $\text{GL}(n, \mathbb{R})$ and $\text{GL}(n, \mathbb{C})$, e.g. the orthogonal group $\text{O}(n)$ and unitary group $\text{U}(n)$, the special linear groups $\text{SL}(n, \mathbb{R})$ and $\text{SL}(n, \mathbb{C})$, and so forth. We saw in Exercise 7.29 that for any locally compact and locally connected Hausdorff space X , the group of homeomorphisms $\text{Homeo}(X)$ 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

$$G \times X \rightarrow X : (g, x) \mapsto g \cdot x$$

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 $G \rightarrow \text{Homeo}(X)$ sending $g \in G$ to the homeomorphism $\varphi_g : X \rightarrow X$ defined by $\varphi_g(x) = g \cdot x$. If G is a discrete group then the converse is also true: every group homomorphism $G \rightarrow \text{Homeo}(X)$ comes from a G -action on X . This is true because as long as the topology of G is discrete, the map $G \times X \rightarrow X : (g, x) \mapsto g \cdot x$ is continuous if and only if the map $X \rightarrow X : x \mapsto g \cdot x$ 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 $p : Y \rightarrow X$, $\text{Aut}(p)$ 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 $\varphi_1 : X \rightarrow X$ 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 $O(n)$ 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 $\mathbb{R}P^n$.

PROPOSITION 17.9. *Regarding $\pi_1(X, x_0)$ as a discrete group, any covering map $p : (Y, y_0) \rightarrow (X, x_0)$ 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 $\text{Aut}(p)$, and the latter acts on Y as explained in Example 17.5.

Alternatively, one can appeal to the uniqueness of the universal cover, so $p : (Y, y_0) \rightarrow (X, x_0)$ is necessarily isomorphic to the specific cover $\tilde{X} = \{\text{paths } x_0 \rightsquigarrow x\} / \sim_{h+}$ 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 $[\gamma] \in \tilde{X}$ is by concatenation:

$$[\alpha] \cdot [\gamma] := [\alpha \cdot \gamma].$$

It is easy to verify that this also defines a group action. □

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

$$(g \cdot \mathcal{U}) \cap \mathcal{U} = \emptyset$$

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 $p : Y \rightarrow X$, the action of $\text{Aut}(p)$ 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 $p : \mathbb{D}^2 \rightarrow X$ with $\deg(p) > 1$.*

PROOF. If $\deg(p) > 1$, then since $\pi_1(\mathbb{D}^2) = 0$, we observe that the cover $p : \mathbb{D}^2 \rightarrow X$ must be regular and therefore has a nontrivial deck transformation group $\text{Aut}(p)$ 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 . \square

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*

$$q : X \rightarrow X/G : x \mapsto G \cdot x$$

is a regular covering map with $\text{Aut}(q) = G$. \square

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 $p : (Y, y_0) \rightarrow (X, x_0)$ such that $p_*\pi_1(Y, y_0) = H$. Since X is reasonable, there exists a universal cover $f : (\tilde{X}, \tilde{x}_0) \rightarrow (X, x_0)$, whose automorphism group is isomorphic to G , so this isomorphism defines a free and properly discontinuous action of G on \tilde{X} . It also defines a free and properly discontinuous action of every subgroup of G on \tilde{X} , and in particular an H -action. Define

$$Y := \tilde{X}/H \quad \text{and} \quad p : Y \rightarrow X : H \cdot \tilde{x} \mapsto f(\tilde{x}).$$

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 $q : (\tilde{X}, \tilde{x}_0) \rightarrow (Y, y_0)$ is now the universal cover of Y , and it fits into the following commutative diagram:

$$\begin{array}{ccc} (\tilde{X}, \tilde{x}_0) & \xrightarrow{f} & (X, x_0) \\ \downarrow q & \nearrow p & \\ (Y, y_0) & & \end{array}$$

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 \tilde{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 $p_*\pi_1(Y, y_0) = H$. \square

18. Manifolds

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 $n \geq 0$ (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

$$\varphi : \mathcal{U} \rightarrow \Omega$$

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 $\{\varphi_\alpha : \mathcal{U}_\alpha \rightarrow \Omega_\alpha\}_{\alpha \in J}$ 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 $\varphi_\alpha : \mathcal{U}_\alpha \rightarrow \Omega_\alpha$ such that Ω_α is either an open ball in $\mathbb{H}^n \setminus \partial\mathbb{H}^n$ 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

$$\partial M := \{p \in M \mid \varphi(p) \in \partial\mathbb{H}^n \text{ for some chart } \varphi\},$$

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 $\partial M = \emptyset$, 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 $\partial M = \emptyset$. You should be aware that some books adopt different conventions for such details, e.g. some authors assume $\partial M = \emptyset$ 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 $\varphi : \mathcal{U} \rightarrow \Omega$, we required²⁴ 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 $\varphi : \mathcal{U} \rightarrow \Omega$ 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 $\varphi : \mathcal{U} \rightarrow \Omega$ 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

²⁴This 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 $\Omega \cap \partial\mathbb{H}^n = \emptyset$, 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 $\varphi : \mathcal{U} \rightarrow \Omega$ 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 $\mathbb{H}^n \setminus \partial\mathbb{H}^n$. For $n \leq 2$, 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 .* \square

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.* \square

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 $\partial M = \emptyset$. 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 $g \geq 0$, $\mathbb{R}\mathbb{P}^2$, \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 $g \geq 1$ 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 $m \geq 1$ 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 , $\mathbb{R}\mathbb{P}^n$, \mathbb{R}^n , $\mathbb{T}^n := S^1 \times \dots \times S^1$, 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 $\varphi_\alpha : \mathcal{U}_\alpha \rightarrow \Omega_\alpha$ and $\varphi_\beta : \mathcal{U}_\beta \rightarrow \Omega_\beta$ 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:

$$\begin{array}{ccc} & \mathcal{U}_\alpha \cap \mathcal{U}_\beta & \\ \swarrow \varphi_\alpha & & \searrow \varphi_\beta \\ \varphi_\alpha(\mathcal{U}_\alpha \cap \mathcal{U}_\beta) & \xrightarrow{\varphi_\beta \circ \varphi_\alpha^{-1}} & \varphi_\beta(\mathcal{U}_\alpha \cap \mathcal{U}_\beta) \end{array}$$

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 $f : M \rightarrow \mathbb{R}$ 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 $\{\varphi_\alpha : \mathcal{U}_\alpha \rightarrow \Omega_\alpha\}_{\alpha \in J}$ 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*).

²⁵Recall 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 $\mathbb{R} \rightarrow \mathbb{R}$ 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, $F : \mathcal{U} \rightarrow \mathbb{R}^k$ is a C^∞ -map and $q \in \mathbb{R}^k$ is a point such that for all $p \in F^{-1}(q)$, the derivative $dF(p) : \mathbb{R}^n \rightarrow \mathbb{R}^k$ 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 $n - k$. \square*

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 $F : \mathbb{R}^{n+1} \rightarrow \mathbb{R} : \mathbf{x} \mapsto |\mathbf{x}|^2$, 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 $\det : \mathbb{R}^{n \times n} \rightarrow \mathbb{R}$. 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 $\mathbf{1}$ 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 $\mathcal{U} \setminus \{x\}$ 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 $n \geq 2$ 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 $\mathbb{R} \times \{0, 1\} \rightarrow X$ 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.* \square

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 $F : \mathbb{R}^n \rightarrow \mathbb{R}^{n-1}$. 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 $\partial M = \emptyset$, 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 $M \cong [0, 1]$. \square

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}_β and a point $p_\beta \in \mathcal{U}_\beta$ that is not in \mathcal{U}_α .*

PROOF. Choose explicit homeomorphisms $\varphi_\alpha : \mathcal{U}_\alpha \rightarrow \mathbb{R}$ and $\varphi_\beta : \mathcal{U}_\beta \rightarrow \mathbb{R}$. The image $\varphi_\beta(\mathcal{W}) \subset \mathbb{R}$ is necessarily 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}_β since this would imply $p_\alpha \in \mathcal{W}$ and thus $t \in \varphi_\alpha(\mathcal{W})$. We claim: $|\varphi_\beta(x_n)| \rightarrow \infty$. 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 $x_n \rightarrow p_\alpha$, 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

$$\varphi_\alpha(\mathcal{W}) = (-\infty, a) \quad \text{and} \quad \varphi_\beta(\mathcal{W}) = (b, \infty),$$

in which case the recipe described above for defining $p_\alpha, p_\beta \in \overline{\mathcal{W}}$ gives

$$p_\alpha = \varphi_\alpha^{-1}(a), \quad p_\beta = \varphi_\beta^{-1}(b).$$

(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 $|\varphi_\beta(x_n)| \rightarrow \infty$ and thus $\varphi_\alpha(x_n) \rightarrow a$, implying $x_n \rightarrow p_\alpha$. 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 $\overline{\mathcal{W}} \rightarrow [b, \infty]$ that sends $p_\alpha \mapsto \infty$, so $\overline{\mathcal{W}}$ has the topology of a compact interval. \square

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 $\varphi_\alpha : \mathcal{U}_\alpha \rightarrow \mathbb{R}$ 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.* \square

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 $\varphi_n : \mathcal{U}_n \rightarrow \mathbb{R}$; 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 $N \geq 2$, 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 $n \geq 2$ 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} . \square

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** $(I, <)$ consists of a set I with a partial order $<$ such that for all pairs of elements $x, y \in I$, at least one of the conditions $x < y$ or $y < x$ 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 (ω_1, \leq) such that for every $x \in \omega_1$, at most countably many elements $y \in \omega_1$ satisfy $y \leq x$.*

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 $x \leq y$ 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 \emptyset as 0, identify the natural number 1 with the set $\{0\} = \{\emptyset\}$, identify 2 with the set $\{0, 1\} = \{\emptyset, \{\emptyset\}\}$, identify

$$3 = \{0, 1, 2\} = \{\emptyset, \{\emptyset\}, \{\{\emptyset\}\}\}$$

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 $(x, s) \leq (y, t)$ whenever either $x \leq y$ or both $x = y$ and $s \leq t$ hold. Writing $x < y$ to mean $x \leq y$ 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. \square*

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 $n \geq 2$, 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

$$\Sigma := \mathbb{H} \amalg \left(\coprod_{a \in \mathbb{R}} X_a \right) / \sim$$

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 .

- Prove that Σ is Hausdorff.
- Prove that Σ is path-connected.
- Prove that every point in Σ has a neighborhood homeomorphic to \mathbb{R}^2 .
- 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

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 $g \geq 0$ and denoted by Σ_g . In particular, we've seen that for each $g \geq 1$, Σ_g is homeomorphic to the g -fold connected sum of copies of \mathbb{T}^2 , and we have also computed its fundamental group

$$\pi_1(\Sigma_g) \cong \left\{ a_1, b_1, \dots, a_g, b_g \mid \prod_{i=1}^g [a_i, b_i] = e \right\},$$

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 $\mathbb{R}\mathbb{P}^2$, e.g. $\mathbb{R}\mathbb{P}^2 \# \mathbb{R}\mathbb{P}^2$ 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 $\#_{i=1}^g \mathbb{R}\mathbb{P}^2$ is homeomorphic to a space obtained from a polygon with $2g$ edges by identifying them in pairs according to the sequence $a_1, a_1, \dots, a_g, a_g$, thus

$$\pi_1(\#_{i=1}^g \mathbb{R}\mathbb{P}^2) \cong \{a_1, \dots, a_g \mid a_1^2 \dots a_g^2 = e\}.$$

EXERCISE 19.3. For $i = 1, \dots, g-1$, let $e_i \in \mathbb{Z}^{g-1}$ denote the i th standard basis vector. Show that there is a well-defined homomorphism $G := \{a_1, \dots, a_g \mid a_1^2 \dots a_g^2 = e\} \rightarrow \mathbb{Z}^{g-1} \oplus \mathbb{Z}_2$ such that

$$a_i \mapsto \begin{cases} (e_i, 0) & \text{for } i = 1, \dots, g-1, \\ (-1, \dots, -1, 1) & \text{for } i = g, \end{cases}$$

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.* \square

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. $\mathbb{R}\mathbb{P}^2 \# \mathbb{T}^2$ is homeomorphic to the connected sum of $\mathbb{R}\mathbb{P}^2$ 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 $[-1, 1] \times S^1$:

$$\Sigma' := \left(\Sigma \setminus (\mathring{\mathbb{D}}^2 \amalg \mathring{\mathbb{D}}^2) \right) \cup_{S^1 \amalg S^1} ([-1, 1] \times S^1).$$

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 $i_1 : \mathbb{D}^2 \rightarrow \Sigma$ 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.

²⁶This 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 $\Sigma = \mathbb{R}\mathbb{P}^2$. 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 $\mathbb{R}\mathbb{P}^2 \# \mathbb{T}^2$ by embedding two small disks disjointly into $\mathbb{M} \subset \mathbb{R}\mathbb{P}^2$, 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 $i_1 : \mathbb{D}^2 \hookrightarrow \mathbb{R}\mathbb{P}^2$ to $i'_1 : \mathbb{D}^2 \hookrightarrow \mathbb{R}\mathbb{P}^2$ through a continuous family of embeddings disjoint from the second disk. This proves that if $\Sigma = \mathbb{R}\mathbb{P}^2$, 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 $\Sigma = \mathbb{R}\mathbb{P}^2$, 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 $g \geq 0$ or $\#_{i=1}^g \mathbb{R}\mathbb{P}^2$ for some $g \geq 1$, 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 $\mathbb{R}\mathbb{P}^2$ by connected sums. (Note that whenever both \mathbb{T}^2 and $\mathbb{R}\mathbb{P}^2$ appear in this collection, Prop. 19.5 allows us to replace \mathbb{T}^2 with two copies of $\mathbb{R}\mathbb{P}^2$, as $\mathbb{R}\mathbb{P}^2 \# \mathbb{R}\mathbb{P}^2$ 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.

$$\Sigma^{(j+1)} \cong \Sigma^{(j)} \setminus \mathring{\mathbb{D}}^2$$

for some embedding $\mathbb{D}^2 \hookrightarrow \Sigma^{(j)} \setminus \partial\Sigma^{(j)}$;

- (2) Attaching a handle (or “cross-handle”) to connect two separate boundary components $\ell_1, \ell_2 \subset \partial\Sigma^{(j)}$, i.e.

$$\Sigma^{(j+1)} \cong \Sigma^{(j)} \cup_{\ell_1 \cup \ell_2} ([-1, 1] \times S^1)$$

for some choice of homeomorphism $\partial([-1, 1] \times S^1) = S^1 \amalg S^1 \rightarrow \ell_1 \amalg \ell_2$;

- (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 $\partial\mathbb{D}^2 = S^1 \rightarrow \ell$;

- (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 $\partial\mathbb{M} \cong S^1 \rightarrow \ell$.

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

$$\Delta^n := \{(t_0, \dots, t_n) \in I^{n+1} \mid t_0 + \dots + t_n = 1\}$$

for each integer $n \geq 0$. 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

$$\Delta_\sigma := \left\{ (t_1, \dots, t_N) \in I^N \mid t_{i_0} + \dots + t_{i_k} = 1 \text{ and } t_j = 0 \text{ for all } v_j \notin \sigma \right\}.$$

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 $\partial\Sigma = \emptyset$, 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.*

²⁷The 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; see 29.

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 $D_1, \dots, D_N \subset \Sigma$ 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 $\Sigma \setminus \bigcup_i \partial D_i$ 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 $\sigma_1, \dots, \sigma_N$. By abuse of notation, we shall also denote by $\sigma_1, \dots, \sigma_N$ the corresponding subsets of Σ homeomorphic to the standard 2-simplex Δ^2 . The latter is homeomorphic to $\mathbb{D}^2 \cong S^2 \setminus \mathring{\mathbb{D}}^2$, thus

$$\Sigma^{(0)} := \sigma_1 \amalg \dots \amalg \sigma_N$$

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

$$\Sigma^{(j+1)} = \Sigma^{(j)} / \sim \quad \text{where} \quad \sim \text{ identifies } \ell_1 \text{ with } \ell_2,$$

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 if some chart maps it to $\mathbb{H}^2 \setminus \partial\mathbb{H}^2$. 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 $\mathbb{H}^2 \setminus \partial\mathbb{H}^2$, then a transition map relating these two charts maps some neighborhood in \mathbb{H}^2 of a point $x \in \mathbb{H}^2 \setminus \partial\mathbb{H}^2$ 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 $f : M \rightarrow N$ 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 $f : M \rightarrow N$ 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 $f : M \rightarrow N$ is a smooth map that has $q \in N$ as a regular value, the preimage $f^{-1}(q) \subset M$ is a smooth submanifold²⁸ of dimension $m - n$. If M has boundary, then one should assume additionally that q is a regular value of the restricted map $f|_{\partial M} : \partial M \rightarrow N$, 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 $f : M \rightarrow N$ is continuous, then every neighborhood of f in $C(M, N)$ with the compact-open topology (cf. Exercise 7.28) contains a smooth map $f_\epsilon : M \rightarrow N$ 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 $f : M \rightarrow N$ 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 $f : M \rightarrow N$.

Hint: If you have a homotopy $H : I \times M \rightarrow N$ 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 $f : M \rightarrow N$, 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 $f : S^2 \rightarrow S^2$ homotopic to the identity is surjective.
 (c) What goes wrong with this discussion if we allow M to be a noncompact manifold? Describe two homotopic maps $f, g : \mathbb{R} \rightarrow S^1$ 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 $S^m \rightarrow S^n$ is homotopic to a constant map.


Hint: What does it mean for a point $q \in S^n$ to be a regular value of $f : S^m \rightarrow S^n$ if $n > m$?

20. Orientations

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:

²⁸A 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 $\varphi : \mathcal{U} \rightarrow \mathbb{R}^m$ 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 $i : K^2 \rightarrow \mathbb{R}^3$ into 3-dimensional Euclidean space from a certain closed 2-manifold K^2 : in differential geometry, one would call $i : K^2 \rightarrow \mathbb{R}^3$ 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 $i : K^2 \rightarrow \mathbb{R}^3$ 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 $\gamma : I \rightarrow K^2$ 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 $\gamma : I \rightarrow \mathbb{T}^2 \subset \mathbb{R}^3$ 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

²⁹Notice 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 .

³⁰The 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 $\mathbb{C} \setminus \{z\}$, and for a suitable choice of base point it is always a generator of $\pi_1(\mathbb{C} \setminus \{z\}) \cong \mathbb{Z}$. For $k \in \mathbb{Z}$, define also the loop

$$\gamma_{z,\epsilon}^k : S^1 \rightarrow \mathbb{C} : e^{i\theta} \mapsto z + \epsilon e^{ki\theta},$$

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 $\alpha : S^1 \rightarrow \mathbb{C} \setminus \{z\}$, the **winding number** (*Windungszahl*) of α about z is an integer characterized uniquely by the condition

$$\text{wind}(\alpha; z) = k \iff \alpha \underset{h}{\sim} \gamma_{z,\epsilon}^k \text{ in } \mathbb{C} \setminus \{z\}.$$

If $\mathcal{U}, \mathcal{V} \subset \mathbb{C}$ are open subsets and $f : \mathcal{U} \rightarrow \mathcal{V}$ 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 $\mathbb{C} \setminus \{w\}$. It follows that there is a well-defined winding number $\text{wind}(f \circ \gamma_{z,\epsilon}; w) \in \mathbb{Z}$, 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 $\mathcal{U} \setminus \{z\}$, so that $f \circ \gamma_{z,\epsilon}$ and $f \circ \gamma_{z,\epsilon'}$ are homotopic in $\mathbb{C} \setminus \{w\}$.

LEMMA 20.1. *In the situation described above, $\text{wind}(f \circ \gamma_{z,\epsilon}; w)$ 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 $\pi_1(B_r(w) \setminus \{w\}) \cong \pi_1(\mathbb{C} \setminus \{w\}) \cong \mathbb{Z}$, and $k := \text{wind}(f \circ \gamma_{z,\epsilon}; w)$ is the unique integer such that $f \circ \gamma_{z,\epsilon}$ is homotopic in $B_r(w) \setminus \{w\}$ to $\gamma_{w,\delta}^k$. Since $\gamma_{z,\epsilon}$ generates $\pi_1(\mathbb{C} \setminus \{z\})$, there is also a unique integer $\ell \in \mathbb{Z}$ such that $f^{-1} \circ \gamma_{w,\delta}$ is homotopic in $\mathbb{C} \setminus \{z\}$ to $\gamma_{z,\epsilon}^\ell$. This implies

$$\gamma_{z,\epsilon} = f^{-1} \circ f \circ \gamma_{z,\epsilon} \underset{h}{\sim} f^{-1} \circ \gamma_{w,\delta}^k \underset{h}{\sim} \gamma_{z,\epsilon}^{k\ell} \text{ in } \mathbb{C} \setminus \{z\},$$

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 $\mathcal{U}_\pm = \{z \in \mathcal{U} \mid \text{wind}(f \circ \gamma_{z,\epsilon}; f(z)) = \pm 1\}$ 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 $\text{wind}(f \circ \gamma_{z,\epsilon}; w)$ 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 $\mathbb{H}^2 \setminus \partial\mathbb{H}^2$, 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 $\mathbb{H}^2 \setminus \partial\mathbb{H}^2$ to $\mathbb{H}^2 \setminus \partial\mathbb{H}^2$.

DEFINITION 20.3. Given open subsets $\mathcal{U}, \mathcal{V} \subset \mathbb{H}^2$, a homeomorphism $f : \mathcal{U} \rightarrow \mathcal{V}$ is called **orientation preserving** (*orientierungserhaltend*) if $\text{wind}(f \circ \gamma_{z,\epsilon}; f(z)) = 1$ for all $z \in \mathbb{H}^2 \setminus \partial\mathbb{H}^2$ and $\epsilon > 0$ sufficiently small. It is called **orientation reversing** (*orientierungsumkehrend*) if $\text{wind}(f \circ \gamma_{z,\epsilon}; f(z)) = -1$ for all $z \in \mathbb{H}^2 \setminus \partial\mathbb{H}^2$ 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 $f : \mathcal{U} \rightarrow \mathcal{V}$ 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 $n \geq 3$, 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 $df(x) : \mathbb{R}^n \rightarrow \mathbb{R}^n$ 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 $n \geq 1$ is a maximal collection of charts $\{\varphi_\alpha : \mathcal{U}_\alpha \rightarrow \Omega_\alpha\}_{\alpha \in J}$ 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 $\epsilon : M \rightarrow \{1, -1\}$, 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 $\pi : J \rightarrow J$ 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, \dots, 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 $\Delta^n = \{(t_0, \dots, t_n) \mid t_0 + \dots + t_n = 1\}$. 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 $n \geq 1$ 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 $\sigma = \{v_0, \dots, v_n\}$, the k th **boundary face** $\partial_{(k)}\sigma$ of σ is the $(n-1)$ -simplex whose vertices include all the v_0, \dots, v_n except v_k . Clearly if the vertices v_0, \dots, v_n come with an ordering, then the vertices of $\partial_{(k)}\sigma$ inherit an ordering from this, though here we

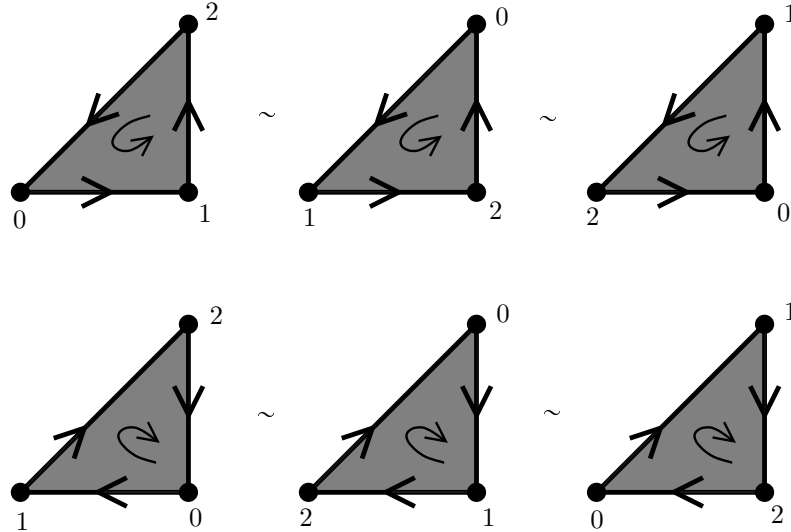


FIGURE 11. The six distinct orderings that define the two possible orientations of a 2-simplex.

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 $n \geq 2$ with vertices v_0, \dots, v_n ordered accordingly, the induced **boundary orientation** of its k th 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 $g \geq 0$.* \square

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 $g \geq 0$. 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 $m \geq 0$ disjoint disks out of the oriented surface Σ_g of genus $g \geq 0$.

- (a) Prove that every compact, orientable, connected surface with boundary is homeomorphic to $\Sigma_{g,m}$ for some values of $g, m \geq 0$.

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

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 $\mathbb{R}^2 \setminus \mathring{\mathbb{D}}^2$, then the fundamental group changes from 0 to \mathbb{Z} , with the boundary of $\mathring{\mathbb{D}}^2$ representing a generator of $\pi_1(\mathbb{R}^2 \setminus \mathring{\mathbb{D}}^2)$, 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 $(\mathbb{R}^2 \setminus \mathring{\mathbb{D}}^2) \times \mathbb{R}$, 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 $\pi_1((\mathbb{R}^2 \setminus \mathring{\mathbb{D}}^2) \times \mathbb{R}) \cong \pi_1(\mathbb{R}^2 \setminus \mathring{\mathbb{D}}^2) \cong \mathbb{Z}$ 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 $\pi_1(\mathbb{R}^3 \setminus \mathring{\mathbb{D}}^3)$ is still zero since $\mathbb{R}^3 \setminus \mathring{\mathbb{D}}^3$ 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 $\gamma : S^1 \rightarrow \Sigma_2$ 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 $\mathbb{D}^2 \rightarrow \Sigma_2$, 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 $k \geq 0$, fix a base point $t_0 \in S^k$ and associate to any pointed space (X, x_0) the set

$$\pi_k(X, x_0) = \{f : (S^k, t_0) \rightarrow (X, x_0)\} / \underset{h_+}{\sim},$$

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 $k \geq 2$, $\pi_k(X, x_0)$ can always be given the structure of an *abelian* group whose identity element is represented by the constant map

$$0 := [(S^k, t_0) \rightarrow (X, x_0) : t \mapsto x_0].$$

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 $S^k \rightarrow S^k$ represents a nontrivial element of $\pi_k(S^k)$ for every $k \geq 1$. This follows from Exercise 19.14, which sketches the notion of the mod 2 mapping degree in order to show that every map $S^k \rightarrow S^k$ homotopic to the identity is surjective (and therefore nonconstant). More generally, one can use the integer-valued mapping degree for maps $S^k \rightarrow S^k$ 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 $f : S^k \rightarrow S^n$, we can approximate f with a map whose image does not contain q and is thus contained in $S^n \setminus \{q\} \cong \mathbb{R}^n$. The latter admits a deformation retraction to any point it contains, so composing the perturbed map $S^k \rightarrow S^n \setminus \{q\}$ with a deformation retraction of $S^n \setminus \{q\}$ 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 $\mathbb{R}^3 \setminus \mathbb{D}^3$ since, by homotopy invariance,

$$\pi_2(\mathbb{R}^3 \setminus \mathbb{D}^3) \cong \pi_2(S^2) \cong \mathbb{Z},$$

with the inclusion $S^2 \hookrightarrow \mathbb{R}^3 \setminus \mathbb{D}^3$ 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 $X_0 = S^1 \times \mathbb{D}^2$, so $X \setminus X_0 = S^1 \times (\mathbb{R}^2 \setminus \mathbb{D}^2)$ admits a deformation retraction to $\partial X_0 = S^1 \times S^1 = \mathbb{T}^2$. By homotopy invariance, we have $\pi_1(X) \cong \pi_1(S^1) \cong \mathbb{Z}$ and $\pi_1(X \setminus X_0) \cong \pi_1(\mathbb{T}^2) \cong \mathbb{Z}^2$, so π_1 does at least partly detect the removal of X_0 from X . But since $X \setminus X_0$ is homotopy equivalent to a surface, there is also an intrinsically 2-dimensional phenomenon going on in this picture, and it seems natural to ask: does $X \setminus X_0$ 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 $\pi_2(X) = \pi_2(X \setminus X_0) = 0$. To see this, we can use the homotopy invariance of π_2 : the spaces X and $X \setminus X_0$ 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 $p : \tilde{Y} \rightarrow Y$ denotes the universal cover of some reasonable space Y , and \tilde{Y} is contractible. Since S^2 is simply connected, any map $f : S^2 \rightarrow Y$ can be lifted to $\tilde{f} : S^2 \rightarrow \tilde{Y}$, but the contractibility of \tilde{Y} then implies that \tilde{f} is homotopic to a constant map. Composing that homotopy with $p : \tilde{Y} \rightarrow Y$ gives a corresponding homotopy of $f = p \circ \tilde{f} : S^2 \rightarrow Y$ 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 role 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 $k \geq 0$, let

$$\Omega_k(X) := \{(M, f)\} / \sim,$$

³¹In 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, $f : M \rightarrow X$ 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 $F : W \rightarrow X$ 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 $M \rightarrow X$, as the homotopy $H : I \times M \rightarrow X$ defines a bordism $(I \times M, H)$.

EXAMPLE 21.9. Recall from Example 21.4 the loop $\gamma : S^1 \rightarrow \Sigma_2$ 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 (\emptyset, \cdot) , where \cdot denotes the unique map $\emptyset \rightarrow X$. Hence $[(S^1, \gamma)] = [(\emptyset, \cdot)] \in \Omega_1(\Sigma_2)$.

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 $f_1 \amalg f_2 : M_1 \amalg M_2 \rightarrow X$ 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

$$0 := [(\emptyset, \cdot)],$$

with $\cdot : \emptyset \rightarrow X$ 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 $F : W \rightarrow X$ with $F|_M = f$.

Referring back to Example 21.7, one can now show that the bordism class represented by the inclusion $\mathbb{T}^2 = \partial \bar{X}_0 \hookrightarrow X \setminus X_0$ is nontrivial in $\Omega_2(X \setminus X_0)$. One way to prove this uses the mod 2 mapping degree (cf. Exercise 19.14) for maps $f : \mathbb{T}^2 \rightarrow \mathbb{T}^2$: 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 $[(\mathbb{T}^2, \text{Id})] \neq 0 \in \Omega_2(\mathbb{T}^2)$ since $\deg_2(\text{Id}) = 1$, 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 $X \setminus X_0$. 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 $F : W \rightarrow X : (s, x) \mapsto f(x)$. Then $\partial W \cong \emptyset \amalg (M \amalg M)$ and $F|_{M \amalg M} = f \amalg f$, hence (W, F) is a bordism between $(M \amalg M, f \amalg f)$ and (\emptyset, \cdot) .³³ \square

³²Note 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.

³³One 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 $\mathbb{H}^n \rightarrow \mathbb{H}^n$ 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 $\partial W \rightarrow M$ 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 is³⁴

$$\Omega_k^{\text{SO}}(X) := \{(M, f)\} / \sim,$$

where M is a closed (but not necessarily connected or nonempty) oriented k -manifold, $f : M \rightarrow X$ 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 $F : W \rightarrow X$ such that

$$\partial W \cong -M_- \amalg M_+$$

and $F|_{M_{\pm}} = f_{\pm}$. The group operation on $\Omega_k^{\text{SO}}(X)$ 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

$$\Omega_0^{\text{SO}}(X) \cong \bigoplus_{\pi_0(X)} \mathbb{Z},$$

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 $\Omega_0^{\text{SO}}(X) \cong \mathbb{Z}^3$. 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 $f : M \rightarrow X$. 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 $\gamma : I_{12} \rightarrow X$, 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 $F : W \rightarrow X$ 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 $M' := M \setminus \{x_1, x_2\}$ 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,

³⁴The ‘‘SO’’ in the notation $\Omega_k^{\text{SO}}(X)$ stands for the group $\text{SO}(k)$, the special orthogonal group. This has to do with the fact that $\text{SO}(k)$ is precisely the subgroup of $\text{O}(k)$ 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

$$\Omega_0(X) \rightarrow \mathbb{Z}_2^3 : [(M, f)] \mapsto (f_1, f_2, f_3)$$

is an isomorphism.

To understand $\Omega_0^{\text{SO}}(X)$, 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 $\epsilon : M \rightarrow \{1, -1\}$ 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 $M = \{x_1, \dots, x_N\}$ 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 $\gamma : I_{12} \rightarrow X_1$ with $\gamma(0) = f(x_1)$ and $\gamma(1) = f(x_2)$, and define $W = I_{12} \amalg I_3 \amalg \dots \amalg I_N$ and $F : W \rightarrow X$ 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 $j = 3, \dots, N$, 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 $M' = M \setminus \{x_1, x_2\}$ with the same orientations on the points x_3, \dots, 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

$$\Omega_0^{\text{SO}}(X) \rightarrow \mathbb{Z}^3 : [(M, f)] \mapsto (f_1, f_2, f_3)$$

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 $X = \{\text{pt}\}$, we abbreviate

$$\Omega_k := \Omega_k(\{\text{pt}\}), \quad \Omega_k^{\text{SO}} := \Omega_k^{\text{SO}}(\{\text{pt}\}),$$

and notice that since there is only one map from each manifold to $\{\text{pt}\}$, 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

$$\Omega_1 = \Omega_1^{\text{SO}} = 0.$$

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

$$\Omega_2^{\text{SO}} = \Omega_3^{\text{SO}} = 0.$$

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 $\mathbb{R}\mathbb{P}^2$. 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 $\mathbb{R}\mathbb{P}^2$, the complex projective space $\mathbb{C}\mathbb{P}^2$, is a closed oriented 4-manifold that never occurs as the boundary of any compact oriented 5-manifold. This implies

$$[\mathbb{R}\mathbb{P}^2] \neq 0 \in \Omega_2, \quad \text{and} \quad [\mathbb{C}\mathbb{P}^2] \neq 0 \in \Omega_4^{\text{SO}}.$$

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 $\{\text{pt}\}$, 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 $\Omega_n^{\text{SO}}(X)$, 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 $n \geq 0$, 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 k th 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

$$\partial_n : C_n(K; G) \rightarrow C_{n-1}(K; G) : \sum_i a_i \sigma_i \mapsto \sum_i a_i (\partial\sigma_i),$$

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 $n \geq 1$ 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 $\text{im } \partial_{n+1} \subset C_n(K; G)$ 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. \square

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 $\text{im } \partial_{n+1}$ is a subgroup of ∂_n for every $n \geq 0$. Since all these groups are abelian and subgroups are therefore normal, we can now consider quotients:

DEFINITION 21.15. The n th **simplicial homology** group of the complex K (with coefficients in G) is

$$H_n^\Delta(K; G) := \ker \partial_n / \text{im } \partial_{n+1}.$$

It is worth comparing this definition to the bordism groups $\Omega_n(X)$ and $\Omega_n^{\text{SO}}(X)$, 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 $f : M \rightarrow X$, 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 condition that two cycles $c, c' \in \ker \partial_n$ represent the same homology class $[c] = [c'] \in H_n^\Delta(K; G)$ if $c - c' \in \text{im } \partial_{n+1}$, 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 $\Omega_n^{\text{SO}}(X)$ 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 eighteen 2-simplices, twenty-seven 1-simplices, and nine vertices labeled as follows:

$$\begin{aligned} S_2 &= \{\sigma_1, \tau_1, \dots, \sigma_9, \tau_9\}, \\ S_1 &= \{a_1, a_2, a_3, b_1, b_2, b_3, \dots, f_1, f_2, f_3, g_1, \dots, g_9\}, \\ S_0 &= \{P_1, P_2, P_3, Q_1, Q_2, Q_3, R_1, R_2, R_3\}. \end{aligned}$$

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

$$\partial \sigma_1 = g_1 - a_1 - d_3, \quad \partial \tau_1 = b_1 + e_3 - g_1, \quad \partial a_1 = P_2 - P_1$$

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 nine generators P_i, Q_i, R_i for $i = 1, 2, 3$, but all nine of them are also homologous to each other since any pair of them can be connected by a path of oriented 1-simplices leading from one to the other, e.g. $\partial a_1 = P_2 - P_1$ implies $[P_1] = [P_2]$, and $\partial e_3 = P_2 - R_2$ implies $[P_2] = [R_2]$. 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 $\Omega_0^{\text{SO}}(\mathbb{T}^2)$ 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

$$\partial(a_3 + d_1 - c_3 - f_1) = 0,$$

so $a_3 + d_1 - c_3 - f_1$ is a 1-cycle, but not a very interesting one, since it is also the boundary of the region filled by the 2-simplices σ_9 and τ_9 : in particular,

$$\partial(\sigma_9 + \tau_9) = (g_9 - c_3 - f_1) + (a_3 + d_1 - g_9) = a_3 + d_1 - c_3 - f_1,$$

hence $[a_3 + d_1 - c_3 - f_1] = 0 \in H_1^\Delta(K; \mathbb{Z})$. 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 $a_1 + a_2 + a_3$, and we easily see that

$$\partial(a_1 + a_2 + a_3) = (P_2 - P_1) + (P_3 - P_2) + (P_1 - P_3) = 0.$$

Another is $b_1 + b_2 + b_3$, 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 six 2-simplices σ_i, τ_i for $i = 1, 2, 3$,

$$\partial(\sigma_1 + \sigma_2 + \sigma_3 + \tau_1 + \tau_2 + \tau_3) = (b_1 + b_2 + b_3) - (a_1 + a_2 + a_3),$$

implying $[a_1 + a_2 + a_3] = [b_1 + b_2 + b_3] \in H_1^\Delta(K; \mathbb{Z})$. Similar reasoning shows that $c_1 + c_2 + c_3$ is yet another 1-cycle representing the same homology class as both of these. 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 any of the three homologous 1-cycles $d_1 + d_2 + d_3$, $e_1 + e_2 + e_3$ or $f_1 + f_2 + f_3$. 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 $\Omega_2^{\text{SO}}(\mathbb{T}^2)$, 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:

$$\partial(\sigma_1 + \tau_1 + \dots + \sigma_9 + \tau_9) = 0.$$

The point is that since the triangulation is oriented, writing down each individual term in this sum would produce a linear combination of 1-simplices 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 eighteen 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

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.

³⁵This is what the conjecture was called in English—one does not translate the word *Hauptvermutung*.

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.

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 $\partial_n : C_n \rightarrow C_{n-1}$ for each $n \in \mathbb{Z}$ such that $\partial_{n-1} \circ \partial_n : C_n \rightarrow C_{n-2}$ 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 $\partial_n : C_n \rightarrow C_{n-1}$ extend uniquely to a homomorphism $\partial : C_* \rightarrow C_*$ 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

$$\partial : C_* \rightarrow C_{*-1}.$$

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 $\text{im } \partial_{n+1} \subset \ker \partial_n$ 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 $\text{im } \partial \subset C_*$ are called **boundaries** (*Ränder*).

DEFINITION 22.2. The **homology** (*Homologie*) of a chain complex (C_*, ∂) is the sequence of abelian groups

$$H_n(C_*, \partial) := \ker \partial_n / \text{im } \partial_{n+1}$$

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** (*Homologieklass*) represented by c . Two cycles $a, b \in \ker \partial_n$ are called **homologous** (*homolog*) if $[a] = [b] \in H_n(C_*, \partial)$, meaning $a - b \in \text{im } \partial_{n+1}$.

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 $f_n : A_n \rightarrow B_n$ for $n \in \mathbb{Z}$ such that the following diagram commutes:

$$(22.1) \quad \begin{array}{ccccccc} \dots & \longrightarrow & A_{n+1} & \xrightarrow{\partial_{n+1}^A} & A_n & \xrightarrow{\partial_n^A} & A_{n-1} & \xrightarrow{\partial_{n-1}^A} & \dots \\ & & \downarrow f_{n+1} & & \downarrow f_n & & \downarrow f_{n-1} & & \\ \dots & \longrightarrow & B_{n+1} & \xrightarrow{\partial_{n+1}^B} & B_n & \xrightarrow{\partial_n^B} & B_{n-1} & \xrightarrow{\partial_{n-1}^B} & \dots \end{array}$$

In other words, a chain map is a homomorphism $f : A_* \rightarrow B_*$ of degree zero satisfying $\partial^B \circ f = f \circ \partial^A$.

PROPOSITION 22.5. Any chain map $f : (A_*, \partial^A) \rightarrow (B_*, \partial^B)$ determines homomorphisms $f_* : H_n(A_*, \partial^A) \rightarrow H_n(B_*, \partial^B)$ 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 $\ker \partial_n^A / \text{im } \partial_{n+1}^A \rightarrow \ker \partial_n^B / \text{im } \partial_{n+1}^B$. This is equally clear, since $a = \partial^A x$ implies $f(a) = f(\partial^A x) = \partial^B f(x)$. \square

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 $n \geq 0$, the **standard n -simplex** is the set

$$\Delta^n = \{(t_0, \dots, t_n) \in I^{n+1} \mid t_0 + \dots + t_n = 1\}.$$

For each $k = 0, \dots, n$, the **k th 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

$$(22.2) \quad \partial_{(k)} \Delta^n \rightarrow \Delta^{n-1} : (t_0, \dots, t_{k-1}, 0, t_{k+1}, \dots, t_n) \mapsto (t_0, \dots, t_{k-1}, t_{k+1}, \dots, t_n).$$

DEFINITION 22.6. Given a topological space X , a **singular n -simplex** in X is a continuous map $\sigma : \Delta^n \rightarrow X$.

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 $\sum_i a_i \sigma_i$ 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 $\sigma : \Delta^n \rightarrow X$, 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 $C_n(X; G) \rightarrow C_{n-1}(X; G)$. 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

$$\partial\sigma := \sum_{k=0}^n (-1)^k (\sigma|_{\partial_{(k)}\Delta^n}) \in C_{n-1}(X; \mathbb{Z}),$$

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

$$\partial : C_n(X; G) \rightarrow C_{n-1}(X; G) : \sum_i a_i \sigma_i \mapsto \sum_i a_i \partial\sigma_i,$$

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

$$H_n(X; G) := H_n(C_*(X; G), \partial).$$

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 $\Omega_0^{\text{SO}}(X)$ 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 $\sigma : \Delta^0 \rightarrow X$ 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 $[0, 1] \rightarrow \Delta^1$ 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 $\sigma : I \rightarrow X$ 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 $\sigma : I \rightarrow X$ 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

³⁶The word “singular” in this context refers to the fact that there is no condition beyond continuity required for the maps $\sigma : \Delta^n \rightarrow X$, 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.

path-component and has two end points with opposite orientations, the conclusion is that both 0-cycles have the same coefficients. \square

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 $\{\text{pt}\}$ is much simpler.

EXERCISE 22.9. Show that for the 1-point space $\{\text{pt}\}$ and any coefficient group G , singular homology satisfies

$$H_n(\{\text{pt}\}; G) \cong \begin{cases} G & \text{for } n = 0, \\ 0 & \text{for } n \neq 0. \end{cases}$$

Hint: For each integer $n \geq 0$, there is exactly one singular n -simplex $\Delta^n \rightarrow \{\text{pt}\}$, so the chain groups $C_n(\{\text{pt}\}; G)$ are all naturally isomorphic to G . What is $\partial : C_n(\{\text{pt}\}; G) \rightarrow C_{n-1}(\{\text{pt}\}; G)$?

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

$$(22.3) \quad \{\text{paths } I \rightarrow X\} \leftrightarrow \mathcal{K}_1(X)$$

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

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

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 $\gamma : I \rightarrow X$ 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 $\gamma : S^1 \rightarrow \Sigma_g$ for $g \geq 2$ whose image separates Σ_g into two surfaces of genus $h \geq 1$ and $k \geq 1$ 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 $\Omega_1^{\text{SO}}(\Sigma_g)$. 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 $h : \pi_1(X) \rightarrow H_1(X; \mathbb{Z})$ is called the **Hurewicz map**. There exists a similar Hurewicz homomorphism $\pi_k(X) \rightarrow H_k(X; \mathbb{Z})$ for every $k \geq 1$, which we will discuss near the end of *Topologie II* if time permits. Note that for $k \geq 2$, $\pi_k(X)$ 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 $n \geq 2$, 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 $\gamma : I \rightarrow X$ with $\gamma(0) = \gamma(1) = x_0$. Identifying the standard 1-simplex

$$\Delta^1 := \{(t_0, t_1) \in \mathbb{R}^2 \mid t_0 + t_1 = 1, t_0, t_1 \geq 0\}$$

with $I := [0, 1]$ via the homeomorphism $\Delta^1 \rightarrow I : (t_0, t_1) \mapsto t_1$, every path $\gamma : I \rightarrow X$ corresponds to a singular 1-simplex $\Delta^1 \rightarrow X$, 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:

- If $\gamma : I \rightarrow X$ satisfies $\gamma(0) = \gamma(1)$, then $\partial \tilde{h}(\gamma) = 0$.
- For any constant path $e : I \rightarrow X$, $\tilde{h}(e) = \partial \sigma$ for some singular 2-simplex $\sigma : \Delta^2 \rightarrow X$.
- For any paths $\alpha, \beta : I \rightarrow X$ with $\alpha(1) = \beta(0)$, the concatenated path $\alpha \cdot \beta : I \rightarrow X$ satisfies $\tilde{h}(\alpha) + \tilde{h}(\beta) - \tilde{h}(\alpha \cdot \beta) = \partial \sigma$ for some singular 2-simplex $\sigma : \Delta^2 \rightarrow X$.

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 $\alpha, \beta : I \rightarrow X$ 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 $h : \pi_1(X) \rightarrow H_1(X; \mathbb{Z}) : [\gamma] \mapsto [\tilde{h}(\gamma)]$.

We call $h : \pi_1(X) \rightarrow H_1(X; \mathbb{Z})$ the **Hurewicz homomorphism**. Notice that since $H_1(X; \mathbb{Z})$ is abelian, $\ker h$ automatically contains the commutator subgroup $[\pi_1(X), \pi_1(X)] \subset \pi_1(X)$ (see Exercise 12.21), thus h descends to a homomorphism on the abelianization of $\pi_1(X)$,

$$\Phi : \pi_1(X) / [\pi_1(X), \pi_1(X)] \rightarrow H_1(X; \mathbb{Z}).$$

We will now show that this is an isomorphism by writing down its inverse. For each point $p \in X$, choose arbitrarily a path $\omega_p : I \rightarrow X$ from x_0 to p , and choose ω_{x_0} in particular to be the constant path. Regarding singular 1-simplices $\sigma : \Delta^1 \rightarrow X$ as paths $\sigma : I \rightarrow X$ 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

$$\tilde{\Psi}(\sigma) := \omega_{\sigma(0)} \cdot \sigma \cdot \omega_{\sigma(1)}^{-1} : I \rightarrow X$$

which begins and ends at the base point x_0 , hence $\tilde{\Psi}(\sigma)$ represents an element of $\pi_1(X)$. Let $\Psi(\sigma)$ denote the equivalence class represented by $\tilde{\Psi}(\sigma)$ in the abelianization $\pi_1(X) / [\pi_1(X), \pi_1(X)]$. This uniquely determines a homomorphism³⁷

$$\Psi : C_1(X; \mathbb{Z}) \rightarrow \pi_1(X) / [\pi_1(X), \pi_1(X)] : \sum_i m_i \sigma_i \mapsto \sum_i m_i \Psi(\sigma_i).$$

- (f) Show that $\Psi(\partial\sigma) = 0$ for every singular 2-simplex $\sigma : \Delta^2 \rightarrow X$, and deduce that Ψ descends to a homomorphism $\Psi : H_1(X; \mathbb{Z}) \rightarrow \pi_1(X) / [\pi_1(X), \pi_1(X)]$.
- (g) Show that $\Psi \circ \Phi$ and $\Phi \circ \Psi$ are both the identity map.
- (h) For a closed surface Σ_g of genus $g \geq 2$, find an example of a nontrivial element in the kernel of the Hurewicz homomorphism $\pi_1(\Sigma_g) \rightarrow H_1(\Sigma_g)$. *Hint: See Exercise 14.13.*

23. Relative homology and long exact sequences

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 $f : X \rightarrow Y$ determines a chain map $f_* : C_*(X; G) \rightarrow C_*(Y; G)$ via the formula $f_*\sigma := f \circ \sigma$ for singular n -simplices $\sigma : \Delta^n \rightarrow X$.*

PROOF. It is straightforward to check that $\partial(f_*\sigma) = f_*(\partial\sigma) \in C_{n-1}(Y; \mathbb{Z})$ for all $\sigma : \Delta^n \rightarrow X$, thus the uniquely determined homomorphism

$$f_* : C_n(X; G) \rightarrow C_n(Y; G) : \sum_i a_i \sigma_i \mapsto \sum_i a_i (f \circ \sigma_i)$$

defines a chain map. □

³⁷Since $\pi_1(X) / [\pi_1(X), \pi_1(X)]$ 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.

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 $f : X \rightarrow Y$ determine group homomorphisms $f_* : H_n(X; G) \rightarrow H_n(Y; G)$ 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 $(\text{Id})_* = 1$. \square*

REMARK 23.3. Recall that in the analogue of Corollary 23.2 for the fundamental group, the map $f : X \rightarrow Y$ 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 $f : X \rightarrow Y$.

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 $f : X \rightarrow Y$ is called a **map of pairs** if $f(A) \subset B$, and in this case we write

$$f : (X, A) \rightarrow (Y, B).$$

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 $f, g : (X, A) \rightarrow (Y, B)$ are **homotopic** if there exists a homotopy $H : I \times X \rightarrow Y$ between f and g such that $H(s, \cdot) : (X, A) \rightarrow (Y, B)$ 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 $f : (X, A) \rightarrow (Y, B)$ and $g : (Y, B) \rightarrow (X, A)$ 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, \emptyset) , 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 $\partial : C_n(X; G) \rightarrow C_{n-1}(X; G)$ 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*),

$$H_n(X, A; G) := H_n(C_*(X, A; G), \partial).$$

The case $A = \emptyset$ 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 $f : (X, A) \rightarrow (Y, B)$ is a map of pairs, then the absolute chain map $f_* : C_*(X; G) \rightarrow C_*(Y; G)$ sends the subgroup $C_*(A; G)$ into $C_*(B; G)$ and thus descends to a chain map

$$f_* : C_*(X, A; G) \rightarrow C_*(Y, B; G),$$

implying the relative version of Corollary 23.2:

THEOREM 23.4. *Maps of pairs $f : (X, A) \rightarrow (Y, B)$ determine group homomorphisms $f_* : H_n(X, A; G) \rightarrow H_n(Y, B; G)$ 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 $[c] \in H_n(X, A; G)$, 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

$$b - c = a + \partial x \quad \text{for some } a \in C_n(A; G), \quad x \in C_{n+1}(X; G).$$

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 $j : X = (X, \emptyset) \hookrightarrow (X, A)$ denote the natural inclusions,³⁸ and consider the sequence of chain maps

$$(23.1) \quad 0 \longrightarrow C_*(A; G) \xrightarrow{i_*} C_*(X; G) \xrightarrow{j_*} C_*(X, A; G) \longrightarrow 0,$$

³⁸Strictly 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 $\emptyset \subset A$.

where the first and last maps are each trivial. The map j_* is obviously surjective, as it is actually just the quotient projection

$$C_*(X; G) \rightarrow C_*(X, G)/C_*(A; G) = C_*(X, A; G).$$

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

$$\dots \rightarrow A_{n-2} \xrightarrow{f_{n-2}} A_{n-1} \xrightarrow{f_{n-1}} A_n \xrightarrow{f_n} A_{n+1} \xrightarrow{f_{n+1}} A_{n+2} \rightarrow \dots$$

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

$$0 \rightarrow A_1 \xrightarrow{f} A_2 \rightarrow 0,$$

and the exactness of the sequence then simply means that $f : A_1 \rightarrow A_2$ 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,

$$0 \rightarrow A_1 \xrightarrow{f_1} A_2 \xrightarrow{f_2} A_3 \rightarrow 0.$$

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*

$$0 \rightarrow A_* \xrightarrow{f} B_* \xrightarrow{g} C_* \rightarrow 0$$

is a short exact sequence of chain maps. Then there exists a natural homomorphism $\partial_ : H_n(C_*, \partial^C) \rightarrow H_{n-1}(A_*, \partial^A)$ for each $n \in \mathbb{Z}$ such that the sequence*

$$(23.2) \quad \begin{aligned} \dots \xrightarrow{\partial_*} H_{n+1}(A_*, \partial^A) \xrightarrow{f_*} H_{n+1}(B_*, \partial^B) \xrightarrow{g_*} H_{n+1}(C_*, \partial^C) \\ \xrightarrow{\partial_*} H_n(A_*, \partial^A) \xrightarrow{f_*} H_n(B_*, \partial^B) \xrightarrow{g_*} H_n(C_*, \partial^C) \\ \xrightarrow{\partial_*} H_{n-1}(A_*, \partial^A) \xrightarrow{f_*} H_{n-1}(B_*, \partial^B) \xrightarrow{g_*} H_{n-1}(C_*, \partial^C) \xrightarrow{\partial_*} \dots \end{aligned}$$

is exact.

The sequence of homology groups in this theorem is called a **long exact sequence** (*lange exakte Sequenz*), and the maps $\partial_* : H_n(C_*, \partial^C) \rightarrow H_{n-1}(A_*, \partial^A)$ 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)** ,

$$(23.3) \quad \dots \rightarrow H_{n+1}(X, A; G) \xrightarrow{\partial_*} H_n(A; G) \xrightarrow{i_*} H_n(X; G) \xrightarrow{j_*} H_n(X, A; G) \xrightarrow{\partial_*} H_{n-1}(A; G) \rightarrow \dots$$

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 $i_* : H_n(A; G) \rightarrow H_n(X; G)$ 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 $n \geq 1$.

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:

$$\begin{array}{ccccccc}
 & & \vdots & & \vdots & & \vdots \\
 & & \downarrow \partial^A & & \downarrow \partial^B & & \downarrow \partial^C \\
 0 & \longrightarrow & A_{n+1} & \xrightarrow{f} & B_{n+1} & \xrightarrow{g} & C_{n+1} \longrightarrow 0 \\
 & & \downarrow \partial^A & & \downarrow \partial^B & & \downarrow \partial^C \\
 0 & \longrightarrow & A_n & \xrightarrow{f} & B_n & \xrightarrow{g} & C_n \longrightarrow 0 \\
 & & \downarrow \partial^A & & \downarrow \partial^B & & \downarrow \partial^C \\
 0 & \longrightarrow & A_{n-1} & \xrightarrow{f} & B_{n-1} & \xrightarrow{g} & C_{n-1} \longrightarrow 0 \\
 & & \downarrow \partial^A & & \downarrow \partial^B & & \downarrow \partial^C \\
 0 & \longrightarrow & A_{n-2} & \xrightarrow{f} & B_{n-2} & \xrightarrow{g} & C_{n-2} \longrightarrow 0 \\
 & & \downarrow \partial^A & & \downarrow \partial^B & & \downarrow \partial^C \\
 & & \vdots & & \vdots & & \vdots
 \end{array}$$

We start by writing down a reasonable candidate for the map $\partial_* : H_n(C_*, \partial^C) \rightarrow H_{n-1}(A_*, \partial^A)$. Given $[c] \in H_n(C_*, \partial^C)$, $c \in C_n$ is necessarily a cycle, and exactness tells us that $g : B_n \rightarrow C_n$ 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

$$f(\partial^A a) = \partial^B (f(a)) = \partial^B \partial^B b = 0$$

³⁹I 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.

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 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

$$\partial_* : H_n(X, A; G) \rightarrow H_{n-1}(A; G).$$

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:

$$(23.4) \quad \partial_*[c] = [\partial c].$$

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 $\Omega_n^{\text{SO}}(X)$ are equivalence classes of maps $f : M \rightarrow X$ 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 $\Delta^k \rightarrow M$ 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 $f : M \rightarrow X$ 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 $[(M, f)] = 0 \in \Omega_n^{\text{SO}}(X)$, it means there exists a compact oriented $(n + 1)$ -manifold W with $\partial W \cong M$ and a map $F : W \rightarrow X$ 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 $n \geq 0$ and $A \in H_n(X; \mathbb{Z})$, there exists a closed n -manifold M , a map $f : M \rightarrow X$ and a number $k \in \mathbb{N}$ such that $kA = f_*[M]$. \square*

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 fundamental class**

$$[M] \in H_n(M, \partial M; \mathbb{Z}).$$

Given a pair (X, A) , any map $f : (M, \partial M) \rightarrow (X, A)$ 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

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 $f_* : H_n(X, A; G) \rightarrow H_n(Y, B; G)$ induced for each $n \in \mathbb{Z}$ by a map of pairs $f : (X, A) \rightarrow (Y, B)$ 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 $A = \emptyset$, in which case the terms “homotopy” and “homotopy equivalence” mean the same thing that they always did.

COROLLARY 24.2. *If $f : (X, A) \rightarrow (Y, B)$ is a homotopy equivalence of pairs, then the induced maps $f_* : H_n(X, A; G) \rightarrow H_n(Y, B; G)$ are isomorphisms.*

PROOF. Suppose $f : (X, A) \rightarrow (Y, B)$ is a homotopy equivalence, so it has a homotopy inverse $g : (Y, B) \rightarrow (X, A)$. 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. \square

The proof of Theorem 24.1 requires another fundamental notion from homological algebra. It should be clear that if $f, g : X \rightarrow Y$ are two non-identical maps, then the induced chain maps $f_*, g_* : C_*(X; G) \rightarrow C_*(Y; G)$ 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 $f, g : (A_*, \partial^A) \rightarrow (B_*, \partial^B)$ is a sequence of homomorphisms $h_n : A_n \rightarrow B_{n+1}$ such that for every $n \in \mathbb{Z}$,

$$f_n - g_n = \partial_{n+1}^B \circ h_n + h_{n-1} \circ \partial_n^A.$$

In other words, a chain homotopy between f and g is a homomorphism $h : A_* \rightarrow B_*$ of degree +1 such that $f - g = \partial^B \circ h + h \circ \partial^A$. We sometimes abuse notation and write

$$h : A_* \rightarrow B_{*+1}$$

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*

$$f_* = g_* : H_n(A_*, \partial^A) \rightarrow H_n(B_*, \partial^B)$$

for all $n \in \mathbb{Z}$.

PROOF. If $h : A_* \rightarrow B_{*+1}$ 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. \square

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 $f, g : (X, A) \rightarrow (Y, B)$, then there also exists a chain homotopy between the induced chain maps $f_*, g_* : C_*(X, A; G) \rightarrow C_*(Y, B; G)$.*

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

$$A = B = \emptyset,$$

as the general case will only require a few extra remarks beyond this. Suppose $H : I \times X \rightarrow Y$ is a homotopy between $f = H(0, \cdot)$ and $g = H(1, \cdot)$. Associate to each singular n -simplex $\sigma : \Delta^n \rightarrow X$ the map

$$h_\sigma : I \times \Delta^n \rightarrow Y : (s, t) \mapsto H(s, \sigma(t)),$$

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 be⁴⁰

$$(24.1) \quad \partial(I \times \Delta^n) = (-\{0\} \times \Delta^n) \cup (\{1\} \times \Delta^n) \cup (-I \times \partial\Delta^n).$$

This relation will be the geometric motivation behind the chain homotopy formula.

The idea now is to define a chain homotopy $h : C_*(X; G) \rightarrow C_{*+1}(Y; G)$ by associating to each singular n -simplex $\sigma : \Delta^n \rightarrow X$ a linear combination of singular $(n + 1)$ -simplices in Y determined by the prism map $h_\sigma : I \times \Delta^n \rightarrow Y$. 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 $\tau_i : \Delta^{n+1} \rightarrow I \times \Delta^n$ 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 $\Delta^n \rightarrow \{1\} \times \Delta^n$, whose attached coefficient we can assume without loss of generality is $+1$;
- (2) Another term for the obvious parametrization $\Delta^n \rightarrow \{0\} \times \Delta^n$, whose attached coefficient must now be -1 for orientation reasons;

⁴⁰One 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.

- (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 $h : C_n(X; G) \rightarrow C_{n+1}(Y; G)$ defined on each singular n -simplex $\sigma : \Delta^n \rightarrow X$ 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 $\tau_i : \Delta^{n+1} \rightarrow I \times \Delta^n$ 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

$$\partial h(\sigma) = g_* \sigma - f_* \sigma - h(\partial \sigma),$$

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 $h : C_*(X; G) \rightarrow C_{*+1}(Y; G)$ 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

$$h_* : C_*(X, A; G) \rightarrow C_{*+1}(Y, B; G)$$

between the relative chain maps $f_*, g_* : C_*(X, A; G) \rightarrow C_*(Y, B; G)$. 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 $n \neq 0$.* \square

COROLLARY 24.8 (via Theorem 22.10). *If X is homotopy equivalent to S^1 , then $H_1(X; \mathbb{Z}) \cong \mathbb{Z}$.* \square

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

$$H_*(X \setminus B, A \setminus B; G) \cong H_*(X, A; G).$$

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 $i : (X \setminus B, A \setminus B) \hookrightarrow (X, A)$ induces isomorphisms*

$$i_* : H_n(X \setminus B, A \setminus B; G) \xrightarrow{\cong} H_n(X, A; G)$$

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 $X \setminus \bar{B}$ 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 $X \setminus B$, i.e. c belongs to the subgroup $C_n(A; G) + C_n(X \setminus B; G) \subset C_n(X; G)$. 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 $c = c_A + c_{X \setminus B}$ for

some $c_A \in C_n(A; G)$ and $C_{X \setminus B} \in C_n(X \setminus B; G)$, then since c is a relative cycle, $\partial c \in C_{n-1}(A; G)$, implying $\partial c_{X \setminus B}$ is also in $C_{n-1}(A; G)$ since ∂c_A must be as well, thus $\partial c_{X \setminus B} \in C_{n-1}(A \setminus B; G)$. This proves that $c_{X \setminus B}$ is a relative n -cycle for the pair $(X \setminus B, A \setminus B)$, so it represents a homology class in $H_n(X \setminus B, A \setminus B; G)$, and obviously

$$i_*[c_{X \setminus B}] = [c]$$

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 ,

$$X = \overset{\circ}{\mathcal{U}} \cup \overset{\circ}{\mathcal{V}}.$$

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 $f : M \rightarrow X$. 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 $m \leq n - 1$. 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 $(\Delta^n, \partial\Delta^n)$,

$$\sum_i \epsilon_i \tau_i \in C_n(\Delta^n; \mathbb{Z}), \quad \partial \sum_i \epsilon_i \tau_i \in C_{n-1}(\partial\Delta^n; \mathbb{Z}),$$

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

$$S : C_n(X; G) \rightarrow C_n(X; G)$$

via the formula

$$S(\sigma) := \sum_i \epsilon_i (\sigma \circ \tau_i)$$

for each $n \geq 0$ and $\sigma : \Delta^n \rightarrow X$. 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. $S : C_*(X; G) \rightarrow C_*(X; G)$ is a chain map.

PROOF. This follows from the relation $\partial S(\sigma) = S(\partial\sigma)$ for each $\sigma : \Delta^n \rightarrow X$, 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. \square

LEMMA 24.11. $S : C_*(X; G) \rightarrow C_*(X; G)$ 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 $\tau_i : \Delta^{n+1} \rightarrow I \times \Delta^n$ are parametrizations of the simplices in the triangulation and the signs $\epsilon_i = \pm 1$ are determined by their orientations. Let

$$\pi : I \times \Delta^n \rightarrow \Delta^n$$

denote the obvious projection map. The desired chain homotopy $h : C_n(X; G) \rightarrow C_{n+1}(X; G)$ is then determined by the formula

$$h(\sigma) = \sum_i \epsilon_i (\sigma \circ \pi \circ \tau_i).$$

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 $S - \mathbb{1} = \partial h + h\partial$. \square

The chain homotopy result implies that our subdivision map $S : C_*(X; G) \rightarrow C_*(X; G)$ has the main property we want, namely it induces the identity homomorphism $H_*(X; G) \rightarrow H_*(X; G)$, 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 $m \geq 1$, 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 $S^m c \in C_*(U; G) + C_*(V; G)$ 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 $X \setminus B$ 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 $X \setminus B$ 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,

$$c = c_A + c_{X \setminus B} \quad \text{for some} \quad c_A \in C_n(A; G), \quad c_{X \setminus B} \in C_n(X \setminus B; G).$$

Having made this assumption, the reason why $i_* : H_n(X \setminus B, A \setminus B; G) \rightarrow H_n(X, A; G)$ 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 $\partial c_{X \setminus B} \in C_n(A; G)$, so that $c_{X \setminus B}$ is a relative n -cycle in $(X \setminus B, A \setminus B)$, thus representing a class $[c_{X \setminus B}] \in H_n(X \setminus B, A \setminus B; G)$ that satisfies

$$i_*[c_{X \setminus B}] = [c].$$

The proof that $i_* : H_n(X \setminus B, A \setminus B; G) \rightarrow H_n(X, A; G)$ is injective uses subdivision in a slightly different way. Suppose $c \in C_n(X \setminus B; G)$ is a relative n -cycle representing a homology class $[c] \in H_n(X \setminus B, A \setminus B; G)$ 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 $X \setminus B$, c is a boundary of some $(n+1)$ -chain in X , modulo one that is contained in A , i.e. we have

$$c = \partial b + a \quad \text{for some } b \in C_{n+1}(X; G) \text{ and } a \in C_n(A; G).$$

Applying ∂ to both sides of this equation gives $\partial c = \partial a$, which implies since c is a relative n -cycle in $(X \setminus B, A \setminus B)$ that $\partial a \in C_n(A \setminus B; G)$, 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 $X \setminus B$ and a as an n -chain in $A \setminus B$, so that the relation $c = \partial b + a$ also implies $[c] = 0 \in H_n(X \setminus B, A \setminus B; G)$. 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 $[c] \in H_n(X \setminus B, A \setminus B; G)$ does not change if we replace c with $S^m c$ for any $m \geq 1$, 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 $c \in C_n(X \setminus B; G)$ and $a \in C_n(A; G)$ this is not new information since we already assumed them to be contained in $X \setminus B$ or A respectively, but for $b \in C_{n+1}(X; G)$ we can now write

$$b = b_A + b_{X \setminus B} \quad \text{for some} \quad b_A \in C_{n+1}(A; G), \quad b_{X \setminus B} \in C_{n+1}(X \setminus B; G).$$

The relation $c = \partial b + a$ thus becomes

$$c = \partial b_{X \setminus B} + (\partial b_A + a),$$

and we observe that since c and $\partial b_{X \setminus B}$ are both n -chains in $X \setminus B$, the same must therefore be true for $\partial b_A + a$, meaning it is actually contained in $A \setminus B$. This proves $[c] = 0 \in H_n(X \setminus B, A \setminus B; G)$. \square

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 $i : (X \setminus B, A \setminus B) \hookrightarrow (X, A)$ induces a **chain homotopy equivalence** (*Kettenhomotopieäquivalenz*)

$$i_* : C_*(X \setminus B, A \setminus B; G) \rightarrow C_*(X, A; G).$$

In case the meaning of this terminology is not obvious, this means there exists a chain map $\psi : C_*(X, A; G) \rightarrow C_*(X \setminus B, A \setminus B; G)$ 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 $U, V \subset X$, the subgroup $C_*(U; G) + C_*(V; G) \subset C_*(X; G)$ is also a chain complex in a natural way. Indeed, the boundary operator on $C_*(X; G)$ maps each of $C_*(U; G)$ and $C_*(V; G)$ to themselves, thus it also preserves their sum.

LEMMA 24.12. *For any subsets $U, V \subset X$ with $X = \overset{\circ}{U} \cup \overset{\circ}{V}$, the inclusion map*

$$j : C_*(U; G) + C_*(V; G) \hookrightarrow C_*(X; G)$$

admits a chain homotopy inverse

$$\rho : C_*(X; G) \rightarrow C_*(U; G) + C_*(V; G)$$

such that $\rho \circ j = \mathbb{1}$, and moreover, there is a chain homotopy $h : C_(X; G) \rightarrow C_{*+1}(X; G)$ of $j \circ \rho$ to the identity such that h vanishes on $C_*(U; G) + C_*(V; G)$.*

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 $C_*(U; G) + C_*(V; G)$ 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 $\sigma : \Delta^n \rightarrow X$ whose boundary is close enough to the boundary of U or 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 $C_*(U; G) + C_*(V; G)$, 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 $h_1 : C_*(X; G) \rightarrow C_{*+1}(X; G)$ denote the chain homotopy provided by Lemma 24.11 for the barycentric subdivision chain map $S : C_*(X; G) \rightarrow C_*(X; G)$, i.e. it satisfies $S - \mathbb{1} = \partial h_1 + h_1 \partial$.

We claim that for all integers $m \geq 0$, the map

$$h_m := h_1 \sum_{k=0}^{m-1} S^k : C_*(X; G) \rightarrow C_{*+1}(X; G)$$

then satisfies

$$(24.2) \quad S^m - \mathbb{1} = \partial h_m + h_m \partial,$$

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 $m - 1 \geq 1$, 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 $\sigma : \Delta^n \rightarrow X$, the iterated subdivision maps S^m can be assumed to satisfy

$$(24.3) \quad S^m(\sigma) \in C_*(\mathcal{U}; G) + C_*(\mathcal{V}; G)$$

if m is large enough, so for each $n \geq 0$ and $\sigma : \Delta^n \rightarrow X$, let $m_\sigma \geq 0$ denote the smallest integer for which (24.3) holds with $m = m_\sigma$. We can then define a homomorphism $h : C_n(X; G) \rightarrow C_{n+1}(X; G)$ for each $n \geq 0$ via

$$h(\sigma) := h_{m_\sigma}(\sigma).$$

Let us see whether this is a chain homotopy. We have

$$\begin{aligned} (\partial h + h \partial)(\sigma) &= \partial h_{m_\sigma}(\sigma) + h_{m_\sigma}(\partial \sigma) + (h - h_{m_\sigma})(\partial \sigma) \\ &= (S^{m_\sigma} - \mathbb{1})(\sigma) + (h - h_{m_\sigma})(\partial \sigma) = ([S^{m_\sigma} + (h - h_{m_\sigma})\partial] - \mathbb{1})(\sigma). \end{aligned}$$

Use this to define $\rho : C_*(X; G) \rightarrow C_*(X; G)$ by

$$\rho(\sigma) := S^{m_\sigma}(\sigma) + (h - h_{m_\sigma})(\partial \sigma),$$

so the relation

$$(24.4) \quad \partial h + h \partial = \rho - \mathbb{1}$$

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 $m_\tau \leq m_\sigma$ 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 $\sigma : \Delta^n \rightarrow X$, $\rho(\sigma)$ 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 $m_\tau \leq m_\sigma$, 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 $h : C_*(X; G) \rightarrow C_{*+1}(X; G)$ vanishes on $C_*(\mathcal{U}; G) + C_*(\mathcal{V}; G)$ since every singular n -simplex $\sigma : \Delta^n \rightarrow X$ with image in either \mathcal{U} or \mathcal{V} satisfies $m_\sigma = 0$, thus $h(\sigma) = h_{m_\sigma}(\sigma) = h_0(\sigma) = 0$. \square

Now we can prove the “chain level” result that implies Theorem 24.9.

LEMMA 24.13. *If $A, B \subset X$ are subsets with $\bar{B} \subset \overset{\circ}{A}$, then the inclusion $i : (X \setminus B, A \setminus B) \hookrightarrow (X, A)$ induces a chain homotopy equivalence $i_* : C_*(X \setminus B, A \setminus B; G) \rightarrow C_*(X, A; G)$.*

PROOF. Consider the quotient chain complex $(C_*(X \setminus B; G) + C_*(A; G))/C_*(A; G)$, 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 $\sigma_i : \Delta^n \rightarrow X$ that have image in $X \setminus B$ but not contained in A . The point here is that while simplices with $\sigma(\Delta^n) \subset A$ are also generators of $C_*(X \setminus B; G) + C_*(A; G)$, they are all equivalent to zero in the quotient. As it happens, the quotient complex $C_*(X \setminus B, A \setminus B; G) = C_*(X \setminus B; G)/C_*(A \setminus B; G)$ can be described in exactly the same way, with the same set of generators: singular simplices that are contained in $X \setminus B$ but not contained in A . Since the obvious inclusion $C_*(X \setminus B; G) \hookrightarrow C_*(X \setminus B; G) + C_*(A; G)$ sends $C_*(A \setminus B; G)$ into $C_*(A; G)$, it follows that this inclusion descends to a chain map of quotient complexes

$$C_*(X \setminus B, A \setminus B; G) \rightarrow (C_*(X \setminus B; G) + C_*(A; G))/C_*(A; G)$$

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

$$(24.5) \quad (C_*(X \setminus B; G) + C_*(A; G))/C_*(A; G) \xrightarrow{j} C_*(X; G)/C_*(A; G) = C_*(X, A; G)$$

induced on these quotients by the obvious inclusion

$$C_*(X \setminus B; G) + C_*(A; G) \xrightarrow{j} C_*(X; G)$$

is a chain homotopy equivalence. Since $X \setminus \bar{B}$ and $\overset{\circ}{A}$ form an open cover of X , Lemma 24.12 provides a chain homotopy inverse for j , namely the map $\rho : C_*(X; G) \rightarrow C_*(X \setminus B; G) + C_*(A; G)$, 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 $h : C_*(X; G) \rightarrow C_{*+1}(X; G)$ that vanishes on $C_*(A; G)$, thus h also descends to the quotient $C_*(X; G)/C_*(A; G)$ as a chain homotopy $h : C_*(X, A; G) \rightarrow C_{*+1}(X, A; G)$ satisfying $j \circ \rho - \mathbb{1} = \partial h + h \partial$ on the quotient complexes. \square

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 $f : V \rightarrow V'$ 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 $|K| \rightarrow |K'|$ which maps each n -simplex in $|K|$ linearly to a k -simplex in $|K'|$ for some $k \leq n$. It is not hard to show that f also naturally induces a chain map $f_* : C_*(K; G) \rightarrow C_*(K'; G)$, 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 $H_*^\Delta(K; G) \rightarrow H_*^\Delta(K'; G)$. 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 $S : C_*(K; G) \rightarrow C_*(K'; G)$ 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 $S_* : H_*^\Delta(K; G) \rightarrow H_*^\Delta(K'; G)$ 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

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** (*Einhangung*) SX of a space X is defined by gluing together two copies of its cone,

$$(25.1) \quad SX = C_+X \cup_X C_-X,$$

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

$$(C_+X, X) \hookrightarrow (SX \setminus \{p_-\}, C_-X \setminus \{p_-\})$$

is a homotopy equivalence of pairs. Indeed, one can define a deformation retraction $H : I \times (SX \setminus \{p_-\}) \rightarrow SX \setminus \{p_-\}$ by pushing points in $C_-X \setminus \{p_-\}$ continuously upward toward X while leaving C_+X fixed, so that $H(1, \cdot)$ is the identity while $H(0, \cdot)$ retracts $SX \setminus \{p_-\}$ to C_+X and $H(s, \cdot)$ preserves $C_-X \setminus \{p_-\}$ for every $s \in I$. The resulting retraction of pairs $(SX \setminus \{p_-\}, C_-X \setminus \{p_-\}) \rightarrow (C_+X, X)$ 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

$$(\mathbb{D}_+^n, S^{n-1}) \hookrightarrow (S^n \setminus \{p_-\}, \mathbb{D}_-^n \setminus \{p_-\}),$$

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:

$$(25.2) \quad \begin{array}{ccc} H_k(X; G) & & H_{k+1}(SX; G) \\ \partial_* \uparrow & & \downarrow \varphi_* \\ H_{k+1}(C_+X, X; G) & \xrightarrow{i_*} & H_{k+1}(SX \setminus \{p_-\}, C_-X \setminus \{p_-\}; G) \xrightarrow{j_*} & H_{k+1}(SX, C_-X; G) \end{array}$$

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

$$\begin{aligned} (SX \setminus \{p_-\}, C_-X \setminus \{p_-\}) &\xrightarrow{j} (SX, C_-X), \\ (SX, \emptyset) &\xrightarrow{\varphi} (SX, C_-X). \end{aligned}$$

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 $k \geq 1$, then ∂_* and φ_* are both also isomorphisms. For the first, consider the long exact sequence of the pair (C_+X, X) :

$$\dots \longrightarrow H_{k+1}(C_+X; G) \longrightarrow H_{k+1}(C_+X, X; G) \xrightarrow{\partial_*} H_k(X; G) \longrightarrow H_k(C_+X; G) \longrightarrow \dots$$

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 $k \geq 1$, so exactness implies that ∂_* is an isomorphism. For φ_* , we instead take an excerpt from the long exact sequence of (SX, C_-X) :

$$\dots \longrightarrow H_{k+1}(C_-X; G) \longrightarrow H_{k+1}(SX; G) \xrightarrow{\varphi_*} H_{k+1}(SX, C_-X; G) \longrightarrow H_k(C_-X; G) \longrightarrow \dots$$

The contractibility of C_-X again makes the first and last terms vanish if $k \geq 1$, 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 $k \geq 1$, the diagram (25.2) defines an isomorphism*

$$S_* = \varphi_*^{-1} \circ j_* \circ i_* \circ \partial_*^{-1} : H_k(X; G) \rightarrow H_{k+1}(SX; G).$$

□

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

$$(25.3) \quad \partial c_+ = -\partial c_- = b$$

and

$$(25.4) \quad S_*[b] = [c_+ + c_-].$$

Note that $c_+ + c_- \in C_{k+1}(SX; G)$ 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 $n \geq 1$, 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 $n \geq 1$:

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 $k = 1, \dots, n-1$, 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. \square

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 $f : \mathbb{D}^n \rightarrow \mathbb{D}^n$ has no fixed point, then we can use it to define a retraction $g : \mathbb{D}^n \rightarrow S^{n-1} = \partial\mathbb{D}^n$. 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 $X \rightarrow Y$ induce homomorphisms $\pi_1(X) \rightarrow \pi_1(Y)$ 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 $f : X \rightarrow A$ 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 , $f_* : H_n(X; G) \rightarrow H_n(A; G)$ is surjective, while $i_* : H_n(A; G) \rightarrow H_n(X; G)$ is injective.

PROOF OF THE BROUWER FIXED POINT THEOREM. Arguing by contradiction, assume a map $f : \mathbb{D}^n \rightarrow \mathbb{D}^n$ without fixed points exists, and therefore also a retraction $g : \mathbb{D}^n \rightarrow S^{n-1}$. We may assume $n \geq 2$ 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

$$g_* : H_{n-1}(\mathbb{D}^n; \mathbb{Z}) \rightarrow H_{n-1}(S^{n-1}; \mathbb{Z}).$$

But this is impossible since $H_{n-1}(\mathbb{D}^n; \mathbb{Z}) \cong H_{n-1}(\{\text{pt}\}; \mathbb{Z}) = 0$ and $H_{n-1}(S^{n-1}; \mathbb{Z}) \cong \mathbb{Z}$. \square

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 $f : \Omega_n \rightarrow \Omega_m$ with $f(0) = 0$. Choose $\epsilon > 0$ small enough so that f maps the ϵ -ball $B_\epsilon^n(0) \subset \Omega_n$ about the origin into the δ -ball $B_\delta^m(0) \subset \Omega_m$ for some $\delta > 0$, where the latter is also small enough so that $B_\delta^m(0) \subset \Omega_m$. Now pick a generator

$$A \in H_{n-1}(B_\epsilon^n(0) \setminus \{0\}; \mathbb{Z}) \cong H_{n-1}(S^{n-1}; \mathbb{Z}) \cong \mathbb{Z}.$$

Since $m \neq n$,

$$H_{n-1}(B_\delta^m(0) \setminus \{0\}; \mathbb{Z}) \cong H_{n-1}(S^{m-1}; \mathbb{Z}) = 0,$$

so restricting f to a map $B_\epsilon^n(0) \setminus \{0\} \rightarrow B_\delta^m(0) \setminus \{0\}$ gives $f_*A = 0 \in H_{n-1}(B_\delta^m(0) \setminus \{0\}; \mathbb{Z})$. But f^{-1} is also defined on $B_\delta^m(0)$, 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 $H_{n-1}(B_\epsilon^n(0) \setminus \{0\}; \mathbb{Z}) \neq 0$. \square

26. Axioms, cells, and the Euler characteristic

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, \emptyset) as usual, one abbreviates

$$h_n(X) := h_n(X, \emptyset).$$

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 $f : (X, A) \rightarrow (Y, B)$ a sequence of homomorphisms

$$f_* : h_n(X, A) \rightarrow h_n(Y, B), \quad n \in \mathbb{Z}$$

with the properties that $(f \circ g)_* = f_* \circ g_*$ whenever the composition of f and g makes sense, and the identity map $\text{Id} : (X, A) \rightarrow (X, A)$ gives rise to the identity homomorphism $\text{Id}_* = \mathbf{1} : h_n(X, A) \rightarrow h_n(X, A)$. 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*

$$\partial_* : h_n(X, A) \rightarrow h_{n-1}(A), \quad n \in \mathbb{Z}.$$

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 $f : (X, A) \rightarrow (Y, B)$ restricts to a continuous map $A \rightarrow B$, so it induces homomorphisms

$f_* : h_n(X, A) \rightarrow h_n(Y, B)$ and $f_* : h_n(A) \rightarrow h_n(B)$, which we would like to fit together with ∂_* into the following commutative diagram for each n :

$$\begin{array}{ccc} h_n(X, A) & \xrightarrow{\partial_*} & h_{n-1}(A) \\ \downarrow f_* & & \downarrow f_* \\ h_n(Y, B) & \xrightarrow{\partial_*} & h_{n-1}(B) \end{array}$$

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 $(X, A) \mapsto h_n(A) := h_n(A, \emptyset)$. 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) $f_* : h_*(X, A) \rightarrow h_*(Y, B)$ depends only on the homotopy class of $f : (X, A) \rightarrow (Y, B)$.
- (EXACTNESS) For the inclusions $i : A \hookrightarrow X$ and $j : (X, \emptyset) \hookrightarrow (X, A)$, the sequence

$$\dots \longrightarrow h_{n+1}(X, A) \xrightarrow{\partial_*} h_n(A) \xrightarrow{i_*} h_n(X) \xrightarrow{j_*} h_n(X, A) \xrightarrow{\partial_*} h_{n-1}(A) \longrightarrow \dots$$

is exact.

- (EXCISION) If $B \subset \bar{B} \subset \overset{\circ}{A} \subset A \subset X$, then the inclusion $(X \setminus B, A \setminus B) \hookrightarrow (X, A)$ induces an isomorphism $h_*(X \setminus B, A \setminus B) \rightarrow h_*(X, A)$.
- (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 $i_*^\alpha : h_*(X_\alpha) \rightarrow h_*(\coprod_{\beta \in J} X_\beta)$ determine an isomorphism

$$\bigoplus_{\alpha \in J} h_*(X_\alpha) \rightarrow h_*\left(\coprod_{\alpha \in J} X_\alpha\right).$$

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 $\partial_* : H_n(X, A) \rightarrow H_{n-1}(A)$ in singular homology satisfy naturality, i.e. for any map $f : (X, A) \rightarrow (Y, B)$ and every $n \in \mathbb{Z}$, the

⁴¹One 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.

diagram

$$\begin{array}{ccc} H_n(X, A) & \xrightarrow{\hat{c}_*} & H_{n-1}(A) \\ \downarrow f_* & & \downarrow f_* \\ H_n(Y, B) & \xrightarrow{\hat{c}_*} & H_{n-1}(B) \end{array}$$

commutes.

- (b) Deduce that for any map $f : (X, A) \rightarrow (Y, B)$, the long exact sequences of (X, A) and (Y, B) in singular homology form the rows of a commutative diagram

$$\begin{array}{cccccccc} \dots & \longrightarrow & H_n(A) & \longrightarrow & H_n(X) & \longrightarrow & H_n(X, A) & \longrightarrow & H_{n-1}(A) & \longrightarrow & \dots \\ & & \downarrow f_* & & \downarrow f_* & & \downarrow f_* & & \downarrow f_* & & \\ \dots & \longrightarrow & H_n(B) & \longrightarrow & H_n(Y) & \longrightarrow & H_n(Y, B) & \longrightarrow & H_{n-1}(B) & \longrightarrow & \dots \end{array}$$

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 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 \tilde{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 $\tilde{H}_n(\{\text{pt}\}) = 0$ 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

$$(26.1) \quad H_0(X; G) \cong \bigoplus_{\pi_0(X)} G.$$

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

$$(26.2) \quad \pi_1(X) \rightarrow H_1(X; \mathbb{Z})$$

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 $\sigma := \{\alpha_0, \dots, \alpha_n\} \subset V$ is defined to be an n -simplex $\sigma \in S$ of the complex $K_{\mathcal{O}}$ if and only if

$$\mathcal{U}_{\alpha_0} \cap \dots \cap \mathcal{U}_{\alpha_n} \neq \emptyset.$$

This easily satisfies the required conditions for a simplicial complex: each vertex $\alpha \in V$ defines a 0-simplex $\{\alpha\} \in S$ since $\mathcal{U}_\alpha \neq \emptyset$, and each face of $\sigma = \{\alpha_0, \dots, \alpha_n\} \in S$ is also a simplex in the complex since every nontrivial subcollection of the sets $\mathcal{U}_{\alpha_0}, \dots, \mathcal{U}_{\alpha_n}$ 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 $n \geq 0$ such that $n + 1$ of the sets in the covering have a nonempty intersection.




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.

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 $\mathcal{U}_{\alpha_0} \cap \dots \cap \mathcal{U}_{\alpha_n}$ 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 $\mathcal{U}_{\alpha_0} \cap \dots \cap \mathcal{U}_{\alpha_n}$ 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 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 $X^0 \subset X^1 \subset \dots \subset X^{n-1} \subset X^n = X$ such that:

- (1) X^0 is a finite discrete set;
- (2) For each $m = 1, \dots, n$, X^m is homeomorphic to a space constructed by attaching finitely many m -disks \mathbb{D}^m to X^{m-1} along maps $\partial\mathbb{D}^m \rightarrow X^{m-1}$.

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** $\varphi_\alpha : S^{m-1} \rightarrow X^{m-1}$ 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 $\varphi_m : \coprod_{\alpha \in \mathcal{K}_m(X)} \partial\mathbb{D}^m \rightarrow X^{m-1}$ denotes the disjoint union of the maps $\varphi_\alpha : S^{m-1} \rightarrow X^{m-1}$, 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

$$e_\alpha^m \cong \mathring{\mathbb{D}}^m \quad \text{for each } \alpha \in \mathcal{K}_m(X),$$

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 $S^{n-1} \cong \partial\Delta^n \rightarrow X^{n-1}$ 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 $\varphi_\alpha : S^{n-1} \rightarrow \{\text{pt}\}$. The resulting space $X = X^n$ is homeomorphic to S^n .

The **cellular homology** of a cell complex $X = \bigcup_{n \geq 0} X^n$ is now defined as follows. Given an abelian coefficient group G , let

$$C_n^{\text{CW}}(X; G) := \bigoplus_{\alpha \in \mathcal{K}_n(X)} G = \left\{ \text{finite sums } \sum_i c_i e_{\alpha_i}^n \mid c_i \in G, \alpha_i \in \mathcal{K}_n(X) \right\}$$

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 $\partial : C_n^{\text{CW}}(X; G) \rightarrow C_{n-1}^{\text{CW}}(X; G)$ 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

$$X_{\beta} := X^{n-1} / (X^{n-1} \setminus e_{\beta}^{n-1}),$$

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 $q : X^{n-1} \rightarrow X_{\beta}$, so composing this with the attaching map $\varphi_{\alpha} : S^{n-1} \rightarrow X^{n-1}$ gives a map between two $(n-1)$ -dimensional spheres

$$q \circ \varphi_{\alpha} : S^{n-1} \rightarrow X_{\beta} \cong S^{n-1}.$$

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 $\mathbb{Z} \rightarrow \mathbb{Z}$ 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:

$$[e_{\beta}^{n-1} : e_{\alpha}^n] := \deg(q \circ \varphi_{\alpha}).$$

Strictly speaking, this definition only makes sense for $n \geq 2$ 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

$$H_n^{\text{CW}}(X; G) := H_n(C_*^{\text{CW}}(X; G), \partial),$$

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

$$H_k^{\text{CW}}(S^n; G) \cong \begin{cases} G & \text{for } k = 0, n, \\ 0 & \text{for all other } k. \end{cases}$$

Indeed, for $n \geq 2$ we can see this without doing any work, because $C_0^{\text{CW}}(S^n; G) \cong C_n^{\text{CW}}(S^n; G) \cong G$ 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 , $H_*^{CW}(X; G) \cong h_*(X)$.*

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 unmanageably 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,

$$(26.3) \quad \sum_{n \in \mathbb{Z}} (-1)^n \dim_{\mathbb{K}} H_n(C_*, \partial) = \sum_{n \in \mathbb{Z}} (-1)^n \dim_{\mathbb{K}} C_n.$$

This follows essentially from the fact that for each $n \in \mathbb{Z}$, writing $Z_n := \ker \partial_n \subset C_n$ and $B_n := \text{im } \partial_{n+1} \subset C_n$, the map $\partial_n : C_n \rightarrow C_{n-1}$ descends to an isomorphism $C_n/Z_n \rightarrow B_{n-1}$, 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 $C_n^{CW}(X; \mathbb{K})$ 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 $H_*(X; \mathbb{K}) \cong H_*^{CW}(X; \mathbb{K})$ 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:

⁴²One 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.

THEOREM. *For any finite cell complex X ,*

$$\chi(X) = \sum_{n=0}^{\infty} (-1)^n (\text{the number of } n\text{-cells}).$$

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 $g \geq 0$: 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.

Second semester (Topologie II)

27. Categories and functors

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 will play a major role in this course from the next lecture onward—is singular homology H_* , which assigns to each space X a whole sequence of abelian groups $H_n(X)$ indexed by the nonnegative integers $n \geq 0$. 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. The first semester of this course did not yet develop the right language to make this notion of a “function” precise, so it is time to do so now.

27.1. Some remarks on set theory. One reason why π_1 cannot actually be called a “function” is that its domain, strictly speaking, is not a set (*Menge*). 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 dealt with this dilemma by coming up with a set of axioms that tell you how to construct new sets from old ones, together with a short list of examples of sets (e.g. the empty set) whose existence clearly needs to be assumed, and insisting that *only* collections of objects that arise by applying the given axioms to the given examples 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 thus 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 standard definitions. This doesn’t make it impossible to define something that intuitively resembles a function on the collection of all topological spaces—it only means that when we define such an object, we are not strictly allowed to call it a “function”. This problem is easy to solve: we shall simply call it something else.

27.2. Categories. Leaving set theory aside, we now introduce some basic notions from category theory. As the examples below should make clear, a category can often be thought of as an answer to the question, “which field of mathematics are we working in?”

DEFINITION 27.1. A **category** (*Kategorie*) \mathcal{C} consists of the following data:

- A collection (i.e. class) $\text{Ob}_{\mathcal{C}}$, whose elements are called the **objects** (*Objekte*) of \mathcal{C} ;
- For each $X, Y \in \text{Ob}_{\mathcal{C}}$ a set $\text{Hom}_{\mathcal{C}}(X, Y)$, which we shall often abbreviate as

$$\text{Hom}(X, Y) := \text{Hom}_{\mathcal{C}}(X, Y)$$

when there is no danger of confusion, whose elements are called the **morphisms from X to Y** (*Morphismen von X nach Y*). For each $X \in \text{Ob}_{\mathcal{C}}$, there is a distinguished⁴³ element $\text{Id}_X \in \text{Hom}(X, X)$ called the **identity morphism** of X ;

- For each $X, Y, Z \in \text{Ob}_{\mathcal{C}}$, a function

$$(27.1) \quad \text{Hom}(X, Y) \times \text{Hom}(Y, Z) \rightarrow \text{Hom}(X, Z) : (f, g) \mapsto g \circ f$$

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, $f \circ \text{Id} = f = \text{Id} \circ f$.

NOTATION. For a category \mathcal{C} , we will often abuse notation and use the symbol \mathcal{C} to indicate not only the category itself but also its collection of objects, hence

$$X \in \mathcal{C} \quad \text{actually means} \quad X \in \text{Ob}_{\mathcal{C}}.$$

A morphism $f \in \text{Hom}(X, Y)$ from X to Y will often be denoted with the same arrow notation that is standard for maps between sets, so

$$f : X \rightarrow Y \quad \text{or} \quad X \xrightarrow{f} Y \quad \text{actually means} \quad f \in \text{Hom}(X, Y).$$

The notation $\text{Hom}(X, Y)$ for a set of morphisms is inspired by Example 27.5 below and similar algebraic examples, in which morphisms are actually homomorphisms respecting given algebraic structures. One also often sees this set denoted by $\text{Mor}(X, Y)$ or $\mathcal{C}(X, Y)$, though we will not use that notation here.

EXAMPLE 27.2. The category Top has $\text{Ob}_{\text{Top}} = \{\text{topological spaces}\}$ and $\text{Hom}(X, Y) = \{f : X \rightarrow Y \mid f \text{ a continuous map}\}$, 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. In accordance with the notation convention described above, the statement

$$X \in \text{Top}$$

thus means that X is a topological space.

EXAMPLE 27.3. The category Set has $\text{Ob}_{\text{Set}} = \{\text{sets}\}$ and $\text{Hom}(X, Y) = \{f : X \rightarrow Y\}$, meaning that morphisms are simply maps between sets, with no continuity requirement 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 $\text{Ob}_{\text{Grp}} = \{\text{groups}\}$, with $\text{Hom}(G, H)$ defined as the set of all group homomorphisms $G \rightarrow H$ for each $G, H \in \text{Grp}$.

⁴³The word “distinguished” appears here because part of the structure of the category \mathcal{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 \mathcal{C} requires each set $\text{Hom}(X, X)$ 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 .

EXAMPLE 27.6. There is a **subcategory** (*Unterkategorie*) \mathbf{Ab} of \mathbf{Grp} whose objects consist of all *abelian* groups, with morphisms defined the same way as in \mathbf{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 \mathbf{hTop} denote the category whose objects are the same as in \mathbf{Top} , but with $\mathbf{Hom}(X, Y)$ defined as the set of homotopy classes of continuous maps $X \rightarrow Y$ and $\text{Id}_X \in \mathbf{Hom}(X, X)$ as the homotopy class of the identity map. The function (27.1) is defined in terms of the usual composition of continuous maps $f : X \rightarrow Y$ and $g : Y \rightarrow Z$ by

$$[g] \circ [f] := [g \circ f].$$

(Exercise: check that this is well defined!) We call \mathbf{hTop} the **homotopy category** of topological spaces.

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 27.3 and 27.4.

DEFINITION 27.8. In any category, a morphism $f \in \mathbf{Hom}(X, Y)$ is called an **isomorphism** (*Isomorphismus*) if there exists a morphism $f^{-1} \in \mathbf{Hom}(Y, X)$ such that $f^{-1} \circ f = \text{Id}_X$ and $f \circ f^{-1} = \text{Id}_Y$. If an isomorphism exists in $\mathbf{Hom}(X, Y)$, we say that the objects X and Y are **isomorphic** (*isomorph*).

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 \mathbf{Top} is a homeomorphism, in \mathbf{Set} it is simply a bijection, in \mathbf{Diff} a diffeomorphism, and in \mathbf{Grp} or \mathbf{Ab} it is the usual notion of group isomorphism. The most interesting case so far is \mathbf{hTop} : two objects in \mathbf{hTop} are isomorphic if and only if they are homotopy equivalent!

The proof of the following is an easy exercise in applying the axioms of a category:

PROPOSITION 27.9. *For any isomorphism $f : X \rightarrow Y$ between two objects of a category, the inverse morphism $f^{-1} : Y \rightarrow X$ is unique.* \square

REMARK 27.10. It is possible to relax Definition 27.1 by allowing $\mathbf{Hom}(X, Y)$ for each $X, Y \in \mathcal{C}$ to be an arbitrary class rather than a set, in which case we are not strictly allowed to call the composition map $\mathbf{Hom}(X, Y) \times \mathbf{Hom}(Y, Z) \rightarrow \mathbf{Hom}(X, Z)$ a “function,” but the definition still makes sense. In this more general framework, the notion described in Definition 27.1 with morphisms forming sets instead of proper classes is called a **locally small** category. All of the categories we deal with in this course will be locally small, and it takes some nontrivial effort to think up an example of one that is not, so we will not worry about this level of generality any further.

27.3. Functors. The next definition gives us a way of relating two categories to each other. As inspiration, you can think of π_1 , a “function” that associates groups to pointed topological spaces, and in fact does so in a way that makes the groups into *topological invariants*. This results mainly from the fact that continuous maps of spaces induce homomorphisms between the corresponding fundamental 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 27.11. Given two categories \mathcal{C} and \mathcal{D} , a **functor** (*Funktor*) $\mathcal{F} : \mathcal{C} \rightarrow \mathcal{D}$ from \mathcal{C} to \mathcal{D} assigns to each object $X \in \mathcal{C}$ an object $\mathcal{F}(X) \in \mathcal{D}$ and to each morphism $f \in \mathbf{Hom}(X, Y)$ between any two objects $X, Y \in \mathcal{C}$ a morphism $\mathcal{F}(f) \in \mathbf{Hom}(\mathcal{F}(X), \mathcal{F}(Y))$ such that:

- (1) $\mathcal{F}(\text{Id}_X) = \text{Id}_{\mathcal{F}(X)}$ for all $X \in \mathcal{C}$;
- (2) $\mathcal{F}(f \circ g) = \mathcal{F}(f) \circ \mathcal{F}(g)$ for all $g \in \text{Hom}(X, Y)$ and $f \in \text{Hom}(Y, Z)$, $X, Y, Z \in \mathcal{C}$.

EXAMPLE 27.12. Denote by Top_* the category whose objects are the **pointed spaces** (X, p) , i.e. a topological space X together with a point $p \in X$, with morphisms defined as continuous **pointed maps**, also known as **base point preserving maps**,

$$\text{Hom}((X, p), (Y, q)) := \{f : X \rightarrow Y \mid f \text{ continuous and } f(p) = q\}.$$

The fundamental group then defines a functor $\pi_1 : \text{Top}_* \rightarrow \text{Grp}$; indeed, it associates to each pointed space (X, p) the group $\pi_1(X, p)$ and to each pointed map $f : (X, p) \rightarrow (Y, q)$ the group homomorphism

$$\pi_1(f) := f_* : \pi_1(X, p) \rightarrow \pi_1(Y, q)$$

such that Id_* is the identity homomorphism and $(f \circ g)_* = f_* \circ g_*$.

EXAMPLE 27.13. There is an obvious functor $\text{Top} \rightarrow \text{hTop}$ that sends each object $X \in \text{Top}$ to itself and sends each continuous map $f : X \rightarrow Y$ to its homotopy class. This is sometimes called a **forgetful functor**, since it is defined by forgetting some (but not all) of the information carried by the morphisms in Top , i.e. it forgets the actual maps $X \rightarrow Y$, but remembers their homotopy classes.

EXAMPLE 27.14. The fundamental group also defines a functor $\pi_1 : \text{hTop}_* \rightarrow \text{Grp}$ where hTop_* is defined to have the same objects as Top_* , but with $\text{Hom}((X, p), (Y, q))$ defined as the set of pointed homotopy classes of maps $(X, p) \rightarrow (Y, q)$. (See Theorem 8.11 in Lecture 8 from last semester.) A slightly fancier way to say this is that the functor $\pi_1 : \text{Top}_* \rightarrow \text{Grp}$ in Example 27.12 is the composition of two functors

$$\text{Top}_* \xrightarrow{\quad \pi_1 \quad} \text{hTop}_* \xrightarrow{\quad \pi_1 \quad} \text{Grp},$$

in which the first is the pointed analogue of the forgetful functor described in Example 27.13. We say in this situation that the functor $\pi_1 : \text{Top}_* \rightarrow \text{Grp}$ **descends** to the (pointed) homotopy category hTop_* .

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 27.15. Given two categories \mathcal{C} and \mathcal{D} , a **contravariant functor** (*kontravarianter Funktor*) $\mathcal{F} : \mathcal{C} \rightarrow \mathcal{D}$ from \mathcal{C} to \mathcal{D} assigns to each $X \in \mathcal{C}$ some $\mathcal{F}(X) \in \mathcal{D}$ and to each $f \in \text{Hom}(X, Y)$ for $X, Y \in \mathcal{C}$ a morphism $\mathcal{F}(f) \in \text{Hom}(\mathcal{F}(Y), \mathcal{F}(X))$ such that

- (1) $\mathcal{F}(\text{Id}_X) = \text{Id}_{\mathcal{F}(X)}$ for all $X \in \mathcal{C}$;
- (2) $\mathcal{F}(f \circ g) = \mathcal{F}(g) \circ \mathcal{F}(f)$ for all $g \in \text{Hom}(X, Y)$ and $f \in \text{Hom}(Y, Z)$, $X, Y, Z \in \mathcal{C}$.

A functor that satisfies the original Definition 27.11 instead of Definition 27.15 can be called **covariant** (*kovariant*) when we want to emphasize that it is not contravariant.

EXAMPLE 27.16. Let $\mathbb{K}\text{-Vect}$ denote the category of vector spaces over a fixed field \mathbb{K} , so $\text{Hom}(V, W) := \text{Hom}_{\mathbb{K}}(V, W)$ is the space of \mathbb{K} -linear maps $V \rightarrow W$. There is a contravariant functor $\Delta : \mathbb{K}\text{-Vect} \rightarrow \mathbb{K}\text{-Vect}$ which sends each vector space V to its dual space $\Delta(V) := V^* := \text{Hom}_{\mathbb{K}}(V, \mathbb{K})$ and sends each morphism $A : V \rightarrow W$ to its transpose $\Delta(A) := A^* : W^* \rightarrow V^*$, 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 $V \rightarrow V$ is the identity $V^* \rightarrow V^*$.

REMARK 27.17. It is possible to avoid Definition 27.15 by instead defining for each category \mathcal{C} the **opposite category** \mathcal{C}^{op} , which has the same collection of objects but reverses the arrows for all morphisms, meaning $\text{Hom}_{\mathcal{C}^{\text{op}}}(X, Y) := \text{Hom}_{\mathcal{C}}(Y, X)$. A contravariant functor $\mathcal{C} \rightarrow \mathcal{D}$ is then the same thing as a covariant functor $\mathcal{C}^{\text{op}} \rightarrow \mathcal{D}$.

EXAMPLE 27.18. One can speak of “functors of multiple variables” in much the same way as with functions. It is not difficult to show for instance that on the category \mathbf{Ab} of abelian groups and homomorphisms,

$$\text{Hom} : \mathbf{Ab} \times \mathbf{Ab} \rightarrow \mathbf{Ab}$$

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 $\text{Hom}(G, H)$ of homomorphisms $G \rightarrow H$.

27.4. Natural transformations. 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 27.19. Given two covariant functors $\mathcal{F}, \mathcal{G} : \mathcal{C} \rightarrow \mathcal{D}$, a **natural transformation** (*natürliche Transformation*) T from \mathcal{F} to \mathcal{G} associates to each $X \in \mathcal{C}$ a morphism $T_X \in \text{Hom}(\mathcal{F}(X), \mathcal{G}(X))$ such that for all $X, Y \in \mathcal{C}$ and $f \in \text{Hom}(X, Y)$, the following diagram commutes:

$$(27.2) \quad \begin{array}{ccc} \mathcal{F}(X) & \xrightarrow{T_X} & \mathcal{G}(X) \\ \downarrow \mathcal{F}(f) & & \downarrow \mathcal{G}(f) \\ \mathcal{F}(Y) & \xrightarrow{T_Y} & \mathcal{G}(Y) \end{array}$$

The statement that T is a natural transformation from \mathcal{F} to \mathcal{G} for two functors $\mathcal{F}, \mathcal{G} : \mathcal{C} \rightarrow \mathcal{D}$ is sometimes written with the notation

$$T : \mathcal{F} \Rightarrow \mathcal{G}, \quad \text{or} \quad \mathcal{F} \xrightarrow{T} \mathcal{G}, \quad \text{or} \quad \begin{array}{ccc} & \mathcal{F} & \\ \curvearrowright & \Downarrow T & \curvearrowleft \\ & \mathcal{G} & \end{array} .$$

A natural transformation between two contravariant functors can be defined analogously.

REMARK 27.20. The meaning of commutative diagrams such as (27.2) in an abstract category-theoretical framework should hopefully be obvious: in the case at hand, the diagram means the relation

$$\mathcal{G}(f) \circ T_X = T_Y \circ \mathcal{F}(f),$$

i.e. it specifies that two specific compositions of morphisms give rise to the same morphism $\mathcal{F}(X) \rightarrow \mathcal{G}(Y)$. A very large portion of the important definitions and results in category theory can be expressed in terms of commutative diagrams, which make sense due to the axioms of a category, without needing to assume that objects are sets or that morphisms are maps between them.

We will see some nice topological examples of natural transformations in the context of bordism theory in §27.7 below. Here is an algebraic example that you may have heard of before:

EXAMPLE 27.21. Consider again the category $\mathbb{K}\text{-Vect}$ of vector spaces over a fixed field \mathbb{K} as in Example 27.16. There is a covariant functor

$$\Delta^2 : \mathbb{K}\text{-Vect} \rightarrow \mathbb{K}\text{-Vect},$$

assigning to each $V \in \mathbb{K}\text{-Vect}$ the dual of its dual space $(V^*)^*$. Let $\text{Id} : \mathbb{K}\text{-Vect} \rightarrow \mathbb{K}\text{-Vect}$ denote the identity functor on $\mathbb{K}\text{-Vect}$, which sends each object and morphism to itself. There is then a natural transformation from Id to Δ^2 that assigns to every $V \in \mathbb{K}\text{-Vect}$ a vector space isomorphism $V \rightarrow (V^*)^*$; see Exercise 27.5.

REMARK 27.22. Whenever a vector space V is finite dimensional, the map $V \rightarrow (V^*)^*$ given by the natural transformation in Example 27.21 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 $\text{Id} : \mathbb{K}\text{-Vect} \rightarrow \mathbb{K}\text{-Vect}$ is a covariant functor while the dualization functor $\Delta : \mathbb{K}\text{-Vect} \rightarrow \mathbb{K}\text{-Vect}$ from Example 27.16 is contravariant, there is no sensible notion of natural transformations from Id to Δ .

27.5. Bordism groups. It would be too ambitious to attempt a serious discussion of bordism theory in this course, but there are two good reasons to introduce the basic definitions now. First, they give us some elegant new topological examples of functors besides π_1 , including some obviously interesting examples of natural transformations. Second, the geometric idea behind bordism groups will give us motivation for the somewhat less straightforward definition of homology groups in the lectures to come.

NOTATION. This is a convenient moment to mention a notational convention that will be in force throughout the semester: we abbreviate the compact unit interval by

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

This will be the meaning of the symbol I in any context that involves homotopies.

For some initial motivation, you can think of π_1 in the following terms: first, elements of $\pi_1(X)$ are represented by base-point preserving maps $\gamma : S^1 \rightarrow X$ defined on a specific closed 1-dimensional manifold, namely the circle S^1 . Two such maps $\gamma, \gamma' : S^1 \rightarrow X$ represent the same element if there exists a pointed homotopy

$$h : S^1 \times I \rightarrow X,$$

between them, so in this situation, the disjoint union $\gamma \amalg \gamma' : S^1 \amalg S^1 \rightarrow X$ of the two maps admits a continuous extension to a map $S^1 \times I \rightarrow X$, whose domain is a specific compact 2-dimensional manifold with boundary naturally homeomorphic to $S^1 \amalg S^1$. This way of describing homotopies ignores base points, but base points are not important for our present purposes: what's important rather is that we are talking about maps into X defined on compact 2-manifolds bounded by closed 1-manifolds. If you take this picture and ask what happens when you allow the domains to be arbitrary compact manifolds of arbitrary dimension, bordism theory is what you get.

For the following definition, recall that an n -dimensional manifold M is called **closed** if it is compact and the $(n - 1)$ -dimensional manifold ∂M defined as the boundary of M is empty. We will generally use the term “manifold” as a synonym for “manifold with boundary,” so all manifolds M are allowed to have a nonempty boundary ∂M , but we shall make no overriding assumptions about whether ∂M is nonempty unless extra words such as “closed” are included. It is useful to note however that for any manifold M , the boundary ∂M is a manifold whose own boundary is always empty:

$$\partial(\partial M) = \emptyset.$$

DEFINITION 27.23. For a space $X \in \text{Top}$ and an integer $n \geq 0$, the n th **unoriented bordism group** $\Omega_n^{\text{O}}(X)$ of X consists of equivalence classes $[(M, \varphi)]$ of pairs (M, φ) in which M is a closed smooth n -manifold and $\varphi : M \rightarrow X$ is a continuous map. We call two such pairs (M, φ) and (N, ψ)

equivalent (or **bordant**) if there exists a **bordism** between them, meaning a pair (W, Φ) in which W is a compact smooth $(n + 1)$ -manifold, $\Phi : W \rightarrow X$ is a continuous map, and there exists a diffeomorphism

$$\partial W \cong M \amalg N$$

such that

$$\Phi|_{\partial W} = \varphi \amalg \psi.$$

We make $\Omega_n^O(X)$ into an abelian group by using disjoint unions to define addition, thus

$$[(M, \varphi)] + [(N, \psi)] := [(M \amalg N, \varphi \amalg \psi)],$$

with the additive identity element defined by

$$0 := [(\emptyset, \cdot)] \in \Omega_n^O(X),$$

where the empty set \emptyset is understood as a smooth manifold of arbitrary dimension, and \cdot denotes the unique map $\emptyset \rightarrow X$.

A few observations are needed before this definition fully makes sense. First, we should check that the bordism relation described above satisfies the conditions of an equivalence relation: for instance, it is reflexive because for any closed n -manifold M and map $\varphi : M \rightarrow X$, the compact $(n + 1)$ -manifold $M \times I$ and map

$$(27.3) \quad M \times I \rightarrow X : (x, t) \mapsto \varphi(x)$$

define a bordism between (M, φ) and itself. The symmetry of the relation is obvious; the most interesting detail is transitivity, which requires some rudimentary knowledge of smooth manifolds and collar neighborhoods, so that two $(n + 1)$ -manifolds with diffeomorphic boundary components can be glued together along those components to form a new $(n + 1)$ -manifold. Since this discussion is not intended as a comprehensive introduction to bordism theory, I will leave that detail to your imagination for now. Once the bordism relation is understood, it is straightforward to check that the addition operation defined via disjoint unions is well defined on equivalence classes. The remaining question to answer is why $\Omega_n^O(X)$ is a group, i.e. why every element has an additive inverse. This also comes from the map (27.3), because there is another way to interpret it: the boundary of $M \times I$ is naturally diffeomorphic to the disjoint union of $M \amalg M$ with \emptyset , which makes $(M \amalg M, \varphi \amalg \varphi)$ bordant to (\emptyset, \cdot) and thus proves

$$[(M, \varphi)] + [(M, \varphi)] = 0 \in \Omega_n^O(X).$$

This not only makes $\Omega_n^O(X)$ a group, but also gives it an especially simple algebraic structure: all of its nontrivial elements have order 2, so the abelian group $\Omega_n^O(X)$ can also be regarded as a vector space over the field \mathbb{Z}_2 .

REMARK 27.24. The domains in Definition 27.23 were all assumed to be *smooth* manifolds rather than just *topological* manifolds, but there is an equally sensible variation on this definition that requires only topological manifolds and replaces the word “diffeomorphism” (in the definition of the bordism relation) with “homeomorphism”. The main reason to include smoothness in the definition is that methods from differential topology make $\Omega_n^O(X)$ easier to compute than its purely topological counterpart. But for our present purposes, this detail will make no difference at all and can safely be ignored.

The following observation makes Ω_n^O into a covariant functor

$$\Omega_n^O : \mathbf{Top} \rightarrow \mathbf{Ab},$$

or equivalently (in light of the fact that all nontrivial elements have order two), a functor $\Omega_n^O : \mathbf{Top} \rightarrow \mathbb{Z}_2\text{-Vect}$. Each continuous map $f : X \rightarrow Y$ induces a map $f_* : \Omega_n^O(X) \rightarrow \Omega_n^O(Y)$ defined by

$$f_*[(M, \varphi)] := [(M, f \circ \varphi)].$$

It is straightforward to check that this map is well defined and is a group homomorphism. It clearly also sends the identity map $X \rightarrow X$ to the identity homomorphism $\Omega_n^O(X) \rightarrow \Omega_n^O(X)$ and satisfies the relation $(f \circ g)_* = f_*g_*$ for any two continuous maps f, g that are composable. In other words: $\Omega_n^O : \mathbf{Top} \rightarrow \mathbf{Ab}$ is a functor.

A less obvious but very useful observation is that $\Omega_n^O : \mathbf{Top} \rightarrow \mathbf{Ab}$ *descends* (in the sense of Example 27.14) to the corresponding homotopy category, and thus also defines a functor

$$\Omega_n^O : \mathbf{hTop} \rightarrow \mathbf{Ab}.$$

This is an immediate consequence of the following result:

PROPOSITION 27.25. *For any two homotopic maps $f, g : X \rightarrow Y$, the induced homomorphisms $f_*, g_* : \Omega_n^O(X) \rightarrow \Omega_n^O(Y)$ are identical.*

PROOF. Assume $H : X \times I \rightarrow Y$ is a homotopy with $H(\cdot, 0) = f$ and $H(\cdot, 1) = g$. Given $[(M, \varphi)] \in \Omega_n^O(X)$, the map

$$M \times I \rightarrow Y : (x, t) \mapsto H(\varphi(x), t)$$

then defines a bordism between $(M, f \circ \varphi)$ and $(M, g \circ \varphi)$, proving $f_*[(M, \varphi)] = g_*[(M, \varphi)] \in \Omega_n^O(Y)$. \square

27.6. Oriented bordism. In case you had hoped for a more interesting group in which not all nontrivial elements have order 2, there is a remedy: one can add a bit of extra data to the domain manifolds that are used to represent bordism classes, namely an orientation. If you know already what it means to equip a smooth manifold with an orientation, then great—if not, then this is not the place to discuss it, though we will give a detailed treatment of orientations for *topological* manifolds later in this semester. For present purposes, it will suffice to take the following facts about orientations on faith:

- (1) Many familiar manifolds such as S^1 and the compact surfaces Σ_g of genus g for each $g \geq 0$ are orientable, but not all manifolds are, e.g. the projective plane \mathbb{RP}^2 and the Klein bottle are not. More generally, no manifold that contains a Möbius band (or equivalently, that is the connected sum of something with \mathbb{RP}^2) can admit an orientation, because the Möbius band contains a loop such that any choice of orientation at one point gets reversed by moving it continuously along the loop.
- (2) For every orientation of a manifold M , there is another orientation called the **opposite orientation**, and if M is connected, then it admits exactly two orientations, which are opposites of each other. For an oriented manifold M , we sometimes denote by $-M$ the same manifold with the opposite orientation.⁴⁴
- (3) For every manifold M with nonempty boundary, an orientation of M naturally determines an orientation of ∂M , called the **boundary orientation**. The opposite orientation of M then determines the opposite boundary orientation, or in symbols,

$$\partial(-M) = -(\partial M).$$

⁴⁴Another popular way of denoting the oriented manifold $-M$ is \bar{M} , especially in certain situations where M comes with a canonical choice of orientation. This is true for instance if M is a *complex* manifold, e.g. the complex projective space $\mathbb{C}\mathbb{P}^n$, which inherits a canonical orientation from its complex structure, and $\bar{\mathbb{C}\mathbb{P}^n}$ then denotes the same real manifold with an orientation opposite to the one determined by the complex structure.

- (4) For any two oriented manifolds M, N with $\dim M = m$ and $\dim N = n$, the Cartesian product $M \times N$ is an $(m + n)$ -manifold that inherits a **product orientation**, which depends in general on the order of the factors, though only if both m and n are odd. In symbols,

$$N \times M \cong (-1)^{mn}(M \times N),$$

meaning that the obvious diffeomorphism $M \times N \xrightarrow{\cong} N \times M$ is orientation reversing if m and n are both odd, and is otherwise orientation preserving.

- (5) For M a 0-manifold (also known as a discrete set with at most countably many points), an orientation is simply a function $M \rightarrow \{1, -1\}$, and for any choice of orientation on the unit interval $I = [0, 1]$, the boundary orientation assigns opposite signs to the two points of $\partial I = \{0, 1\}$.

DEFINITION 27.26. The n th **oriented bordism group**

$$\Omega_n^{\text{SO}}(X)$$

of a space X is defined by modifying Definition 27.23 as follows: the manifold M in each representative (M, φ) is equipped with an orientation, and the manifold W in an **oriented bordism** (W, Φ) between (M, φ) and (N, ψ) is also oriented and equipped with an orientation-preserving diffeomorphism

$$\partial W \cong -M \amalg N,$$

where ∂W is assumed to carry the boundary orientation.

Let us clarify why reversing the orientation of either M or N in the oriented bordism relation is the right thing to do. For any oriented manifold M , assigning the product orientation to $M \times I$ and then the boundary orientation to $\partial(M \times I)$ gives a natural *orientation-preserving* diffeomorphism

$$\partial(M \times I) \cong -M \amalg M.$$

The trivial homotopy (27.3) thus implies again that the bordism relation satisfies $(M, \varphi) \sim (M, \varphi)$, but in the oriented setting, it does *not* imply $[(M, \varphi)] + [(M, \varphi)] = 0$, so that elements of $\Omega_n^{\text{SO}}(X)$ do not need to have order two. Instead, the additive inverse of any given $[(M, \varphi)] \in \Omega_n^{\text{SO}}(X)$ is obtained by reversing the orientation of M ,

$$-[(M, \varphi)] = [(-M, \varphi)] \in \Omega_n^{\text{SO}}(X).$$

REMARK 27.27. The letters “O” and “SO” appearing in the notation $\Omega_n^{\text{O}}(X)$ and $\Omega_n^{\text{SO}}(X)$ refer to the orthogonal group $\text{O}(n)$ and special orthogonal group $\text{SO}(n)$ respectively, which makes some sense if you recall that $\text{SO}(n)$ is precisely the subgroup of $\text{O}(n)$ consisting of transformations $\mathbb{R}^n \rightarrow \mathbb{R}^n$ that preserve orientation. A fuller explanation of this notation would be too much of a digression for now, but suffice it to say there also exist other versions of bordism groups corresponding to other families of Lie groups that act linearly on Euclidean space, in which the manifold M in representatives (M, φ) of bordism classes is equipped with extra structure respected by those group actions.

The following easy computation (see Exercise 27.7) demonstrates that, indeed, elements of $\Omega_n^{\text{SO}}(X)$ need not have order 2 in general.

PROPOSITION 27.28. *For any space X , there are natural isomorphisms*

$$\Omega_0^{\text{O}}(X) \cong \bigoplus_{\pi_0(X)} \mathbb{Z}_2, \quad \text{and} \quad \Omega_0^{\text{SO}}(X) \cong \bigoplus_{\pi_0(X)} \mathbb{Z},$$

i.e. $\Omega_0^{\text{O}}(X)$ is a vector space over \mathbb{Z}_2 with a canonical basis in bijective correspondence with the set $\pi_0(X)$ of path-components of X , and $\Omega_0^{\text{SO}}(X)$ is a free abelian group with the same basis. \square

Beyond the case $n = 0$, computations of $\Omega_n^O(X)$ and $\Omega_n^{SO}(X)$ are generally doable, but too difficult to attempt before learning about homology and cohomology, which will be the main objectives of this semester's course. One sees a revealing symptom of the difficulty when one tries to compute either of these groups for the simplest possible nonempty topological space, namely a one-point space.

NOTATION. We will frequently denote by

$$\{*\} \in \mathbf{Top}$$

a topological space consisting of only one point, with the point in this space denoted by

$$* \in \{*\}.$$

The symbol $* \in X$ is sometimes also used to denote the base point of a pointed space $X \in \mathbf{Top}_*$, if it has not been given any other name.

REMARK 27.29. We sometimes abuse terminology by speaking of “the” one-point space, but of course one-point spaces are not unique, since the one element in the space can literally be anything, e.g. the sets $\{1\}$ and $\{2\}$ are not identical since $1 \neq 2$, and they are also different from the set whose only element is the banana you ate for breakfast this morning. In light of the fact that set theory has no way of defining a “set of all things,” the collection of all possible one-point spaces forms a proper class rather than a set. However, one does have a strong form of uniqueness up to isomorphism in the category \mathbf{Top} or \mathbf{Top}_* : there exists a unique homeomorphism between any two one-point spaces, and this is why referring to them all as “the” one-point space does not do any harm.

For a one-point space $\{*\}$ and any given manifold M , there is only one possible map $M \rightarrow \{*\}$, so elements of the groups⁴⁵

$$\Omega_n^O := \Omega_n^O(\{*\}), \quad \Omega_n^{SO} := \Omega_n^{SO}(\{*\})$$

can be regarded simply as equivalence classes $[M]$ of closed n -manifolds, and the information encoded in these groups is therefore a coarse version of the classification of closed n -manifolds, subject to an equivalence relation in which boundaries of compact manifolds are equated with the empty set. The classification problem up to homeomorphism or diffeomorphism is well understood for manifolds of dimension at most two, but already from dimension three upward, complete classifications are not known, and the problem is not generally considered tractable. From this perspective, it seems slightly surprising that Ω_n^O and Ω_n^{SO} can in fact be computed, and the answers are often not difficult to write down, but proving them usually takes quite a bit of work. By the end of this semester, we will at least be able to fill in all the gaps in the following special case:

PROPOSITION 27.30. *The group $\Omega_2^O = \Omega_2^O(\{*\})$ is isomorphic to \mathbb{Z}_2 , and its unique nontrivial element is the bordism class of the projective plane \mathbb{RP}^2 .*

PROOF SKETCH. The nontriviality of $[\mathbb{RP}^2] \in \Omega_2^O(\{*\})$ means that \mathbb{RP}^2 is not diffeomorphic to the boundary of any compact 3-manifold. If you take this on faith for a moment, the rest of the computation follows easily from the classification of surfaces, as described in Lecture 19 from last semester. Indeed, every closed and *orientable* surface can be presented as the smooth boundary of a compact region in \mathbb{R}^3 , and thus represents the trivial element in Ω_2^O . The closed, connected

⁴⁵The symbols Ω_n^{SO} and Ω_n^O now each have two possible interpretations, either as functors $\mathbf{Top} \rightarrow \mathbf{Ab}$ or as the groups obtained by plugging a one-point space into these functors. It depends on the context.

and non-orientable surfaces, in turn, are all homeomorphic (and in fact also diffeomorphic⁴⁶) to connected sums of N copies of $\mathbb{R}\mathbb{P}^2$ for some $N \in \mathbb{N}$. A convenient fact to use in this situation is that for any two closed manifolds M, N of the same dimension n , there exists a compact $(n + 1)$ -manifold whose boundary is diffeomorphic to the disjoint union of M, N and the connected sum $M \# N$. This can be proved with a picture, and I will leave it as an exercise, but if you need a hint, try looking up some information on *handle attachment* in geometric topology—the key trick is to “attach a 1-handle” to $(M \amalg N) \times I$. With this understood, one now sees that every closed and connected surface is bordant to some disjoint union of copies of $\mathbb{R}\mathbb{P}^2$, and therefore so is every closed and disconnected surface.

So, why is $\mathbb{R}\mathbb{P}^2$ not the boundary of any compact 3-manifold? This is harder to explain, but it will follow easily from some computations of homological invariants carried out later in this course. In particular, the Poincaré duality isomorphism implies that the Euler characteristic (an integer-valued invariant that is defined for a wide class of topological spaces including all compact manifolds) of every closed odd-dimensional manifold is zero. If $\mathbb{R}\mathbb{P}^2$ were the boundary of some compact 3-manifold Y , then by gluing Y to a copy of itself along the boundary, one would obtain a closed 3-manifold

$$X := Y \cup_{\mathbb{R}\mathbb{P}^2} Y$$

whose Euler characteristic $\chi(X)$ satisfies $\chi(X) = 2\chi(Y) - \chi(\mathbb{R}\mathbb{P}^2) = 2\chi(Y) - 1$, and therefore could not be zero. \square

27.7. More examples of natural transformations. The bordism groups provide us with some examples of natural transformations that are quite easy to write down. Proving the required naturality property, i.e. that the required diagrams commute, is a straightforward exercise in each case.

EXAMPLE 27.31. For every space X and $n \geq 0$, there is an obvious *forgetful* homomorphism

$$\Omega_n^{\text{SO}} \rightarrow \Omega_n^{\text{O}} : [(M, \varphi)] \mapsto [(M, \varphi)]$$

defined by forgetting the orientation of the manifold M . Regarding both Ω_n^{SO} and Ω_n^{O} as covariant functors $\text{Top} \rightarrow \text{Ab}$, this defines a natural transformation from Ω_n^{SO} to Ω_n^{O} .

EXAMPLE 27.32. Since S^1 is a closed orientable 1-manifold, one can associate to any pointed space (X, p) a map

$$(27.4) \quad \pi_1(X, p) \xrightarrow{h} \Omega_1^{\text{SO}}(X) : [\gamma] \mapsto [(S^1, \gamma)],$$

defined by regarding representatives of elements in $\pi_1(X, p)$ as maps $\gamma : S^1 \rightarrow X$. This map is well defined because a homotopy between two maps $\gamma, \gamma' : S^1 \rightarrow X$ gives rise to a bordism between (S^1, γ) and (S^1, γ') . The lemma below shows that $h : \pi_1(X, p) \rightarrow \Omega_1^{\text{SO}}(X)$ is also a group homomorphism; it is a variation on what is known in homology theory as the *Hurewicz homomorphism*, and we will later see another version of it in that context.

LEMMA 27.33. *For any two base-point preserving loops $\alpha, \beta : S^1 \rightarrow X$ and their concatenation $\alpha \cdot \beta : S^1 \rightarrow X$, $(S^1, \alpha \cdot \beta)$ is bordant to $(S^1 \amalg S^1, \alpha \amalg \beta)$.*

PROOF. See Exercise 27.8. \square

⁴⁶It is a nontrivial fact that for $n \leq 3$ (though emphatically not for $n \geq 4$), every topological n -manifold admits a smooth structure, and two smooth n -manifolds are homeomorphic if and only if they are diffeomorphic. For closed surfaces, the easiest way to prove this is probably by reproving the standard classification of surfaces in the smooth category. In fact, this is easier than working only with topological surfaces and continuous maps, because Riemannian geometry makes the existence of triangulations on *smooth* manifolds easier to prove.

Continuing with Example 27.32, the following “naturality” property of the map $h : \pi_1(X, p) \rightarrow \Omega_1^{\text{SO}}(X)$ is nearly immediate from the definitions: for any pointed map $f : (X, p) \rightarrow (Y, q)$, the diagram

$$\begin{array}{ccc} \pi_1(X, p) & \xrightarrow{h} & \Omega_1^{\text{SO}}(X) \\ \downarrow f_* & & \downarrow f_* \\ \pi_1(Y, q) & \xrightarrow{h} & \Omega_1^{\text{SO}}(Y) \end{array}$$

commutes. This *almost* amounts to the statement that h defines a natural transformation from π_1 to Ω_1^{SO} , though before we can say this in precise terms, we have a minor bookkeeping issue to deal with, as $\pi_1 : \mathbf{Top}_* \rightarrow \mathbf{Grp}$ and $\Omega_1^{\text{SO}} : \mathbf{Top} \rightarrow \mathbf{Ab}$ are not functors between exactly the same pairs of categories, strictly speaking. The distinction between \mathbf{Grp} and \mathbf{Ab} is easy to erase since the latter is a subcategory of the former, i.e. we can equally well regard Ω_1^{SO} as a functor $\mathbf{Top} \rightarrow \mathbf{Grp}$. For the distinction between \mathbf{Top}_* and \mathbf{Top} , the obvious thing to do is define Ω_1^{SO} as a functor $\mathbf{Top}_* \rightarrow \mathbf{Grp}$ by composing the usual $\Omega_1^{\text{SO}} : \mathbf{Top} \rightarrow \mathbf{Grp}$ with the obvious *forgetful* functor $\mathbf{Top}_* \rightarrow \mathbf{Top}$, replacing each pointed space (X, p) with the unpointed space X . With this understood, the commuting diagram above shows that h defines a natural transformation from π_1 to Ω_1^{SO} if both are regarded as functors $\mathbf{Top}_* \rightarrow \mathbf{Grp}$. For a slightly different variation, we could observe that since $\Omega_1^{\text{SO}}(X)$ is abelian, the map $h : \pi_1(X, p) \rightarrow \Omega_1^{\text{SO}}(X)$ always vanishes on the commutator subgroup

$$[\pi_1(X, p), \pi_1(X, p)] \subset \pi_1(X, p),$$

and thus descends to a well-defined homomorphism on the abelianization of the fundamental group,

$$\pi_1(X, p) / [\pi_1(X, p), \pi_1(X, p)] \xrightarrow{h} \Omega_1^{\text{SO}}(X).$$

By now you should be unsurprised to learn that *abelianization* can also be regarded as a functor

$$\mathbf{Grp} \xrightarrow{\text{ab}} \mathbf{Ab} : G \mapsto \text{ab}(G) := G / [G, G],$$

and we can then also regard h as a natural transformation from $\text{ab} \circ \pi_1 : \mathbf{Top}_* \rightarrow \mathbf{Ab}$ to $\Omega_{\text{SO}}^1 : \mathbf{Top}_* \rightarrow \mathbf{Ab}$.

EXAMPLE 27.34. Let $\Omega_n^\bullet : \mathbf{Top} \rightarrow \mathbf{Ab}$ denote either the unoriented or oriented bordism functor. For any two spaces X, Y and integers $m, n \geq 0$, one can define a product operation

$$\begin{aligned} \Omega_m^\bullet(X) \otimes \Omega_n^\bullet(Y) &\xrightarrow{\times} \Omega_{m+n}^\bullet(X \times Y), \\ [(M, \varphi)] \otimes [(N, \psi)] &\mapsto [(M, \varphi)] \times [(N, \psi)] := [(M \times N, \varphi \times \psi)]. \end{aligned}$$

I will leave it as an exercise to convince yourself that this operation is well defined, and to clarify precisely what it means to say that it is *natural*: in particular, for each fixed pair of integers $m, n \geq 0$, this product defines a natural transformation between two functors from the *product category* $\mathbf{Top} \times \mathbf{Top}$ to \mathbf{Ab} .

27.8. Exercises.

EXERCISE 27.1. Prove Proposition 27.9 (isomorphisms have unique inverses).

EXERCISE 27.2. Verify the claim in Example 27.18 that $\text{Hom} : \mathbf{Ab} \times \mathbf{Ab} \rightarrow \mathbf{Ab}$ defines a contravariant functor in its first variable and a covariant functor in its second variable.

EXERCISE 27.3. Suppose \mathcal{A} is a category whose objects form a set X , such that for each pair $x, y \in X$, the set of morphisms $\text{Hom}(x, y)$ 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 $\text{Hom}(x, y) \neq \emptyset$.

- (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 \mathcal{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 $f : X \rightarrow Y$ need to have in order for it to define a functor from \mathcal{A} to \mathcal{B} ?

EXERCISE 27.4. In any category \mathcal{C} , each object X has an **automorphism group** (also called **isotropy group**) $\text{Aut}(X)$, consisting of all the isomorphisms in $\text{Hom}(X, X)$. A **groupoid** is a category in which all morphisms are also isomorphisms.

- (a) Show that if \mathcal{G} is a groupoid and \mathbf{Grp} denotes the usual category of groups with homomorphisms, there exists a contravariant functor from \mathcal{G} to \mathbf{Grp} that assigns to each object X of \mathcal{G} its automorphism group $\text{Aut}(X)$. How does this functor act on morphisms $X \rightarrow Y$? Could you alternatively define it as a *covariant* functor? Conclude either way that whenever X and Y are isomorphic objects in \mathcal{G} (meaning there exists an isomorphism in $\text{Hom}(X, Y)$), the groups $\text{Aut}(X)$ and $\text{Aut}(Y)$ are isomorphic.
- (b) Given a topological space X and two points x, y , let $\text{Hom}(x, y)$ denote the set of homotopy classes (with fixed end points) of paths $[0, 1] \rightarrow X$ from x to y , and define a composition function $\text{Hom}(x, y) \times \text{Hom}(y, z) \rightarrow \text{Hom}(x, z) : (\alpha, \beta) \mapsto \alpha \cdot \beta$ by the usual notion of concatenation of paths. Show that this notion of morphisms defines a groupoid whose objects are the points in X .⁴⁷ In this case, what are the automorphism groups $\text{Aut}(x)$ and the isomorphisms $\text{Aut}(y) \rightarrow \text{Aut}(x)$ given by the functor in part (a)?

EXERCISE 27.5. Consider the category $\mathbb{K}\text{-Vect}$ of vector spaces over a fixed field \mathbb{K} .

- (a) Show that there is a covariant functor Δ^2 from $\mathbb{K}\text{-Vect}$ to itself, assigning to each $V \in \mathbb{K}\text{-Vect}$ the dual of its dual space $(V^*)^*$. Describe how this functor acts on morphisms.
- (b) Construct a natural transformation from the identity functor $\text{Id} : \mathbb{K}\text{-Vect} \rightarrow \mathbb{K}\text{-Vect}$ to Δ^2 that assigns to every $V \in \mathbb{K}\text{-Vect}$ a linear injection $V \rightarrow (V^*)^*$, which is an isomorphism whenever V is finite dimensional.

EXERCISE 27.6. 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 $V \rightarrow \bar{V} : v \mapsto \bar{v}$ denotes the identity map, then scalar multiplication on \bar{V} is defined so as to make this map complex antilinear, giving the formula

$$\lambda \bar{v} := \overline{\lambda v} \in \bar{V} \quad \text{for } \lambda \in \mathbb{C}, v \in V.$$

- (a) Show that there is a covariant functor $\kappa : \mathbb{C}\text{-Vect} \rightarrow \mathbb{C}\text{-Vect}$ that sends each $V \in \mathbb{C}\text{-Vect}$ to its conjugate \bar{V} , and describe how this functor acts on morphisms.
- (b) Show that if T is a natural transformation from $\text{Id} : \mathbb{C}\text{-Vect} \rightarrow \mathbb{C}\text{-Vect}$ to $\kappa : \mathbb{C}\text{-Vect} \rightarrow \mathbb{C}\text{-Vect}$, then T assigns to each $V \in \mathbb{C}\text{-Vect}$ the zero map $V \rightarrow \bar{V}$.

Hint: What does the naturality of T imply about the specific morphism $V \rightarrow V : v \mapsto iv$?

Comment: The map $V \rightarrow \bar{V} : v \mapsto \bar{v}$ is always a real-linear isomorphism, but it is not complex linear and is thus not a morphism in $\mathbb{C}\text{-Vect}$. 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 $\text{Id} \rightarrow \kappa$ is a symptom of the fact that there is generally no canonical way to define such isomorphisms.

EXERCISE 27.7. Prove Proposition 27.28 on the computation of $\Omega_0^0(X)$ and $\Omega_0^{\text{SO}}(X)$ for any space X .

⁴⁷It is called the **fundamental groupoid** of X .

EXERCISE 27.8. Prove Lemma 27.33, showing that the natural map $h : \pi_1(X, p) \rightarrow \Omega_1^{\text{SO}}(X)$ is a group homomorphism.

Hint: You are looking for an oriented bordism (Σ, Φ) in which Σ is a compact surface with three boundary components—the simplest surface of this kind is a so-called “pair of pants,” which has the topology of a disk with two holes cut out. Assuming Σ is a pair of pants, try to define $\Phi : \Sigma \rightarrow X$ by first thinking about which subset of Σ should be mapped to the base point of X . If you know anything about Morse theory, there is a relatively simple Morse-theoretic picture that will almost immediately lead to the construction you need: it involves the gradient flow of a Morse function $f : \Sigma \rightarrow \mathbb{R}$ that is constant on each boundary component and has exactly one critical point of index 1 in the interior.

28. Axioms for homology theories

We will not yet define any specific homology theory in this lecture, but we shall introduce the standard set of axioms satisfied by homology theories, and demonstrate their usefulness in computations. Along the way, we encounter a fundamental tool from homological algebra: exact sequences.

28.1. The category of R -modules. The bordism theories Ω_n^{O} and Ω_n^{SO} in the previous lecture were defined as functors from **Top** to the category **Ab** of abelian groups, though we saw that the groups $\Omega_n^{\text{O}}(X)$ can also be regarded as vector spaces over \mathbb{Z}_2 . For homology theory, it is also possible to work entirely in the category **Ab**, but it is sometimes profitable to generalize this to a category that includes both abelian groups and vector spaces as special cases. This generalization does not require any extra effort, so we might as well work in the more general setting from the beginning.

NOTATION. For the rest of this course, unless otherwise noted, the symbol

$$R$$

will always denote a fixed commutative ring with unit, the choice of which will often not matter. We then denote by

$$R\text{-Mod}$$

the **category of modules over R** , whose morphisms are the R -module homomorphisms. For two modules $G, H \in R\text{-Mod}$, we will denote the set of R -module homomorphisms $G \rightarrow H$ (which is also an R -module) by

$$\text{Hom}_R(G, H) := \text{Hom}_{R\text{-Mod}}(G, H)$$

whenever there is a need to specify R , but the abbreviated notation

$$\text{Hom}(G, H) := \text{Hom}_R(G, H)$$

can also be used when the context is clear. Similarly, we can denote the tensor product of two R -modules by $G \otimes_R H$ whenever R needs to be specified, but we will otherwise abbreviate it as

$$G \otimes H := G \otimes_R H.$$

A trivial R -module⁴⁸ is denoted by

$$0 \in R\text{-Mod}.$$

For our purposes, abelian groups will be the most important special case of R -modules (see Example 28.1 below), and for that reason, we will sometimes abuse terminology and use the word “group” in places where the word “module” would be more appropriate.

⁴⁸As with one-point spaces, there is not a unique trivial R -module, but there is a unique R -module isomorphism between any two of them.

EXAMPLE 28.1. All abelian groups $G \in \mathbf{Ab}$ can equivalently be regarded as modules over the commutative ring \mathbb{Z} , with scalar multiplication nx for $n \in \mathbb{Z}$ and $x \in G$ determined in the obvious way by the addition operation. Group homomorphisms are then automatically also \mathbb{Z} -module homomorphisms, and in this sense, the categories \mathbf{Ab} and $\mathbb{Z}\text{-Mod}$ are completely equivalent.

EXAMPLE 28.2. If R is a field \mathbb{K} , then an R -module is the same thing as a vector space over \mathbb{K} , and $R\text{-Mod}$ is in this case equivalent to the category $\mathbb{K}\text{-Vect}$ of vector spaces.

In this course, we will in practice almost exclusively be interested in the special cases where R is either \mathbb{Z} or a field (most often either \mathbb{Z}_2 , \mathbb{Q} , \mathbb{R} or \mathbb{K}), and the category of R -modules will thus serve mainly as a single umbrella that encompasses both abelian groups and vector spaces.

One subtlety worth noting is that for any choice of the ring R , an R -module can always also be regarded as an abelian group, just by forgetting its scalar multiplication while keeping the addition operation, but doing this changes the definitions of tensor products $G \otimes H$ and the set of homomorphisms $\text{Hom}(G, H)$. For instance, if $G, H \in \mathbb{R}\text{-Vect}$ are real vector spaces, then they are also abelian groups and thus \mathbb{Z} -modules $G, H \in \mathbb{Z}\text{-Mod}$, but their tensor product in the sense of real vector spaces satisfies the relation

$$rg \otimes h = g \otimes rh \in G \otimes_{\mathbb{R}} H \quad \text{for all} \quad g \in G, h \in H, r \in \mathbb{R},$$

whereas the tensor product $G \otimes_{\mathbb{Z}} H$ in the sense of abelian groups only satisfies this when $r \in \mathbb{Z}$. Similarly, every \mathbb{R} -linear map $G \rightarrow H$ is also a homomorphism of abelian groups, but the converse is quite false.

DEFINITION 28.3. A **basis** of an R -module G is a subset $\mathcal{B} \subset G$ such that every element $g \in G$ can be written in the form

$$g = \sum_{b \in \mathcal{B}} g_b b$$

for some coefficients $g_b \in G$ that are uniquely determined by g , at most finitely-many of which are nonzero. An R -module is called **free** if it admits a basis.

A choice of basis $\mathcal{B} \subset G$ for a free R -module is equivalent to a choice of R -module isomorphism

$$\bigoplus_{b \in \mathcal{B}} R \xrightarrow{\cong} G,$$

so for instance, an abelian group (i.e. \mathbb{Z} -module) is free if and only if it is isomorphic to a direct sum of copies of \mathbb{Z} . Obviously, not every abelian group G has this property, e.g. it is never true if G is finite. On the other hand, a standard argument in linear algebra (using Zorn's lemma for the infinite-dimensional case) shows that every vector space admits a basis, so when R is a field, all R -modules are free. This basic fact is one of the key advantages of having the freedom to work with vector spaces instead of just abelian groups.

28.2. Exact sequences and splittings. In homological algebra, exact sequences play a role comparable to that of Cauchy sequences in analysis; that is to say, the entire subject would be impossible without them.

By a **sequence** (*Sequenz*) of R -modules, we mean a linearly ordered collection of modules A_n for $n \in \mathbb{Z}$, together with R -module homomorphisms $\alpha_n : A_n \rightarrow A_{n+1}$. Depending on the context in which sequences arise, we can allow n to vary over any contiguous subset of the integers, which may be unbounded, or bounded above and/or below, so the sequence itself may have finitely or infinitely many terms, with or without a starting or end point. Let us call A_n an **interior** term of the sequence if the sequence also includes both A_{n-1} and A_{n+1} , thus giving rise to a three-term subsequence

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

In this situation, we say that the sequence is **exact** (*exakt*) at the term A_n if

$$\text{im } \alpha_{n-1} = \ker \alpha_n.$$

We do not define exactness for non-interior terms, i.e. terms that are at the beginning or end of the sequence. An **exact sequence** (*exakte Sequenz*) of R -modules is a sequence that is exact at all of its interior terms.

EXAMPLE 28.4. A sequence of the form $0 \rightarrow A \xrightarrow{f} B \rightarrow 0$ is exact if and only if f is an isomorphism.

EXAMPLE 28.5. An exact sequence with five terms that begins and ends with trivial modules

$$0 \rightarrow A \xrightarrow{f} B \xrightarrow{g} C \rightarrow 0$$

is called a **short exact sequence** (*kurze exakte Sequenz*). Exactness means in this case that f is injective, g is surjective, and $\text{im } f = \ker g$. A popular class of examples is the sequence

$$0 \rightarrow A \hookrightarrow B \xrightarrow{q} B/A \rightarrow 0$$

for any submodule $A \subset B$, where q denotes the quotient projection. Another is

$$(28.1) \quad 0 \rightarrow A \xrightarrow{i} A \oplus C \xrightarrow{p} C \rightarrow 0$$

for any two modules A and C , with the obvious inclusion map $i(a) := (a, 0)$ and projection map $p(a, c) := c$.

DEFINITION 28.6. A short exact sequence $0 \rightarrow A \rightarrow B \rightarrow C \rightarrow 0$ is said to **split**, and is then called a **split exact sequence**, if there exists an isomorphism $B \cong A \oplus C$ identifying it with the sequence in (28.1).

In the category of abelian groups, there are easy examples of short exact sequences that do not split, e.g. writing $q: \mathbb{Z} \rightarrow \mathbb{Z}/2\mathbb{Z} =: \mathbb{Z}_2$ for the quotient projection,

$$0 \rightarrow \mathbb{Z} \xrightarrow{\cdot 2} \mathbb{Z} \xrightarrow{q} \mathbb{Z}_2 \rightarrow 0$$

is such an example, since \mathbb{Z} is not isomorphic to $\mathbb{Z} \oplus \mathbb{Z}_2$. The next result, whose proof is a straightforward exercise, gives a useful practical criterion for short exact sequences to split, and its corollary implies in particular that they *always* split if R is a field.

THEOREM 28.7. *The following conditions on a short exact sequence $0 \rightarrow A \xrightarrow{f} B \xrightarrow{g} C \rightarrow 0$ are equivalent:*

- (i) *The sequence splits;*
- (ii) *The injective homomorphism $f: A \rightarrow B$ admits a left-inverse $B \rightarrow A$;*
- (iii) *The surjective homomorphism $g: B \rightarrow C$ admits a right-inverse $C \rightarrow B$.*

□

COROLLARY 28.8. *If C is a free R -module, then every short exact sequence $0 \rightarrow A \rightarrow B \rightarrow C \rightarrow 0$ splits.*

PROOF. Use a basis of C to define a right-inverse for the surjective map $B \rightarrow C$. □

Here is another popular application of exactness whose proof is an easy exercise.

THEOREM 28.9. *For an exact sequence of the form*

$$\dots \rightarrow A_n \xrightarrow{f_n} B_n \rightarrow C_n \rightarrow A_{n+1} \xrightarrow{f_{n+1}} B_{n+1} \rightarrow C_{n+1} \rightarrow \dots,$$

the following conditions are equivalent:

- (i) The modules C_n are trivial for every n ;
- (ii) The maps $f_n : A_n \rightarrow B_n$ are isomorphisms for every n .

□

28.3. Relative bordism groups. For a first real-life example of an exact sequence that arises naturally in topology, we can generalize the previous lecture’s discussion and define **relative bordism groups**

$$\Omega_n^O(X, A)$$

for every so-called **pair of spaces** (X, A) , meaning a space X together with a choice of subset $A \subset X$. Given two pairs of spaces (X, A) and (Y, B) , a **map of pairs**

$$f : (X, A) \rightarrow (Y, B) \quad \text{or} \quad (X, A) \xrightarrow{f} (Y, B)$$

is a continuous map $f : X \rightarrow Y$ such that $f(A) \subset B$, thus if we assign subspace topologies to A and B , the restriction $f|_A$ becomes a continuous map $A \rightarrow B$. Let us focus the discussion for now on *unoriented* bordism theory; the oriented case is completely analogous. Elements of $\Omega_n^O(X, A)$ are equivalence classes $[(M, \varphi)]$ in which M is a compact smooth n -manifold that is allowed to have nonempty boundary, and φ is a map of pairs

$$(M, \partial M) \xrightarrow{\varphi} (X, A).$$

Two such pairs (M, φ) and (N, ψ) are equivalent if there is a **relative bordism** between them: this means a pair (W, Φ) consisting of a compact smooth $(n + 1)$ -manifold W equipped with a smooth embedding

$$M \amalg N \hookrightarrow \partial W,$$

and a map of pairs

$$(W, \partial W \setminus (M \amalg N)) \xrightarrow{\Phi} (X, A)$$

such that $\Phi|_{M \amalg N} = \varphi \amalg \psi$. Note that while the domain of $\varphi : M \rightarrow X$ in this definition is allowed to have nonempty boundary, it may also be closed, thus the definition still makes sense if $A = \emptyset$ and just reproduces the so-called **absolute** bordism groups defined in the previous lecture,

$$\Omega_n^O(X, \emptyset) = \Omega_n^O(X).$$

The group structure of $\Omega_n^O(X, A)$ is again defined via disjoint unions, and there is a straightforward way of associating to each map of pairs $f : (X, A) \rightarrow (Y, B)$ a group homomorphism

$$\Omega_n^O(X, A) \xrightarrow{f_*} \Omega_n^O(Y, B),$$

so that Ω_n^O becomes a functor

$$\mathbf{Top}^{\text{rel}} \xrightarrow{\Omega_n^O} \mathbf{Ab},$$

defined on the category $\mathbf{Top}^{\text{rel}}$ of pairs of spaces, whose morphisms are maps of pairs. We can identify \mathbf{Top} with the subcategory of $\mathbf{Top}^{\text{rel}}$ whose objects are pairs of the form (X, \emptyset) , and then interpret $\Omega_n^O : \mathbf{Top}^{\text{rel}} \rightarrow \mathbf{Ab}$ as an extension of the previously-defined functor $\Omega_n^O : \mathbf{Top} \rightarrow \mathbf{Ab}$.

For any pair (X, A) and $n \geq 1$, there is also a group homomorphism

$$\begin{aligned} \Omega_n^O(X, A) &\xrightarrow{\hat{c}_*} \Omega_{n-1}^O(A), \\ [(M, \varphi)] &\longmapsto [(\partial M, \varphi|_{\partial M})], \end{aligned}$$

which is well defined because if (W, Φ) is a relative bordism between two representatives (M, φ) and (N, ψ) , then restricting Φ to the compact n -manifold obtained by removing the interiors of M and N from ∂W defines an absolute bordism between $(\partial M, \varphi|_{\partial M})$ and $(\partial N, \psi|_{\partial N})$. One can

interpret ∂_* as a natural transformation between two functors $\mathbf{Top}^{\text{rel}} \rightarrow \mathbf{Ab}$, the details of which I will leave to the reader. What I really want to point out about ∂_* is the following:

THEOREM 28.10. *Given a pair of spaces (X, A) , let $i : A \hookrightarrow X$ and $j : (X, \emptyset) \hookrightarrow (X, A)$ denote the obvious inclusions. Then the sequence of abelian groups*

$$\begin{aligned} \dots \longrightarrow \Omega_n^{\text{O}}(A) \xrightarrow{i_*} \Omega_n^{\text{O}}(X) \xrightarrow{j_*} \Omega_n^{\text{O}}(X, A) \xrightarrow{\partial_*} \Omega_{n-1}^{\text{O}}(A) \xrightarrow{i_*} \Omega_{n-1}^{\text{O}}(X) \longrightarrow \dots \\ \dots \longrightarrow \Omega_1^{\text{O}}(X, A) \xrightarrow{\partial_*} \Omega_0^{\text{O}}(A) \xrightarrow{i_*} \Omega_0^{\text{O}}(X) \xrightarrow{j_*} \Omega_0^{\text{O}}(X, A) \longrightarrow 0 \end{aligned}$$

is exact.

COROLLARY 28.11 (via Theorem 28.9). *For a pair of spaces (X, A) , the map $\Omega_n^{\text{O}}(A) \rightarrow \Omega_n^{\text{O}}(X)$ induced by the inclusion $A \hookrightarrow X$ is an isomorphism for every $n \geq 0$ if and only if $\Omega_n^{\text{O}}(X, A) = 0$ for every $n \geq 0$. \square*

We will later see an analogue of Theorem 28.10 in singular homology that plays a major role in that theory, and whose proof requires some elementary but non-obvious ideas from homological algebra. It's worth noting that the proof of Theorem 28.10, by comparison, is much more direct and straightforward; see Exercise 28.2.

28.4. The Eilenberg-Steenrod axioms. In the early history of homology, multiple packages of invariants were proposed that were easier to compute than the bordism groups, while seeming to measure similar topological information. The resulting theories differ in the details of their definitions—some of them drastically—but turn out to be naturally isomorphic if one restricts them to a “nice” class of spaces, which in practice includes all of the spaces that one is typically interested in, such as manifolds. Eventually, singular homology settled into a special role as the “standard” homology theory that everyone needs to learn, but in fact, one usually doesn't need to know its precise definition in order to use it. What's much more important are the *formal* properties that it satisfies, which are common to all homology theories, and were codified in the middle of the 20th century as a set of axioms due to Eilenberg and Steenrod [ES52], with a bit of extra input from Milnor [Mil62].

DEFINITION 28.12. Fix as usual a commutative ring R with unit. An **axiomatic homology theory** h_* valued in the category of R -modules is a collection $\{h_n\}_{n \in \mathbb{Z}}$ of covariant functors

$$\mathbf{Top}^{\text{rel}} \xrightarrow{h_n} R\text{-Mod} : (X, A) \mapsto h_n(X, A)$$

defined for each $n \in \mathbb{Z}$, which also determine functors $h_n : \mathbf{Top} \rightarrow R\text{-Mod}$ by defining

$$h_n(X) := h_n(X, \emptyset).$$

For a map of pairs $f : (X, A) \rightarrow (Y, B)$, the R -module homomorphism induced by the functor h_n is denoted by

$$h_n(X, A) \xrightarrow{f_*} h_n(Y, B).$$

The data of a homology theory also includes natural transformations ∂_* from the functor $\mathbf{Top}^{\text{rel}} \rightarrow R\text{-Mod} : (X, A) \mapsto h_n(X, A)$ to the functor $\mathbf{Top}^{\text{rel}} \rightarrow R\text{-Mod} : (X, A) \mapsto h_{n-1}(A)$ for each $n \in \mathbb{Z}$, and we require the following axioms:

- (HOMOTOPY) For any two homotopic maps of pairs $f, g : (X, A) \rightarrow (Y, B)$, the induced morphisms $f_*, g_* : h_n(X, A) \rightarrow h_n(Y, B)$ are identical. (See Remark 28.14 below for the notion of a *homotopy of maps of pairs*.)

- (EXACTNESS) For all pairs (X, A) with inclusion maps $i : A \hookrightarrow X$ and $j : (X, \emptyset) \hookrightarrow (X, A)$, the sequence

$$\dots \longrightarrow h_{n+1}(X, A) \xrightarrow{\partial_*} h_n(A) \xrightarrow{i_*} h_n(X) \xrightarrow{j_*} h_n(X, A) \xrightarrow{\partial_*} h_{n-1}(A) \longrightarrow \dots$$

is exact.

- (EXCISION) For any pair (X, A) and any subset $B \subset A$ such that there exists a continuous function $u : X \rightarrow I$ equal to 0 on B and 1 on $X \setminus A$, the map induced by the inclusion $(X \setminus B, A \setminus B) \hookrightarrow (X, A)$ is an isomorphism

$$h_n(X \setminus B, A \setminus B) \xrightarrow{\cong} h_n(X, A) \quad \text{for every } n \in \mathbb{Z}.$$

- (DIMENSION) For any one-point space $\{*\}$, $h_n(\{*\}) = 0$ for all $n \neq 0$. The group $h_0(\{*\})$ is then called the **coefficient group** of the homology theory.⁴⁹
- (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 map determined by the induced homomorphisms

$$i_*^\alpha : h_n(X_\alpha) \rightarrow h_n\left(\coprod_{\beta \in J} X_\beta\right)$$

is an isomorphism

$$\bigoplus_{\alpha \in J} i_*^\alpha : \bigoplus_{\alpha \in J} h_n(X_\alpha) \xrightarrow{\cong} h_n\left(\coprod_{\beta \in J} X_\beta\right).$$

You should be able to convince yourself without much trouble that the bordism functors $\Omega_n^O : \mathbf{Top}^{\text{rel}} \rightarrow \mathbf{Ab} = \mathbb{Z}\text{-Mod}$ and their oriented counterparts Ω_n^{SO} each satisfy four out of the five Eilenberg-Steenrod axioms; see in particular Exercises 28.2 and 28.3. They do *not* satisfy the dimension axiom: this follows from Proposition 27.30 in the case of unoriented bordism theory, and there is a similar result for the oriented theory involving complex (instead of real) projective spaces. We call h_* a **generalized homology theory** if it satisfies all of the Eilenberg-Steenrod axioms except for dimension. In some contexts, the word “generalized” is removed, so that homology theories are typically assumed to satisfy four axioms instead of five, and those which also satisfy the dimension axiom are called **ordinary** homology theories. We will generally assume the dimension axiom in this semester and will not make use of any theories that don’t satisfy it, but some of the results we prove about homology theories will be equally valid for generalized theories, since they do not depend on the dimension axiom.

A few further comments on the axioms are in order.

REMARK 28.13. The original list in [ES52] included three additional axioms at the beginning of the list, but the first two of these are equivalent to the statement that the h_n are functors, and the third simply requires ∂_* to be a natural transformation.

REMARK 28.14. The following definition is hopefully intuitive: a **homotopy** between two maps of pairs $f, g : (X, A) \rightarrow (Y, B)$ is a homotopy $H : X \times I \rightarrow Y$ between f and g such that $H(\cdot, t)$ is also a map of pairs $(X, A) \rightarrow (Y, B)$ for every $t \in I$, so in other words, H satisfies the condition

$$H(A \times I) \subset B.$$

⁴⁹There is a slightly awkward semantic issue in this definition: strictly speaking, what we are calling “ $\{*\}$ ” is not a unique space, but simply *any* choice of space that happens to contain only one element. It follows that the coefficient group $h_0(\{*\})$ 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 $P_0 \rightarrow P_1$, which induces a canonical isomorphism $h_0(P_0) \rightarrow h_0(P_1)$, and the coefficient group of a homology theory is unique in this sense.

We could also have chosen to hide the homotopy axiom by calling the h_n functors

$$\mathbf{hTop}^{\text{rel}} \xrightarrow{h_n} R\text{-Mod}$$

instead of $\mathbf{Top}^{\text{rel}} \rightarrow R\text{-Mod}$, where $\mathbf{hTop}^{\text{rel}}$ denotes the **homotopy category of pairs of spaces**, having the same objects as $\mathbf{Top}^{\text{rel}}$, but with homotopy classes of maps of pairs as morphisms. Note that a homotopy of maps of pairs $(X, \emptyset) \rightarrow (Y, \emptyset)$ is just a homotopy of maps $X \rightarrow Y$, making \mathbf{hTop} naturally a subcategory of $\mathbf{hTop}^{\text{rel}}$.

REMARK 28.15. 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 28.4), thus Eilenberg and Steenrod did not need it, because they were mainly concerned with computations for compact polyhedra—compactness precludes infinite disjoint unions.

REMARK 28.16. One often sees the excision axiom stated under a weaker hypothesis on the sets $B \subset A \subset X$, namely that the closure of B is contained in the interior of A . You might find it a challenge to think up an example in which that hypothesis is satisfied but the one we stated is not, and I don't encourage you to try, because within the class of spaces that are typically considered interesting to study, the two are fully equivalent; moreover, in all interesting situations I'm aware of, it is as easy to verify the stronger hypothesis as the weaker one. Singular homology does satisfy excision under the weaker hypothesis, but the existence of a function $u : X \rightarrow I$ separating B from $X \setminus A$ is a more natural condition from other points of view, especially in homotopy-theoretic reformulations of homology. The hypotheses originally stated in [ES52] also required B to be open, which is another detail that makes no meaningful difference for the class of spaces typically of interest.

REMARK 28.17. The reason the dimension axiom has the name that it does is that if it were not included in the list of axioms, then for every homology theory h_* , one could use arbitrary degree shifts to define new homology theories such as k_* with $k_n(X, A) := h_{n+1}(X, A)$. The dimension axiom prevents this, in the hope that the value of the subscript n in $h_n(X, A)$ will then have some geometric meaning. The reason for calling $h_0(\{*\})$ a “coefficient group” will become clearer when we write down concrete examples of homology theories.

REMARK 28.18. It is sometimes useful to expand the definition and allow an axiomatic homology theory to be a functor $\mathcal{C} \rightarrow R\text{-Mod}$ defined on a suitable subcategory \mathcal{C} of $\mathbf{Top}^{\text{rel}}$, so that we need not define $h_*(X, A)$ for *all* pairs (X, A) , but only a subclass. One useful example 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. Others include the categories of *polyhedra* and *CW-complexes*, which we'll have more to say about in future lectures. 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 $(X, \emptyset) \hookrightarrow (X, A)$ —are actually morphisms in the category \mathcal{C} . In [ES52], this concern motivates the definition of the notion of an *admissible* category of pairs, though we have no need to reproduce that definition here.

28.5. Reduced homology. Assume h_* is a collection of functors as in Definition 28.12 satisfying at least the homotopy and exactness axioms. For technical reasons that will become clearer in the next section, it is sometimes useful to replace the groups $h_n(X)$ with certain subgroups $\tilde{h}_n(X) \subset h_n(X)$ called **reduced homology** groups. To define them, we denote by

$$X \xrightarrow{\epsilon} \{*\}$$

the unique map from any given space X to a one-point space. We then use the induced homomorphisms $\epsilon_* : h_n(X) \rightarrow h_n(\{*\})$ to define

$$\tilde{h}_n(X) := \ker \left(h_n(X) \xrightarrow{\epsilon_*} h_n(\{*\}) \right) \subset h_n(X).$$

Observe that if h_* satisfies the dimension axiom, then $\tilde{h}_n(X) = h_n(X)$ for all $n \neq 0$. If $n = 0$ or the dimension axiom is not satisfied, then we can typically expect $\tilde{h}_n(X)$ and $h_n(X)$ to be different, and the best way to relate them to each other is through a split exact sequence. Indeed, observe that the map $\epsilon : X \rightarrow \{*\}$ is not only trivially surjective, but also admits a right-inverse, defined by choosing any embedding

$$\{*\} \xrightarrow{i} X.$$

It then follows from functoriality that the homomorphism $\epsilon_* : h_n(X) \rightarrow h_n(\{*\})$ likewise is surjective and admits a right-inverse, thus by Theorem 28.7,

$$0 \rightarrow \tilde{h}_n(X) \hookrightarrow h_n(X) \xrightarrow{\epsilon_*} h_n(\{*\}) \rightarrow 0$$

is a split exact sequence, implying the existence of an isomorphism

$$h_n(X) \cong \tilde{h}_n(X) \oplus h_n(\{*\}).$$

If h_* satisfies the dimension axiom and has coefficient group $G = h_0(\{*\})$, this becomes

$$h_n(X) \cong \begin{cases} \tilde{h}_n(X) & \text{if } n \neq 0, \\ \tilde{h}_n(X) \oplus G & \text{if } n = 0. \end{cases}$$

One should keep in mind however that this isomorphism is not generally canonical: it depends on the choice of inclusion $i : \{*\} \hookrightarrow X$, which determines the splitting of the exact sequence relating $\tilde{h}_n(X)$ and $h_n(X)$.

Let us clarify why \tilde{h}_n for each $n \in \mathbb{Z}$ is naturally also a functor $\mathbf{Top} \rightarrow R\text{-Mod}$.

PROPOSITION 28.19. *The homomorphisms $f_* : h_n(X) \rightarrow h_n(Y)$ induced by any continuous map $f : X \rightarrow Y$ send $\tilde{h}_n(X)$ into $\tilde{h}_n(Y)$.*

PROOF. Denote $\epsilon^X : X \rightarrow \{*\}$ and $\epsilon^Y : Y \rightarrow \{*\}$ for the unique maps, and notice that $\epsilon^Y \circ f = \text{Id} \circ \epsilon^X$, thus the following diagram commutes.

$$\begin{array}{ccc} h_n(X) & \xrightarrow{f_*} & h_n(Y) \\ \downarrow \epsilon_*^X & & \downarrow \epsilon_*^Y \\ h_n(\{*\}) & \xrightarrow{\mathbb{1}} & h_n(\{*\}) \end{array}$$

This implies that $f_*(\ker \epsilon_*^X) \subset \ker \epsilon_*^Y$. □

The next result reveals the main advantage of using \tilde{h}_* in place of h_* for certain applications.

PROPOSITION 28.20. *If X is a contractible space, then $\tilde{h}_n(X) = 0$ for every n .*

PROOF. Contractibility implies that the map $\epsilon : X \rightarrow \{*\}$ is a homotopy equivalence, thus by the homotopy axiom, $\epsilon_* : h_n(X) \rightarrow h_n(\{*\})$ is an isomorphism, and its kernel $\tilde{h}_n(X)$ is therefore trivial. □

REMARK 28.21. If h_* also satisfies the dimension axiom, then we also have $h_n(X) = 0$ for all $n \neq 0$ whenever X is contractible, but $h_0(X)$ is typically nontrivial, as it is isomorphic to the coefficient group. As a consequence, some of the standard applications of reduced homology can

also be carried out with unreduced homology, but only if the degree 0 groups are excluded from consideration.

The relative version of reduced homology is defined in a trivial way: we set

$$\tilde{h}_n(X, A) := h_n(X, A) \quad \text{whenever} \quad A \neq \emptyset.$$

This seemingly naive definition is justified by the following considerations. Note first that the functors $\tilde{h}_n : \mathbf{Top} \rightarrow R\text{-Mod}$ now extend to pairs as functors $\mathbf{Top}^{\text{rel}} \rightarrow R\text{-Mod}$; here there is nothing to check since the existence of a map of pairs $(X, A) \rightarrow (Y, B)$ with $A \neq \emptyset$ implies $B \neq \emptyset$, so that both reduced relative homology groups match the unreduced case. Next, observe that for any space X , the relative homology groups $h_n(X, X)$ all vanish; this follows from the exactness axiom and Theorem 28.9, as we have an exact sequence

$$\dots \rightarrow h_n(X) \xrightarrow{1} h_n(X) \rightarrow h_n(X, X) \xrightarrow{\partial_*} h_{n-1}(X) \xrightarrow{1} h_{n-1}(X) \rightarrow \dots$$

It follows that $\tilde{h}_n(X, A)$ for $A \neq \emptyset$ is in fact the kernel of the map

$$h_n(X, A) \xrightarrow{\epsilon_*} h_n(\{*\}, \{*\}) = 0$$

induced by the unique map of pairs $\epsilon : (X, A) \rightarrow (\{*\}, \{*\})$. Moreover, the naturality of the connecting homomorphisms ∂_* gives a commutative diagram

$$\begin{array}{ccc} h_{n+1}(X, A) & \xrightarrow{\partial_*} & h_n(A) \\ \downarrow \epsilon_* & & \downarrow \epsilon_* \\ h_{n+1}(\{*\}, \{*\}) & \xrightarrow{\partial_*} & h_n(\{*\}) \end{array}$$

Since the term $h_{n+1}(\{*\}, \{*\})$ is trivial, this diagram proves that the image of $\partial_* : h_{n+1}(X, A) \rightarrow h_n(A)$ is always in the subgroup $\tilde{h}_n(A)$. We can therefore write down a well-defined sequence of homomorphisms

$$\dots \rightarrow \tilde{h}_{n+1}(X, A) \xrightarrow{\partial_*} \tilde{h}_n(A) \xrightarrow{i_*} \tilde{h}_n(X) \xrightarrow{j_*} \tilde{h}_n(X, A) \xrightarrow{\partial_*} \tilde{h}_{n-1}(A) \rightarrow \dots$$

using the usual inclusions $i : A \hookrightarrow X$ and $j : (X, \emptyset) \hookrightarrow (X, A)$. It is not immediately obvious whether this sequence is exact, but consider the commutative diagram

$$\begin{array}{cccccccc} & & 0 & & 0 & & 0 & & 0 & & \\ & & \downarrow & & \downarrow & & \downarrow & & \downarrow & & \\ \dots & \longrightarrow & \tilde{h}_n(A) & \xrightarrow{i_*} & \tilde{h}_n(X) & \xrightarrow{j_*} & \tilde{h}_n(X, A) & \xrightarrow{\partial_*} & \tilde{h}_{n-1}(A) & \longrightarrow & \dots \\ & & \downarrow & & \downarrow & & \parallel & & \downarrow & & \\ \dots & \longrightarrow & h_n(A) & \xrightarrow{i_*} & h_n(X) & \xrightarrow{j_*} & h_n(X, A) & \xrightarrow{\partial_*} & h_{n-1}(A) & \longrightarrow & \dots \\ & & \downarrow \epsilon_* & & \downarrow \epsilon_* & & \downarrow \epsilon_* & & \downarrow \epsilon_* & & \\ \dots & \longrightarrow & h_n(\{*\}) & \xrightarrow{=} & h_n(\{*\}) & \longrightarrow & 0 & \longrightarrow & h_{n-1}(\{*\}) & \longrightarrow & \dots \\ & & \downarrow & & \downarrow & & \downarrow & & \downarrow & & \\ & & 0 & & 0 & & 0 & & 0 & & \end{array}$$

Here the bottom two nontrivial rows are exact due to the exactness axiom, and all columns in the diagram are short exact sequences by construction. The rest is algebra:

PROPOSITION 28.22. *Assume the following diagram of R -modules commutes, all its columns are exact sequences, and the bottom two nontrivial rows are also exact sequences:*

$$\begin{array}{ccccccccc}
 & & 0 & & 0 & & 0 & & 0 & & 0 & & & & \\
 & & \downarrow & & \downarrow & & \downarrow & & \downarrow & & \downarrow & & & & \\
 \dots & & A_{n+2} & & A_{n+1} & & A_n & & A_{n-1} & & A_{n-2} & & \dots & & \\
 & & \downarrow \iota_{n+2} & & \downarrow \iota_{n+1} & & \downarrow \iota_n & & \downarrow \iota_{n-1} & & \downarrow \iota_{n-2} & & & & \\
 \dots & \longrightarrow & B_{n+2} & \xrightarrow{g_{n+2}} & B_{n+1} & \xrightarrow{g_{n+1}} & B_n & \xrightarrow{g_n} & B_{n-1} & \xrightarrow{g_{n-1}} & B_{n-2} & \longrightarrow & \dots & & \\
 & & \downarrow \epsilon_{n+2} & & \downarrow \epsilon_{n+1} & & \downarrow \epsilon_n & & \downarrow \epsilon_{n-1} & & \downarrow \epsilon_{n-2} & & & & \\
 \dots & \longrightarrow & C_{n+2} & \xrightarrow{h_{n+2}} & C_{n+1} & \xrightarrow{h_{n+1}} & C_n & \xrightarrow{h_n} & C_{n-1} & \xrightarrow{h_{n-1}} & C_{n-2} & \longrightarrow & \dots & & \\
 & & \downarrow & & \downarrow & & \downarrow & & \downarrow & & \downarrow & & & & \\
 & & 0 & & 0 & & 0 & & 0 & & 0 & & & &
 \end{array}$$

Then the top nontrivial row can be endowed uniquely with maps $f_n : A_n \rightarrow A_{n-1}$ such that the diagram still commutes, and these make that row into an exact sequence.

PROOF. The method behind this proof is commonly known as *diagram chasing*, and we will later see several other examples of it. The basic idea is straightforward: at every step, we examine a particular term in the diagram, consider what is already known about the maps going into and out of that term, and then deduce whatever we can from given conditions such as exactness. In typical situations, whatever can be deduced tells you which term to examine in the next step.

If $f_n : A_n \rightarrow A_{n-1}$ 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 $g_n \iota_n(a) \in \ker \epsilon_{n-1} = \text{im } \iota_{n-1}$. 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 $\dots A_{n+1} \xrightarrow{f_{n+1}} A_n \xrightarrow{f_n} A_{n-1} \rightarrow \dots$ 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 commutativity 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 28.23. *For any pair of spaces (X, A) and any homology theory h_* , there is an exact sequence of reduced homology groups*

$$\dots \longrightarrow \tilde{h}_{n+1}(X, A) \xrightarrow{\partial_*} \tilde{h}_n(A) \xrightarrow{i_*} \tilde{h}_n(X) \xrightarrow{j_*} \tilde{h}_n(X, A) \xrightarrow{\partial_*} \tilde{h}_{n-1}(A) \longrightarrow \dots,$$

where $i : A \hookrightarrow X$ and $j : (X, \emptyset) \hookrightarrow (X, A)$ are the obvious inclusions and $\partial_* : \tilde{h}_n(X, A) \rightarrow \tilde{h}_{n-1}(A)$ is the same map as the usual connecting homomorphism $h_n(X, A) \rightarrow h_{n-1}(A)$. □

28.6. Suspension isomorphisms. The following general construction leads easily to a complete computation of $h_*(S^n)$ for any axiomatic homology theory. We assume in this section that h_* is a generalized homology theory, so it satisfies all the conditions in Definition 28.12 except possibly the dimension axiom.⁵⁰

Recall that for an arbitrary space X , the **suspension** (*Einhängung*) of X is a space ΣX 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,

$$\Sigma X := C_+X \cup_X C_-X.$$

THEOREM 28.24. *For every space X , integer $k \in \mathbb{Z}$ and generalized homology theory h_* , the diagram (28.2) below gives rise to a natural isomorphism*

$$\Sigma_* := \varphi_*^{-1} \circ j_* \circ i_* \circ \partial_*^{-1} : \tilde{h}_k(X) \rightarrow \tilde{h}_{k+1}(\Sigma X).$$

PROOF. Let

$$p_+ \in C_+X \subset \Sigma X \quad \text{and} \quad p_- \in C_-X \subset \Sigma X$$

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

$$(28.2) \quad \begin{array}{ccc} \tilde{h}_k(X) & & \tilde{h}_{k+1}(\Sigma X) \\ \partial_* \uparrow & & \downarrow \varphi_* \\ \tilde{h}_{k+1}(C_+X, X) & \xrightarrow{i_*} \tilde{h}_{k+1}(\Sigma X \setminus \{p_-\}, C_-X \setminus \{p_-\}) & \xrightarrow{j_*} \tilde{h}_{k+1}(\Sigma X, C_-X) \end{array}$$

in which three of the maps are determined by the obvious inclusions of pairs,

$$\begin{aligned} (C_+X, X) &\xrightarrow{i} (\Sigma X \setminus \{p_-\}, C_-X \setminus \{p_-\}), \\ (\Sigma X \setminus \{p_-\}, C_-X \setminus \{p_-\}) &\xrightarrow{j} (\Sigma X, C_-X), \\ (\Sigma X, \emptyset) &\xrightarrow{\varphi} (\Sigma X, C_-X). \end{aligned}$$

⁵⁰In fact, the additivity axiom is also not strictly necessary for this discussion, since by Exercise 28.4, it follows from the other axioms in the case of finite disjoint unions.

The first of these is a homotopy equivalence, as there exists a deformation retraction of the pair $(\Sigma X \setminus \{p_-\}, C_- X \setminus \{p_-\})$ to $(C_+ X, X)$, thus i_* is an isomorphism by the homotopy axiom. Since one can easily define a function $u : \Sigma X \rightarrow I$ that vanishes at p_- and equals 1 on $C_+ X$, the excision axiom implies that j_* is also an isomorphism. For the other two maps, we consider the exact sequences provided by Theorem 28.23 for the pairs $(\Sigma X, C_- X)$ and $(C_+ X, X)$, that is

$$\dots \longrightarrow \tilde{h}_{k+1}(C_- X) \longrightarrow \tilde{h}_{k+1}(\Sigma X) \xrightarrow{\varphi_*} \tilde{h}_{k+1}(\Sigma X, C_- X) \longrightarrow \tilde{h}_k(C_- X) \longrightarrow \dots$$

and

$$\dots \longrightarrow \tilde{h}_{k+1}(C_+ X) \longrightarrow \tilde{h}_{k+1}(C_+ X, X) \xrightarrow{\hat{\varphi}_*} \tilde{h}_k(X) \longrightarrow \tilde{h}_k(C_+ X) \longrightarrow \dots$$

The contractibility of $C_\pm X$ implies via Proposition 28.20 that

$$\tilde{h}_k(C_\pm X) \cong \tilde{h}_k(\{*\}) = 0, \quad \text{and} \quad \tilde{h}_{k+1}(C_\pm X) \cong \tilde{h}_{k+1}(\{*\}) = 0,$$

thus the exactness of these two sequences implies that φ_* and ∂_* are both isomorphisms.

The *naturality* of the map $\Sigma_* : \tilde{h}_k(X) \rightarrow \tilde{h}_{k+1}(\Sigma X)$ has a precise meaning, because the suspension operation can be understood as a functor $\Sigma : \mathbf{Top} \rightarrow \mathbf{Top}$, and the statement is then that Σ_* defines a natural transformation between two functors $\mathbf{Top} \rightarrow R\text{-Mod}$, namely \tilde{h}_k and $\tilde{h}_{k+1} \circ \Sigma$. This follows in a straightforward way using the naturality of the homomorphisms ∂_* ; the details are an exercise. \square

28.7. Homology groups of spheres. Recall that the suspension of a sphere is also a sphere, but one dimension higher:

$$\Sigma(S^n) \cong S^{n+1}.$$

This fact and Theorem 28.24 make possible an inductive computation of $h_*(S^n)$ for every axiomatic homology theory and every $n \geq 0$, using the fact that S^0 is the disjoint union of two one-point spaces. Here is the statement; the proof is Exercise 28.6.

THEOREM 28.25. *Assume h_* is an axiomatic homology theory with coefficient group $h_0(\{*\}) = G$. Then for each pair of integers $k \in \mathbb{Z}$ and $n \geq 1$,*

$$h_k(S^n) \cong \begin{cases} G & \text{if } k = 0 \text{ or } k = n, \\ 0 & \text{otherwise.} \end{cases}$$

\square

28.8. Exercises.

EXERCISE 28.1 (*). Prove Theorem 28.7 on split exact sequences, and Theorem 28.9 on long exact sequences with every third term vanishing.

EXERCISE 28.2. Prove that the sequence of relative and absolute bordism groups in Theorem 28.10 is exact. Here are a couple of hints:

- For exactness at $\Omega_n^O(X)$: $j_*[(M, \varphi)] = 0$ means $\varphi : M \rightarrow X$ can be extended over a compact $(n+1)$ -manifold W with $M \subset \partial W$ such that the extension maps $\partial W \setminus M$ into A . In this situation, M is a *closed* manifold—what does that imply about $\partial W \setminus M$?
- For exactness at $\Omega_n^O(X, A)$: $\partial_*[(M, \varphi)] = 0$ means that $\varphi|_{\partial M} : \partial M \rightarrow A$ extends to a map $V \rightarrow A$ on some compact n -manifold V whose boundary is identified with ∂M . Build a closed n -manifold out of V and M .

EXERCISE 28.3. Prove that the bordism theories Ω_*^O and Ω_*^{SO} satisfy the homotopy, excision, and additivity axioms.

Hint for excision: Suppose $u : X \rightarrow I$ is a function as specified in the excision axiom, and $\varphi : (M, \partial M) \rightarrow (X, A)$ is a map of pairs so that (M, φ) represents a bordism class. By standard

results about smooth manifolds, the function $u \circ \varphi : M \rightarrow I$ can be perturbed to a function $v : M \rightarrow I$ such that for some $r \in (0, 1)$, both $v^{-1}(r) \subset M$ and $(v|_{\partial M})^{-1}(r) \subset \partial M$ are smooth submanifolds.

EXERCISE 28.4. Assume h_* is a collection of functors $h_n : \mathbf{Top}^{\text{rel}} \rightarrow R\text{-Mod}$ for $n \in \mathbb{Z}$ satisfying the exactness and excision axioms of Eilenberg-Steenrod. Given two spaces X, Y and the natural inclusions $i^X : X \hookrightarrow X \amalg Y$ and $i^Y : Y \hookrightarrow X \amalg Y$, show that the map

$$i_*^X \oplus i_*^Y : h_n(X) \oplus h_n(Y) \rightarrow h_n(X \amalg Y) : (x, y) \mapsto i_*^X x + i_*^Y y$$

is an isomorphism, and deduce that h_* also satisfies the additivity axiom for all *finite* disjoint unions.

Hint: Apply exactness and excision to the pairs $(X \amalg Y, X)$ and $(X \amalg Y, Y)$.

EXERCISE 28.5 (*). Assume h_* is an axiomatic homology theory with coefficient group $h_0(\{*\}) = G$. For any two spaces X and Y with maps $\epsilon^X : X \rightarrow \{*\}$ and $\epsilon^Y : Y \rightarrow \{*\}$, show that the natural isomorphism $h_n(X \amalg Y) \cong h_n(X) \oplus h_n(Y)$ identifies $\tilde{h}_n(X \amalg Y)$ with $\ker(\epsilon_*^X \oplus \epsilon_*^Y) \subset h_n(X; G) \oplus h_n(Y; G)$. Then apply this in the case $X = Y = \{*\}$ to identify $\tilde{h}_0(\{*\} \amalg \{*\})$ with the kernel of the map

$$\mathbb{1} \oplus \mathbb{1} : G \oplus G \rightarrow G : (g, h) \mapsto g + h,$$

which is isomorphic to G .

EXERCISE 28.6 (*). Given an axiomatic homology theory h_* with coefficient group G , use Theorem 28.24, Exercise 28.5 and an inductive argument to derive a general formula for $\tilde{h}_k(S^n)$ for all $k \in \mathbb{Z}$ and $n \geq 0$, and then deduce from it Theorem 28.25.

EXERCISE 28.7. One of the most popular simple applications of homology is the Brouwer fixed point theorem, which states that for the closed disk $\mathbb{D}^n \subset \mathbb{R}^n$ of any dimension $n \in \mathbb{N}$, every continuous map $f : \mathbb{D}^n \rightarrow \mathbb{D}^n$ has a fixed point.

- (a) Deduce the Brouwer fixed point theorem from the following statement: For each $n \in \mathbb{N}$, the disk \mathbb{D}^n does not admit any retraction to its boundary $\partial \mathbb{D}^n = S^{n-1}$.

Hint: If $f : \mathbb{D}^n \rightarrow \mathbb{D}^n$ has no fixed points, then there is a unique line through x and $f(x)$ for every $x \in \mathbb{D}^n$.

- (b) Assuming the existence of an axiomatic homology theory h_* with a nontrivial coefficient group, deduce from the computation of $h_*(S^{n-1})$ that retractions $\mathbb{D}^n \rightarrow S^{n-1}$ cannot exist.

EXERCISE 28.8 (*). The subject of this exercise is a standard tool in homological algebra known as the **five-lemma**.

- (a) Suppose the following diagram commutes and that both of its rows are exact, meaning $\text{im } f = \ker g$, $\text{im } g' = \ker h'$ and so forth:

$$\begin{array}{ccccccccc} A & \xrightarrow{f} & B & \xrightarrow{g} & C & \xrightarrow{h} & D & \xrightarrow{i} & E \\ \downarrow \alpha & & \downarrow \beta & & \downarrow \gamma & & \downarrow \delta & & \downarrow \varepsilon \\ A' & \xrightarrow{f'} & B' & \xrightarrow{g'} & C' & \xrightarrow{h'} & D' & \xrightarrow{i'} & E' \end{array}$$

Prove that if α, β, δ and ε are all isomorphisms, then so is γ .

- (b) Here is an application: given a homology theory h_* and a map of pairs $f : (X, A) \rightarrow (Y, B)$, show that if any two of the induced maps $f_* : h_n(X) \rightarrow h_n(Y)$, $f_* : h_n(A) \rightarrow h_n(B)$ and $f_* : h_n(X, A) \rightarrow h_n(Y, B)$ are isomorphisms for every n , then so is the third.

- (c) Prove that every homology theory h_* also satisfies a relative version of the additivity axiom, involving disjoint unions of pairs of spaces

$$\coprod_{\beta \in J} (X_\beta, A_\beta) := \left(\coprod_{\beta \in J} X_\beta, \coprod_{\beta \in J} A_\beta \right).$$

29. Simplicial homology

As mental preparation for the definition of singular homology, it will be helpful to start with a different theory that is similar but more restrictive. Simplicial homology requires strictly more data for its definition than just a topological space, and thus can only be defined on spaces that are “nice” enough to admit such data. What it lacks in generality, it makes up for in computability and geometric transparency. One can think of simplicial homology as a combinatorial variant of bordism theory, one that is based on simpler building blocks than manifolds, and can thus be studied without any understanding of the (generally difficult) problem of classifying manifolds.

29.1. Simplicial complexes and polyhedra. The spaces on which simplicial homology is defined are called polyhedra, and they are much more restrictive than arbitrary topological spaces, but nonetheless include most of the typical examples of interest, e.g. all smooth manifolds. Intuitively, a polyhedron is a space that can be constructed by gluing together “triangles” of various dimensions, and the resulting decomposition of a polyhedron into “triangular” pieces is therefore known as a *triangulation*. The first necessary step is to define the n -dimensional generalization of a triangle.

DEFINITION 29.1. For an integer $n \geq 0$, the **standard n -simplex** is the topological space

$$\Delta^n := \{(t_0, \dots, t_n) \in \mathbb{R}^{n+1} \mid t_0 + \dots + t_n = 1\},$$

endowed with the subspace topology as a subset of \mathbb{R}^{n+1} . The $n+1$ standard basis vectors of \mathbb{R}^{n+1} are called the **vertices** (*Eckpunkte*) of Δ^n , and for arbitrary subsets $J \subset \{0, \dots, n\}$, the sets of the form

$$\{(t_0, \dots, t_n) \in \Delta^n \mid t_j = 0 \text{ for all } j \in J\}$$

are called the **faces** (*Seiten* or *Facetten*) of Δ^n ; these include in particular the $n+1$ **boundary faces** (*Seitenflächen*)

$$\partial_{(j)} \Delta^n := \{(t_0, \dots, t_n) \in \Delta^n \mid t_j = 0\}, \quad j = 0, \dots, n.$$

This definition 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. Observe that every face of Δ^n is homeomorphic to Δ^k for some $k \leq n$, and since the coordinates of \mathbb{R}^{n+1} come with a canonical ordering, there is even a canonical choice of homeomorphism. For instance, the boundary faces $\partial_{(j)} \Delta^n$ are all homeomorphic to Δ^{n-1} , and the canonical homeomorphisms take the form

$$(29.1) \quad \Delta^{n-1} \xrightarrow{\cong} \partial_{(j)} \Delta^n : (t_0, \dots, t_{n-1}) \mapsto (t_0, \dots, t_{j-1}, 0, t_j, \dots, t_{n-1}).$$

We will make frequent use of these canonical homeomorphisms to identify each face of a standard simplex with another standard simplex.

In order to explain how copies of Δ^n for various $n \geq 0$ can be glued together to form a polyhedron, we need to define simplicial complexes, which are fundamentally combinatorial objects.

DEFINITION 29.2. A **simplicial complex** (*Simplizialkomplex*) $K = (V, S)$ consists of two sets V and S , called the sets of **vertices** (*Eckpunkte*) and **simplices** (*Simplizes*) respectively, where the elements of S are 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

$$\dim K := \sup_{\sigma \in S} \dim \sigma \in \{0, 1, 2, \dots, \infty\},$$

where we write $\dim \sigma = n$ whenever σ is an n -simplex.

DEFINITION 29.3. A **subcomplex** $K' \subset K$ of a simplicial complex $K = (V, S)$ is a simplicial complex $K' = (V', S')$ such that $V' \subset V$ and $S' \subset S$.

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 $V \rightarrow I$, 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 $\sigma = \{v_0, \dots, v_n\}$ in K , we define the set

$$|\sigma| := \left\{ t \in I^V \mid \sum_{v \in \sigma} t_v = 1 \text{ and } t_v = 0 \text{ for all } v \notin \sigma \right\}.$$

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 useful choice here. We will instead let the topology of $|K|$ be determined by that of the individual simplices:

DEFINITION 29.4. Given a simplicial complex $K = (V, S)$, the topology of its polyhedron $|K| \subset I^V$ is defined such that a subset $\mathcal{U} \subset |K|$ is open if and only if $\mathcal{U} \cap |\sigma|$ is an open subset of $|\sigma|$ for every simplex $\sigma \in S$.

In other words, $|K|$ is equipped with the strongest⁵¹ 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

⁵¹For 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 29.1).

was already said for the case where K is finite, and you should then prove the following proposition as an exercise:

PROPOSITION 29.5. *For any simplicial complex $K = (V, S)$ and any space X , a map $f : |K| \rightarrow X$ is continuous if and only if $f|_{|\sigma|} : |\sigma| \rightarrow X$ is continuous for every simplex $\sigma \in S$. \square*

DEFINITION 29.6. A topological space X is a **polyhedron** (*Polyeder*) if it is homeomorphic to the polyhedron $|K|$ of some simplicial complex K . A choice of such a homeomorphism $X \cong |K|$ is called a **triangulation** (*Triangulierung*) or **simplicial decomposition** of the space X .

REMARK 29.7. The definition of the term *triangulation* given above is perhaps stricter than some other sensible definitions of this term that one could imagine. What everyone can agree upon is that a triangulation of X should decompose X as a union of compact subsets, each of which is homeomorphic to a standard simplex, such that the intersection of any two of them is a common face of both; this includes the case where one of them is a face of the other, but also cases in which their interiors are disjoint. Definition 29.6 does decompose X in this way, but having a specific choice of homeomorphism $X \cong |K|$ is actually a lot more information, and it is debateable whether this amount of information is truly necessary for most of the important applications of triangulations. It will be useful for our purposes, however, when we want to write down precise relations between the simplicial and singular homologies of a triangulated space.

DEFINITION 29.8. For each integer $n \geq 0$, the **n -skeleton** (*n -Skelett* or *n -Gerüst*) of a simplicial complex $K = (V, S)$ is the subcomplex $K^n = (V, S^n)$ of K whose set of simplices $S^n \subset S$ consists of all $\sigma \in S$ with $\dim \sigma \leq n$. Similarly, the **n -skeleton** of a polyhedron X with triangulation $X \cong |K|$ is the subspace $X^n \subset X$ formed by the polyhedron of the n -skeleton K^n of K .

This definition presents a polyhedron X as the union of a nested sequence of subspaces, its skeleta of various dimensions,

$$X^0 \subset X^1 \subset X^2 \subset \dots \subset \bigcup_{n=0}^{\infty} X^n = X,$$

each of which is also a polyhedron. In particular, 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 29.9. 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 29.10. 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} .

EXAMPLE 29.11. Suppose $V = \{v_0, \dots, v_n\}$ for some $n \geq 2$ and S is defined to consist of all the one-element subsets $\{v_i\}$ plus the 1-simplices $\{v_i, v_{i+1}\}$ for $i = 0, \dots, n - 1$ and $\{v_n, v_0\}$. Then $|K|$ is a 1-dimensional polyhedron homeomorphic to S^1 .

EXAMPLE 29.12. 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 29.13. 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 $\{1, \dots, n\}$ for n sufficiently large, thus you can try to picture $|K|$ as the union of an infinite nested sequence of simplices $\Delta^0 \subset \Delta^1 \subset \Delta^2 \subset \dots$, where each Δ^k is a boundary face of Δ^{k+1} .

DEFINITION 29.14. Given two simplicial complexes $K_1 = (V_1, S_1)$ and $K_2 = (V_2, S_2)$, a **simplicial map** (*simpliciale Abbildung*) from K_1 to K_2 is a function $f : V_1 \rightarrow V_2$ such that $f(\sigma) \in S_2$ for every $\sigma \in S_1$.

Note that a simplicial map $K_1 \rightarrow K_2$ 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 $k \leq n$. There is a natural way to turn any simplicial map into a continuous map of the polyhedra $|K_1| \rightarrow |K_2|$. 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 formal⁵² 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

$$f : |K_1| \rightarrow |K_2| : \sum_{v \in V_1} t_v e_v \mapsto \sum_{v \in V_1} t_v e_{f(v)} \in I^{V_2}.$$

In other words, for each simplex $\sigma \in S_1$, f maps $|\sigma|$ onto $|f(\sigma)|$ via the restriction of the obvious linear map $\mathbb{R}^\sigma \rightarrow \mathbb{R}^{f(\sigma)}$ that sends basis vectors e_v to $e_{f(v)}$ for $v \in \sigma$. We have thus defined a functor

$$\mathbf{Simp} \rightarrow \mathbf{Top} : K \mapsto |K|,$$

where \mathbf{Simp} is the category of simplicial complexes with morphisms defined to be simplicial maps. Notice that $f : |K_1| \rightarrow |K_2|$ always maps the n -skeleton of $|K_1|$ into the n -skeleton of $|K_2|$ for every $n \geq 0$.

Since we will often be concerned mainly with compact manifolds, the following result enables us to restrict attention to finite simplicial complexes:

PROPOSITION 29.15. *A simplicial complex $K = (V, S)$ is finite if and only if its polyhedron $|K|$ is compact.*

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 29.15 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 $\sigma_1, \sigma_2, \dots \in S$ 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 $\mathcal{U}_i \cap |\sigma| = \emptyset$ 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, \dots . This proves that the infinite subset $\{x_1, x_2, \dots\} \subset |K|$ is discrete, hence $|K|$ cannot be compact. \square

⁵²The 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.

29.2. The category of chain complexes. In absolute bordism theory, crucial roles are played by the words “closed” and “boundary”: elements are represented by maps defined on *closed* manifolds rather than manifolds that are noncompact or have boundary, and the equivalence relation arises from the fact that certain manifolds are the boundaries of others. One trivial and yet important detail here is the fact that for a compact manifold M with boundary, its boundary ∂M is always a closed manifold: the compactness of ∂M is automatic since $\partial M \subset M$ is a closed subset, but being a boundary also means that ∂M cannot have any boundary points of its own.

In the usual constructions of homology theories—which do not require any knowledge of manifolds—there is an algebraic device that gives useful meaning to the words *closed* and *boundary*, and the fact that boundaries have no boundary of their own is then encoded by a simple algebraic equation, taking the form “ $\partial^2 = 0$ ”.

DEFINITION 29.16. A **chain complex** (*Kettenkomplex*) of R -modules is a sequence of R -modules taking the form

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

and satisfying the relation

$$(29.2) \quad \partial_n \circ \partial_{n+1} = 0$$

for every $n \in \mathbb{Z}$.

Let’s add some helpful terminology and notation to the definition above. The collection of R -modules C_n forming a chain complex can be packaged together as a single R -module

$$C_* := \bigoplus_{n \in \mathbb{Z}} C_n,$$

and writing $\partial : C_* \rightarrow C_*$ for the homomorphism determined uniquely by the maps $\partial_n : C_n \rightarrow C_{n-1}$ for all n , the defining relation (29.2) is then written succinctly as

$$\partial^2 = 0.$$

The chain complex itself can then be denoted by (C_*, ∂) , often abbreviated simply as C_* . We call ∂ the **boundary map** or **boundary operator** (*Randoperator*) of the complex. An element $c \in C_*$ is said to be **homogeneous** (*homogen*) if it belongs to the specific submodule $C_n \subset C_*$ for some $n \in \mathbb{Z}$, which is then called the **degree** (*Grad*) of c , sometimes written as

$$|c| := n \quad \text{for } c \in C_n,$$

and the homogeneous elements of degree n are also called the **n -chains** (*n -Ketten*) of the complex. We say that $c \in C_*$ is **closed** (*geschlossen*) if it satisfies

$$\partial c = 0,$$

and the closed n -chains are called the **n -cycles** (*n -Zykel*) of the complex. Further, $c \in C_*$ is a **boundary** (*Rand*) if it satisfies

$$c = \partial a \quad \text{for some } a \in C_*,$$

and the n -cycles that are also boundaries are called the **n -boundaries**. The relation $\partial^2 = 0$ is equivalent to the condition that all boundaries are also cycles, in other words, $\text{im } \partial_{n+1}$ is always a submodule of $\ker \partial_n$.

REMARK 29.17. For the boundary map $\partial : C_* \rightarrow C_*$ of a chain complex, one sometimes abuses notation and writes

$$\partial : C_* \rightarrow C_{*-1}$$

to emphasize the fact that ∂ sends n -chains to $(n-1)$ -chains for each n . A fancier way to say this is that C_* is naturally a \mathbb{Z} -graded R -module, and the boundary map is a homomorphism of **degree** -1 .

DEFINITION 29.18. The **homology** $H_*(C_*) = H_*(C_*, \partial)$ of the chain complex C_* is the collection of quotient modules

$$H_n(C_*) := \ker(\partial_n) / \operatorname{im}(\partial_{n+1}).$$

Their direct sum is denoted by

$$H_*(C_*) = \bigoplus_{n \in \mathbb{Z}} H_n(C_*).$$

Given a chain complex C_* , elements $[c] \in H_n(C_*)$ are called **homology classes** of degree n : their representatives $c \in C_n$ are n -cycles, and two such n -cycles c, c' represent the same homology class if and only if $c' - c$ is a boundary, in which case we say that they are **homologous**.

DEFINITION 29.19. Given two chain complexes (A_*, ∂^A) and (B_*, ∂^B) , a **chain map** (*Kettenabbildung*) from (A_*, ∂^A) to (B_*, ∂^B) is a collection of homomorphisms $f_n : A_n \rightarrow B_n$ for $n \in \mathbb{Z}$ such that the following diagram commutes:

$$(29.3) \quad \begin{array}{ccccccc} \dots & \longrightarrow & A_{n+1} & \xrightarrow{\partial_{n+1}^A} & A_n & \xrightarrow{\partial_n^A} & A_{n-1} & \xrightarrow{\partial_{n-1}^A} & \dots \\ & & \downarrow f_{n+1} & & \downarrow f_n & & \downarrow f_{n-1} & & \\ \dots & \longrightarrow & B_{n+1} & \xrightarrow{\partial_{n+1}^B} & B_n & \xrightarrow{\partial_n^B} & B_{n-1} & \xrightarrow{\partial_{n-1}^B} & \dots \end{array}$$

In other words, a chain map is a homomorphism $f : A_* \rightarrow B_*$ that maps n -chains to n -chains for each $n \in \mathbb{Z}$ and satisfies $\partial^B \circ f = f \circ \partial^A$.

It is easy to check that the composition of two chain maps is also a chain map, and so is the identity map on any chain complex, thus we can define a category

$$\operatorname{Ch}(R\text{-Mod}) \quad \text{often abbreviated as} \quad \operatorname{Ch},$$

whose objects are chain complexes of R -modules, with chain maps as morphisms. The following easy observation then produces a functor

$$\operatorname{Ch}(R\text{-Mod}) \xrightarrow{H_n} R\text{-Mod}$$

for each $n \in \mathbb{Z}$, sending each chain complex to its homology in degree n and each chain map to the induced homomorphism between homologies.

PROPOSITION 29.20. Any chain map $f : (A_*, \partial^A) \rightarrow (B_*, \partial^B)$ determines homomorphisms $f_* : H_n(A_*, \partial^A) \rightarrow H_n(B_*, \partial^B)$ 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 $\ker \partial_n^A / \operatorname{im} \partial_{n+1}^A \rightarrow \ker \partial_n^B / \operatorname{im} \partial_{n+1}^B$. This is equally clear, since $a = \partial^A x$ implies $f(a) = f(\partial^A x) = \partial^B f(x)$. \square

29.3. Ordered simplicial homology. We now describe the first of two versions of the so-called **simplicial chain complex** (*simplizialer Kettenkomplex*) of a simplicial complex $K = (V, S)$, the homology of which will be the **simplicial homology** (*simpliziale Homologie*) of K . We will see later that with a bit of care, simplicial homology can be defined as a collection of functors on the subcategory of **Top** consisting of all polyhedra, without needing to specify how each polyhedron is triangulated. For now, however, the definition of the simplicial homology groups will depend explicitly on a simplicial complex, and thus gives us functors $\mathbf{Simp} \rightarrow R\text{-Mod}$.

The first version of the simplicial chain complex is algebraically simpler than the second, while the second will be easier to interpret geometrically. In practice, we will eventually be able to choose freely between them, because (for slightly nontrivial reasons) their homologies turn out to be naturally isomorphic.

REMARK 29.21. For readers who have seen the definition of simplicial homology in the first semester of these notes (cf. Lecture 21): the complex defined in §29.4 below is cosmetically different from the one that was defined there, but is easily seen to be isomorphic to it (see Remark 29.24). 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. For the rest of this lecture, and in fact for most of the rest of this course, you should assume that

$$G \in R\text{-Mod}$$

is an arbitrary choice of R -module, which will typically play the role of the coefficient group in whichever version of homology is under discussion. We will include G in the notation for homology in situations where the choice of coefficient group matters, but omit it whenever this choice plays no important role.

Given a simplicial complex $K = (V, S)$, define the set

$$\mathcal{K}_n^o(K) := \left\{ (v_0, \dots, v_n) \in V^{\times(n+1)} \mid \text{there exists a } \sigma \in S \text{ with } v_i \in \sigma \text{ for all } i = 0, \dots, n \right\}$$

for each $n \geq 0$. 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, \dots, v_n are all distinct, though if they are, then it means $\{v_0, \dots, v_n\} \in S$ is an n -simplex of the complex K , and the ordered tuple (v_0, \dots, 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 with

$$C_n^o(K) = \bigoplus_{\sigma \in \mathcal{K}_n^o(K)} G$$

for each $n \geq 0$, so that n -chains can be written uniquely as finite sums $\sum_i a_i \sigma_i$ with coefficients $a_i \in G$ attached to canonical generators $\sigma_i \in \mathcal{K}_n^o(K)$. In particular, if the coefficient module G is taken to be the ring R itself, then $C_n^o(K)$ is the free R -module over the set $\mathcal{K}_n^o(K)$; in the case $R = \mathbb{Z}$, it is thus a free abelian group. Linearity and the formula

$$(29.4) \quad \partial(v_0, \dots, v_n) := \sum_{k=0}^n (-1)^k (v_0, \dots, v_{k-1}, v_{k+1}, \dots, v_n)$$

uniquely determine a boundary map $\partial : C_n^o(K) \rightarrow C_{n-1}^o(K)$ on this complex for each $n \geq 1$, and we define $C_n^o(K)$ to be trivial for each $n < 0$, so that $\partial : C_n^o(K) \rightarrow C_{n-1}^o(K)$ is necessarily trivial

for each $n \leq 0$. It is a straightforward exercise in sign cancellations to verify that ∂ satisfies $\partial^2 = 0$. The resulting homology groups

$$H_*^o(K) := H_*^o(K; G) := H_*(C_*^o(K; G), \partial)$$

will be called the **ordered simplicial homology** of K with coefficients in G .

In order to view ordered simplicial homology as a functor, we associate to each simplicial map $f : K_1 \rightarrow K_2$ and each $n \geq 0$ the unique R -module homomorphism

$$f_* : C_n^o(K_1) \rightarrow C_n^o(K_2)$$

determined by linearity and the formula

$$f_*(v_0, \dots, v_n) := (f(v_0), \dots, f(v_n)).$$

It is straightforward to check that this defines a chain map $C_*^o(K_1) \rightarrow C_*^o(K_2)$, and thus gives us a functor

$$C_*^o : \text{Simp} \rightarrow \text{Ch}(R\text{-Mod}).$$

Composing this with the algebraic homology functors $H_n : \text{Ch}(R\text{-Mod}) \rightarrow R\text{-Mod}$ gives us functors

$$H_n^o : \text{Simp} \rightarrow R\text{-Mod};$$

in particular, simplicial maps $f : K_1 \rightarrow K_2$ induce R -module homomorphisms $f_* : H_n^o(K_1) \rightarrow H_n^o(K_2)$ for every n .

29.4. Oriented simplicial homology. The second version of the simplicial chain complex has a similar but smaller set of generators, because it excludes tuples (v_0, \dots, 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 result makes this possible; its proof is an exercise.

LEMMA 29.22. *For each $n \geq 1$, the boundary map $\partial : C_n^o(K; \mathbb{Z}) \rightarrow C_{n-1}^o(K; \mathbb{Z})$ defined via (29.4) preserves the subgroup of $C_*^o(K; \mathbb{Z})$ generated by all elements of the form*

$$(29.5) \quad (v_0, \dots, v_n) \in C_n^o(K; \mathbb{Z}) \quad \text{with } v_i = v_j \text{ for some } i \neq j$$

or of the form

$$(29.6) \quad (v_0, \dots, v_n) - (-1)^{|\tau|} (v_{\tau(0)}, \dots, v_{\tau(n)}) \in C_n^o(K; \mathbb{Z})$$

for arbitrary $(v_0, \dots, v_n) \in \mathcal{K}_n^o(K)$ and permutations $\tau \in S_{n+1}$, where $(-1)^{|\tau|} = \pm 1$ denotes the sign of the permutation. \square

DEFINITION 29.23. An **orientation** (*Orientierung*) of an n -simplex $\sigma \in S$ for $n \geq 1$ 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, \dots, v_n can be written with the notation

$$\pm[v_0, \dots, v_n],$$

which is understood to mean the simplex $\{v_0, \dots, v_n\}$ with orientation determined by the ordering v_0, \dots, 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 Lemma 29.22, 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 submodule generated by products of arbitrary coefficients $g \in G$ with elements of the form (29.5) or (29.6); this is sometimes called the group of **degenerate chains**. For each generator (v_0, \dots, v_n) of $C_n^o(K; \mathbb{Z})$, we shall denote the equivalence class that it represents in the quotient complex by

$$[v_0, \dots, v_n] \in C_n^\Delta(K; \mathbb{Z}).$$

This means

$$[v_0, \dots, v_n] = 0 \quad \text{if } v_i = v_j \text{ for some } i \neq j,$$

whereas if the vertices v_0, \dots, v_n are all distinct, then $[v_0, \dots, v_n]$ 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

$$[v_0, \dots, v_i, \dots, v_j, \dots, v_n] = -[v_0, \dots, v_j, \dots, v_i, \dots, v_n]$$

for each pair $i \neq j$ in $\{0, \dots, n\}$. The boundary map $\partial : C_n^\Delta(K) \rightarrow C_{n-1}^\Delta(K)$ is thus determined by the formula

$$(29.7) \quad \partial[v_0, \dots, v_n] = \sum_{k=0}^n (-1)^k [v_0, \dots, v_{k-1}, v_{k+1}, \dots, v_n],$$

and Lemma 29.22 guarantees that this formula is independent of the order in which the vertices are written. We will denote the resulting **oriented simplicial homology** by

$$H_*^\Delta(K) := H_*(C_*^\Delta(K)).$$

One checks easily that the chain maps $f_* : C_*^o(K_1) \rightarrow C_*^o(K_2)$ induced by any simplicial map $f : K_1 \rightarrow K_2$ descend to the quotient as chain maps

$$f_* : C_*^\Delta(K_1) \rightarrow C_*^\Delta(K_2),$$

thus giving a functor

$$C_*^\Delta : \text{Simp} \rightarrow \text{Ch}(R\text{-Mod}),$$

which composes with the algebraic homology functor to produce functors

$$H_n^\Delta : \text{Simp} \rightarrow R\text{-Mod}$$

for each $n \geq 0$.

Notice moreover that since the chain complex $C_*^\Delta(K)$ is defined as a quotient of $C_*^o(K)$, the quotient projection

$$C_*^o(K) \rightarrow C_*^\Delta(K) : (v_0, \dots, v_n) \mapsto [v_0, \dots, v_n]$$

is also a chain map, and thus induces a natural sequence of homomorphisms

$$H_n^o(K) \rightarrow H_n^\Delta(K), \quad n \geq 0.$$

The word “natural” is meant here in its technical sense, as the map from ordered to oriented simplicial homology can be seen as a natural transformation between two functors $\text{Simp} \rightarrow R\text{-Mod}$.

We will later see in fact that on the homology level (though not on the level of chain complexes), these maps are always isomorphisms. This fact, however, requires a lengthier discussion, and we do not need it just yet.

REMARK 29.24. While it is not so obvious from the definition above, $C_n^\Delta(K) = C_n^\Delta(K; G)$ for a simplicial complex K can be identified with a complex of the form

$$C_n^\Delta(K) \cong \bigoplus_{\sigma \in \mathcal{K}_n^\Delta(K)} G,$$

where the set $\mathcal{K}_n^\Delta(K)$ of generators consists of all n -simplices $\sigma = \{v_0, \dots, v_n\}$ in the complex K . This perspective gives $C_n^\Delta(K)$ a similar formal structure to that of $C_n^o(K)$, so that for instance $C_n^\Delta(K; \mathbb{Z})$ is also a free abelian group, but with a smaller and more manageable set of generators than $C_n^o(K; \mathbb{Z})$. Indeed, each generator of $C_n^o(K)$ is an ordered tuple (v_0, \dots, v_n) of vertices in a simplex, but the generators of $C_n^\Delta(K)$ are instead actual n -simplices $\{v_0, \dots, v_n\}$ of the complex K , meaning that the vertices v_0, \dots, v_n are required to be distinct, and the order in which they are written does not matter. Some choices are required, however, before $C_n^\Delta(K)$ can be presented in this way: if one makes an arbitrary choice of orientation for each simplex $\{v_0, \dots, v_n\}$ of K and writes its vertices in an order consistent with the chosen orientation, then the resulting oriented n -simplex $[v_0, \dots, v_n]$ can be used as a generator of $C_n^\Delta(K)$, and there is no need to consider other permutations of the vertices v_0, \dots, v_n . Writing down $\partial : C_n^\Delta(K) \rightarrow C_{n-1}^\Delta(K)$ then requires taking some care with signs, to account for the fact that the arbitrarily chosen orientations of the $(n-1)$ -simplices of K may or may not agree with the orientations of the boundary faces appearing in the usual formula for $\partial[v_0, \dots, v_n]$. The result is essentially the definition of $H_n^\Delta(K)$ 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 (see e.g. Figure 15). The alternative formulation as a quotient complex shows why it does not actually depend on the choices of orientations.

Comparing the ordered and oriented simplicial chain complexes, the oriented complex has a more obvious geometric interpretation, because its generators are in bijective correspondence with actual simplices. 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, the *ordered* complex is the one that admits a straightforward relationship with the singular homology of a polyhedron.

29.5. Exercises.

EXERCISE 29.1 (*). Prove Proposition 29.5: For any simplicial complex $K = (V, S)$ and any space X , a map $f : |K| \rightarrow X$ is continuous if and only if $f|_{|\sigma|} : |\sigma| \rightarrow X$ is continuous for every simplex $\sigma \in S$.

EXERCISE 29.2 (*). Prove Lemma 29.22, which establishes that the definition of ∂ on the oriented simplicial chain complex makes sense.

Hint: One does not really need to examine all possible tuples (v_0, \dots, v_n) and all of their permutations. It suffices to check cases where $v_k = v_{k+1}$ for some k , and permutations that interchange two neighboring elements.

EXERCISE 29.3. Figure 16 shows a simplicial complex $K = (V, S)$ whose associated polyhedron $|K|$ is homeomorphic to the Klein bottle. There are nine vertices labeled P_i, Q_i, R_i for $i = 1, 2, 3$, twenty-seven 1-simplices labeled by letters $a_i, b_i, c_i, d_i, e_i, f_i$ for $i = 1, 2, 3$ and g_i for $i = 1, \dots, 9$, and eighteen 2-simplices labeled σ_i, τ_i for $i = 1, \dots, 9$. The picture also shows a choice of orientation for each of the 2-simplices (circular arrows represent a cyclic ordering of the vertices) and 1-simplices (arrows point from the first vertex to the last). If we additionally endow each 0-simplex with the




FIGURE 15. The picture shows a simplicial complex K with polyhedron $|K| \cong \mathbb{T}^2$, 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 $\partial : C_n^\Delta(K; \mathbb{Z}) \rightarrow C_{n-1}^\Delta(K; \mathbb{Z})$ gives e.g. $\partial\sigma_1 = g_1 - a_1 - d_3$, $\partial\tau_1 = b_1 + e_3 - g_1$, $\partial a_1 = P_2 - P_1$, $\partial a_2 = P_3 - P_2$, $\partial a_3 = P_1 - P_3$, 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 eighteen 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 $H_0^\Delta(K; \mathbb{Z}) \cong \mathbb{Z} \cong H_0(\mathbb{T}^2; \mathbb{Z})$.

positive orientation, every letter in the picture can be regarded as representing an oriented simplex, and thus a generator of the oriented simplicial chain complex $C_*^\Delta(K; \mathbb{Z})$.

- (a) Write down the 1-chains $\partial\sigma_i, \partial\tau_i \in C_1^\Delta(K; \mathbb{Z})$ explicitly for each $i = 1, \dots, 9$.
- (b) Prove that $H_2^\Delta(K; \mathbb{Z}_2) \cong \mathbb{Z}_2$, and write down a specific cycle in $C_2^\Delta(K; \mathbb{Z}_2)$ that generates it.
- (c) Prove that $H_2^\Delta(K; \mathbb{Z}) = 0$.




FIGURE 16. The Klein bottle as a polyhedron.

- (d) Show that the 1-cycle $d_1 + d_2 + d_3$ represents a nontrivial homology class $[d_1 + d_2 + d_3]$ in both $H_1^\Delta(K; \mathbb{Z})$ and $H_1^\Delta(K; \mathbb{Z}_2)$, but satisfies $2[d_1 + d_2 + d_3] = 0 \in H_1^\Delta(K; \mathbb{Z})$ and $[d_1 + d_2 + d_3] = 0 \in H_1^\Delta(K; \mathbb{Q})$.

EXERCISE 29.4 (*). The following computations may give you a hint as to why $h_0(\{*\})$ is called the *coefficient group* of an axiomatic homology theory h_* . In the simplicial context, let $\{*\}$ denote a simplicial complex that has exactly one vertex.

- (a) Prove that $H_0^\Delta(\{*\}; G) \cong G$ and $H_n^\Delta(\{*\}; G) = 0$ for all $n \neq 0$.
Hint: This is nearly trivial.
- (b) Prove that $H_0^o(\{*\}; G) \cong G$ and $H_n^o(\{*\}; G) = 0$ for all $n \neq 0$.
Remark: This is slightly less trivial than part (a), but not difficult.

30. Triangulated manifolds and subdivision

I claimed in the previous lecture that simplicial homology can be viewed as a combinatorial variant of bordism theory. To see what I mean by this, we need to talk about manifolds with triangulations.

30.1. Triangulated manifolds. In this lecture, we will not need any knowledge of smoothness, so the word “manifold” means *topological* manifold, i.e. a second countable Hausdorff space that is locally homeomorphic to a finite-dimensional vector space or half-space. It should be assumed that all manifolds M may have nonempty boundary ∂M unless stated otherwise.

DEFINITION 30.1. An n -dimensional **triangulated manifold** is a topological n -manifold M equipped with a triangulation $M \cong |K|$ that identifies ∂M with the polyhedron of a subcomplex $K' \subset K$.

The following is a consequence of the local Euclidean structure of manifolds:

PROPOSITION 30.2. *If $M \cong |K|$ is a triangulated n -dimensional manifold, then the associated simplicial complex K is n -dimensional, and every $(n - 1)$ -simplex σ in K is a boundary face of either one or two n -simplices, where the former is the case if and only if σ belongs to the subcomplex triangulating ∂M . \square*

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 a compact n -manifold into an n -chain in the simplicial chain complex of K .

It is easiest to explain how this works in $H^\Delta(K; \mathbb{Z}_2)$. Using \mathbb{Z}_2 as a coefficient group has the advantage that for any n -simplex $\sigma = \{v_0, \dots, v_n\}$ of K , we have

$$[v_0, \dots, v_n] = -[v_0, \dots, v_n] \in C_n^\Delta(K; \mathbb{Z}_2),$$

so that all choices of ordering for the vertices v_0, \dots, v_n produce the same element, and there is thus no need to worry about orientations. Given a compact triangulated n -manifold $M \cong |K|$, we can define an oriented simplicial n -chain by

$$(30.1) \quad c_M := \sum_{\sigma} \mathbf{v}_{\sigma} \in C_n^\Delta(K; \mathbb{Z}_2),$$

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 an arbitrary order. Note that this definition would not make sense if M were not compact, but according to Proposition 29.15, compactness implies that K is a *finite* simplicial complex, so that the sum in the definition of c_M is finite. If $\partial M \neq \emptyset$, then the subcomplex $K' \subset K$ triangulating ∂M similarly defines a simplicial $(n - 1)$ -chain

$$c_{\partial M} \in C_{n-1}^\Delta(K'; \mathbb{Z}_2) \subset C_{n-1}^\Delta(K; \mathbb{Z}_2),$$

where we are regarding $C_{n-1}^\Delta(K'; \mathbb{Z}_2)$ as a submodule of $C_{n-1}^\Delta(K; \mathbb{Z}_2)$, which makes sense because the canonical generators of $C_{n-1}^\Delta(K'; \mathbb{Z}_2)$ (i.e. the $(n - 1)$ -simplices of K') are also $(n - 1)$ -simplices of K and thus generators of $C_{n-1}^\Delta(K; \mathbb{Z}_2)$. If $\partial M = \emptyset$, then the recipe above defines the trivial $(n - 1)$ -chain, and we can therefore sensibly write

$$c_{\partial M} = 0 \in C_{n-1}^\Delta(K; \mathbb{Z}_2) \quad \text{if } \partial M = \emptyset.$$

PROPOSITION 30.3. *The chains $c_M \in C_n^\Delta(K; \mathbb{Z}_2)$ and $c_{\partial M} \in C_{n-1}^\Delta(K'; \mathbb{Z}_2) \subset C_{n-1}^\Delta(K; \mathbb{Z}_2)$ in the situation above satisfy*

$$\partial c_M = c_{\partial M}.$$

PROOF. By Proposition 30.2, applying ∂ to the right hand side of (30.1) produces exactly two copies of each $(n-1)$ -simplex of K that is not in K' , 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 $c_{\partial M}$. \square

The proposition implies in particular that whenever M is a *closed* triangulated n -manifold, the n -chain c_M is a cycle, and thus represents a homology class

$$[M] := [c_M] \in H_n^\Delta(K; \mathbb{Z}_2).$$

We call this the (simplicial) **fundamental class** of M , and refer to c_M as a **fundamental cycle**. In the case $\partial M \neq \emptyset$, we will see when we discuss *relative* simplicial homology that c_M still represents a relative homology class for the triangulated pair of spaces $(M, \partial M)$, thus the terms *fundamental cycle* and *fundamental class* remain appropriate.

Fundamental cycles and classes can also be defined in ordered simplicial homology, but this requires some choices.

DEFINITION 30.4. 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. Since we are only talking about compact manifolds in this lecture, our simplicial complexes are always finite, 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 again that $M \cong |K|$ is a compact triangulated n -manifold, and let $K' \subset K$ denote the subcomplex whose polyhedron is identified with ∂M . Working with \mathbb{Z}_2 coefficients, any choice of admissible ordering for K determines an ordered simplicial n -chain of the form

$$c_M := \sum_{\sigma} \mathbf{v}_{\sigma} \in C_n^o(K; \mathbb{Z}_2),$$

in which the sum ranges again 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 $\partial M \neq \emptyset$, the admissible ordering on K restricts to an admissible ordering on K' , and thus similarly determines an ordered simplicial $(n-1)$ -chain

$$c_{\partial M} \in C_{n-1}^o(K'; \mathbb{Z}_2) \subset C_{n-1}^o(K; \mathbb{Z}_2),$$

and we take $c_{\partial M}$ to be $0 \in C_{n-1}^o(K; \mathbb{Z}_2)$ if $\partial M = \emptyset$. It is easy to verify that the analogue of Proposition 30.3 also holds in this situation, and we thus have

$$\partial c_M = c_{\partial M} \in C_{n-1}^o(K; \mathbb{Z}_2).$$

It is clear from the construction that the natural chain map $C_*^o(K; \mathbb{Z}_2) \rightarrow C_*^\Delta(K; \mathbb{Z}_2)$ sends each ordered fundamental cycle to the oriented fundamental cycle, so when M is closed, it therefore sends an ordered simplicial fundamental class $[M] \in H_n^o(K; \mathbb{Z}_2)$ to the oriented simplicial fundamental class $[M] \in H_n^\Delta(K; \mathbb{Z}_2)$. Once we've proved that the natural map $H_n^o(K; \mathbb{Z}_2) \rightarrow H_n^\Delta(K; \mathbb{Z}_2)$ is an isomorphism, we will be able to deduce from this that the class $[M] \in H_n^\Delta(K; \mathbb{Z}_2)$ is independent

of choices, even though the fundamental cycle $c_M \in C_n^o(K; \mathbb{Z}_2)$ that represents it does depend on the choice of admissible ordering for the complex.

The following result is a worthwhile exercise in the computation of simplicial homology.

THEOREM 30.5. *For any closed and connected triangulated n -manifold $M \cong |K|$, $H_n^\Delta(K; \mathbb{Z}_2)$ is isomorphic to \mathbb{Z}_2 , and its unique nontrivial element is the fundamental class $[M]$. \square*

30.2. Oriented triangulations. In order to extend the construction of fundamental cycles from \mathbb{Z}_2 to integer coefficients, we need triangulations with a bit of extra structure.

DEFINITION 30.6. Suppose $n \geq 1$ 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 30.6 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 (29.7) 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 30.7. In addition to being consistent with our usual formulas for boundary operators on chain complexes, there is some geometric motivation behind Definition 30.6. 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 30.6) of the boundary faces of σ , which form a triangulation of $\partial|\sigma|$.

DEFINITION 30.8. 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. 30.2) 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 30.8. (The trouble is with the 1-simplices labeled d_1, d_2, d_3 .) 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 30.9. The triangulation of S^{n-1} described in Example 29.10 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 cancellation 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 Proposition 30.10 below.

We now consider whether a version of the fundamental cycle $c_M \in C_n^\Delta(K; \mathbb{Z}_2)$ for a compact triangulated n -manifold $M \cong |K|$ with boundary $\partial M \cong |K'|$ can also be defined with integer coefficients. Indeed, suppose that an orientation has been chosen for each of the n -simplices σ of K , and consider an n -chain of the form

$$(30.2) \quad c_M := \sum_{\sigma} \mathbf{v}_{\sigma} \in C_n^\Delta(K; \mathbb{Z}),$$

where as usual the sum ranges over the set of all n -simplices $\sigma = \{v_0, \dots, v_n\}$ in K , and $\mathbf{v}_{\sigma} = [v_0, \dots, v_n]$ is defined by ordering the vertices in accordance with the chosen orientation. Since each $(n-1)$ -simplex of the subcomplex $K' \subset K$ triangulating ∂M is a boundary face of a unique n -simplex, the chosen orientations of the n -simplices determine boundary orientations of the $(n-1)$ -simplices of K' , which we can use to define an $(n-1)$ -chain

$$c_{\partial M} \in C_{n-1}^\Delta(K'; \mathbb{Z}) \subset C_{n-1}^\Delta(K; \mathbb{Z}).$$

The formula

$$\partial c_M = c_{\partial M}$$

is then satisfied if and only if the chosen orientations of the n -simplices satisfy the condition in Definition 30.8: indeed, this condition means that all contributions to ∂c_M from $(n-1)$ -simplices not in ∂M appear in cancelling pairs, while each $(n-1)$ -simplex in ∂M appears exactly once with the correct sign. In the case $\partial M = \emptyset$, $c_M \in C_n^\Delta(K; \mathbb{Z})$ is then a cycle and thus represents an integral fundamental class

$$[M] := [c_M] \in H_n^\Delta(K; \mathbb{Z}).$$

We summarize:

PROPOSITION 30.10. *For any compact triangulated n -manifold $M \cong |K|$ with an oriented triangulation, the induced triangulation of the boundary $\partial M \cong |K'|$ admits a unique orientation for which each $(n-1)$ -simplex of K' is oriented as the boundary of an oriented n -simplex of K . The resulting fundamental cycles in $C_n^\Delta(K; \mathbb{Z})$ as constructed above then satisfy the relation $\partial c_M = c_{\partial M}$.*

PROOF. The discussion preceding the statement showed that if the triangulation of M is oriented and the orientations of its n -simplices are used in defining $c_M \in C_n^\Delta(K; \mathbb{Z})$ and (via boundary orientations) $c_{\partial M} \in C_{n-1}^\Delta(K'; \mathbb{Z}) \subset C_{n-1}^\Delta(K; \mathbb{Z})$, then $\partial c_M = c_{\partial M}$. One detail not yet addressed is that the boundary orientations on the $(n-1)$ -simplices of K' really do define an oriented triangulation of ∂M : this follows from the relation

$$\partial c_{\partial M} = \partial(\partial c_M) = 0,$$

which means that the two contributions to $\partial c_{\partial M}$ from each $(n-2)$ -simplex in ∂M cancel each other. \square

Here is another worthwhile computational exercise:

THEOREM 30.11. *For any closed and connected n -manifold $M \cong |K|$ with an oriented triangulation, $H_n^\Delta(K; \mathbb{Z})$ is isomorphic to \mathbb{Z} , and its fundamental class $[M]$ is a generator. \square*

For an analogue in ordered simplicial homology with integer coefficients, we can again choose an admissible ordering for K , but we need to be aware that the resulting ordering of the vertices (v_0, \dots, v_n) of each simplex σ might not be consistent with the chosen orientation of σ . We can account for this by including appropriate signs in the formula: we define

$$c_M := \sum_{\sigma} \epsilon_{\sigma} \mathbf{v}_{\sigma} \in C_n^{\circ}(K; \mathbb{Z}),$$

where for each n -simplex $\sigma = \{v_0, \dots, v_n\}$ of our oriented triangulation, with vertices arranged in increasing order, we set $\epsilon_\sigma = \pm 1$ so that $\epsilon_\sigma[v_0, \dots, v_n]$ defines the chosen orientation. Defining

$$c_{\partial M} \in C_{n-1}^o(K'; \mathbb{Z}) \subset C_{n-1}^o(K; \mathbb{Z})$$

in the same manner via the admissible ordering and boundary orientation, the same arguments as before prove $\partial c_M = c_{\partial M}$.

30.3. Triangulated bordism. With triangulated manifolds in hand, the similarity between homology and bordism theory can be made more explicit. Suppose

$$X \cong |K|$$

is a space (but not necessarily a manifold) triangulated by a simplicial complex K , and suppose

$$M \cong |L|$$

is a closed triangulated n -manifold with an oriented triangulation. Any simplicial map $\varphi : L \rightarrow K$ induces a continuous map $\varphi : M \rightarrow X$, so that the pair (M, φ) represents an element of the oriented bordism group $\Omega_n^{\text{SO}}(X)$. A corresponding simplicial homology class can be defined by

$$\varphi_*[M] \in H_n^\Delta(K; \mathbb{Z}),$$

using the integral fundamental class $[M] = [c_M] \in H_n^\Delta(L; \mathbb{Z})$ defined via the oriented triangulation of M . Further, suppose there is an oriented bordism (W, Φ) between (M, φ) and another such pair (N, ψ) , equipped with the additional data of an oriented triangulation: more precisely, $W \cong |L'|$ is a compact $(n+1)$ -manifold with an oriented triangulation, $\Phi : L' \rightarrow K$ is a simplicial map inducing the continuous map $\Phi : W \rightarrow X$, and there is a homeomorphism $\partial W \cong M \amalg (-N)$ that identifies the subcomplex triangulating M with K and $\Phi|_M$ with φ . The minus sign in front of N means that we consider $N \subset \partial W$ to be equipped with an oriented triangulation whose n -simplices carry the *opposite* of the boundary orientations they inherit from the oriented $(n+1)$ -simplices triangulating W . With this understood, we can write $\psi := \Phi|_N : N \rightarrow X$ and obtain a simplicial homology class

$$\psi_*[N] \in H_n^\Delta(K; \mathbb{Z})$$

in the same manner as $\varphi_*[M]$, but the triangulated bordism (W, Φ) tells us more: the orientation reversal on N gives the relation

$$\partial c_W = c_{\partial W} = c_M - c_N \in C_n^\Delta(L'; \mathbb{Z}),$$

and plugging this into the chain map $\Phi_* : C_*^\Delta(L'; \mathbb{Z}) \rightarrow C_*^\Delta(K; \mathbb{Z})$, we have

$$\partial(\Phi_* c_W) = \Phi_*(c_M - c_N) = \varphi_* c_M - \psi_* c_N \in C_n^\Delta(K; \mathbb{Z}),$$

implying

$$\varphi_*[M] = \psi_*[N] \in H_n^\Delta(K; \mathbb{Z}).$$

This matches what happens in bordism theory: two simplicial homology classes represented by closed triangulated manifolds with simplicial maps are the same whenever there is a *triangulated* bordism between them. We will see when we study fundamental classes in singular homology that the entire discussion makes sense in that context as well, but without any need for triangulations.

REMARK 30.12. Orientations were needed for all the triangulations in the discussion above because we were working with integer coefficients. If we did not have orientations, the entire discussion would still make sense after uniformly replacing the coefficient group \mathbb{Z} by \mathbb{Z}_2 , and $H_n^\Delta(X; \mathbb{Z}_2)$ thus becomes the combinatorial variant of the *unoriented* bordism group $\Omega_n^{\text{O}}(X)$.

30.4. Barycentric subdivision. I would now like to describe a specific triangulation of the standard n -simplex Δ^n . One can reasonably ask why it is worth bothering to triangulate a simplex: after all, Δ^n is already a polyhedron in a trivial way. But the point of the following construction is to decompose Δ^n into n -simplices that are strictly *smaller*, and iterating the process will then produce triangulations whose individual n -simplices are as small as we like. This will come in handy when we need to prove the formal properties of singular homology, and it also has some important theoretical consequences for simplicial homology, including one ingredient in the proof that $H_*^s(K)$ and $H_*^{\Delta}(K)$ really are invariants of the *polyhedron* $|K|$, and not just of the particular simplicial complex K that is used for triangulating it.

For each $n \geq 0$, the point

$$b_n := \left(\frac{1}{n+1}, \dots, \frac{1}{n+1} \right) \in \Delta^n \subset \mathbb{R}^{n+1}$$

is called the **barycenter** of the standard n -simplex; you should imagine it as the center of mass of Δ^n . The following inductive procedure uniquely determines a decomposition of Δ^n for each $n \geq 0$ into smaller pieces $\delta^n \subset \Delta^n$ that are homeomorphic to Δ^n :

- (1) For $n = 0$, the one-point space Δ^0 cannot be decomposed any further, so its triangulation consists only of a single 0-simplex.
- (2) If the triangulation of Δ^{n-1} has already been defined, then using the canonical identification of the boundary face $\partial_{(k)}\Delta^n$ for each $k = 0, \dots, n$ with Δ^{n-1} , each $(n-1)$ -simplex $\delta^{n-1} \subset \partial_{(k)}\Delta^n$ in its triangulation determines an n -simplex $\delta^n \subset \Delta^n$ as the convex hull of δ^{n-1} and the barycenter b_n .

For a more precise description of barycentric subdivision, we should specify an abstract simplicial complex K along with a homeomorphism $|K| \cong \Delta^n$ defining the triangulation of Δ^n . It is most natural to define K so that its vertices are points in Δ^n , and since Δ^n is a subset of the vector space \mathbb{R}^{n+1} , the following condition becomes relevant:

DEFINITION 30.13. Given a vector space V of dimension at least n , a set of n points in V is said to be in **general position** if they are not contained in any $(n-2)$ -dimensional plane.

For example, three points in a vector space of dimension at least 2 are in general position if they are not colinear. In general, for a given set of points $v_0, \dots, v_n \in V$ with $\dim V \geq n+1$, the unique linear map $\mathbb{R}^{n+1} \rightarrow V$ sending the standard basis of \mathbb{R}^{n+1} to the vectors v_0, \dots, v_n restricts to $\Delta^n \subset \mathbb{R}^{n+1}$ as an *embedding* $\Delta^n \hookrightarrow V$ if and only if the points v_0, \dots, v_n are in general position.

The abstract simplicial complex K arising from the barycentric subdivision of Δ^n can now be described as follows. One first triangulates the standard 0-simplex with the simplicial complex whose only vertex is the one point in $\Delta^0 \subset \mathbb{R}$. Then inductively, having defined triangulations of the boundary faces $\partial_{(k)}\Delta^n \cong \Delta^{n-1}$ via complexes whose vertices are all identified with points in $\partial_{(k)}\Delta^n$, each n -simplex of K is defined to have vertices b_n, v_1, \dots, v_n , where $v_1, \dots, v_n \in \partial_{(k)}\Delta^n$ are the vertices of an $(n-1)$ -simplex in the complex triangulating $\partial_{(k)}\Delta^n \cong \Delta^{n-1}$ for some $k = 0, \dots, n$. The homeomorphism identifying $|K|$ with Δ^n sends each simplex $|\sigma| \subset |K|$ into Δ^n via the restriction of the unique linear map that sends each vertex to itself. That this actually defines a homeomorphism $|K| \cong \Delta^n$ follows from the following proposition, whose proof is Exercise 30.2:

PROPOSITION 30.14. *For the abstract simplicial complex K described above, whose vertices are points in Δ^n :*

- (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 $v_0, \dots, v_k \in \Delta^n$ for some simplex $\{v_0, \dots, v_k\}$ of K .*

□

It is worth pausing a moment now to draw pictures of the barycentric subdivisions of Δ^n for $n = 1, 2$: the subdivision of Δ^1 will have only two 1-simplices, and since Δ^2 has three boundary faces, its subdivision has six 2-simplices. In general, the number of n -simplices in the subdivision of Δ^n will be $(n + 1)!$. You'll find a picture of the subdivision of Δ^3 in [Hat02], among other places; if you ever find a convincing picture of the case $n = 4$, let me know.

The triangulation defined above can be turned into an integral fundamental cycle

$$c_{\Delta^n} \in C_n^\Delta(K; \mathbb{Z}) \quad \text{or} \quad c_{\Delta^n} \in C_n^o(K; \mathbb{Z})$$

after making some additional choices, namely of orientations and/or admissible orderings. 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, \dots, 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 $\pm[v_1, \dots, v_n]$ is its chosen orientation, define the ordering and orientation of the n -simplex $\{b_n, v_1, \dots, v_n\}$ in Δ^n to be given by

$$(b_n, v_1, \dots, v_n) \quad \text{and} \quad \pm(-1)^k[b_n, v_1, \dots, v_n]$$

respectively. Following our usual prescriptions to define fundamental cycles as simplicial n -chains in $C_n^o(K; \mathbb{Z})$ or $C_n^\Delta(K; \mathbb{Z})$, the barycentric subdivisions of Δ^n and its boundary faces are then related via the formula

$$\partial c_{\Delta^n} = \sum_{k=0}^n (-1)^k c_{\partial_{(k)}\Delta^n} \quad \text{in} \quad C_{n-1}^o(K; \mathbb{Z}) \text{ or } C_{n-1}^\Delta(K; \mathbb{Z}),$$

where as usual, $c_{\partial_{(k)}\Delta^n}$ is defined by identifying $\partial_{(k)}\Delta^n$ with Δ^{n-1} and is then regarded as an element of $C_{n-1}^\Delta(K; \mathbb{Z})$ since the vertices of the subdivision of $\partial_{(k)}\Delta^n$ are also vertices of the subdivision of Δ^n .

Since we can now subdivide a standard simplex into smaller simplices of the same dimension, we can also subdivide *any* polyhedron. Indeed, assuming $K = (V, S)$ is an arbitrary simplicial complex, for each simplex $\sigma \in S$ of K , the corresponding subset $|\sigma| \subset |K|$ has a well-defined barycenter $b_\sigma \in |\sigma|$. We can then construct a new simplicial complex $K' = (V', S')$ whose vertices are points in the polyhedron $|K|$, including all the vertices of K plus all the barycenters of its simplices, and such that the simplices of K' correspond to simplices in the barycentric subdivisions of the individual simplices of K . The unique linear map $\mathbb{R}^{V'} \rightarrow \mathbb{R}^V$ that sends the basis vector $e_v \in \mathbb{R}^{V'}$ corresponding to each vertex $v \in V'$ to the location of that vertex in $|K| \subset \mathbb{R}^V$ now restricts to a homeomorphism

$$|K'| \xrightarrow{\cong} |K|,$$

and we can therefore sensibly call K' the **barycentric subdivision** of the simplicial complex K .

A natural question now arises: what relation is there between the simplicial homologies of K and K' ? Their simplicial chain complexes are obviously not the same; in general, the chain complex for the subdivision K' has many more generators than that of K . But the polyhedra of these two complexes are the same, and it turns out that simplicial homology recognizes this fact. The result is best stated in terms of a concrete chain map

$$C_*^\Delta(K) \xrightarrow{S} C_*^\Delta(K'),$$

which can be defined by associating to each oriented n -simplex of K the n -chain of K' determined by its fundamental cycle. In the next lecture we will prove:

THEOREM 30.15. *The map $S_* : H_*^\Delta(K) \rightarrow H_*^\Delta(K')$ induced by the chain map described above is an isomorphism.*

30.5. Exercises.

EXERCISE 30.1. Prove Theorems 30.5 and 30.11 on the computation of $H_n^\Delta(K; \mathbb{Z}_2)$ and $H_n^\Delta(K; \mathbb{Z})$ for a closed and connected triangulated n -manifold $M \cong |K|$, in the second case with an oriented triangulation. Show moreover that if the triangulation does not admit any orientation, then $H_n^\Delta(K; \mathbb{Z}) = 0$.

EXERCISE 30.2. Prove Proposition 30.14, showing that the simplicial complex K with vertices in Δ^n defined via the barycentric subdivision algorithm actually defines a triangulation of Δ^n . *Hint: Argue inductively on n . Given any point $p \in \Delta^n$ distinct from the barycenter b_n , draw a straight line from b_n through p . What can you say about the point where this line exits Δ^n ?*

31. Chain homotopy and simplicial approximation

I owe you an explanation of why Theorem 30.15 is true, but in this lecture I also want to sketch a deep application of this theorem, showing that the isomorphism class of the oriented simplicial homology $H_*^\Delta(K)$ of a finite simplicial complex K depends only on its polyhedron $|K|$. In other words, simplicial homology is a *topological* invariant, not just an invariant of abstract simplicial complexes. For concreteness, we shall work with the oriented rather than the ordered simplicial chain complex, which is not a loss of generality since we will also show in the next lecture that $H_*^o(K) \cong H_*^\Delta(K)$. A few tricky details will be omitted, and we will make up for this later in the semester by deriving a second proof of the topological invariance of $H_*^\Delta(K)$ from cellular homology. But several of the ideas discussed in this lecture will also be useful for other purposes, when we develop singular homology and study its applications.

In the early history of homology theory, it was widely believed that the topological invariance of simplicial homology should be deduced from Theorem 30.15 in combination with a result called the *Hauptvermutung*, which conjectured that any two triangulations of the same polyhedron could be made identical up to homeomorphism after sufficiently many iterations of the barycentric subdivision algorithm. At some point, the invariance of simplicial homology was proven by other means, and the Hauptvermutung remained an open question until it was, ironically, shown to be *false* in the 1960's. Theorem 30.15 can be viewed nonetheless as an important ingredient in a proof that $H_*^\Delta(K)$ depends only on $|K|$.

To state the main result properly, let

$$\text{Cpct}^\Delta \subset \text{Top}$$

denote the subcategory whose objects consist of all compact polyhedra, with arbitrary continuous maps as morphisms. We should clarify: a compact topological space X is an object of Cpct^Δ if and only if it is homeomorphic to the polyhedron $|K|$ of some finite simplicial complex, but the actual complex K and homeomorphism $X \cong |K|$ are *not* considered to be part of the data defining an object of Cpct^Δ . In general, a polyhedron has infinitely many distinct choices of possible triangulations, and without choosing specific triangulations, there is no canonical way to define what it means for a map between two polyhedra to be simplicial. This is one of a few reasons why we allow *all* continuous maps as morphisms in Cpct^Δ , rather than just simplicial maps.

THEOREM 31.1. *There exists for each integer $n \geq 0$ and each R -module G a functor*

$$H_n^\Delta = H_n^\Delta(\cdot; G) : \text{Cpct}^\Delta \rightarrow R\text{-Mod}$$

that assigns to each compact polyhedron X the simplicial homology $H_n^\Delta(K; G)$ of some finite simplicial complex K whose polyhedron $|K|$ is homeomorphic to X .

Since homeomorphisms are isomorphisms in the category Cpct^Δ , this result implies:

COROLLARY 31.2. *If K and K' are two finite simplicial complexes with homeomorphic polyhedra $|K| \cong |K'|$, then their simplicial homologies $H_n^\Delta(K)$ and $H_n^\Delta(K')$ are isomorphic. \square*

Notice what Theorem 31.1 does not say: we are not claiming that the functor $H_n^\Delta : \text{Cpct}^\Delta \rightarrow R\text{-Mod}$ is unique or canonical, and in fact, some arbitrary choices will need to be made in order to define it at the end of this lecture. The need for choices, however, does not detract from the power of the theorem: the mere fact that H_n^Δ is a functor on the category Cpct^Δ , whose morphisms are arbitrary continuous maps, is enough to deduce useful consequences such as Corollary 31.2.

31.1. The homotopy question. In preparation for proving Theorem 30.15, let us consider a slightly different question about the functoriality of simplicial homology. Suppose $f, g : K \rightarrow L$ are two simplicial maps between simplicial complexes such that the induced continuous maps of polyhedra $|K| \rightarrow |L|$ are homotopic. Does it follow that the induced homomorphisms

$$f_*, g_* : H_*^\Delta(K) \rightarrow H_*^\Delta(L)$$

are identical? We've seen that bordism theory has a homotopy invariance property of this type, and a similar property is also incorporated into the Eilenberg-Steenrod axioms for homology theories.

The assumption in the present context is that there exists a continuous map

$$I \times |K| \xrightarrow{h} |L|$$

with $h(0, \cdot) = f$ and $h(1, \cdot) = g$. In order to make something useful out of this in simplicial homology, it would seem natural to impose an extra condition and require h to be a simplicial map, but here we encounter an obstacle: it is not obvious whether $I \times |K|$ has a natural triangulation, which would be needed in order for the notion of a simplicial map to make sense. The polyhedron $|K|$ is a union of simplices $|\sigma| \cong \Delta^n$ of various dimensions $n \geq 0$, and this decomposes $I \times |K|$ into "prism-shaped" subsets of the form

$$I \times \Delta^n \cong \Delta^1 \times \Delta^n.$$

If we can find a sufficiently natural way of triangulating $\Delta^1 \times \Delta^n$, we will obtain from this a triangulation of $I \times |K|$ and thus be able to speak of simplicial homotopies $h : I \times |K| \rightarrow |L|$ between f and g .

31.2. Triangulating products of simplices. Let us frame the question a bit more generally: Is there a natural way to triangulate $\Delta^m \times \Delta^n$ for every pair of integers $m, n \geq 0$? This product of simplices is a manifold of dimension $m + n$ with boundary

$$\begin{aligned} \partial(\Delta^m \times \Delta^n) &= (\partial\Delta^m \times \Delta^n) \cup (\Delta^m \times \partial\Delta^n) \\ &= \left(\bigcup_{k=0}^m \partial_{(k)} \Delta^m \times \Delta^n \right) \cup \left(\bigcup_{k=0}^n \Delta^m \times \partial_{(k)} \Delta^n \right). \end{aligned}$$

Notice that each term in the union on the second line is canonically homeomorphic to a product of the form $\Delta^k \times \Delta^\ell$ for $k \leq m$ and $\ell \leq n$ with $k + \ell = m + n - 1$. This suggests an inductive condition that would be natural to require on our triangulations: if we assume that suitable triangulations of $\Delta^k \times \Delta^\ell$ have already been constructed for all $k + \ell < m + n$, then we would like our triangulation of $\Delta^m \times \Delta^n$ to reproduce these triangulations when restricted to its smooth boundary faces. We shall now describe a direct construction that produces this result.

Denote the standard basis of $\mathbb{R}^{m+n+2} = \mathbb{R}^{m+1} \times \mathbb{R}^{n+1}$ by

$$(e_0, 0), \dots, (e_m, 0), (0, f_0), \dots, (0, f_n) \in \mathbb{R}^{m+1} \times \mathbb{R}^{n+1},$$

so we can regard e_0, \dots, e_m as the vertices of Δ^m and f_0, \dots, f_n as the vertices of Δ^n . The triangulation of $\Delta^m \times \Delta^n$ we construct has vertex set

$$V := \{(e_i, f_j) \in \Delta^m \times \Delta^n \mid i \in \{0, \dots, m\} \text{ and } j \in \{0, \dots, n\}\},$$

and its k -simplices for $k = 0, \dots, m+n$ will be the convex hulls of certain $(k+1)$ -tuples of these vertices; in this way, the triangulation $\Delta^m \times \Delta^n \cong |K|$ will be uniquely determined once we have specified a suitable abstract simplicial complex $K = (V, S)$. To specify which subsets should be the vertices of a simplex in K , endow the set $\{0, \dots, m\} \times \{0, \dots, n\}$ with the total order such that $(i, j) \leq (i', j')$ if and only if $i \leq i'$ and $j \leq j'$, so *strict* inequality $(i, j) < (i', j')$ means additionally that $i < i'$ or $j < j'$. For $k = 0, \dots, m+n$, the k -simplices σ of K are then defined as

$$\sigma = \{(e_{i_0}, f_{j_0}), \dots, (e_{i_k}, f_{j_k})\} \subset \Delta^m \times \Delta^n,$$

for all possible strictly increasing sequences

$$(i_0, j_0) < \dots < (i_k, j_k) \in \{0, \dots, m\} \times \{0, \dots, n\}.$$

By this definition, we observe that the $(m+n)$ -simplices all correspond to sequences $(i_0, j_0) < \dots < (i_{m+n}, j_{m+n})$ that begin with $(i_0, j_0) = (0, 0)$ and end with $(i_{m+n}, j_{m+n}) = (m, n)$, thus all of them contain the two specific vertices (e_0, f_0) and (e_m, f_n) . Boundary faces σ of these $(m+n)$ -simplices come in three types, corresponding to sequences $(i_0, j_0) < \dots < (i_{m+n-1}, j_{m+n-1})$ that satisfy the following conditions:

- (1) The sequence j_0, \dots, j_{m+n-1} takes every value in $\{0, \dots, n\}$ but i_0, \dots, i_{m+n-1} misses exactly one value $i \in \{0, \dots, m\}$.
- (2) The sequence i_0, \dots, i_{m+n-1} takes every value in $\{0, \dots, m\}$ but j_0, \dots, j_{m+n-1} misses exactly one value $j \in \{0, \dots, n\}$.
- (3) There are two consecutive terms of the form $(i, j), (i+1, j+1)$.

In the first two cases, the $m+n$ vertices of σ all lie in one of the convex sets

$$\partial_{(i)} \Delta^m \times \Delta^n \quad \text{or} \quad \Delta^m \times \partial_{(j)} \Delta^n.$$

As observed above, the union of these sets for all $i = 0, \dots, m$ and $j = 0, \dots, n$ is $\partial(\Delta^m \times \Delta^n)$, and these boundary faces thus determine an $(m+n-1)$ -dimensional subcomplex $K' \subset K$ in which the convex hull of the vertices of each simplex is contained in $\partial(\Delta^m \times \Delta^n)$. It is easy to check that *all* simplices of K with convex hull contained in $\partial(\Delta^m \times \Delta^n)$ are of this form, because for any two points $p, q \in \partial(\Delta^m \times \Delta^n)$ that do not both belong to the same one of the $m+n+2$ convex subsets mentioned above, the line segment from p to q passes through the interior of $\Delta^m \times \Delta^n$. In particular, boundary faces of the third type in the list above do not belong to the subcomplex K' .

Since the vertices $(e_i, f_j) \in V$ are all points in $\Delta^m \times \Delta^n \subset \mathbb{R}^{m+n+2}$ and the latter is a convex set, the unique linear map $\mathbb{R}^V \rightarrow \mathbb{R}^{m+n+2}$ sending $e_v \mapsto v$ for each $v \in V$ restricts to the polyhedron $|K| \subset \mathbb{R}^V$ as a map

$$(31.1) \quad |K| \rightarrow \Delta^m \times \Delta^n.$$

Exercise 31.1 shows that this map is a homeomorphism, and thus defines a triangulation of $\Delta^m \times \Delta^n$; moreover, restricting it to the subcomplex K' formed by vertices contained in $\partial(\Delta^m \times \Delta^n)$ gives a homeomorphism $|K'| \cong \partial(\Delta^m \times \Delta^n)$.

In order to define fundamental cycles $c_{\Delta^m \times \Delta^n}$ in $H_{m+n}^o(K; \mathbb{Z})$ and $H_{m+n}^\Delta(K; \mathbb{Z})$ from our triangulation, we need to endow it with an orientation and choose an admissible ordering for the simplicial complex K . The latter is easy, as the total order on $\{0, \dots, k\} \times \{0, \dots, \ell\}$ determines a total order on the set of all vertices of K . In order to define a suitable orientation, let $\mathbf{S}(m, n)$ denote the set of all strictly increasing sequences $(i_0, j_0) < \dots < (i_{m+n}, j_{m+n})$ of $m+n+1$

elements in $\{0, \dots, m\} \times \{0, \dots, n\}$, and write $\sigma_{\mathbf{s}}$ for the $(m+n)$ -simplex of K determined by each $\mathbf{s} \in \mathbf{S}(m, n)$. Denote by $\mathbf{s}_0 \in \mathbf{S}(m, n)$ the specific sequence

$$(0, 0) < (1, 0) < \dots < (m, 0) < (m, 1) < \dots < (m, n),$$

and define the **parity** $|\mathbf{s}| \in \mathbb{Z}_2$ of any element $\mathbf{s} \in \mathbf{S}(m, n)$ 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:

$$(i-1, j) < (i, j) < (i, j+1) \quad \rightsquigarrow \quad (i-1, j) < (i-1, j+1) < (i, j+1).$$

LEMMA 31.3. *The parity $|\mathbf{s}| \in \mathbb{Z}_2$ of elements $\mathbf{s} \in \mathbf{S}(m, n)$ is independent of choices.*

PROOF SKETCH. We can interpret $(-1)^{|\mathbf{s}|} \in \{1, -1\}$ as the sign of a permutation of $m+n$ elements, which include m copies of the letter R (for “right”) and n copies of the letter U (for “up”). □

The chosen orientation and admissible ordering for K determine a fundamental cycle

$$c_{\Delta^m \times \Delta^n} \in C_{m+n}^{\Delta}(K; \mathbb{Z}) \quad \text{or} \quad C_{m+n}^o(K; \mathbb{Z}).$$

For each $m, n \geq 1$, the relation $\partial c_{\Delta^m \times \Delta^n} = c_{\partial(\Delta^m \times \Delta^n)}$ then becomes the following formula under the usual identification between boundary faces and simplices of one dimension lower:

$$(31.2) \quad \partial c_{\Delta^m \times \Delta^n} = \sum_{i=0}^m (-1)^i c_{\partial_{(i)} \Delta^m \times \Delta^n} + (-1)^m \sum_{j=0}^n (-1)^j c_{\Delta^m \times \partial_{(j)} \Delta^n}.$$

REMARK 31.4. If we regard $\partial \Delta^m \times \Delta^n$ and $\Delta^m \times \partial \Delta^n$ as compact topological $(m+n-1)$ -manifolds with matching boundary $\partial \Delta^m \times \partial \Delta^n$ and endow both with the obvious oriented triangulations and admissible orderings that they inherit from $\Delta^m \times \Delta^n$, the formula in (31.2) takes the slightly prettier form

$$\partial c_{\Delta^m \times \Delta^n} = c_{\partial \Delta^m \times \Delta^n} + (-1)^m c_{\Delta^m \times \partial \Delta^n}.$$

When we introduce the homological cross product later in this semester, the singular homology version of this relation will take the form

$$\partial(c_{\Delta^m} \times c_{\Delta^n}) = \partial c_{\Delta^m} \times c_{\Delta^n} + (-1)^m c_{\Delta^m} \times \partial c_{\Delta^n},$$

which is written in terms of the obvious fundamental cycle $c_{\Delta^k} \in C_k(\Delta^k; \mathbb{Z})$ for the standard simplex of each dimension with its trivial triangulation, and a bilinear product operation

$$C_*(X) \otimes C_*(Y) \rightarrow C_*(X \times Y) : A \otimes B \mapsto A \times B$$

that relates the singular chain complexes of any two spaces X, Y and sends $C_m(X) \otimes C_n(Y)$ in general to $C_{m+n}(X \times Y)$. We will have plenty to say about this product later, but the detail I want to comment on right now is the sign $(-1)^m$ 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^m and Δ^n 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 necessary when writing $\partial c_{\Delta^m} \times c_{\Delta^n}$, since the three objects ∂ , c_{Δ^m} and c_{Δ^n} appear here in the same order as on the left hand side, but writing $c_{\Delta^m} \times \partial c_{\Delta^n}$ exchanges the order of ∂ and c_{Δ^m} , and since ∂ has odd degree, a sign change must then result if and only if c_{Δ^m} also has odd degree, meaning m 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.

31.3. From simplicial homotopies to chain homotopies. Let us identify the unit interval I with the standard 1-simplex via the homeomorphism

$$I \xrightarrow{\cong} \Delta^1 \subset \mathbb{R}^2 : t \mapsto (1 - t, t).$$

The oriented triangulation of $\Delta^1 \times \Delta^n$ constructed in the previous section for each $n \geq 0$ yields an oriented triangulation of $I \times \Delta^n$ whose restriction to the smooth faces of

$$\partial(I \times \Delta^n) = (\{1\} \times \Delta^n) \cup (\{0\} \times \Delta^n) \cup \left(\bigcup_{k=0}^n I \times \partial_{(k)} \Delta^n \right)$$

matches the trivial triangulation of Δ^n and the constructed trivialization of $I \times \Delta^{n-1}$.

Now consider again the polyhedron $|K|$ discussed in §31.1 above, and choose an admissible ordering for the underlying simplicial complex K . The ordering determines an identification of each n -simplex of $|K|$ for $n = 0, 1, 2, \dots$ with the standard n -simplex, and applying the triangulation algorithm then gives a triangulation of $I \times |K|$, whose underlying simplicial complex we shall denote in the following by K_I . For this triangulation, the two inclusions

$$|K| \xrightarrow{\iota_j} I \times |K| : p \mapsto (j, p) \quad \text{for } j = 0, 1$$

are both simplicial maps. Now suppose $\sigma = [v_0, \dots, v_n]$ is an n -simplex of K , equipped with the orientation it inherits from the admissible ordering, and let $|\sigma| \subset |K|$ denote the corresponding subset homeomorphic to Δ^n in the polyhedron. Our triangulation determines an oriented triangulation of the $(n + 1)$ -dimensional manifold $I \times \Delta^n \cong I \times |\sigma| \subset I \times |K|$, thus a fundamental cycle

$$c_{I \times |\sigma|} \in C_{n+1}^\Delta(K_I; \mathbb{Z}),$$

for which the formula (31.2) specializes to this situation as

$$\partial c_{I \times |\sigma|} = (\iota_1)_* \sigma - (\iota_0)_* \sigma - c_{I \times \partial|\sigma|}.$$

Here $c_{I \times \partial|\sigma|}$ is an abbreviation for the signed sum of fundamental cycles of the induced triangulation of $I \times \Delta^{n-1}$ for each boundary face $\Delta^{n-1} \cong \partial_{(k)} \Delta^n$ of $|\sigma| \cong \Delta^n$.

Since $I \times |K|$ is now a polyhedron, we can sensibly impose an extra condition on the homotopy $h : I \times |K| \rightarrow |L|$, and require it to be a simplicial map, i.e. a **simplicial homotopy** between f and g . With this assumption in place, there is a unique homomorphism

$$C_n^\Delta(K) \xrightarrow{h_\#} C_{n+1}^\Delta(L)$$

defined for each $n \geq 0$ via linearity and the formula

$$h_\#(\sigma) := h_* c_{I \times |\sigma|}.$$

for oriented n -simplices $\sigma \in [v_0, \dots, v_n]$ of K . Since $h_* : C_*(K_I) \rightarrow C_*(L)$ is a chain map, the formula above for $\partial c_{I \times |\sigma|}$ implies

$$\partial h_\#(\sigma) = g_* \sigma - f_* \sigma - h_\#(\partial \sigma),$$

so that $h_\#$ satisfies the so-called **chain homotopy relation**

$$\partial h_\# + h_\# \partial = g_* - f_*.$$

A brief algebraic digression is now in order.

DEFINITION 31.5. Given two chain maps $f, g : (A_*, \partial^A) \rightarrow (B_*, \partial^B)$, a **chain homotopy** (*Kettenhomotopie*) from f to g is a homomorphism $h : A_* \rightarrow B_*$ that satisfies $h(A_n) \subset B_{n+1}$ for each $n \in \mathbb{Z}$ and the chain homotopy relation

$$\partial^B h + h \partial^A = g - f.$$

We say that f and g are **chain homotopic** if there exists a chain homotopy from f to g .

One easily checks that the notion of chain homotopy defines an equivalence relation between chain maps, and moreover, if f_0 and f_1 are chain homotopic and have well-defined compositions of chain maps $f_j \circ g$, then $f_0 \circ g$ and $f_1 \circ g$ are also chain homotopic; a similar statement applies to compositions of the form $g \circ f_j$. The upshot is that there is a well-defined **homotopy category of chain complexes**

$$\mathbf{hCh}(R\text{-Mod}) \quad \text{abbreviated as} \quad \mathbf{hCh},$$

in which the objects are chain complexes of R -modules and the morphisms are chain homotopy classes of chain maps. An isomorphism in the category \mathbf{hCh} is called a **chain homotopy equivalence** (*Kettenhomotopieäquivalenz*), so a chain map $f : A_* \rightarrow B_*$ is a chain homotopy equivalence if and only if it admits a **chain homotopy inverse** $g : B_* \rightarrow A_*$, meaning a chain map such that the compositions $g \circ f$ and $f \circ g$ are each chain homotopic to identity maps.

There are two convincing reasons why the category \mathbf{hCh} is important to define: the first is that simplicial homotopies between simplicial maps give rise to chain homotopies between the induced chain maps, as shown above—and we will see later that chain homotopies in the singular chain complex similarly arise from arbitrary homotopies between continuous maps. The second reason is the following easy result, which tells us that the algebraic homology functors $H_n : \mathbf{Ch} \rightarrow R\text{-Mod}$ descend to the homotopy category as functors $\mathbf{hCh} \rightarrow R\text{-Mod}$.

PROPOSITION 31.6. *If $f, g : A_* \rightarrow B_*$ are chain homotopic chain maps, then for each $n \in \mathbb{Z}$, the homomorphisms $f_*, g_* : H_n(A_*) \rightarrow H_n(B_*)$ they induce on homology are identical.*

PROOF. Given $[a] \in H_n(A_*)$, the representative $a \in A_n$ is a cycle, so the chain homotopy relation gives $g(a) - f(a) = \partial^B h(a) + h \partial^A a = \partial^B h(a)$, implying $[f(a)] = [g(a)] \in H_n(B_*)$. \square

Putting all of this together implies:

COROLLARY 31.7. *If $f, g : |K| \rightarrow |L|$ are simplicial maps related by a simplicial homotopy, then for each $n \in \mathbb{Z}$, the induced maps $f_*, g_* : H_n^\Delta(K) \rightarrow H_n^\Delta(L)$ are identical.* \square

We have used oriented simplicial homology in this discussion for the sake of concreteness, but the discussion also makes sense for ordered simplicial homology.

31.4. Subdivision defines a chain homotopy equivalence. Now that the notion of a chain homotopy equivalence has been defined, we can explain the real reason behind Theorem 30.15. Assume K is a simplicial complex and K' is the complex defined from K by barycentric subdivision, giving rise to the chain map

$$S : C_*^\Delta(K) \rightarrow C_*^\Delta(K')$$

described in the previous lecture. A special class of chain maps in the other direction can be defined as follows. By definition, every vertex v in the complex K' is the barycenter of a particular simplex σ_v in the polyhedron $|K|$; note that this includes the vertices of K' that are also vertices of K , since the latter are also 0-simplices of K . For each vertex v of K' , let $w(v)$ denote an arbitrary choice of a vertex of the simplex σ_v in K that has v as its barycenter. One can check that this defines a simplicial map

$$K' \xrightarrow{\pi} K : v \mapsto w_v,$$

and we call it a **projection** since it necessarily sends each vertex of K' that is also a vertex of K to itself. The following result now implies Theorem 30.15:

THEOREM 31.8. *For any choice of projection $\pi : K' \rightarrow K$, the induced chain map $\pi_* : C_*^\Delta(K') \rightarrow C_*^\Delta(K)$ is a chain homotopy inverse of $S : C_*^\Delta(K) \rightarrow C_*^\Delta(K')$, implying in particular that the latter is a chain homotopy equivalence.*

A complete proof of this theorem can be found e.g. in [ES52, Theorem VI.7.1]; here we shall content ourselves with a brief sketch. One can verify directly that $\pi_* S : C_*^\Delta(K) \rightarrow C_*^\Delta(K)$ is the identity map. It then remains to show that

$$S\pi_* : C_*^\Delta(K') \rightarrow C_*^\Delta(K')$$

is chain homotopic to the identity. The details of this chain homotopy would require too much of a digression, but there is a geometric construction in the background that is worth understanding: in a different context, the same construction will later give us a relatively straightforward construction of a chain homotopy for the natural subdivision operator on singular homology.

The construction is yet another triangulation of the prism $I \times \Delta^n$, one that interpolates between the trivial triangulation of $\{0\} \times \Delta^n$ and the barycentric subdivision of $\{1\} \times \Delta^n$. Like the other explicit triangulations we've discussed, it decomposes $I \times \Delta^n$ into convex regions determined by sets of vertices in general position, and it 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 $n \geq 1$ 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 $I \times \partial_{(k)} \Delta^n \cong I \times \Delta^{n-1}$, one with vertices v_0, \dots, v_n and $(1, b_n)$.

Exercise 31.2 implies that this defines an oriented triangulation of $I \times \Delta^n$.

Intuitively, the three pieces of the boundary

$$\partial(I \times \Delta^n) = (\{1\} \times \Delta^n) \cup (\{0\} \times \Delta^n) \cup (I \times \partial \Delta^n)$$

with their induced triangulations now correspond to the three terms on the right hand side of a chain homotopy relation

$$\partial h_\# = S\pi_* - \mathbb{1} - h_\# \partial$$

for some chain homotopy $h_\# : C_n^\Delta(K') \rightarrow C_{n+1}^\Delta(K')$. In the simplicial context, it is not so straightforward to make this intuition precise, but we will return to this subject in the near future in the context of singular homology, where the definition of the corresponding chain homotopy is more straightforward.

31.5. Simplicial approximation. The last major ingredient needed for a proof of Theorem 31.1 is a result that relates the categories Cpct^Δ and Simp :

THEOREM 31.9 (simplicial approximation). *If $X \cong |K|$ and $Y \cong |L|$ are compact polyhedra and $f : X \rightarrow Y$ is a continuous map, then after finitely-many iterations of barycentric subdivision to replace the triangulation of X with a finer triangulation $X \cong |K'|$, f is homotopic to a map $g : X \rightarrow Y$ that arises from a simplicial map $K' \rightarrow L$. Moreover, for every $x \in X$, $g(x)$ is contained in the smallest simplex of Y containing $f(x)$.*

We might have naively hoped for the theorem to state that every continuous map between polyhedra with fixed choices of triangulations is homotopic to a simplicial map—but there are easy counterexamples to that statement. For instance, every non-surjective map $S^1 \rightarrow S^1$ has its

image in a contractible space $S^1 \setminus \{*\} \cong \mathbb{R}$ and is thus homotopic to a constant, implying that every map $S^1 \rightarrow S^1$ homotopic to the identity is surjective. But if we choose two triangulations $S^1 \cong |K|$ and $S^1 \cong |L|$ such that L has strictly more vertices than K , then no simplicial map $K \rightarrow L$ can be surjective, and the identity $S^1 \rightarrow S^1$ therefore cannot be homotopic to any simplicial map with respect to these particular triangulations. Of course, this problem goes away if we are also allowed to replace K with a triangulation that has arbitrarily many vertices, e.g. by iterated barycentric subdivision.

We refer to [Hat02, §2.C] for a detailed proof of Theorem 31.9, but the following explains the basic idea.

SKETCH OF THE PROOF OF THEOREM 31.9. For each vertex $v \in X$, define the so-called **open star** of v as the open neighborhood

$$\text{st } v \subset X$$

of v formed by the union of the interiors of all simplices in X that have v as a vertex. Figure 17 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 31.3). The collection of all open stars of vertices defines an open covering of any polyhedron. Now given $f : X \rightarrow Y$ 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 17 again)

$$\text{st } v \subset f^{-1}(\text{st } w_v).$$

Having associated to each $v \in X$ some $w_v \in Y$ with this property, there is a unique simplicial map $g : X \rightarrow Y$ that satisfies $g(v) = w_v$: indeed, for every simplex $\{v_0, \dots, v_n\}$ of X , Exercise 31.3 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 $h : I \times X \rightarrow Y$ from f to g can be defined by choosing $h(\cdot, x) : I \rightarrow Y$ for every $x \in X$ to be the linear path from $f(x)$ to $g(x)$ in the smallest simplex containing $f(x)$. \square

31.6. Simplicial homology as a topological invariant. Here is a sketch of a proof of Theorem 31.1.

By the axiom of choice, we can associate to every compact polyhedron $X \in \text{Cpct}^\Delta$ a specific choice of finite simplicial complex K_X and triangulation $X \cong |K_X|$; having done this, define

$$H_n^\Delta(X) := H_n^\Delta(K_X).$$

For each continuous map $f : X \rightarrow Y$ between compact polyhedra, we can apply the simplicial approximation theorem to find a sufficiently fine subdivision K'_X of K_X and a simplicial map $g : K'_X \rightarrow K_Y$ for which the associated continuous map $g : X \rightarrow Y$ is homotopic to f . Writing $S_* : H_n^\Delta(K_X) \rightarrow H_n^\Delta(K'_X)$ for the isomorphism defined via iterated barycentric subdivision, the homomorphism $f_* : H_n^\Delta(X) \rightarrow H_n^\Delta(Y)$ induced by f can then be defined by

$$\begin{array}{c}
 \xrightarrow{f_*} \\
 H_n^\Delta(X) = H_n^\Delta(K_X) \xrightarrow{\cong} H_n^\Delta(K'_X) \xrightarrow{g_*} H_n^\Delta(K_Y) = H_n^\Delta(Y) .
 \end{array}$$

The result of §31.3 on simplicial homotopies can be used in showing that the map $f_* : H_n^\Delta(X) \rightarrow H_n^\Delta(Y)$ defined in this way is independent of choices. Putting all this together produces a functor $H_n^\Delta : \text{Cpct}^\Delta \rightarrow R\text{-Mod}$.




FIGURE 17. A map $f : X \rightarrow Y$ 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 31.9 will then produce a simplicial map $g : X \rightarrow Y$ 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} .

31.7. Exercises.

EXERCISE 31.1. Prove that the map $|K| \rightarrow \Delta^m \times \Delta^n$ described in (31.1) is a homeomorphism. *Hint: This is probably not the only possible approach, but here an inductive argument as in Exercise 30.2 is also possible. Use the fact that certain points are contained in all the n -simplices.*

EXERCISE 31.2. Let L denote the $(n + 1)$ -dimensional abstract simplicial complex formed by the sets of vertices in $I \times \Delta^n$ described in §31.4, 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 30.2 to show that L defines a triangulation of $I \times \Delta^n$ for which the subcomplex L' triangulates $\partial(I \times \Delta^n)$.

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 $n \geq 0$.

EXERCISE 31.3. Given vertices v_0, \dots, v_k in a polyhedron X , show that $\bigcap_{i=0}^k \text{st } v_i \neq \emptyset$ if and only if X contains a simplex whose vertices are v_0, \dots, v_k .

32. Acyclic models and relative homology

I want to tie up a few loose ends regarding simplicial homology before we move on to singular homology in the next lecture. One important topic is the reason why the ordered simplicial homology $H_*^o(K)$ and its oriented counterpart $H_*^\Delta(K)$ are isomorphic: we will prove this using the method of acyclic models, which will also be quite useful in our later discussion of products in singular homology and cohomology. We also take this opportunity to introduce *relative* simplicial homology, and explain the general algebraic mechanism that leads to long exact sequences of homology groups.

Several results in this lecture will apply equally well to the ordered and oriented versions of simplicial homology, and the following notational convention will allow us to talk about both at the same time:

$$H_*^\bullet := H_*^o \text{ or } H_*^\Delta,$$

$$C_*^\bullet := C_*^o \text{ or } C_*^\Delta.$$

32.1. Reduced simplicial homology. We discussed the reduced version \tilde{h}_* of an axiomatic homology theory h_* in Lecture 28. A reduced version of simplicial homology can be defined analogously, after observing that the one-point space $\{*\}$ is a polyhedron, whose underlying simplicial complex consists only of a single vertex. We shall also denote this one-point simplicial complex by $\{*\}$, and let

$$K \xrightarrow{\epsilon} \{*\}$$

denote the unique simplicial map from any given simplicial complex K to the one-point complex. The reduced (ordered or oriented) simplicial homology is then defined by

$$\tilde{H}_n^\bullet(K) = \tilde{H}_n^\bullet(K; G) := \ker \left(H_n^\bullet(K; G) \xrightarrow{\epsilon_*} H_n^\bullet(\{*\}; G) \right).$$

As with axiomatic homology, we can always choose a right-inverse of $\epsilon : K \rightarrow \{*\}$, which in this context must be a simplicial map $\{*\} \hookrightarrow K$, and the induced homomorphism on homology gives rise to a splitting of the short exact sequence

$$0 \rightarrow \tilde{H}_n^\bullet(K) \hookrightarrow H_n^\bullet(K) \xrightarrow{\epsilon_*} H_n^\bullet(\{*\}) \rightarrow 0,$$

and thus an isomorphism $H_n^\bullet(K) \cong \tilde{H}_n^\bullet(K) \oplus H_n^\bullet(\{*\})$. We recall from Exercise 29.4 that $H_n^\bullet(\{*\}; G)$ is trivial for $n \neq 0$ and is naturally isomorphic to the coefficient group G for $n = 0$, so the result is

$$H_n^\bullet(K; G) \cong \begin{cases} \tilde{H}_n^\bullet(K; G) \oplus G & \text{for } n = 0, \\ \tilde{H}_n^\bullet(K; G) & \text{for } n \neq 0. \end{cases}$$

Working through Exercise 29.4 also leads to the following observation: In both versions of the simplicial chain complex for $\{*\}$, the boundary map $\partial_1 : C_1^\bullet(\{*\}) \rightarrow C_0^\bullet(\{*\})$ at degree 1 is trivial. Indeed, this is immediate in the oriented chain complex because $C_1^\Delta(\{*\})$ is trivial due to the lack of 1-simplices, while in the ordered chain complex, $C_1^o(\{*\})$ has only a single generator $(*, *)$ determined by the unique vertex $* \in \{*\}$, which satisfies

$$\partial(*, *) = (*) - (*) = 0.$$

Since $\epsilon_* : C_*^\bullet(K) \rightarrow C_*^\bullet(\{*\})$ is a chain map, the relation $\epsilon_* \partial = \partial \epsilon_*$ then implies that the composition of $\partial_1 : C_1^\bullet(K) \rightarrow C_0^\bullet(K)$ with the so-called **augmentation** $\epsilon_* : C_0^\bullet(K) \rightarrow C_0^\bullet(\{*\}) = G$ is trivial, leading to the so-called **augmented chain complex**

$$\dots \longrightarrow C_2^\bullet(K; G) \xrightarrow{\partial_2} C_1^\bullet(K; G) \xrightarrow{\partial_1} C_0^\bullet(K; G) \xrightarrow{\epsilon_*} G \longrightarrow 0 \longrightarrow 0 \longrightarrow \dots,$$

in which we use the natural isomorphism $C_0^\bullet(\{*\}; G) \cong G$ to replace $C_0^\bullet(\{*\}; G)$ by the coefficient group G , and the map $\epsilon_* : C_0^\bullet(K; G) \rightarrow G$ can then be expressed via the direct formula

$$\epsilon_* \left(\sum_i a_i \sigma_i \right) = \sum_i a_i$$

for any finite linear combination of generators σ_i with coefficients $a_i \in G$. We shall denote the augmented chain complex by $\tilde{C}_*^\bullet(K) = \tilde{C}_*^\bullet(K; G)$, with chain groups

$$\tilde{C}_n^\bullet(K; G) := \begin{cases} C_n^\bullet(K; G) & \text{for } n \neq -1, \\ G & \text{for } n = -1, \end{cases}$$

and boundary map $\partial : \tilde{C}_*^\bullet(K) \rightarrow \tilde{C}_*^\bullet(K)$ matching that of the usual chain complex $C_*^\bullet(K)$ except at degree 0, where it is defined to be the augmentation $\epsilon_* : C_0^\bullet(K; G) \rightarrow G$. The following result is a near immediate consequence of the definitions.

PROPOSITION 32.1. *There is a natural isomorphism*

$$H_*(\tilde{C}_*^\bullet(K; G)) \xrightarrow{\cong} \tilde{H}_*(K; G)$$

that takes the form $[c] \mapsto [c]$ for cycles $c \in \tilde{C}_n^\bullet(K; G)$ of degree $n \geq 0$. □

32.2. The cone of a simplicial complex. The point of defining reduced simplicial homology is to have a version of simplicial homology that vanishes in *all* degrees for certain contractible polyhedra that arise in applications. Here is a popular class of examples.

DEFINITION 32.2. The **cone** of a simplicial complex $K = (V, S)$ is the simplicial complex $CK = (CV, CS)$ with vertices

$$CV := V \cup \{*\}$$

and simplices

$$CS := S \cup \{\{v_0, \dots, v_n, *\} \mid \{v_0, \dots, v_n\} \in S\},$$

where $*$ denotes an *extra* vertex that is assumed to be not an element of the original vertex set V .

The polyhedron $|CK|$ of a cone complex CK has an obvious identification with the topological cone $C|K|$ of the original polyhedron $|K|$, in which the summit of the cone corresponds to the extra vertex $* \in CV$.

DEFINITION 32.3. A chain complex A_* is called **chain contractible** if the identity map $\mathbb{1} : A_* \rightarrow A_*$ is chain homotopic to the trivial chain map $0 : A_* \rightarrow A_*$.

If A_* is chain contractible, then looking at induced maps $H_*(A_*) \rightarrow H_*(A_*)$, we find that the identity map and the zero map on $H_*(A_*)$ must be identical, which is only possible if $H_*(A_*) = 0$. A chain complex with the latter property is said to be **acyclic**, in other words, A_* has no cycles other than those which are trivial in the sense of being boundaries. Chain contractible complexes are thus acyclic; one can view this as an algebraic counterpart to the topological fact that contractible spaces have trivial reduced homology according to the axioms.

LEMMA 32.4. *For any simplicial complex K , the augmented simplicial chain complex $\tilde{C}_*^o(CK; \mathbb{Z})$ of its cone is chain contractible.*

PROOF. Let us write down a proof for the ordered chain complex, from which a proof for the oriented complex can be obtained just by changing round brackets into square brackets. For each $n \geq 0$, we can specify a homomorphism $h_\# : \tilde{C}_n^o(CK; \mathbb{Z}) \rightarrow \tilde{C}_{n+1}^o(CK; \mathbb{Z})$ by saying how it is defined on an arbitrary generator $(v_0, \dots, v_n) \in C_n^o(CK; \mathbb{Z}) = \tilde{C}_n^o(CK; \mathbb{Z})$, so we define

$$\tilde{C}_n^o(CK; \mathbb{Z}) \xrightarrow{h_\#} \tilde{C}_{n+1}^o(CK; \mathbb{Z}) : (v_0, \dots, v_n) \mapsto (*, v_0, \dots, v_n),$$

and we extend it to $n = -1$ by specifying its value on the generator $1 \in \mathbb{Z} = \tilde{C}_{-1}^o(CK; \mathbb{Z})$, namely

$$\tilde{C}_{-1}^o(CK; \mathbb{Z}) \xrightarrow{h_\#} \tilde{C}_0^o(CK; \mathbb{Z}) : 1 \mapsto (*).$$

We then have

$$\partial h_\#(v_0, \dots, v_n) = (v_0, \dots, v_n) - h_\# \partial(v_0, \dots, v_n) \quad \text{and} \quad \partial h_\#(1) = \epsilon_*(*) = 1 = 1 - h_\# \partial(1)$$

since $\tilde{C}_{-2}^o(CK; \mathbb{Z})$ is trivial by definition and thus $\partial(1) = 0$. This establishes the chain homotopy relation $\partial h_\# + h_\# \partial = \mathbb{1} = \mathbb{1} - 0$. \square

EXAMPLE 32.5. For some $n \geq 1$, suppose K is a simplicial complex containing only a single n -simplex and all its faces, so $|K| \cong \Delta^n$. Then K can be identified with the cone of a complex K' with $|K'| \cong \Delta^{n-1}$, and the lemma above therefore implies that $\tilde{C}_*(K; \mathbb{Z})$ is chain contractible.

32.3. Natural chain homotopy equivalences. Recall that for any simplicial complex K and any choice of coefficients, the quotient projection $(v_0, \dots, v_n) \mapsto [v_0, \dots, v_n]$ determines a natural chain map

$$C_*^o(K) \xrightarrow{\Psi_K} C_*^\Delta(K).$$

Here the word *natural* carries a precise meaning that will be important to clarify: it means that for any other simplicial complex L with a simplicial map $f : L \rightarrow K$, the diagram

$$\begin{array}{ccc} C_*^o(L) & \xrightarrow{\Psi_L} & C_*^\Delta(L) \\ \downarrow f_* & & \downarrow f_* \\ C_*^o(K) & \xrightarrow{\Psi_K} & C_*^\Delta(K) \end{array}$$

commutes. As a special case, suppose $L \subset K$ is a subcomplex and $f : L \hookrightarrow K$ is the inclusion map: the chain maps $f_* : C_*^o(L) \rightarrow C_*^o(K)$ and $f_* : C_*^\Delta(L) \rightarrow C_*^\Delta(K)$ are then likewise inclusions of subcomplexes, and naturality then implies firstly that Ψ_K sends the subcomplex $C_*^o(L) \subset C_*^o(K)$ into the subcomplex $C_*^\Delta(L) \subset C_*^\Delta(K)$, and secondly that Ψ_L is simply the restriction of Ψ_K to $C_*^o(L)$.

With this observation as motivation, let us say more generally that for a specific simplicial complex K , a chain map $\Psi : C_*^o(K) \rightarrow C_*^\Delta(K)$ is **natural** if for every subcomplex $L \subset K$, Ψ sends $C_*^o(L)$ into $C_*^\Delta(L)$. In the same manner, one can define the notion of a *natural* chain map in the other direction $C_*^\Delta(K) \rightarrow C_*^o(K)$, or between each of $C_*^\Delta(K)$ or $C_*^o(K)$ and itself. Such chain

maps can always be interpreted as natural transformations between two functors to the category of chain complexes, defined on a category that has subcomplexes of K as objects and inclusion maps as morphisms.

The following result explains why the natural homomorphism $H_*^o(K) \rightarrow H_*^\Delta(K)$ induced by the chain map Ψ_K is always an isomorphism, thus making the ordered and oriented versions of simplicial homology interchangeable in practice.

THEOREM 32.6. *For every simplicial complex K , the natural chain map $\Psi_K : C_*^o(K) \rightarrow C_*^\Delta(K)$ is a chain homotopy equivalence.*

The theorem will follow from three lemmas, each of which should be understood to hold for an arbitrary simplicial complex K :

LEMMA 32.7. *There exists a natural chain map $\Phi : C_*^\Delta(K; \mathbb{Z}) \rightarrow C_*^o(K; \mathbb{Z})$ that is determined in degree 0 by the formula*

$$\Phi([v]) := (v) \quad \text{for all vertices } v \text{ of } K,$$

and moreover, natural chain maps with this property are unique up to chain homotopy.

LEMMA 32.8. *Natural chain maps $C_*^o(K; \mathbb{Z}) \rightarrow C_*^o(K; \mathbb{Z})$ matching the identity map in degree 0 are unique up to chain homotopy.*

LEMMA 32.9. *Natural chain maps $C_*^\Delta(K; \mathbb{Z}) \rightarrow C_*^\Delta(K; \mathbb{Z})$ matching the identity map in degree 0 are unique up to chain homotopy.*

Notice that the statements of the last two lemmas only involve uniqueness, not existence; the existence is clear in both cases because the identity map is a chain map that satisfies the required properties. This trivial observation is used in the following proof.

PROOF OF THEOREM 32.6. Writing $\Psi := \Psi_K$, the uniqueness up to chain homotopy in Lemmas 32.8 and 32.9 implies that if we are working with integer coefficients, $\Phi \circ \Psi$ and $\Psi \circ \Phi$ are both chain homotopic to the identity, so that the chain map Φ from Lemma 32.7 is a chain homotopy inverse for Ψ . The validity of this result extends to arbitrary coefficients for relatively trivial algebraic reasons explained in Remark 32.10 below. \square

REMARK 32.10. Here is why in the proof of Theorem 32.6, it suffices to consider chain complexes with integer coefficients. The three lemmas above provide chain maps between chain complexes with integer coefficients, but the resulting formulas for these maps on the canonical generators of $C_*^o(K; \mathbb{Z})$ and $C_*^\Delta(K; \mathbb{Z})$ determine via linearity chain maps on $C_*^o(K; G)$ and $C_*^\Delta(K; G)$ for any coefficient group G . The same applies to chain homotopies, e.g. if $h : C_*^o(K; \mathbb{Z}) \rightarrow C_{*+1}^o(K; \mathbb{Z})$ is a chain homotopy between $\Phi \circ \Psi_K : C_*^o(K; \mathbb{Z}) \rightarrow C_*^o(K; \mathbb{Z})$ and the identity, then it determines via linearity a chain homotopy $h : C_*^o(K; G) \rightarrow C_{*+1}^o(K; G)$ between $\Phi \circ \Psi_K : C_*^o(K; G) \rightarrow C_*^o(K; G)$ and the identity for any choice of coefficients G .

The proofs of Lemmas 32.7, 32.8 and 32.9 are very similar, and are based on an idea known as the method of **acyclic models**. We shall carry out the details only for Lemma 32.7.

PROOF OF LEMMA 32.7. For the entirety of this proof, we assume

$$G := \mathbb{Z}$$

and omit the coefficient group from the notation. We shall prove by induction on the degree $n \geq 0$ that it is possible to construct homomorphisms $\Phi_L : C_n^\Delta(L) \rightarrow C_n^o(L)$ for every subcomplex $L \subset K$ such that Φ_L is the restriction of Φ_K to $C_n^\Delta(L) \subset C_n^\Delta(K)$ and the chain map relation $\Phi_L \partial = \partial \Phi_L$ is satisfied. It would of course suffice to construct Φ_K such that it sends $C_n^\Delta(L) \rightarrow C_n^o(L)$ for every

subcomplex $L \subset K$ and then define Φ_L as the restriction, but in practice, we shall do things the other way around, and define Φ_L first for a special class of subcomplexes such that the definition of Φ_K is then uniquely determined.

The beginning of the induction is to define $\Phi_K : C_0^\Delta(L) \rightarrow C_0^o(L) : [v] \mapsto (v)$ as specified in the statement of the lemma.

For a given $n \geq 1$, we then assume that $\Phi_K : C_k^\Delta(K) \rightarrow C_k^o(K)$ has already been defined for every $k \leq n-1$ such that it sends $C_k^\Delta(L)$ to $C_k^o(L)$ for every subcomplex $L \subset K$ and satisfies $\Phi_K \partial = \partial \Phi_K$. The idea for the inductive step is now to first define $\Phi_L : C_n^\Delta(L) \rightarrow C_n^o(L)$ for a specific class of “model” subcomplexes $L \subset K$, which will determine $\Phi_K : C_n^\Delta(K) \rightarrow C_n^o(K)$ via the naturality condition. The model complexes are defined as follows: For any n -simplex $\sigma = \{v_0, \dots, v_n\}$ of K , let $L_\sigma \subset K$ denote the subcomplex that contains only σ and all its faces. Note that since $n \geq 1$, L_σ can be identified with the cone of an $(n-1)$ -dimensional complex as in Example 32.5, so Lemma 32.4 implies that both versions of the augmented simplicial chain complex for L_σ are acyclic; this is why L_σ is called an “acyclic model”. Now, there is only one generator $\sigma = [v_0, \dots, v_n] \in C_n^\Delta(L_\sigma)$, so $\Phi_{L_\sigma} : C_n^\Delta(L_\sigma) \rightarrow C_n^o(L_\sigma)$ will be determined as soon as we choose a value for $\tau := \Phi_{L_\sigma}(\sigma) \in C_n^o(L_\sigma)$, which must be required to satisfy

$$\partial\tau = \partial\Phi_{L_\sigma}(\sigma) = \Phi_{L_\sigma}(\partial\sigma) \in C_{n-1}^o(L_\sigma).$$

The right hand side of this expression has already been defined due to the inductive hypothesis. Moreover, it is a cycle in the augmented chain complex $\tilde{C}_*^o(L_\sigma)$, since

$$\partial\Phi_{L_\sigma}(\partial\sigma) = \Phi_{L_\sigma}(\partial^2\sigma) = 0 \in \tilde{C}_{n-2}^o(L_\sigma),$$

where we should clarify that in the case $n = 1$, the operator ∂ acting on 0-chains is actually the augmentation $\epsilon_* : C_0^o(L_\sigma; \mathbb{Z}) \rightarrow \mathbb{Z}$. Since $\tilde{C}_*^o(L_\sigma)$ is acyclic, it follows that $\Phi_{L_\sigma}(\partial\sigma)$ is also a boundary, and we can therefore define $\Phi_{L_\sigma}(\sigma)$ to be any choice of element $\tau \in C_n^o(L_\sigma)$ such that $\partial\tau = \Phi_{L_\sigma}(\partial\sigma)$.

Having made such choices and defined $\Phi_{L_\sigma} : C_n^\Delta(L_\sigma) \rightarrow C_n^o(L_\sigma)$ for the model subcomplex $L_\sigma \subset K$ corresponding to each n -simplex σ of K , we observe now that there is a unique definition of $\Phi_K : C_n^\Delta(K) \rightarrow C_n^o(K)$ that has the correct restriction to all of these subcomplexes, and it automatically satisfies both the chain map relation and the naturality condition.

The construction of Φ_K beyond degree zero involved some arbitrary choices, so it remains to show that any other natural chain map Φ'_K that matches Φ_K in degree zero is chain homotopic to it. We shall use a similar inductive argument to construct homomorphisms $h_K : C_k^\Delta(K) \rightarrow C_{k+1}^o(K)$ that satisfy the chain homotopy relation $\partial h_K + h_K \partial = \Phi'_K - \Phi_K$, and here as well it will be convenient to impose a naturality condition, namely that h_K has a well-defined restriction $h_L : C_k^\Delta(L) \rightarrow C_{k+1}^o(L)$ for every subcomplex $L \subset K$. To start the induction, it suffices to define $h_K : C_0^\Delta(K) \rightarrow C_1^o(K)$ as the trivial homomorphism since $\Phi_K = \Phi'_K$ on $C_0^\Delta(K)$. Now assume that h_K and its restrictions h_L satisfying the chain homotopy relation have already been defined on chains of degree $k \leq n-1$ for some $n \geq 1$. For each n -simplex $\sigma = \{v_0, \dots, v_n\}$ of K , we again consider the corresponding model subcomplex $L_\sigma \subset K$, and define $h_{L_\sigma} : C_n^\Delta(L_\sigma) \rightarrow C_{n+1}^o(L_\sigma)$ so that it sends the unique generator $\sigma = [v_0, \dots, v_n] \in C_n^\Delta(L_\sigma)$ to some element $\tau := h_{L_\sigma}(\sigma) \in C_{n+1}^o(L_\sigma)$ satisfying

$$\partial\tau = \partial h_{L_\sigma}(\sigma) = -h_{L_\sigma}(\partial\sigma) + \Phi'_{L_\sigma}(\sigma) - \Phi_{L_\sigma}(\sigma) \in C_n^o(L_\sigma).$$

This is possible due to acyclicity, since the inductive hypothesis implies that the right hand side is a cycle:

$$\begin{aligned} \partial(-h_{L_\sigma}(\partial\sigma) + \Phi'_{L_\sigma}(\sigma) - \Phi_{L_\sigma}(\sigma)) &= (-\partial h_{L_\sigma} + \Phi'_{L_\sigma} - \Phi_{L_\sigma})(\partial\sigma) \\ &= h_{L_\sigma}(\partial^2\sigma) = 0. \end{aligned}$$

Having extended h_{L_σ} to degree n for each of the model subcomplexes $L_\sigma \subset K$, there is again a unique definition of $h_K : C_n^\Delta(K) \rightarrow C_{n+1}^\circ(K)$ that has the correct restriction to each of these subcomplexes, and it automatically satisfies the chain homotopy relation. \square

The proofs of Lemmas 32.8 and 32.9 are similar, but shorter since one only needs to construct chain homotopies, the existence of suitable chain maps being obvious. Lemma 32.8 also requires the knowledge that $\tilde{C}_*^\circ(L_\sigma)$ is an acyclic chain complex for each of the model complexes $L_\sigma \subset K$, while for Lemma 32.9, one must instead use the fact that $\tilde{C}_*^\Delta(L_\sigma)$ is acyclic.

32.4. Relative homology. In §28.3 we saw that there is a relative version of bordism theory defined for pairs of spaces $(X, A) \in \mathbf{Top}^{\text{rel}}$, with long exact sequences that relate the relative bordism groups of (X, A) to the absolute bordism groups of X and A . Something similar is true in all versions of homology theory; let's discuss briefly how it works in simplicial homology.

A **simplicial pair** (K, L) is a simplicial complex K together with a subcomplex $L \subset K$, and a **map of simplicial pairs** $f : (K, L) \rightarrow (K', L')$ is a simplicial map $f : K \rightarrow K'$ that sends L into L' and thus also defines a simplicial map $L \rightarrow L'$. Let us denote by $\mathbf{Simp}^{\text{rel}}$ the category whose objects are simplicial pairs and whose morphisms are maps of simplicial pairs. We can identify the category \mathbf{Simp} of simplicial complexes with the subcategory

$$\mathbf{Simp} \subset \mathbf{Simp}^{\text{rel}}$$

consisting of pairs of the form (K, \emptyset) . The following definition makes sense because for any subcomplex $L \subset K$, the generators of $C_*^\bullet(L)$ are also generators of $C_*^\bullet(K)$, thus making the chain complex $C_*^\bullet(L)$ into a subcomplex of $C_*^\bullet(K)$.

DEFINITION 32.11. The (ordered or oriented) **relative simplicial homology** of a simplicial pair (K, L) with coefficients in G is defined in each degree $n \in \mathbb{Z}$ as the homology of the quotient chain complex $C_*^\bullet(K; G)/C_*^\bullet(L; G)$, thus

$$H_n^\bullet(K, L) = H_n^\bullet(K, L; G) := H_n(C_*^\bullet(K, L; G)), \quad \text{where}$$

$$C_*^\bullet(K, L) = C_*^\bullet(K, L; G) := C_*^\bullet(K; G) / C_*^\bullet(L; G).$$

Relative simplicial homology defines functors

$$H_n^\bullet : \mathbf{Simp}^{\text{rel}} \rightarrow R\text{-Mod}$$

in a straightforward way: any map of simplicial pairs $f : (K, L) \rightarrow (K', L')$ induces a chain map $f_* : C_*^\bullet(K) \rightarrow C_*^\bullet(K')$ that also sends $C_*^\bullet(L)$ to $C_*^\bullet(L')$ and thus descends to the quotients as a chain map $f_* : C_*^\bullet(K, L) \rightarrow C_*^\bullet(K', L')$, inducing maps

$$H_n^\bullet(K, L) \xrightarrow{f_*} H_n^\bullet(K', L')$$

for each n . In keeping with the identification of \mathbf{Simp} with a subcomplex of $\mathbf{Simp}^{\text{rel}}$, we observe that $H_n(K, \emptyset)$ is the same thing as $H_n(K)$.

Elements $[c] \in H_n^\bullet(K, L)$ can be represented by **relative n -cycles**

$$c \in C_n^\bullet(K) \quad \text{such that} \quad \partial c \in C_{n-1}^\bullet(L).$$

Here, the condition $\partial c \in C_{n-1}^\bullet(L)$ means that the image of c under the quotient projection $C_n^\bullet(K) \rightarrow C_n^\bullet(K, L)$ is a cycle, and we understand $[c] \in H_n^\bullet(K, L)$ to mean the homology class represented by that cycle. Two relative n -cycles $a, b \in C_n^\bullet(K)$ then represent the same relative homology class in $H_n^\bullet(K, L)$ if and only if $a - b = \partial c + d$ for some $c \in C_{n+1}^\bullet(K)$ and $d \in C_n^\bullet(L)$. There is a natural homomorphism defined for each $n \geq 1$ by

$$H_n^\bullet(K, L) \xrightarrow{\hat{c}_*} H_{n-1}^\bullet(L) : [c] \mapsto [\partial c].$$

Note that, in spite of appearances, the class $[\partial c] \in H_{n-1}^\bullet(L)$ in this expression need not be trivial, because c is an n -chain in K , but might not be an n -chain in L .

THEOREM 32.12. *Given a simplicial pair (K, L) , let $i : L \hookrightarrow K$ and $j : (K, \emptyset) \hookrightarrow (K, L)$ denote the obvious inclusion maps. Then the sequence*

$$\begin{aligned} \dots \longrightarrow H_{n+1}^\bullet(K, L) &\xrightarrow{\hat{c}_*} H_n^\bullet(L) \xrightarrow{i_*} H_n^\bullet(K) \xrightarrow{j_*} H_n^\bullet(K, L) \xrightarrow{\hat{c}_*} H_{n-1}^\bullet(L) \xrightarrow{i_*} H_{n-1}^\bullet(K) \longrightarrow \dots \\ &\longrightarrow H_0^\bullet(L) \xrightarrow{i_*} H_0^\bullet(K) \xrightarrow{j_*} H_0^\bullet(K, L) \longrightarrow 0 \end{aligned}$$

is exact.

It is not hard to verify the exactness of the sequence in this theorem explicitly, but there is also an underlying algebraic phenomenon that deserves more attention. Since $C_*^\bullet(K, L)$ is a quotient, every simplicial pair (K, L) gives rise to an obvious short exact sequence

$$0 \rightarrow C_*^\bullet(L) \xrightarrow{i_*} C_*^\bullet(K) \xrightarrow{j_*} C_*^\bullet(K, L) \rightarrow 0,$$

in which each term is a chain complex and the maps between them are chain maps. The inclusion $C_*^\bullet(L) \hookrightarrow C_*^\bullet(K)$ of chain complexes is in fact the chain map i_* induced by the inclusion $i : L \hookrightarrow K$ of simplicial complexes, and since $j : (K, \emptyset) \hookrightarrow (K, L)$ is actually the identity map, the quotient projection $C_*^\bullet(K) \rightarrow C_*^\bullet(K)/C_*^\bullet(L)$ can similarly be understood as the chain map $j_* : C_*^\bullet(K) \rightarrow C_*^\bullet(K, L)$ induced by j . Algebraically, it turns out that short exact sequences of chain complexes and chain maps *always* give rise to long exact sequences relating their homology groups:

PROPOSITION 32.13. *Suppose $0 \rightarrow A_* \xrightarrow{f} B_* \xrightarrow{g} C_* \rightarrow 0$ is a short exact sequence of chain complexes and chain maps. Then for each $n \in \mathbb{Z}$ there exists a so-called **connecting homomorphism** $\partial_* : H_n(C_*) \rightarrow H_{n-1}(A_*)$ such that the sequence*

$$\begin{aligned} \dots \xrightarrow{\hat{c}_*} H_{n+1}(A_*) &\xrightarrow{f_*} H_{n+1}(B_*) \xrightarrow{g_*} H_{n+1}(C_*) \\ &\xrightarrow{\hat{c}_*} H_n(A_*) \xrightarrow{f_*} H_n(B_*) \xrightarrow{g_*} H_n(C_*) \\ &\xrightarrow{\hat{c}_*} H_{n-1}(A_*) \xrightarrow{f_*} H_{n-1}(B_*) \xrightarrow{g_*} H_{n-1}(C_*) \xrightarrow{\hat{c}_*} \dots \end{aligned}$$

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

$$\begin{array}{ccccccc} 0 & \longrightarrow & A_* & \xrightarrow{f} & B_* & \xrightarrow{g} & C_* \longrightarrow 0 \\ & & \downarrow \alpha & & \downarrow \beta & & \downarrow \gamma \\ 0 & \longrightarrow & A'_* & \xrightarrow{f'} & B'_* & \xrightarrow{g'} & C'_* \longrightarrow 0 \end{array}$$

in which all maps are chain maps and the bottom row is also exact, and we denote the resulting connecting homomorphisms by $\partial'_* : H_n(C'_*) \rightarrow H_{n-1}(A'_*)$. Then the diagram

$$\begin{array}{ccccccccccc} \dots & \longrightarrow & H_{n+1}(C_*) & \xrightarrow{\hat{c}_*} & H_n(A_*) & \xrightarrow{f_*} & H_n(B_*) & \xrightarrow{g_*} & H_n(C_*) & \xrightarrow{\hat{c}_*} & H_{n-1}(A_*) \longrightarrow \dots \\ & & \downarrow \gamma_* & & \downarrow \alpha_* & & \downarrow \beta_* & & \downarrow \gamma_* & & \downarrow \alpha_* \\ \dots & \longrightarrow & H_{n+1}(C'_*) & \xrightarrow{\hat{c}'_*} & H_n(A'_*) & \xrightarrow{f'_*} & H_n(B'_*) & \xrightarrow{g'_*} & H_n(C'_*) & \xrightarrow{\hat{c}'_*} & H_{n-1}(A'_*) \longrightarrow \dots \end{array}$$

also commutes.

The proof of this result is by “diagram chasing,” which we already saw examples of in Proposition 28.22 and Exercise 28.8 (the five-lemma). Let’s do the first step, which is to write down a

reasonable candidate for the map $\partial_* : H_n(C_*) \rightarrow H_{n-1}(A_*)$. We are given a commuting diagram of the form

$$\begin{array}{ccccccc}
 & & \vdots & & \vdots & & \vdots \\
 & & \downarrow & & \downarrow & & \downarrow \\
 0 & \longrightarrow & A_n & \xrightarrow{f} & B_n & \xrightarrow{g} & C_n \longrightarrow 0 \\
 & & \downarrow \partial & & \downarrow \partial & & \downarrow \partial \\
 0 & \longrightarrow & A_{n-1} & \xrightarrow{f} & B_{n-1} & \xrightarrow{g} & C_{n-1} \longrightarrow 0 \\
 & & \downarrow \partial & & \downarrow \partial & & \downarrow \partial \\
 0 & \longrightarrow & A_{n-2} & \xrightarrow{f} & B_{n-2} & \xrightarrow{g} & C_{n-2} \longrightarrow 0 \\
 & & \downarrow \partial & & \downarrow \partial & & \downarrow \partial \\
 & & \vdots & & \vdots & & \vdots
 \end{array}$$

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 $\partial_* : H_n(C_*) \rightarrow H_{n-1}(A_*)$ 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.

Similarly, it is not hard to see why the “functoriality” aspect of the statement 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 complexes or long exact sequences of R -modules, with morphisms defined in each case via commutative diagrams, such that Proposition 32.13 produces a functor from the former category to the latter. See Exercise 32.3 for a precise formulation in these terms. Exercise 32.2 shows moreover that applying Proposition 32.13 to the short exact sequence $0 \rightarrow C_*^*(L) \rightarrow C_*^*(K) \rightarrow C_*^*(K, L) \rightarrow 0$ for a simplicial pair (K, L) produces the same connecting homomorphism as in the statement of Theorem 32.12.

32.5. Exercises.

EXERCISE 32.1. Carry out the rest of the details of the diagram chase to prove the exactness of the sequence in Proposition 32.13.

EXERCISE 32.2 (*). Show that for any simplicial pair (K, L) , the connecting homomorphisms $\partial_* : H_n^*(K, L) \rightarrow H_{n-1}^*(L)$ that arise by plugging the short exact sequence of simplicial chain

complexes $0 \rightarrow C_*^\bullet(L) \hookrightarrow C_*^\bullet(K) \rightarrow C_*^\bullet(K, L) \rightarrow 0$ into Proposition 32.13 are given by the formula $\partial_*[c] = \partial c$ for any relative n -cycle $c \in C_n^\bullet(K)$ in (K, L) .

EXERCISE 32.3. Consider the categories **Short** and **Long**, defined as follows. Objects in **Short** are short exact sequences of chain complexes $0 \rightarrow A_* \xrightarrow{f} B_* \xrightarrow{g} C_* \rightarrow 0$ of R -modules, with a morphism from this object to another object $0 \rightarrow A'_* \xrightarrow{f'} B'_* \xrightarrow{g'} C'_* \rightarrow 0$ 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:

$$\begin{array}{ccccccccc} 0 & \longrightarrow & A_* & \xrightarrow{f} & B_* & \xrightarrow{g} & C_* & \longrightarrow & 0 \\ & & \downarrow \alpha & & \downarrow \beta & & \downarrow \gamma & & \\ 0 & \longrightarrow & A'_* & \xrightarrow{f'} & B'_* & \xrightarrow{g'} & C'_* & \longrightarrow & 0 \end{array}$$

The objects in **Long** are long exact sequences of \mathbb{Z} -graded R -modules $\dots \rightarrow C_{n+1} \xrightarrow{\delta} A_n \xrightarrow{F} B_n \xrightarrow{G} C_n \xrightarrow{\delta} A_{n-1} \rightarrow \dots$, with morphisms from this to another object $\dots \rightarrow C'_{n+1} \xrightarrow{\delta'} A'_n \xrightarrow{F'} B'_n \xrightarrow{G'} C'_n \xrightarrow{\delta'} A'_{n-1} \rightarrow \dots$ 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:

$$\begin{array}{cccccccccccc} \dots & \longrightarrow & C_{n+1} & \xrightarrow{\delta} & A_n & \xrightarrow{F} & B_n & \xrightarrow{G} & C_n & \xrightarrow{\delta} & A_{n-1} & \longrightarrow & \dots \\ & & \downarrow \gamma & & \downarrow \alpha & & \downarrow \beta & & \downarrow \gamma & & \downarrow \alpha & & \\ \dots & \longrightarrow & C'_{n+1} & \xrightarrow{\delta'} & A'_n & \xrightarrow{F'} & B'_n & \xrightarrow{G'} & C'_n & \xrightarrow{\delta'} & A'_{n-1} & \longrightarrow & \dots \end{array}$$

- Show that there is a covariant functor $\mathbf{Simp}^{\text{rel}} \rightarrow \mathbf{Short}$ assigning to each simplicial pair (K, L) its short exact sequence of (ordered or oriented) simplicial chain complexes.
- Show that there is also a covariant functor $\mathbf{Short} \rightarrow \mathbf{Long}$ 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 $\mathbf{Simp}^{\text{rel}} \rightarrow \mathbf{Long}$.)

33. Singular homology

33.1. Definitions. The immediate disadvantage of simplicial homology is that its definition requires strictly more data than just a topological space: we need to have a triangulation of that space, and it takes considerable effort to see why different triangulations of the same space produce isomorphic homologies. The definition of singular homology resembles that of simplicial homology, but it explicitly removes the need for a triangulation. The price to be paid for this is that the resulting chain complex seems absurdly large: so large, in fact, that one might find it surprising at first that it is ever possible to explicitly compute the singular homology of a space. I advise you not to think too much about this when you first read the definition, as we will subsequently discuss some properties that make computations of singular homology quite a reasonable task.

DEFINITION 33.1. A **singular n -simplex** (*singulärer n -Simplex*) in a topological space X is defined to be a continuous map $\sigma : \Delta^n \rightarrow X$. Let

$$\mathcal{K}_n(X) := \{ \sigma : \Delta^n \rightarrow X \mid \sigma \text{ is continuous} \}$$

denote the set of all singular n -simplices in X . The **singular chain complex** (*singulärer Kettenkomplex*) $C_*(X) = C_*(X; G)$ of X with coefficients in the R -module G is defined such that

$C_n(X) = 0$ for all $n < 0$, and for $n \geq 0$,

$$C_n(X) := \bigoplus_{\sigma \in \mathcal{K}_n(X)} G.$$

The boundary map $\partial : C_n(X) \rightarrow C_{n-1}(X)$ for $n \geq 1$ is uniquely determined by linearity and the formula

$$\partial\sigma = \sum_{k=0}^n (-1)^k (\sigma|_{\partial_{(k)}\Delta^n}),$$

where the identification (29.1) is used in order to view each term in the summation as a singular $(n-1)$ -simplex $\sigma|_{\partial_{(k)}\Delta^n} : \Delta^{n-1} \rightarrow X$, making the linear combination an element of $C_{n-1}(X; \mathbb{Z})$. The homology groups of this chain complex form the **singular homology** of X with coefficients in G ,

$$H_n(X) = H_n(X; G) := H_n(C_*(X; G)).$$

There is a fairly obvious way to make

$$C_* : \mathbf{Top} \rightarrow \mathbf{Ch}(R\text{-Mod})$$

into a functor: any continuous map $f : X \rightarrow Y$ between spaces induces a unique chain map

$$f_* : C_*(X) \rightarrow C_*(Y)$$

determined by linearity and the formula

$$f_*(\sigma) := f \circ \sigma$$

for singular simplices $\sigma : \Delta^n \rightarrow X$. Composing this functor with $H_n : \mathbf{Ch}(R\text{-Mod}) \rightarrow R\text{-Mod}$ makes singular homology itself into a collection of functors

$$H_n : \mathbf{Top} \rightarrow R\text{-Mod},$$

meaning in particular that continuous maps $f : X \rightarrow Y$ induce homomorphisms $f_* : H_n(X) \rightarrow H_n(Y)$ for every $n \geq 0$.

For a pair of spaces $(X, A) \in \mathbf{Top}^{\text{rel}}$, there is a similarly straightforward extension of the definitions above to the notion of **relative singular homology**

$$H_n(X, A) = H_n(X, A; G) := H_n(C_*(X, A; G)), \quad \text{where}$$

$$C_*(X, A) = C_*(X, A; G) := C_*(X; G) / C_*(A; G),$$

which makes sense because singular simplices in A are also singular simplices in X , making $C_*(A)$ naturally a subcomplex of $C_*(X)$. We have $H_n(X, \emptyset) = H_n(X)$ for all spaces X , and relative singular homology can thus be regarded as an extension of the functor $H_n : \mathbf{Top} \rightarrow R\text{-Mod}$ over the larger category $\mathbf{Top}^{\text{rel}}$, in which maps of pairs $f : (X, A) \rightarrow (Y, B)$ induce chain maps $f_* : C_*(X) \rightarrow C_*(Y)$ that descend to the quotients as chain maps $f_* : C_*(X, A) \rightarrow C_*(Y, B)$ and thus induce homomorphisms $f_* : H_n(X, A) \rightarrow H_n(Y, B)$ for all n . As with relative simplicial homology, we can represent relative singular homology classes $[c] \in H_n(X, A)$ via **relative cycles** $c \in C_n(X)$, which are assumed to satisfy $\partial c \in C_{n-1}(A)$, and writing them in this way gives rise to an obvious connecting homomorphism

$$H_n(X, A) \xrightarrow{\partial_*} H_{n-1}(A) : [c] \mapsto [\partial c].$$

33.2. Non-axiomatic properties. Before we get to the purely “formal” (i.e. axiomatic) properties of singular homology, let us discuss a few features it has that other axiomatic homology theories do not.

THEOREM 33.2. *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

$$\sum_i a_i x_i \in C_0(X)$$

of generators $x_i \in X$ with coefficients $a_i \in G$. 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

$$(33.1) \quad I \xrightarrow{\cong} \Delta^1 : t \mapsto (1-t, t).$$

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 (33.1) matches boundary points via the correspondence

$$\partial I \ni 0 \mapsto \partial_{(1)} \Delta^1 \subset \partial \Delta^1, \quad \partial I \ni 1 \mapsto \partial_{(0)} \Delta^1 \subset \partial \Delta^1,$$

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 $\partial : C_n(X) \rightarrow C_{n-1}(X)$, 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 (33.1) to identify singular 1-simplices in X with paths $\gamma : I \rightarrow X$ and likewise identify singular 0-simplices with points $x \in X$ in the canonical way, then the operator $\partial : C_1(X) \rightarrow C_0(X)$ is now determined by the formula

$$(33.2) \quad \partial \gamma = \gamma(1) - \gamma(0).$$

This tells you why two 0-cycles of the form $m x, m y \in C_0(X)$ 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 (33.1) between I and Δ^1 also gives rise to a natural homomorphism

$$(33.3) \quad h : \pi_1(X, p) \rightarrow H_1(X; \mathbb{Z})$$

sending the homotopy class of the loop $\gamma : I \rightarrow X$ to the homology class that it represents when regarded as a singular 1-chain with integer coefficients; note that by (33.2), this 1-chain is a cycle because $\gamma : I \rightarrow X$ has the same start and end point. The map (33.3) 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 33.3. *If X is path-connected, then the Hurewicz map (33.3) descends to the abelianization of $\pi_1(X) := \pi_1(X, p)$ as an isomorphism*

$$\pi_1(X) / [\pi_1(X), \pi_1(X)] \xrightarrow{\cong} H_1(X; \mathbb{Z}).$$

One can prove Theorem 33.3 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.

The third property I want to mention is a relationship between simplicial and singular homology. Suppose $K = (V, S)$ is a simplicial complex, with polyhedron $|K|$. There is then a natural chain map

$$C_*^o(K) \rightarrow C_*(|K|)$$

defined by associating to each generator (v_0, \dots, v_n) in degree n of the ordered simplicial chain complex the unique singular n -simplex $\sigma : \Delta^n \rightarrow |K|$ that extends to a linear map $\mathbb{R}^{n+1} \rightarrow \mathbb{R}^V$ sending the standard basis of \mathbb{R}^{n+1} to the vectors $e_{v_0}, \dots, e_{v_n} \in \mathbb{R}^V$. As usual, the word "natural" has a precise meaning here, and the map $C_*^o(K) \rightarrow C_*(|K|)$ can be described as a natural transformation between two functors $\mathbf{Simp} \rightarrow \mathbf{Ch}(R\text{-Mod})$. Letting chain maps descend to maps between homology groups, we obtain natural homomorphisms

$$H_n^o(K) \rightarrow H_n(|K|).$$

In light of the natural isomorphisms $H_n^o(K) \xrightarrow{\cong} H_n^\Delta(K)$, we also obtain from this natural homomorphisms $H_n^\Delta(K) \rightarrow H_n(|K|)$, and we will see when we study cellular homology that the latter is also an isomorphism.

One useful application of this relationship is a construction of fundamental cycles in singular homology: If $M \cong |K|$ is a compact triangulated n -manifold, with a choice of admissible ordering for the underlying simplicial complex, then feeding the resulting fundamental cycle $c_M \in C_n^o(K; \mathbb{Z}_2)$ into the natural chain map $C_n^o(K; \mathbb{Z}_2) \rightarrow C_n(M; \mathbb{Z}_2)$ produces a singular fundamental cycle

$$c_M \in C_n(M; \mathbb{Z}_2) \quad \text{such that} \quad \partial c_M = c_{\partial M} \in C_{n-1}(\partial M; \mathbb{Z}_2) \subset C_{n-1}(M; \mathbb{Z}_2).$$

If the triangulation is also oriented, then this can all also be done with integer coefficients, producing an integral fundamental cycle

$$c_M \in C_n(M; \mathbb{Z}) \quad \text{such that} \quad \partial c_M = c_{\partial M} \in C_{n-1}(\partial M; \mathbb{Z}) \subset C_{n-1}(M; \mathbb{Z}).$$

33.3. The axioms. Here is the main result of this lecture.

THEOREM 33.4. *For any R -module G , the functors $H_n(\cdot; G) : \mathbf{Top}^{\text{rel}} \rightarrow R\text{-Mod}$ and connecting homomorphisms $\partial_* : H_n(X, A; G) \rightarrow H_{n-1}(A; G)$ defined for all $(X, A) \in \mathbf{Top}^{\text{rel}}$ and $n \in \mathbb{Z}$ satisfy the axioms of a homology theory (in the sense of Eilenberg-Steenrod) with coefficient group G .*

Let's first dispense with the axioms that are easy exercises. Since a one-point space $\{*\}$ admits only one singular n -simplex $\sigma : \Delta^n \rightarrow \{*\}$ for each $n \geq 0$, a computation completely analogous to Exercise 29.4 shows that

$$H_n(\{*\}; G) \cong \begin{cases} 0 & \text{if } n \neq 0, \\ G & \text{if } n = 0. \end{cases}$$

This shows that $H_* := H_*(\cdot; G)$ satisfies the *dimension* axiom, and moreover, the isomorphism $H_0(\{*\}; G) \cong G$ is canonical. The *additivity* axiom is a similarly straightforward consequence of the definitions.

Now for the interesting part.

PROPOSITION 33.5 (the homotopy axiom). *For any two homotopic maps of pairs $f, g : (X, A) \rightarrow (Y, B)$, the induced homomorphisms $f_*, g_* : H_n(X, A) \rightarrow H_n(Y, B)$ on singular homology are identical.*

PROOF. We consider first the case of absolute homology, so assume $h : I \times X \rightarrow Y$ is a homotopy between $f := h(0, \cdot)$ and $g := h(1, \cdot)$. For each $n \geq 0$, there is a unique homomorphism $h_\# : C_n(X) \rightarrow C_{n+1}(Y)$ determined by linearity and the following formula for $h_\#(\sigma) \in C_{n+1}(Y; \mathbb{Z})$ on an arbitrary singular n -simplex $\sigma : \Delta^n \rightarrow X$: we use the maps

$$I \times \Delta^n \xrightarrow{\text{Id} \times \sigma} I \times X \xrightarrow{h} Y,$$

together with the integral fundamental cycle $c_{I \times \Delta^n} \in C_{n+1}(I \times \Delta^n; \mathbb{Z})$ arising from the oriented triangulation of $I \times \Delta^n \cong \Delta^1 \times \Delta^n$ described in §31.2, to define

$$h_\#(\sigma) := h_*(\text{Id} \times \sigma)_* c_{I \times \Delta^n} \in C_{n+1}(Y; \mathbb{Z}).$$

One now deduces from the formula for $\partial c_{I \times \Delta^n}$ that $h_\# : C_*(X) \rightarrow C_{*+1}(Y)$ is a chain homotopy between f_* and g_* .

Extending this result to the setting of a homotopy $h : (I \times X, I \times A) \rightarrow (Y, B)$ between two maps of pairs $f, g : (X, A) \rightarrow (Y, B)$ requires only the extra observation that since $h(I \times A) \subset B$, the chain homotopy $h_\#$ constructed above descends to the quotient as a chain homotopy $C_*(X, A) \rightarrow C_{*+1}(Y, B)$ between the two chain maps $f_*, g_* : C_*(X, A) \rightarrow C_*(Y, B)$. \square

PROPOSITION 33.6 (the exactness axiom). *For any pair of spaces $(X, A) \in \text{Top}^{\text{rel}}$ with inclusion maps $i : A \hookrightarrow X$ and $j : (X, \emptyset) \hookrightarrow (X, A)$, the sequence*

$$\dots \longrightarrow H_n(A) \xrightarrow{i_*} H_n(X) \xrightarrow{j_*} H_n(X, A) \xrightarrow{\hat{c}_*} H_{n-1}(A) \longrightarrow \dots \longrightarrow H_0(X, A) \longrightarrow 0$$

is exact.

PROOF. This is a straightforward consequence of Proposition 32.13 and the obvious short exact sequence of chain complexes

$$0 \longrightarrow C_*(A) \xrightarrow{i_*} C_*(X) \xrightarrow{j_*} C_*(X, A) \longrightarrow 0,$$

one only needs to check that the connecting homomorphism produced by the diagram chase in the proof of Proposition 32.13 is the specific map $H_n(X, A) \rightarrow H_{n-1}(A) : [c] \mapsto [\partial c]$. \square

Recall that for the *excision* axiom formulated in Lecture 28, the hypothesis was that $B \subset A \subset X$ and there exists a continuous function $u : X \rightarrow I$ that “separates” B from $X \setminus A$ in the sense that $u|_B \equiv 0$ and $u|_{X \setminus A} \equiv 1$. In singular homology, it suffices to work with a slightly weaker variant of this hypothesis.

PROPOSITION 33.7 (the excision axiom). *Assume $B \subset A \subset X$ such that the closure of B is contained in the interior of A . Then the inclusion of pairs $i : (X \setminus B, A \setminus B) \hookrightarrow (X, A)$ induces an isomorphism $i_* : H_n(X \setminus B, A \setminus B) \xrightarrow{\cong} H_n(X, A)$ for every n .*

The proof requires a bit of preparation. For each $n \geq 0$, there is a unique homomorphism

$$S : C_n(X) \rightarrow C_n(X)$$

that is determined by linearity and the formula

$$S(\sigma) := \sigma_* c_{\Delta^n} \in C_n(X; \mathbb{Z}),$$

where $\sigma : \Delta^n \rightarrow X$ is an arbitrary singular n -simplex and $c_{\Delta^n} \in C_n(\Delta^n; \mathbb{Z})$ is the integral fundamental cycle defined via barycentric subdivision of Δ^n . We will refer to this as the **subdivision operator** on the singular chain complex.

LEMMA 33.8. *For each $m \in \mathbb{N}$, the m th iterate $S^m : C_*(X) \rightarrow C_*(X)$ of the subdivision operator is a chain map, and there exists a chain homotopy $h_m : C_*(X) \rightarrow C_{*+1}(X)$ between S^m and the identity map. Moreover, for any subspace $A \subset X$, both S^m and h_m preserve the subcomplex $C_*(A) \subset C_*(X)$.*

PROOF. We prove the statement first for $m = 1$. For a given singular n -simplex $\sigma : \Delta^n \rightarrow X$, the chain map relation $\partial S(\sigma) = \sigma_* \partial c_{\Delta^n} = \sigma_* c_{\partial \Delta^n} = S(\partial \sigma)$ follows from the inductive nature of the barycentric subdivision algorithm described in §30.4. To see why S is then chain homotopic to the identity, one uses the oriented triangulation of $I \times \Delta^n$ described in §31.4 and its integral fundamental cycle $c_{I \times \Delta^n} \in C_{n+1}(I \times \Delta^n; \mathbb{Z})$ to define, for each $n \geq 0$, a homomorphism $h_1 : C_n(X) \rightarrow C_{n+1}(X)$ that is given on each generator $\sigma : \Delta^n \rightarrow X$ by

$$h_1(\sigma) = (\text{pr}_2)_*(\text{Id} \times \sigma)_* c_{I \times \Delta^n} \in C_{n+1}(X; \mathbb{Z}),$$

with $\text{pr}_2 : I \times X \rightarrow X$ denoting the projection to the second factor. Since the triangulation of $I \times \Delta^n$ restricts to $\partial(I \times \Delta^n)$ as the trivial triangulation of $\{0\} \times \Delta^n$, the barycentric subdivision of $\{1\} \times \Delta^n$, and the $(n-1)$ -dimensional case of the same triangulation on each face of $I \times \partial \Delta^n$, we have

$$\partial h_1(\sigma) = (\text{pr}_2)_*(\text{Id} \times \sigma)_* \partial c_{I \times \Delta^n} = S(\sigma) - \sigma - h_1(\partial \sigma),$$

and thus the chain homotopy relation $\partial h_1 + h_1 \partial = S - \mathbb{1}$. It is clear from the construction that both S and h_1 preserve $C_*(A) \subset C_*(X)$ for any $A \subset X$.

For arbitrary $m \in \mathbb{N}$, it is now obvious that S^m is also a chain map and is chain homotopic to $\mathbb{1}^m = \mathbb{1}$, but we need to check that there is a chain homotopy h_m that preserves subcomplexes $C_*(A) \subset C_*(X)$. One can see this by writing down an inductive definition of h_m , for which various choices are possible, e.g. $h_m := h_{m-1}S + h_1$ does the job. \square

Taking m large enough, the operator S^m can be applied in principle to replace any singular chain $c \in C_n(X)$ with a chain $S^m c \in C_n(X)$ whose constituent singular simplices are as “small” we we like: in particular, if X is covered by the interiors of two subsets

$$X = \mathring{U} \cup \mathring{V}, \quad \mathcal{U}, \mathcal{V} \subset X,$$

then for any given chain $c \in C_n(X)$, taking $m \in \mathbb{N}$ sufficiently large makes the n -chain $S^m c \in C_n(X)$ *decomposable* with respect to this covering, meaning

$$S^m c = u + v \quad \text{for some} \quad u \in C_n(\mathcal{U}), \quad v \in C_n(\mathcal{V}),$$

because every singular n -simplex in the finite linear combination forming $S^m c$ can be assumed to have its image entirely inside either \mathcal{U} or \mathcal{V} . Moreover, if $c \in C_n(X)$ is a cycle, then $S^m c \in C_n(X)$ is also a cycle, and the chain homotopy relation

$$S^m c - c = \partial h_m c + h_m \partial c = \partial h_m c$$

shows that c and $S^m c$ represent the same singular homology class. A relative version of this observation will be used in the proof below, and we can now see the significance of the condition $B \subset \mathring{A}$: it means that the interiors of A and $X \setminus B$ form an open covering of X .

PROOF OF PROPOSITION 33.7. Given any class $[c] \in H_n(X, A)$ represented by a relative n -cycle $c \in C_n(X)$, we observe that for each $m \in \mathbb{N}$, the chain $S^m c \in C_n(X)$ satisfies

$$\partial(S^m c) = S^m(\partial c) \in C_{n-1}(A),$$

since the subdivision operator $S : C_*(X) \rightarrow C_*(X)$ preserves the subcomplex $C_*(A) \subset C_*(X)$, hence $S^m c$ is also a relative n -cycle. Moreover, the chain homotopy relation $S^m c - c = \partial h_m c + h_m \partial c$ implies $[S^m c] = [c] \in H_n(X, A)$, since $\partial c \in C_{n-1}(A)$ implies $h_m \partial c \in C_n(A)$. With this in mind, since the interiors of A and $X \setminus B$ cover X , we can assume without loss of generality after replacing c by $S^m c$ for some $m \in \mathbb{N}$ sufficiently large that the chain c can be decomposed as

$$c = c_A + c_{X \setminus B} \quad \text{for some} \quad c_A \in C_n(A), \quad c_{X \setminus B} \in C_n(X \setminus B).$$

Having made this assumption, the fact that $c \in C_n(X, A)$ is a relative n -cycle means $\partial c \in C_{n-1}(A)$ and therefore also $\partial c_{X \setminus B} \in C_n(A)$, so that $c_{X \setminus B}$ is a relative n -cycle in $(X \setminus B, A \setminus B)$, thus representing a class $[c_{X \setminus B}] \in H_n(X \setminus B, A \setminus B)$ that satisfies

$$i_*[c_{X \setminus B}] = [c].$$

This proves that $i_* : H_n(X \setminus B, A \setminus B) \rightarrow H_n(X, A)$ is surjective.

To show that $i_* : H_n(X \setminus B, A \setminus B) \rightarrow H_n(X, A)$ is injective, suppose $c \in C_n(X \setminus B)$ is a relative n -cycle in $(X \setminus B, A \setminus B)$ representing a class $[c] \in H_n(X \setminus B, A \setminus B)$ with $i_*[c] = 0 \in H_n(X, A)$, which means that if c is viewed as an n -chain in X , we have

$$c = \partial b + a \quad \text{for some } b \in C_{n+1}(X) \text{ and } a \in C_n(A).$$

By applying S^m to both sides for m sufficiently large, we can assume without loss of generality that b decomposes as

$$b = b_A + b_{X \setminus B} \quad \text{for some} \quad b_A \in C_n(A), \quad b_{X \setminus B} \in C_n(X \setminus B).$$

We then have $c - \partial b_{X \setminus B} = \partial b_A + a$, in which the left hand side is a chain in $X \setminus B$ and the right hand side is a chain in A , implying that the right hand side is also a chain in $A \setminus B$, and the relation $c = \partial b_{X \setminus B} + (\partial b_A + a)$ therefore implies $[c] = 0 \in H_n(X \setminus B, A \setminus B)$. \square

The proof of Theorem 33.4 is now complete, and now that we have established the existence of at least one axiomatic homology theory with any given choice of coefficient group, there are many immediate corollaries, e.g. the Brouwer fixed point theorem (cf. Exercise 28.7). In particular, the computation of $h_*(S^n)$ carried out in Lecture 28 can now be considered a valid computation of singular homology, giving

$$H_k(S^n; G) \cong \begin{cases} G & \text{if } k = 0 \text{ or } k = n, \\ 0 & \text{otherwise.} \end{cases}$$

33.4. Chain-level excision. When we study singular cohomology later in this semester, it will be useful to have a stronger variant of the excision property, one that applies to the singular *chain complex* rather than just to its homology:

THEOREM 33.9. *Assume $B \subset A \subset X$ such that the closure of B is contained in the interior of A . Then the inclusion of pairs $i : (X \setminus B, A \setminus B) \hookrightarrow (X, A)$ induces a chain homotopy equivalence $i_* : C_*(X \setminus B, A \setminus B) \rightarrow C_*(X, A)$.*

This obviously implies Proposition 33.7, and we will see in the following that it also more-or-less follows from it, for somewhat nontrivial reasons. The usefulness of Theorem 33.9 will be that it almost immediately implies that similar statements hold after applying certain standard algebraic operations to chain complexes, such as the one that turns homology into cohomology. One gets a hint of this from the following easy observation, which is based on the same trick as Remark 32.10:

LEMMA 33.10. *If Theorem 33.9 holds for singular chain complexes with coefficients in \mathbb{Z} , then it also holds for arbitrary choices of coefficients.*

PROOF. Under the stated hypotheses on $B \subset A \subset X$, assume it is known that the map $i_* : C_*(X \setminus B, A \setminus B; \mathbb{Z}) \rightarrow C_*(X, A; \mathbb{Z})$ is a chain homotopy equivalence, which means there exists a chain map $g : C_*(X, A; \mathbb{Z}) \rightarrow C_*(X \setminus B, A \setminus B; \mathbb{Z})$, a chain homotopy h_1 between $i_*g : C_*(X, A; \mathbb{Z}) \rightarrow C_*(X, A; \mathbb{Z})$ and the identity, and a chain homotopy h_2 between $gi_* : C_*(X \setminus B, A \setminus B; \mathbb{Z}) \rightarrow C_*(X \setminus B, A \setminus B; \mathbb{Z})$ and the identity. For any coefficient module G , linearity and the definitions of these maps on the generators of the singular chain complex (i.e. on singular simplices) uniquely determine similar chain maps and chain homotopies that relate $C_*(X \setminus B, A \setminus B; G)$ and $C_*(X, A; G)$ in the same manner. \square

With the lemma in mind, our goal is now to prove that Theorem 33.9 holds in the special case $G = \mathbb{Z}$. We will deduce this from some general results about chain complexes in the next subsection.

33.5. Chain contractions and mapping cones. Recall that a chain complex C_* is called *chain contractible* if there exists a chain homotopy of the identity map $C_* \rightarrow C_*$ to the trivial chain map $0 : C_* \rightarrow C_*$.

LEMMA 33.11. *A chain complex of R -modules C_* is chain contractible if and only if there is a splitting of C_n for each $n \in \mathbb{Z}$ into submodules $C_n = A_n \oplus B_n$ such that $C_n \xrightarrow{\partial} C_{n-1}$ vanishes on A_n and maps B_n isomorphically to A_{n-1} .*

PROOF. Assume $h : C_* \rightarrow C_*$ satisfies $h(C_n) \subset C_{n+1}$ for every $n \in \mathbb{Z}$ and $\partial h + h\partial = \mathbf{1}$. We observe that the homomorphisms ∂h and $h\partial$ in this case are complementary projections, since $\partial^2 = 0$ implies

$$(\partial h)^2 = \partial(h\partial)h = \partial(\mathbf{1} - \partial h)h = \partial h, \quad \text{and} \quad (h\partial)^2 = h(\partial h)\partial = h(\mathbf{1} - h\partial)\partial = h\partial.$$

We therefore obtain a splitting $C_* = A_* \oplus B_*$ with $A_* := \text{im}(\partial h)$ and $B_* := \text{im}(h\partial)$, and setting $A_n := A_* \cap C_n$ and $B_n := B_* \cap C_n$ for each $n \in \mathbb{Z}$ gives $C_n = A_n \oplus B_n$. Since $\partial(\partial h) = 0$, ∂ vanishes on A_* ; moreover, the definitions of the projections imply that ∂h is the identity map on A_* while $h\partial$ is the identity map on B_* , showing that for each $n \in \mathbb{Z}$, one obtains an inverse of $B_n \xrightarrow{\partial} A_{n-1}$ by composing $A_{n-1} \xrightarrow{h} C_n$ with the projection $C_n \xrightarrow{h\partial} B_n$.

Conversely, if splittings $C_n = A_n \oplus B_n$ with the stated properties are given, then defining $h : C_n \rightarrow C_{n+1}$ for each $n \in \mathbb{Z}$ to be trivial on B_n and an inverse of $B_{n+1} \xrightarrow{\partial} A_n$ on A_n gives a chain contraction. \square

We can now clarify the advantage of focusing on the case $G = \mathbb{Z}$ in the proof of Theorem 33.9: it is the fact that chain complexes over \mathbb{Z} are *free* abelian groups, thus making the following result applicable.

LEMMA 33.12. *A chain complex C_* of free abelian groups is acyclic if and only if it is chain contractible.*

PROOF. Assume the chain complex C_* is acyclic and $C_n \subset C_*$ is a free abelian group for each $n \in \mathbb{Z}$. By a basic result in algebra (see e.g. [Lan02, §III.7]), subgroups of free abelian groups are also free, and this applies in particular to the subgroups

$$Z_n := \ker \left(C_n \xrightarrow{\partial_n} C_{n-1} \right).$$

Acyclicity means that the map

$$C_n \xrightarrow{\partial_n} Z_{n-1}$$

is surjective for every n , and we therefore have a short exact sequence

$$0 \rightarrow Z_n \hookrightarrow C_n \xrightarrow{\partial} Z_{n-1} \rightarrow 0,$$

which splits since Z_{n-1} is free. This splitting identifies C_n with $Z_n \oplus Z_{n-1}$ so that ∂ becomes the projection $Z_n \oplus Z_{n-1} \rightarrow Z_{n-1}$, and chain contractibility now follows from Lemma 33.11. \square

REMARK 33.13. Lemma 33.12 also holds under the hypothesis that each $C_n \subset C_*$ is a free R -module if the underlying ring R is a principal ideal domain: the key detail is that under this assumption, submodules of free R -modules are also free, thus providing the splitting of short exact sequences used in the proof. This fact about principal ideal domains depends in general on Zorn's lemma, so it may seem a bit abstract, but one can also avoid using it if one is willing to assume the chain complex is *bounded* above or below, which also suffices for our present purposes; see Exercise 33.2.

What we need next is a way to deduce that something is a chain homotopy equivalence from the fact that some other complex is chain contractible. The right tool for this is the mapping cone.

A quick digression on simplicial complexes will provide some useful motivation. For a simplicial pair (K, L) , the cone CL of L also contains L itself as a subcomplex, so we can define the **cone** of the pair (K, L) as the simplicial complex

$$\text{cone}(K, L) := CL \cup_L K,$$

in which the subcomplexes $L \subset CL$ and $L \subset K$ are identified with each other. The vertices of $\text{cone}(K, L)$ thus consist of the vertices of K plus one extra vertex labelled $*$, while its simplices consist of the simplices of K plus, for each $n \geq 0$ and each n -simplex $\{v_0, \dots, v_n\}$ of L , the $(n+1)$ -simplex $\{*, v_0, \dots, v_n\}$. Topologically, the polyhedron $|\text{cone}(K, L)|$ is a space obtained by attaching $|K|$ to the cone $C|L|$ along $|L|$, thus making the inclusion $|L| \hookrightarrow |\text{cone}(K, L)|$ homotopic to a constant map.

The augmented chain complex $\tilde{C}_*^\Delta(\text{cone}(K, L); \mathbb{Z})$ contains two types of generators. First, since $K \subset \text{cone}(K, L)$ is a subcomplex, there are the generators corresponding to simplices of K , in addition to $1 \in \mathbb{Z} = \tilde{C}_{-1}^\Delta(K; \mathbb{Z})$, making $\tilde{C}_*^\Delta(K; \mathbb{Z})$ a subcomplex of $\tilde{C}_*^\Delta(\text{cone}(K, L); \mathbb{Z})$. Secondly, each oriented simplex $[v_0, \dots, v_n]$ of L gives rise to a generator $[*, v_0, \dots, v_n]$ of $\tilde{C}_*^\Delta(\text{cone}(K, L); \mathbb{Z})$, defining for each $n \geq -1$ a homomorphism

$$\tilde{C}_n^\Delta(L; \mathbb{Z}) \xrightarrow{j} \tilde{C}_{n+1}^\Delta(\text{cone}(K, L); \mathbb{Z})$$

such that $j[v_0, \dots, v_n] := [*, v_0, \dots, v_n]$ and, for the case $n = -1$, $j(1) := [*]$. The map j identifies every n -chain in $\tilde{C}_*^\Delta(L; \mathbb{Z})$ with an $(n+1)$ -chain in $\tilde{C}_*^\Delta(\text{cone}(K, L); \mathbb{Z})$, but j is not a chain map and its image is not a subcomplex: instead, we have

$$\partial j[v_0, \dots, v_n] = \partial[*, v_0, \dots, v_n] = [v_0, \dots, v_n] - j(\partial[v_0, \dots, v_n])$$

and, using $\partial = \epsilon_*$ for the degree 0 part of the augmented chain complex, $\partial j(1) = \epsilon_*[*] = 1 = 1 - j(\partial(1))$. The result is a direct sum decomposition

$$\tilde{C}_n^\Delta(\text{cone}(K, L); \mathbb{Z}) \cong \tilde{C}_{n-1}^\Delta(L; \mathbb{Z}) \oplus \tilde{C}_n^\Delta(K; \mathbb{Z})$$

for each $n \geq -1$ such that the boundary map on $\tilde{C}_*^\Delta(\text{cone}(K, L); \mathbb{Z})$ decomposes in block form as

$$\partial = \begin{pmatrix} -\partial^L & 0 \\ i_* & \partial^K \end{pmatrix},$$

where ∂^L and ∂^K denote the boundary maps on the augmented simplicial chain complexes of L and K respectively, and $i : L \hookrightarrow K$ is the inclusion. One can take this as topological motivation for the following algebraic definition.

DEFINITION 33.14. The **mapping cone** of a chain map $f : (A_*, \partial^A) \rightarrow (B_*, \partial^B)$ is the chain complex $(\text{cone}(f)_*, \partial)$ with

$$\text{cone}(f)_n := A_{n-1} \oplus B_n \quad \text{and} \quad \partial := \begin{pmatrix} -\partial^A & 0 \\ f & \partial^B \end{pmatrix}.$$

REMARK 33.15. The literature contains a variety of alternative versions of Definition 33.14 with slightly different sign conventions.

It is straightforward to check that for any chain map $f : A_* \rightarrow B_*$, one obtains a short exact sequence of chain complexes

$$0 \longrightarrow B_* \xrightarrow{i} \text{cone}(f)_* \xrightarrow{\pi} A_*[-1] \longrightarrow 0,$$

where we denote by $A_*[-1]$ the chain complex A_* with its grading shifted so that $A_*[-1]_n := A_{n-1}$, and the maps i and π are the obvious inclusion and projection respectively,

$$B_n \xrightarrow{i} A_{n-1} \oplus B_n, \quad A_{n-1} \oplus B_n \xrightarrow{\pi} A_{n-1}.$$

Plugging this into Proposition 32.13 thus gives a long exact sequence that relates the homology groups of A_* , B_* and the cone, and by inspection of the usual diagram chase, one finds that the connecting homomorphism $H_n(A_*[-1]) = H_{n-1}(A_*) \xrightarrow{\hat{c}_*} H_{n-1}(B_*)$ in this case is simply the map induced on homology by the chain map $f : A_* \rightarrow B_*$, so the long exact sequence takes the form

$$(33.4) \quad \dots \longrightarrow H_n(A_*) \xrightarrow{f_*} H_n(B_*) \xrightarrow{i_*} H_n(\text{cone}(f)_*) \xrightarrow{\pi_*} H_{n-1}(A_*) \xrightarrow{f_*} H_{n-1}(B_*) \longrightarrow \dots$$

The exactness of this sequence implies:

PROPOSITION 33.16. *A chain map $f : A_* \rightarrow B_*$ induces isomorphisms $H_n(A_*) \rightarrow H_n(B_*)$ for all $n \in \mathbb{Z}$ if and only if its mapping cone $\text{cone}(f)_*$ is acyclic.* \square

The following is a chain-level analogue of Proposition 33.16.

THEOREM 33.17. *A chain map $f : A_* \rightarrow B_*$ is a chain homotopy equivalence if and only if its mapping cone $\text{cone}(f)_*$ is chain contractible.*

PROOF. Suppose $\text{cone}(f)_*$ admits a chain contraction, so for each $n \in \mathbb{Z}$, there is a homomorphism

$$h = \begin{pmatrix} \alpha & \beta \\ \gamma & \delta \end{pmatrix} : A_{n-1} \oplus B_n = \text{cone}(f)_n \rightarrow \text{cone}(f)_{n+1} = A_n \oplus B_{n+1}$$

satisfying $h\partial + \partial h = \mathbb{1}$, which amounts to the four equations

$$\begin{aligned} -\partial^A \alpha - \alpha \partial^A + \beta f &= \mathbb{1}, \\ f \alpha + \partial^B \gamma - \gamma \partial^A + \delta f &= 0, \\ -\partial^A \beta + \beta \partial^B &= 0, \\ f \beta + \partial^B \delta + \delta \partial^B &= \mathbb{1} \end{aligned}$$

for the maps $\alpha : A_{n-1} \rightarrow A_n$, $\beta : B_n \rightarrow A_n$, $\gamma : A_{n-1} \rightarrow B_{n+1}$ and $\delta : B_n \rightarrow B_{n+1}$. The third equation makes β a chain map $B_* \rightarrow A_*$, the first makes $-\alpha$ a chain homotopy between $\beta \circ f$ and the identity $A_* \rightarrow A_*$, and the fourth makes δ a chain homotopy between $f \circ \beta$ and the identity $B_* \rightarrow B_*$, proving that β is a chain homotopy inverse of f .

We will not need the converse in the application below, so we include here only a sketch of its proof, adapted from an argument in [Bro94, Prop. 0.7].

As a preparatory observation, note that for any two chain complexes (A_*, ∂^A) and (B_*, ∂^B) , there is a chain complex $(\text{Hom}(A_*, B_*), \partial)$ whose degree n part consists of the homomorphisms

$\varphi : A_* \rightarrow B_*$ that satisfy $\varphi(A_k) \subset B_{k+n}$ for all $k \in \mathbb{Z}$, with the boundary operator ∂ given on homogeneous elements $\varphi \in \text{Hom}(A_*, B_*)$ of degree $|\varphi|$ by

$$\partial\varphi := \partial^B \circ \varphi - (-1)^{|\varphi|} \varphi \circ \partial^A.$$

The choice of sign convention used here is motivated by a convention that we will later use for defining tensor products of chain complexes, and it ensures for instance that the obvious evaluation map

$$\text{Hom}(A_*, B_*) \otimes A_* \rightarrow B_* : \varphi \otimes a \mapsto \varphi(a)$$

is a chain map. This detail is unimportant for now; one can easily check in any case that $(\text{Hom}(A_*, B_*), \partial)$ as defined above is a chain complex. Moreover, the 0-cycles in $\text{Hom}(A_*, B_*)$ are precisely the chain maps from A_* to B_* , and two such cycles are homologous if and only if they are chain homotopic.

Next, we have two claims whose proofs are both straightforward exercises:

Claim 1: For any fixed chain complex C_* , there exists a covariant functor $\text{Ch} \rightarrow \text{Ch}$ that sends each chain complex A_* to the chain complex $\text{Hom}(C_*, A_*)$ and sends each chain map $f : A_* \rightarrow B_*$ to the chain map

$$\text{Hom}(C_*, f) : \text{Hom}(C_*, A_*) \rightarrow \text{Hom}(C_*, B_*) : \varphi \mapsto f \circ \varphi,$$

and moreover, the chain homotopy class of $\text{Hom}(C_*, f)$ depends only on the chain homotopy class of f .

Claim 2: For any chain map $f : A_* \rightarrow B_*$ and any third chain complex C_* , there is a natural isomorphism between the chain complexes $\text{Hom}(C_*, \text{cone}(f)_*)$ and $\text{cone}(\text{Hom}(C_*, f))_*$.

With these ingredients in place, suppose $f : A_* \rightarrow B_*$ is a chain homotopy equivalence, and abbreviate $C_* := \text{cone}(f)_*$. Claim 1 implies that $\text{Hom}(C_*, f) : \text{Hom}(C_*, A_*) \rightarrow \text{Hom}(C_*, B_*)$ is then also a chain homotopy equivalence, and by claim 2, its mapping cone is naturally isomorphic to $\text{Hom}(C_*, C_*)$, implying via Proposition 33.16 that $\text{Hom}(C_*, C_*)$ is acyclic. The vanishing of $H_0(\text{Hom}(C_*, C_*))$ means that every chain map $C_* \rightarrow C_*$ is chain homotopic to zero: since this applies in particular to the identity map $C_* \rightarrow C_*$, it follows that C_* is chain contractible. \square

Theorem 33.9 in the case $G = \mathbb{Z}$ is an immediate consequence of the following:

COROLLARY 33.18. *For two chain complexes A_*, B_* of free abelian groups, a chain map $f : A_* \rightarrow B_*$ is a chain homotopy equivalence if and only if the induced maps $f_* : H_n(A_*) \rightarrow H_n(B_*)$ are isomorphisms for all $n \in \mathbb{Z}$.*

PROOF. If $f_* : H_n(A_*) \rightarrow H_n(B_*)$ is an isomorphism for every n , then by Proposition 33.16, $\text{cone}(f)_*$ is acyclic. Since the chain groups A_{n-1} and B_n are free abelian groups, the same holds for $\text{cone}(f)_n = A_{n-1} \oplus B_n$, and it then follows via Lemma 33.12 that $\text{cone}(f)_*$ is also chain contractible. The result now follows from Theorem 33.17. \square

33.6. Exercises.

EXERCISE 33.1. 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 (33.3) defines a natural transformation from π_1 to $H_1(\cdot; \mathbb{Z})$.

EXERCISE 33.2. A chain complex C_* is said to be **bounded below** or **bounded above** if $C_n = 0$ for all $n \in \mathbb{Z}$ sufficiently small or sufficiently large respectively, e.g. all of the chain complexes that we have used for defining topological invariants so far have been bounded below,

since they satisfy $C_n = 0$ for $n < 0$. Show that if C_* is an acyclic chain complex of R -modules that is bounded above or below and the modules $C_n \subset C_*$ are all free, then C_* is chain contractible.

Hint: Construct a chain contraction inductively by degree, as in the method of acyclic models. The assumption that C_ is bounded above or below gives you a place to start the induction.*

EXERCISE 33.3 (*). Prove that for any two subsets $U, V \subset X$ with $X = \overset{\circ}{U} \cup \overset{\circ}{V}$, the obvious inclusion

$$C_*(U) + C_*(V) \hookrightarrow C_*(X)$$

is a chain homotopy equivalence.

EXERCISE 33.4 (*). The **reduced** version of singular homology

$$\tilde{H}_n(X) = \tilde{H}_n(X; G) := \ker \left(H_n(X; G) \xrightarrow{\epsilon_*} H_n(\{*\}; G) \right)$$

is defined in terms of the unique map $\epsilon : X \rightarrow \{*\}$. Show that $\tilde{H}_*(X)$ can also be identified in a natural way with the homology of an **augmented** singular chain complex $\tilde{C}_*(X; G)$, taking the form

$$\dots \longrightarrow C_2(X; G) \xrightarrow{\partial} C_1(X; G) \xrightarrow{\partial} C_0(X; G) \xrightarrow{\epsilon_*} G \longrightarrow 0 \longrightarrow 0 \longrightarrow \dots,$$

where the **augmentation** map $\epsilon_* : C_0(X; G) \rightarrow G$ takes the form $\epsilon_* \sum_i a_i \sigma_i := \sum_i a_i$.

34. Pairs, triples, and the Mayer-Vietoris sequence

In this lecture we discuss three further properties that are common to all axiomatic homology theories, though some of them are a bit easier to understand in the specific example of singular homology.

34.1. Triples. A **triple** (X, A, B) of spaces consists of a space X and two nested subsets $B \subset A \subset X$. In this situation, there are two obvious inclusion maps

$$(A, B) \xrightarrow{i} (X, B) \xrightarrow{j} (X, A),$$

and in singular homology there are also natural homomorphisms

$$(34.1) \quad H_n(X, A) \xrightarrow{\hat{c}_*} H_{n-1}(A, B) : [c] \mapsto [\partial c],$$

where $c \in C_n(X)$ denotes a relative n -cycle in (X, A) , which makes $\partial c \in C_{n-1}(A)$ an $(n-1)$ -cycle in A and therefore also a relative $(n-1)$ -cycle in (A, B) .

THEOREM 34.1. *For any axiomatic homology theory h_* , there exist natural homomorphisms $\partial_* : h_n(X, A) \rightarrow h_{n-1}(A, B)$ associated to every triple of spaces (X, A, B) such that the sequence*

$$\dots \longrightarrow h_n(A, B) \xrightarrow{i_*} h_n(X, B) \xrightarrow{j_*} h_n(X, A) \xrightarrow{\hat{c}_*} h_{n-1}(A, B) \longrightarrow \dots \longrightarrow h_0(X, A) \longrightarrow 0$$


is exact, and in the case of singular homology, ∂_ is the map defined in (34.1).*

PROOF. Let us first give a proof that applies only to singular homology, since it is quicker and more intuitive. It is simply a matter of applying Proposition 32.13 with the short exact sequence of relative chain complexes

$$0 \longrightarrow C_*(A, B) \xrightarrow{i_*} C_*(X, B) \xrightarrow{j_*} C_*(X, A) \longrightarrow 0$$

and then inspecting the diagram chase to check that the resulting connecting homomorphism is the map specified in (34.1).

A different argument is needed in order to produce the same result for an arbitrary axiomatic homology theory h_* . It is based on the following “braid” diagram:
(34.2)



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) , as provided by the exactness axiom. The fourth strand is the sequence

$$(34.3) \quad \dots \rightarrow h_{n+1}(X, A) \xrightarrow{\partial} h_n(A, B) \xrightarrow{i} h_n(X, B) \xrightarrow{j} h_n(X, A) \xrightarrow{\partial} h_{n-1}(A, B) \rightarrow \dots,$$

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. The rest is an exercise in diagram chasing; see Exercise 34.1. \square

34.2. Good pairs. 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 the *reduced* homology $\tilde{H}_*(X/A)$ of the quotient space. 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 34.2. *For any homology theory h_* , any space X and a point $x \in X$, the inclusion of pairs $(X, \emptyset) \hookrightarrow (X, \{x\})$ induces an isomorphism*

$$\tilde{h}_*(X) \xrightarrow{\cong} \tilde{h}_*(X, \{x\}) = h_*(X, \{x\}).$$

PROOF. This is immediate from the long exact sequence of $(X, \{x\})$ in reduced homology since $\tilde{h}_*(\{x\}) = 0$. \square

Now, observe that the quotient projection $q : X \rightarrow X/A$ is also a map of pairs $(X, A) \rightarrow (X/A, A/A)$ and thus induces a morphism $h_*(X, A) \rightarrow h_*(X/A, A/A)$. 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

$$(X \setminus A, A \setminus A) \xrightarrow{q} ((X/A) \setminus (A/A), (A/A) \setminus (A/A)),$$

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 $B \subset A \subset X$ that we remove from A needs

to satisfy a strict condition separating it from $X \setminus A$, and this condition usually does not hold if $B = A$. 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 adapted from [Hat02].

DEFINITION 34.3. 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 $V \subset X$ such that there exists a continuous function $u : X \rightarrow I$ with $u|_A \equiv 0$ and $u|_{X \setminus V} \equiv 1$.

REMARK 34.4. We have formulated Definition 34.3 to be compatible with the hypothesis of the excision axiom as stated in Lecture 28, but if h_* is singular homology, for which excision holds under a weaker hypothesis, then one can also use a weaker variant of Definition 34.3 in which the existence of the separating function $u : X \rightarrow I$ does not need to be assumed explicitly. That is how the condition is stated in [Hat02].

EXAMPLE 34.5. $(\mathbb{D}^n, \partial\mathbb{D}^n)$ 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 34.6. 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 34.7 below; see Exercise 34.2.

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 34.7. *If (X, A) is a good pair, then for every axiomatic homology theory h_* , the natural quotient map $q : (X, A) \rightarrow (X/A, A/A)$ induces an isomorphism*

$$q_* : h_*(X, A) \xrightarrow{\cong} h_*(X/A, A/A),$$

implying via Lemma 34.2 that there is a natural isomorphism $h_*(X, A) \cong \tilde{h}_*(X/A)$.

PROOF. Fix a neighborhood $V \subset X$ of A such that A is a deformation retract of V and there exists a function $u : X \rightarrow I$ that vanishes on A and equals 1 outside of V . The deformation retraction implies that 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

$$0 = h_k(V, A) \rightarrow h_k(X, A) \rightarrow h_k(X, V) \rightarrow h_{k-1}(V, A) = 0,$$

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 $h_*(V/A, A/A) \cong \tilde{h}_*(V/A) = 0$, 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

$$\begin{array}{ccccc} h_*(X, A) & \xrightarrow{i_*} & h_*(X, V) & \xleftarrow{k_*} & h_*(X \setminus A, V \setminus A) \\ \downarrow q_* & & \downarrow q_* & & \downarrow q_* \\ h_*(X/A, A/A) & \xrightarrow{j_*} & h_*(X/A, V/A) & \xleftarrow{\ell_*} & h_*((X/A) \setminus (A/A), (V/A) \setminus (A/A)), \end{array}$$

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. The rightmost map labeled q_* in this diagram is also an isomorphism since it is induced by the map

$$(X \setminus A, V \setminus A) \xrightarrow{q} ((X/A) \setminus (A/A), (V/A) \setminus (A/A)),$$

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

The following simple example will appear frequently when we compute the homology of CW-complexes.

EXAMPLE 34.8. 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

$$h_k(\mathbb{D}^n, \partial \mathbb{D}^n) \cong \tilde{h}_k(\mathbb{D}^n / \partial \mathbb{D}^n) \cong \tilde{h}_k(S^n) \cong \begin{cases} G & \text{if } k = n, \\ 0 & \text{otherwise,} \end{cases}$$

where $G := h_0(\{*\})$ 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 $h_k(\mathbb{D}^n, \partial \mathbb{D}^n) \rightarrow \tilde{h}_{k-1}(S^{n-1})$ is an isomorphism.)

34.3. Mayer-Vietoris sequences in simplicial homology. It is time to discuss the analogue in homology of the Seifert-van Kampen theorem.

The problem is as follows: we are given a space $X = A \cup B$, and we would like to compute the homology $h_*(X)$ in terms of $h_*(A)$, $h_*(B)$ and $h_*(A \cap B)$. Of the versions of homology that we've discussed so far, the problem is simplest to solve in simplicial homology, so let's start with that. As in Lecture 32, we use the notation $H_*^\bullet(K)$ and $C_*^\bullet(K)$ to denote either the ordered or oriented versions of simplicial homology.

Suppose $K = A \cup B$ is a simplicial complex containing $A, B \subset K$ as subcomplexes, in which case $A \cap B$ is automatically also a subcomplex of K . For clarity: the notation $K = A \cup B$ means that every simplex of K is a simplex of A or a simplex of B , and $A \cap B$ denotes the subcomplex of K consisting of simplices that belong to both of the subcomplexes A and B . The inclusions

$$(34.4) \quad \begin{array}{ccc} & A & \\ i^A \nearrow & & \searrow j^A \\ A \cap B & & K \\ i^B \searrow & & \nearrow j^B \\ & B & \end{array}$$

are all simplicial maps, and the induced chain maps can be assembled into a short exact sequence of simplicial chain complexes

$$(34.5) \quad 0 \longrightarrow C_*^\bullet(A \cap B) \xrightarrow{(i_*^A, -i_*^B)} C_*^\bullet(A) \oplus C_*^\bullet(B) \xrightarrow{j_*^A \oplus j_*^B} C_*^\bullet(K) \longrightarrow 0.$$

Note in particular that the map $j_*^A \oplus j_*^B$ is surjective since every generator of $C_*^\bullet(K)$ is also a generator of $C_*^\bullet(A)$ or $C_*^\bullet(B)$; this is worth pointing out because it is the detail that will become more complicated when we adapt this discussion to singular homology, but let's come back to that in a moment. Plugging (34.5) into Proposition 32.13 gives the **Mayer-Vietoris sequence** of simplicial homology groups, in which close inspection of the usual diagram chase reveals the connecting homomorphism to be of the form stated below:

THEOREM 34.9. *For any simplicial complex $K = A \cup B$ that is the union of two subcomplexes $A, B \subset K$, with simplicial inclusion maps written as in (34.4), there exists a long exact sequence*

$$\begin{aligned} \dots \longrightarrow H_{n+1}^\bullet(K) \xrightarrow{\hat{c}_*} H_n^\bullet(A \cap B) \xrightarrow{(i_*^A, -i_*^B)} H_n^\bullet(A) \oplus H_n^\bullet(B) \xrightarrow{j_*^A \oplus j_*^B} H_n^\bullet(K) \xrightarrow{\hat{c}_*} H_{n-1}^\bullet(A \cap B) \\ \longrightarrow \dots \longrightarrow H_0^\bullet(A \cap B) \xrightarrow{(i_*^A, -i_*^B)} H_0^\bullet(A) \oplus H_0^\bullet(B) \xrightarrow{j_*^A \oplus j_*^B} H_0^\bullet(K) \longrightarrow 0, \end{aligned}$$

with connecting homomorphisms $\partial_* : H_n^\bullet(K) \rightarrow H_{n-1}^\bullet(A \cap B)$ characterized by the formula

$$\partial_*[a + b] = [\partial a]$$

for any simplicial n -cycle of the form $a + b \in C_n^\bullet(K)$ with $a \in C_n^\bullet(A)$ and $b \in C_n^\bullet(B)$. □

REMARK 34.10. The placement of the minus signs in (34.5) and in the exact sequence of Theorem 34.9 is not universally standard; other reasonable conventions on this are possible.

REMARK 34.11. One easily checks that the connecting homomorphisms in Theorem 34.9 (and therefore the entire Mayer-Vietoris sequence) are natural in the sense that for any simplicial map $f : K \rightarrow K'$ such that $K' = A' \cup B'$ for subcomplexes $A', B' \subset K'$ with $f(A) \subset A'$ and $f(B) \subset B'$, the diagram

$$\begin{array}{ccc} H_n^\bullet(K) & \xrightarrow{\hat{c}_*} & H_{n-1}^\bullet(A \cap B) \\ \downarrow f_* & & \downarrow f_* \\ H_n^\bullet(K') & \xrightarrow{\hat{c}_*} & H_{n-1}^\bullet(A' \cap B') \end{array}$$

commutes.

34.4. Mayer-Vietoris sequences in singular homology. Consider next an arbitrary topological space $X = A \cup B$ that is the union of two subsets $A, B \subset X$, with inclusion maps denoted by

(34.6)

The obvious analogue of the short exact sequence (34.5) for the singular chain complex is not generally exact, because the map $j_*^A \oplus j_*^B : C_*(A) \oplus C_*(B) \rightarrow C_*(X)$ is not surjective: there can be singular simplices $\sigma : \Delta^n \rightarrow X$ with image not fully contained in either A or B . We've seen this problem before, namely in the proof of the excision property for singular homology, and we know how to solve it: subdivision. But in order to make this solution possible, we will have to assume some conditions on the subsets $A, B \subset X$. Let's first observe what can be proved easily: the obvious analogue of (34.5) does become exact if we replace the final term $C_*(X)$ with the subcomplex

$$C_*(A + B) := C_*(A) + C_*(B) \subset C_*(X),$$

producing the short exact sequence

(34.7) $0 \longrightarrow C_*(A \cap B) \xrightarrow{(i_*^A, -i_*^B)} C_*(A) \oplus C_*(B) \xrightarrow{J^A \oplus J^B} C_*(A + B) \longrightarrow 0,$

where we denote by J^A and J^B the inclusions of chain complexes

$$C_*(A) \xleftarrow{J^A} C_*(A + B) \xleftarrow{J^B} C_*(B).$$

There is of course a long exact sequence of homology groups associated to the short exact sequence above, but in order for that long exact sequence to be useful as a way of computing $H_*(X)$, we need to be able to deduce the latter from the subcomplex $C_*(A + B) \subset C_*(X)$. Subdivision will make this possible in certain situations, in particular by establishing the conditions of Definition 34.12 below.

NOTATION. We have left the coefficient group G unspecified in this discussion; for situations where the choice of coefficients matters, we shall denote

$$C_*(A + B; G) := C_*(A; G) + C_*(B; G) \subset C_*(X; G).$$

DEFINITION 34.12. For a space X , two subsets $A, B \subset X$ are said to form an **excisive couple** in X (for singular homology) if the inclusion of subcomplexes $C_*(A + B; \mathbb{Z}) \hookrightarrow C_*(A \cup B; \mathbb{Z}) \subset C_*(X; \mathbb{Z})$ induces an isomorphism on homology,

$$H_*(C_*(A + B; \mathbb{Z})) \xrightarrow{\cong} H_*(A \cup B; \mathbb{Z}).$$

REMARK 34.13. We have formulated Definition 34.12 so that the condition is maximally easy to check, but several seemingly stronger conditions are equivalent to it. Indeed, since $C_*(A \cup B; \mathbb{Z})$ and $C_*(A + B; \mathbb{Z})$ are chain complexes of free abelian groups, the defining property of an excisive couple implies via Corollary 33.18 that the inclusion $C_*(A + B; \mathbb{Z}) \hookrightarrow C_*(A \cup B; \mathbb{Z})$ is also a chain homotopy equivalence. From this, it follows as in Remark 32.10 and Lemma 33.10 that $C_*(A + B; G) \hookrightarrow C_*(A \cup B; G)$ is also a chain homotopy equivalence for arbitrary coefficient groups G , and the homological condition in Definition 34.12 therefore also holds with an arbitrary choice of coefficients. With this in mind, we shall mostly go back to omitting the coefficients from the notation from here on.

Here is the most common situation in which one encounters excisive couples:

PROPOSITION 34.14. *If $A, B \subset X$ are subsets such that $A \cup B = \overset{\circ}{A} \cup \overset{\circ}{B}$ is the union of their interiors, then they form an excisive couple. In particular, this holds whenever A and B are both open in X .*

PROOF. We can replace X by $A \cup B$ and thus assume without loss of generality that the interiors of A and B cover X . The result is then immediate from Exercise 33.3, which states that the inclusion $C_*(A + B) \hookrightarrow C_*(X)$ is in this case a chain homotopy equivalence. In fact, one can first use barycentric subdivision to prove that $A, B \subset X$ forms an excisive couple; the crucial detail is that since every point lies in at least one of either $\overset{\circ}{A}$ or $\overset{\circ}{B}$, applying barycentric subdivision sufficiently many times to any given chain produces one whose constituent singular simplices are each fully contained in either A or B . This result implies the stronger statement in Exercise 33.3 due to the general properties of chain complexes and mapping cones discussed in §33.5. \square

For an excisive couple $A, B \subset X$ with $X = A \cup B$, we can define homomorphisms

$$H_n(X) \xrightarrow{\hat{c}_*} H_{n-1}(A \cap B) : [a + b] \mapsto [\partial a] = -[\partial b], \quad \text{where } a \in C_n(A) \text{ and } b \in C_n(B).$$

Indeed, the defining property of an excisive couple implies that every homology class $[c] \in H_n(X)$ can be represented by the sum of a chain a in A with a chain b in B , and the condition that $a + b$ is a cycle then means $\partial a = -\partial b$, so that the $(n - 1)$ -cycle ∂a necessarily lies in $A \cap B$. One can deduce from the injectivity of the map $H_n(C_*(A + B)) \rightarrow H_n(X)$ that the homology class $[\partial a] \in H_{n-1}(A \cap B)$ does not depend on the choice of representation $[c] = [a + b]$; see Exercise 34.3(a). Applying Proposition 32.13 and Exercise 34.3(b), we have:

THEOREM 34.15. *For any space $X = A \cup B$ that is the union of two subsets forming an excisive couple for singular homology, the sequence*

$$\begin{aligned} \dots \longrightarrow H_{n+1}(X) \xrightarrow{\hat{c}_*} H_n(A \cap B) \xrightarrow{(i_*^A, -i_*^B)} H_n(A) \oplus H_n(B) \xrightarrow{j_*^A \oplus j_*^B} H_n(X) \xrightarrow{\hat{c}_*} H_{n-1}(A \cap B) \\ \longrightarrow \dots \longrightarrow H_0(A \cap B) \xrightarrow{(i_*^A, -i_*^B)} H_0(A) \oplus H_0(B) \xrightarrow{j_*^A \oplus j_*^B} H_0(X) \longrightarrow 0 \end{aligned}$$

is exact. Moreover, the connecting homomorphisms are natural in the sense that for any map $f : X \rightarrow X'$ such that $X' = A' \cup B'$ with $A', B' \subset X'$ an excisive couple and $f(A) \subset A'$ and $f(B) \subset B'$, the diagram

$$\begin{array}{ccc} H_n(X) & \xrightarrow{\hat{c}_*} & H_{n-1}(A \cap B) \\ \downarrow f_* & & \downarrow f_* \\ H_n(X') & \xrightarrow{\hat{c}_*} & H_{n-1}(A' \cap B') \end{array}$$

commutes. □

Note that the naturality part of the statement follows from the functoriality of the short \rightarrow long exact sequence construction, cf. Proposition 32.13 and Exercise 32.3.

We will have more use for excisive couples later, when we discuss product structures in relative homology and cohomology. The following result will be important in that discussion, and also provides a useful hint about how to generalize beyond singular homology to the axiomatic setting. Observe that for any two subsets $A, B \subset X$, we have a nested sequence of subcomplexes

$$C_*(A + B) = C_*(A) + C_*(B) \subset C_*(A \cup B) \subset C_*(X)$$

and therefore also quotient projections

$$C_*(X) \rightarrow \frac{C_*(X)}{C_*(A + B)} \rightarrow \frac{C_*(X)}{C_*(A \cup B)} = C_*(X, A \cup B),$$

which are chain maps.

PROPOSITION 34.16. *For a space X with two subsets $A, B \subset X$, the following conditions are equivalent.*

- (i) $A, B \subset X$ form an excisive couple for singular homology;
- (ii) The quotient projection

$$\frac{C_*(X; \mathbb{Z})}{C_*(A + B; \mathbb{Z})} \rightarrow C_*(X, A \cup B; \mathbb{Z})$$

induces isomorphisms of homology groups;

- (iii) The map $H_*(A, A \cap B; \mathbb{Z}) \rightarrow H_*(A \cup B, B; \mathbb{Z})$ induced by the inclusion $(A, A \cap B) \hookrightarrow (A \cup B, B)$ is an isomorphism;
- (iv) The map $H_*(B, A \cap B; \mathbb{Z}) \rightarrow H_*(A \cup B, A; \mathbb{Z})$ induced by the inclusion $(B, A \cap B) \hookrightarrow (A \cup B, A)$ is an isomorphism.

Moreover, if any of these conditions hold, then the same conditions also hold with \mathbb{Z} replaced by an arbitrary coefficient group G .

PROOF. We shall prove that the four statements are equivalent using an arbitrary coefficient group G instead of \mathbb{Z} . The fact that everything then follows from the special case $G = \mathbb{Z}$ will then be a consequence of Remark 34.13.

To see that conditions (i) and (ii) are equivalent, consider the commutative diagram

$$\begin{array}{ccccccc}
 0 & \longrightarrow & C_*(A + B) & \hookrightarrow & C_*(X) & \longrightarrow & \frac{C_*(X)}{C_*(A + B)} \longrightarrow 0 \\
 & & \downarrow & & \downarrow \mathbb{1} & & \downarrow \\
 0 & \longrightarrow & C_*(A \cup B) & \hookrightarrow & C_*(X) & \longrightarrow & C_*(X, A \cup B) \longrightarrow 0,
 \end{array}$$

in which both rows are short exact sequences of chain complexes and all arrows are either inclusions or quotient projections. Transforming both rows via Proposition 32.13 into long exact sequences of homology groups then produces a diagram of the form

$$\begin{array}{ccccccccc}
 \dots & \longrightarrow & H_n(C_*(A + B)) & \longrightarrow & H_n(X) & \longrightarrow & H_n\left(\frac{C_*(X)}{C_*(A + B)}\right) & \longrightarrow & H_{n-1}(C_*(A + B)) & \longrightarrow & H_{n-1}(X) & \longrightarrow & \dots \\
 & & \downarrow & & \downarrow \mathbb{1} & & \downarrow & & \downarrow & & \downarrow \mathbb{1} & & \\
 \dots & \longrightarrow & H_n(A \cup B) & \longrightarrow & H_n(X) & \longrightarrow & H_n(X, A \cup B) & \longrightarrow & H_{n-1}(A \cup B) & \longrightarrow & H_{n-1}(X) & \longrightarrow & \dots
 \end{array}$$

Each of (i) and (ii) amounts to a condition in which every third vertical arrow in this last diagram is an isomorphism, so if either holds, then the five-lemma implies that the other one does as well.

We next prove that conditions (i) and (iii) are equivalent; since condition (i) is symmetric with respect to the interchange of A and B , it will follow that both are also equivalent to (iv). We observe first that the inclusion $C_*(A) \hookrightarrow C_*(A + B)$ descends to an isomorphism of quotient complexes

$$C_*(A, A \cap B) = \frac{C_*(A)}{C_*(A \cap B)} \cong \frac{C_*(A + B)}{C_*(B)}.$$

The diagram

$$\begin{array}{ccccc}
 \frac{C_*(A)}{C_*(A \cap B)} & \xlongequal{\quad} & C_*(A, A \cap B) & \xrightarrow{i_*} & C_*(A \cup B, B) & \xlongequal{\quad} & \frac{C_*(A \cup B)}{C_*(B)} \\
 & & & \searrow \cong & & \nearrow \Phi & \\
 & & & & \frac{C_*(A + B)}{C_*(B)} & &
 \end{array}$$

now shows that the chain map induced by the inclusion $i : (A, A \cap B) \hookrightarrow (A \cup B, B)$ induces isomorphisms on homology if and only if the same is true of the chain map Φ , which is the result of letting the inclusion $C_*(A + B) \hookrightarrow C_*(A \cup B)$ descend to the quotient by $C_*(B)$. Another argument using long exact sequences and the five-lemma shows that Φ induces isomorphisms on homology if and only if the inclusion $C_*(A + B) \hookrightarrow C_*(A \cup B)$ does; the necessary argument is closely analogous to the proof that (i) \Leftrightarrow (ii), so we leave the details as an exercise. \square

34.5. Axiomatic Mayer-Vietoris sequences. The singular homology version of the Mayer-Vietoris sequence is a sufficient tool for all applications that we will encounter, nonetheless, we give here a quick sketch of how it can be generalized to the axiomatic setting.

Observe first that conditions (iii) and (iv) in Proposition 34.16 give a hint as to the motivation for the word “excisive”: if we write $X' := A \cup B$, $A' := B$ and $B' := B \setminus (A \cap B)$, then the inclusion $(A, A \cap B) \hookrightarrow (A \cup B, B)$ becomes $(X' \setminus B', A' \setminus B') \hookrightarrow (X', A')$ and condition (iii) thus becomes the statement that excision holds in singular homology for the triple (X', A', B') . Statements of this form make sense in any axiomatic homology theory.

DEFINITION 34.17. Given an axiomatic homology theory h_* and a space X with two subsets $A, B \subset X$, we will call (A, B) an **excisive couple** in X for the theory h_* if the obvious inclusions of pairs induce isomorphisms

$$h_*(A, A \cap B) \xrightarrow{\cong} h_*(A \cup B, B) \quad \text{and} \quad h_*(B, A \cap B) \xrightarrow{\cong} h_*(A \cup B, A).$$

EXAMPLE 34.18. The excision axiom for h_* implies that (A, B) will be an excisive couple in X whenever there exists a continuous function $u : A \cup B \rightarrow I$ that equals 0 outside of B and 1 outside of A . Urysohn's lemma guarantees this on well-behaved spaces X that are covered by the interiors of A and B .

If $X = A \cup B$ and (A, B) is an excisive couple for h_* , then using the inclusion map $(X, \emptyset) \hookrightarrow (X, B)$ and the connecting homomorphism $h_n(A, A \cap B) \rightarrow h_{n-1}(A \cap B)$ from the long exact sequence of the pair $(A, A \cap B)$, the diagram

$$(34.8) \quad \begin{array}{ccc} h_n(X) & \xrightarrow{\partial_*^A} & h_{n-1}(A \cap B) \\ \downarrow & & \uparrow \\ h_n(X, B) & \xleftarrow{\cong} & h_n(A, A \cap B) \end{array}$$

determines a homomorphism $\partial_* : h_n(X) \rightarrow h_{n-1}(A \cap B)$ for each $n \in \mathbb{Z}$. This requires only the first of the two isomorphisms in Definition 34.17; if one instead uses the second isomorphism and interchanges the roles of A and B , one analogously obtains another homomorphism $h_n(X) \rightarrow h_{n-1}(A \cap B)$, which turns out to differ from the one defined above only by a sign. This is the content of Lemma 34.20 below, and its proof makes use of the following alternative characterization of excisive couples:

PROPOSITION 34.19. For two subsets $A, B \subset X$, (A, B) is an excisive couple for h_* if and only if the inclusions of pairs

$$(A, A \cap B) \xleftarrow{k^A} (A \cup B, A \cap B) \xleftarrow{k^B} (B, A \cap B)$$

give rise to isomorphisms

$$h_n(A, A \cap B) \oplus h_n(B, A \cap B) \xrightarrow[\cong]{k_*^A \oplus k_*^B} h_n(A \cup B, A \cap B)$$

for all n .


PROOF. See Exercise 34.9 or [ES52, P. 34, Theorem 14.2]. □

LEMMA 34.20. If $X = A \cup B$ with (A, B) an excisive couple for the homology theory h_* , then for each $n \in \mathbb{Z}$, the two maps $\partial_*^A, \partial_*^B : h_n(X) \rightarrow h_{n-1}(A \cap B)$ defined in terms of natural inclusions and connecting homomorphisms from long exact sequences of pairs via the diagram

$$\begin{array}{ccccc} h_{n-1}(A \cap B) & \xleftarrow{\partial_*^A} & h_n(X) & \xrightarrow{\partial_*^B} & h_{n-1}(A \cap B) \\ \uparrow & & \swarrow & & \uparrow \\ h_n(B, A \cap B) & \xrightarrow{\cong} & h_n(X, A) & \xleftarrow{\cong} & h_n(X, B) \xleftarrow{\cong} h_n(A, A \cap B) \end{array}$$

satisfy $\partial_*^A + \partial_*^B = 0$.

PROOF. The maps ∂_*^A and ∂_*^B are obtained by following the highest and lowest rightward paths respectively from $h_n(X)$ to $h_{n-1}(A \cap B)$ in the diagram



where the maps ∂_*^a , ∂_*^b and ∂_* are all connecting homomorphisms from long exact sequences of pairs, and all other maps in the diagram are induced by inclusions of pairs of spaces. Note that while this diagram commutes, it does not imply that all paths from $h_n(X)$ to $h_{n-1}(A \cap B)$ give the same map, and in particular it does not imply $\partial_*^A = \partial_*^B$; the latter would follow if f_*^A and f_*^B were invertible, but in general they are not. By Proposition 34.19, however, each $x \in h_n(X)$ determines unique elements $a \in h_n(A, A \cap B)$ and $b \in h_n(B, A \cap B)$ such that

$$k_*^A a + k_*^B b = \varphi_* x \in h_n(X, A \cap B).$$

The compositions $f_*^A k_*^A$ and $f_*^B k_*^B$ each come from long exact sequences of triples, and are therefore trivial, thus we have

$$\varphi_*^A x = f_*^A \varphi_* x = f_*^A (k_*^A a + k_*^B b) = f_*^A k_*^B b \in h_n(B, A \cap B),$$

implying $\partial_*^A x = \partial_*^B b$, and a similar argument shows $\partial_*^B x = \partial_*^A a$. Since the horizontal maps in the middle of the diagram belong to the long exact sequence of the pair $(X, A \cap B)$, it follows that

$$0 = \partial_* \varphi_* x = \partial_* (k_*^A a + k_*^B b) = \partial_*^A a + \partial_*^B b = (\partial_*^A + \partial_*^B) x.$$

□

Since the two homomorphisms $\partial_*^A, \partial_*^B : h_n(X) \rightarrow h_{n-1}(A \cap B)$ described above only differ by a sign, it does not matter which one we use for the purposes of exactness in the theorem below. We chose in (34.8) to define $\partial_* := \partial_*^B$, because this convention is consistent with the version of the Mayer-Vietoris sequence derived for singular homology in §34.4, but this is not really an important detail.

THEOREM 34.21. *For any axiomatic homology theory h_* and any space $X = A \cup B$ that is the union of two subsets forming an excisive couple (A, B) for h_* , the sequence*

$$\dots \longrightarrow h_{n+1}(X) \xrightarrow{\partial_*} h_n(A \cap B) \xrightarrow{\begin{pmatrix} i_*^A & -i_*^B \end{pmatrix}} h_n(A) \oplus h_n(B) \xrightarrow{j_*^A \oplus j_*^B} h_n(X) \xrightarrow{\partial_*} h_{n-1}(A \cap B) \longrightarrow \dots$$

is exact, and it is natural with respect to maps $f : X \rightarrow X'$ such that $X' = A' \cup B'$ with (A', B') an excisive couple and $f(A) \subset A'$ and $f(B) \subset B'$.

The proof of the theorem is a diagram chase, which we will mostly leave as an exercise, but here is a useful hint. All terms appearing in the sequence of Theorem 34.21 also appear in the

where we denote the inclusions of pairs

$$\begin{array}{ccccc}
 & & (A, C) & & \\
 & \nearrow^{i^A} & & \searrow_{j^A} & \\
 (A \cap B, C \cap D) & & & & (X, Y) \\
 & \searrow_{i^B} & & \nearrow^{j^B} & \\
 & & (B, D) & &
 \end{array}$$

See Exercise 34.8 for a proof of the exactness of this sequence of singular homology.

34.7. Exercises.

EXERCISE 34.1. Deduce via the following steps that the sequence (34.3) appearing as the fourth strand in the braid diagram (34.2) is exact:

- 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.
- Prove that $j \circ i = 0$ by factoring it through the group $h_*(A, A)$, which is always zero. (Why?)
- Use a purely algebraic diagram-chasing argument to prove that the kernel of each map in the sequence (34.3) is contained in the image of the previous one.

EXERCISE 34.2. Show that for the pair (X, A) in Example 34.6, $H_1(X/A; \mathbb{Z}) \not\cong H_1(X, A; \mathbb{Z})$.
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 $\prod_{n \in \mathbb{N}} \mathbb{Z}$.

EXERCISE 34.3 (*). Assume $X = A \cup B$ is any space with subsets (A, B) forming an excisive couple for singular homology.

- Prove that the definition of the homomorphism

$$H_n(X) \xrightarrow{\partial_*} H_{n-1}(A \cap B) : [a + b] \mapsto [\partial a] = -[\partial b], \quad \text{where } a \in C_n(A) \text{ and } b \in C_n(B)$$

is independent of choices.

- Prove that the map ∂_* in part (a) is in fact the connecting homomorphism in the long exact sequence that arises by plugging the short exact sequence (34.7) into Proposition 32.13.

EXERCISE 34.4. Derive from the Mayer-Vietoris sequence a simple proof that there is an isomorphism $\tilde{h}_n(X) \cong \tilde{h}_{n+1}(\Sigma X)$ for every axiomatic homology theory h_* , every $n \in \mathbb{Z}$ and every space X , where ΣX denotes the suspension of X .

EXERCISE 34.5. Use Mayer-Vietoris sequences to compute $H_*(\mathbb{T}^2; \mathbb{Z})$ by decomposing the torus \mathbb{T}^2 into a union of open subsets each homotopy equivalent to S^1 .

Hint: There is a useful algebraic trick for turning any long exact sequence

$$\dots \longrightarrow A \xrightarrow{\alpha} B \xrightarrow{\beta} C \xrightarrow{\gamma} D \xrightarrow{\delta} E \longrightarrow \dots$$

into a short exact sequence with a specific term in the middle, e.g.

$$0 \longrightarrow \operatorname{coker}(\alpha) \xrightarrow{\beta} C \xrightarrow{\gamma} \ker(\delta) \longrightarrow 0,$$

where $\text{coker}(\alpha) := B/\text{im}(\alpha)$. If this short exact sequence splits, one obtains from it a formula for C .

EXERCISE 34.6. Use Mayer-Vietoris sequences to compute $H_*(X; \mathbb{Z})$ and $H_*(X; \mathbb{Z}_2)$, where X is

- (a) The projective plane $\mathbb{R}P^2$.
- (b) The Klein bottle.

Hint: $\mathbb{R}P^2$ 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 $\mathbb{R}P^2 \# \mathbb{R}P^2$.

EXERCISE 34.7. 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{D}^n from each of X and Y and then gluing $X \setminus \mathring{D}^n$ and $Y \setminus \mathring{D}^n$ together along an identification of their boundary spheres.

- (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 $H_k(X \setminus \mathring{D}^n)$, and then see what happens when you glue $X \setminus \mathring{D}^n$ and $Y \setminus \mathring{D}^n$ 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 $n \geq 2$, and without orientability we still have $H_{n-1}(X \# Y; \mathbb{Z}_2) \cong H_{n-1}(X; \mathbb{Z}_2) \oplus H_{n-1}(Y; \mathbb{Z}_2)$. (One can deduce both results from the properties of fundamental classes in singular homology, which we will discuss later.) 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.

EXERCISE 34.8. Here is a way to derive the relative Mayer-Vietoris sequence in singular homology. Assume (X, Y) , (A, C) and (B, D) are pairs of spaces such that $X = A \cup B$ and $Y = C \cup D$. Defining the quotient chain complex

$$C_*(A + B, C + D) := C_*(A + B) / C_*(C + D) = \frac{C_*(A) + C_*(B)}{C_*(C) + C_*(D)},$$

we notice that the inclusion $C_*(A + B) \hookrightarrow C_*(X)$ descends to a chain map

$$C_*(A + B, C + D) \rightarrow C_*(X, Y).$$


- (a) Show that if (A, B) is an excisive couple in X and (C, D) is an excisive couple in Y , then the chain map $C_*(A + B, C + D) \rightarrow C_*(X, Y)$ induces an isomorphism on homology.

Hint: Five-lemma!

- (b) Under the same assumptions as in part (a), derive the exact sequence (34.9) in singular homology $h_* := H_*$ from a short exact sequence of chain complexes $0 \rightarrow C_*(A \cap B, C \cap D) \rightarrow C_*(A, C) \oplus C_*(B, D) \rightarrow C_*(A + B, C + D) \rightarrow 0$.

EXERCISE 34.9. Prove Proposition 34.19. You may find the following diagram helpful, in which both diagonals are long exact sequences of triples, and the two vertical arrows are the maps that

are required to be isomorphisms in the definition of an excisive couple:



EXERCISE 34.10. Fill in the details of the proof of Theorem 34.21 on the exactness of the axiomatic Mayer-Vietoris sequence.

35. Mapping tori and maps between spheres

The two topics of this 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.

35.1. 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 $f : X \rightarrow X$, the **mapping torus** (*Abbildungstorus*) of f is defined to be the quotient space

$$X_f := (X \times I) / \sim, \quad \text{where } (x, 0) \sim (f(x), 1) \text{ for all } x \in X.$$

We can regard X itself as a subspace of X_f via the inclusion map⁵³

$$i : X \hookrightarrow X_f : x \mapsto [(x, 1)].$$

THEOREM 35.1. *For any map $f : X \rightarrow X$ and its mapping torus X_f , every axiomatic homology theory h_* admits a long exact sequence*

$$\dots \longrightarrow h_{k+1}(X_f) \longrightarrow h_k(X) \xrightarrow{1-f_*} h_k(X) \xrightarrow{i_*} h_k(X_f) \longrightarrow h_{k-1}(X) \longrightarrow \dots$$

Let's do an example before talking about the proof.

EXAMPLE 35.2. For each $n \in \mathbb{N}$, the n -torus $\mathbb{T}^n = S^1 \times \dots \times S^1$ is the mapping torus of the identity map $\text{Id} : \mathbb{T}^{n-1} \rightarrow \mathbb{T}^{n-1}$, so the exact sequence of the mapping torus includes segments of the form

$$\dots \longrightarrow h_k(\mathbb{T}^{n-1}) \xrightarrow{0} h_k(\mathbb{T}^{n-1}) \xrightarrow{i_*} h_k(\mathbb{T}^n) \xrightarrow{\Phi} h_{k-1}(\mathbb{T}^{n-1}) \xrightarrow{0} h_{k-1}(\mathbb{T}^{n-1}) \longrightarrow \dots$$

⁵³Of course there are also natural maps $X \rightarrow X_f : x \mapsto [(x, t)]$ 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 $f(x) = f(y)$.

The triviality of the two maps $\mathbb{1} - \text{Id}_* = 0$ here means that we actually have a short exact sequence

$$(35.1) \quad 0 \longrightarrow h_k(\mathbb{T}^{n-1}) \xrightarrow{i_*} h_k(\mathbb{T}^n) \xrightarrow{\Phi} h_{k-1}(\mathbb{T}^{n-1}) \longrightarrow 0.$$

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

$$H_k(\mathbb{T}^2; \mathbb{Z}) \cong H_k(S^1; \mathbb{Z}) \oplus H_{k-1}(S^1; \mathbb{Z})$$

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 (35.1) 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 $r \geq 0$, the **rank** (*Rang*) of the group, and these ranks satisfy $\text{rank } H_k(\mathbb{T}^n; \mathbb{Z}) = \text{rank } H_k(\mathbb{T}^{n-1}; \mathbb{Z}) + \text{rank } H_{k-1}(\mathbb{T}^{n-1}; \mathbb{Z})$, so they are precisely the numbers in *Pascal's triangle*, i.e. the familiar binomial coefficients:

$$\text{rank } H_k(\mathbb{T}^n; \mathbb{Z}) = \binom{n}{k} \quad \text{for } 0 \leq k \leq n, \quad H_k(\mathbb{T}^n; \mathbb{Z}) = 0 \text{ for } k > n.$$

Theorem 35.1 gives rise to a similarly straightforward computation of the homology of the Klein bottle; see Exercise 35.1.

To prove the theorem, we shall first state a more general result that implies it. Given two spaces X, Y and maps $f, g : X \rightarrow Y$, define the space

$$Z := ((X \times I) \amalg Y) / \sim \quad \text{where } (x, 0) \sim f(x) \text{ and } (x, 1) \sim g(x) \text{ for all } x \in X.$$

This space comes with a natural inclusion

$$i : Y \hookrightarrow Z,$$

and the special case with $X = Y$ and $g = \text{Id}$ reproduces the mapping torus X_f of $f : X \rightarrow X$. Theorem 35.1 follows immediately from the next statement:

THEOREM 35.3. *Given $f, g : X \rightarrow Y$ and the space Z described above, there exists a long exact sequence*

$$\dots \longrightarrow h_{k+1}(Z) \longrightarrow h_k(X) \xrightarrow{g_* - f_*} h_k(Y) \xrightarrow{i_*} h_k(Z) \longrightarrow h_{k-1}(X) \longrightarrow \dots$$

for every axiomatic homology theory h_* .

REMARK 35.4. 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 $j : M \rightarrow X$. This gives rise to a map $\tilde{j} : M \times I \rightarrow X \times I : (x, t) \mapsto (j(x), t)$, 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{j} 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 $i_* : H_n(Y; \mathbb{Z}) \rightarrow H_n(Z; \mathbb{Z})$.

PROOF OF THEOREM 35.1. Consider the map of pairs $q : (X \times I, X \times \partial I) \rightarrow (Z, Y)$ defined as the composition of the two maps

$$(X \times I, X \times \partial I) \hookrightarrow ((X \times I) \amalg Y, X \times \partial I) \rightarrow (Z, Y),$$

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

$$(35.2) \quad \begin{array}{ccccccc} \dots & \longrightarrow & h_{k+1}(X \times I) & \xrightarrow{\beta_*} & h_{k+1}(X \times I, X \times \partial I) & \xrightarrow{\partial_*} & h_k(X \times \partial I) & \xrightarrow{\alpha_*} & h_k(X \times I) & \xrightarrow{\beta_*} & \dots \\ & & \downarrow q_* & & \downarrow q_* & & \downarrow q_* & & \downarrow q_* & & \\ \dots & \longrightarrow & h_{k+1}(Z) & \longrightarrow & h_{k+1}(Z, Y) & \xrightarrow{\partial_*^-} & h_k(Y) & \xrightarrow{i_*} & h_k(Z) & \longrightarrow & \dots \end{array}$$

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 $\beta : (X \times I, \emptyset) \hookrightarrow (X \times I, X \times \partial I)$ are the obvious inclusions. Since $X \times \partial I = X \times \{0, 1\} \cong X \amalg X$, the additivity axiom gives an isomorphism

$$(35.3) \quad j_*^0 \oplus j_*^1 : h_k(X) \oplus h_k(X) \xrightarrow{\cong} h_k(X \times \partial I),$$

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 $(\alpha \circ j^i)_* : h_k(X) \rightarrow h_k(X \times I)$ for $i = 0, 1$ are both the same isomorphism. It follows that

$$\alpha_* \circ (j_*^0 \oplus j_*^1) = (\alpha_* \circ j_*^0) \oplus (\alpha_* \circ j_*^1) : h_k(X) \oplus h_k(X) \rightarrow h_k(X \times I)$$

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

$$(35.4) \quad \Psi : h_k(X) \xrightarrow{\cong} \ker \alpha_* : c \mapsto (j_*^0 \oplus j_*^1)(-c, c) = j_*^1 c - j_*^0 c.$$

Exactness of the top row now implies $\beta_* = 0$, and the connecting homomorphism $\partial_* : H_k(X \times I, X \times \partial I) \rightarrow H_{k-1}(X \times \partial I)$ is thus injective. This makes ∂_* an isomorphism onto its image, which is $\ker \alpha_*$.

Now observe that for the map $q : X \times \partial I \rightarrow Y$, the compositions $q \circ j^i : X \rightarrow Y$ for $i = 0, 1$ are the maps f and g respectively, thus we have

$$(35.5) \quad q_* \circ \Psi : h_k(X) \rightarrow h_k(Y) : c \mapsto g_* c - f_* c = (g_* - f_*)c.$$

On the other hand, the map

$$(X \times I)/(X \times \partial I) \xrightarrow{q} Z/Y$$

determined by $q : (X \times I, X \times \partial I) \rightarrow (Z, Y)$ is a homeomorphism and thus induces an isomorphism

$$q_* : \tilde{h}_*((X \times I)/(X \times \partial I)) \xrightarrow{\cong} \tilde{h}_*(Z/Y),$$

and since both pairs are *good* in the sense of Definition 34.3, Theorem 34.7 implies that the map $q_* : h_*(X \times I, X \times \partial I) \rightarrow h_*(Z, Y)$ is also an isomorphism. We can put all of this information together to produce a commutative diagram

$$\begin{array}{ccccc} h_{k+1}(Z, Y) & \xrightarrow{q_*^{-1}} & h_{k+1}(X \times I, X \times \partial I) & \xrightarrow{\partial_*} & \ker \alpha_* & \xrightarrow{\Psi^{-1}} & h_k(X) \\ & & & & \downarrow q_* & \swarrow g_* - f_* & \\ & & & & h_k(Y) & & \end{array}$$

in which all the horizontal maps on the top row are isomorphisms. The composition of these maps therefore gives an isomorphism $h_{k+1}(Z, Y) \rightarrow h_k(X)$ that we can use to replace $h_{k+1}(Z, Y)$ by $h_k(X)$ in the bottom row of (35.2); the original connecting homomorphism $\partial_*^- : h_{k+1}(Z, Y) \rightarrow$

$h_k(Y)$ then gets replaced by the map $g_* - f_* : h_k(X) \rightarrow h_k(Y)$, producing an exact sequence as in the statement of the theorem. \square

One last comment about mapping tori: the usefulness of the exact sequence in Theorem 35.1 depends heavily on how easy it is to compute the homomorphism $f_* : h_*(X) \rightarrow h_*(X)$. This is not always easy, but sometimes it is, particularly in cases where $h_*(X)$ is relatively simple.

35.2. Degrees of maps between spheres. For the second topic today, let's talk about the homomorphism $f_* : H_n(S^n) \rightarrow H_n(S^n)$ induced by a map $f : S^n \rightarrow S^n$. For $n \geq 1$, 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 $\mathbb{Z} \rightarrow \mathbb{Z}$ or $\mathbb{Z}_2 \rightarrow \mathbb{Z}_2$ admit very simple characterizations. In order to accommodate the case $n = 0$ under the same umbrella, we can use reduced homology since $\tilde{H}_0(S^0; \mathbb{Z}) \cong \mathbb{Z}$ and $\tilde{H}_0(S^0; \mathbb{Z}_2) \cong \mathbb{Z}_2$.

DEFINITION 35.5. The **mapping degree** (*Abbildungsgrad*)

$$\deg(f) \in \mathbb{Z}$$

of a map $f : S^n \rightarrow S^n$ is the unique integer k such that the homomorphism $f_* : \tilde{H}_n(S^n; \mathbb{Z}) \rightarrow \tilde{H}_n(S^n; \mathbb{Z})$ 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 $f_* : \tilde{H}_n(S^n; \mathbb{Z}_2) \rightarrow \tilde{H}_n(S^n; \mathbb{Z}_2)$ is the map $c \mapsto kc$.

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 $f : M \rightarrow M$ on a closed connected manifold M that need not be orientable, and we will touch upon this in the next lecture.

PROPOSITION 35.6. *The integer-valued degree for maps $S^n \rightarrow S^n$ has the following properties.*

- (1) *If $f, g : S^n \rightarrow S^n$ are homotopic, then $\deg(f) = \deg(g)$.*
- (2) *For any $f, g : S^n \rightarrow S^n$, $\deg(f \circ g) = \deg(f) \cdot \deg(g)$.*
- (3) *The identity map $S^n \rightarrow S^n$ has $\deg(\text{Id}) = 1$.*
- (4) *If f is constant, then $\deg(f) = 0$.*
- (5) *The degree of any map $f : S^1 \rightarrow S^1$ matches its **winding number** (*Windungszahl*), i.e. it is the unique $k \in \mathbb{Z}$ such that any continuous function $\theta : [0, 1] \rightarrow \mathbb{R}$ with $f(e^{2\pi i t}) = e^{2\pi i \theta(t)}$ satisfies $\theta(1) - \theta(0) = k$.*

PROOF. The first three properties are immediate from the homotopy invariance of $\tilde{H}_*(\cdot; \mathbb{Z})$ and the fact that it is a functor. For the fourth, observe that any constant map $f : S^n \rightarrow S^n$ can be factored as $i \circ \epsilon$ for the unique map $\epsilon : S^n \rightarrow \{*\}$ and a suitable inclusion $i : \{*\} \hookrightarrow S^n$, thus $f_* : \tilde{H}_n(S^n; \mathbb{Z}) \rightarrow \tilde{H}_n(S^n; \mathbb{Z})$ factors through $\tilde{H}_n(\{*\}; \mathbb{Z}) = 0$. Finally, the fifth property follows from standard facts about $\pi_1(S^1)$ and the natural isomorphism $\tilde{H}_1(S^1; \mathbb{Z}) = H_1(S^1; \mathbb{Z}) \cong \pi_1(S^1)$. \square

Recall that the suspension $\Sigma X = C_+ X \cup_X C_- X$ of a space X can be regarded as a functor $\text{Top} \rightarrow \text{Top}$ sending objects X to ΣX , where maps $f : X \rightarrow Y$ are transformed to maps

$$\Sigma f : \Sigma X \rightarrow \Sigma Y : [(x, t)] \mapsto [(f(x), t)].$$

In particular, any map $f : S^n \rightarrow S^n$ gives rise to a map $\Sigma f : S^{n+1} \rightarrow S^{n+1}$ using the identification $\Sigma S^n \cong S^{n+1}$.

PROPOSITION 35.7. *For any $f : S^n \rightarrow S^n$, $\deg(f) = \deg(\Sigma f)$.*

PROOF. A useful detail you may have learned if you worked out Exercise 34.4 is that the isomorphism $\tilde{H}_{n+1}(\Sigma X; \mathbb{Z}) \rightarrow \tilde{H}_n(X; \mathbb{Z})$ can always be constructed as the connecting homomorphism in a Mayer-Vietoris exact sequence for ΣX . Given a map $f : S^n \rightarrow S^n$, the naturality of this connecting homomorphism produces a commuting diagram

$$\begin{array}{ccc} \tilde{H}_{n+1}(S^{n+1}; \mathbb{Z}) & \xrightarrow{\partial_*} & \tilde{H}_n(S^n; \mathbb{Z}) \\ \downarrow (\Sigma f)_* & & \downarrow f_* \\ \tilde{H}_{n+1}(S^{n+1}; \mathbb{Z}) & \xrightarrow{\partial_*} & \tilde{H}_n(S^n; \mathbb{Z}) \end{array}$$

where the two maps labeled ∂_* are the same isomorphism. Now if $(\Sigma f)_*c = kc$ for some nontrivial $c \in \tilde{H}_{n+1}(S^{n+1}; \mathbb{Z})$, it follows that $\partial_*(\Sigma f)_*c = k\partial_*c = f_*\partial_*c$, where $\partial_*c \in \tilde{H}_n(S^n; \mathbb{Z})$ is also nontrivial, hence $\deg(\Sigma f) = k = \deg(f)$. \square

PROPOSITION 35.8. *If $f : S^n \rightarrow S^n$ is the restriction to the unit sphere $S^n \subset \mathbb{R}^{n+1}$ of an orthogonal linear transformation $\mathbf{A} \in O(n+1)$, then $\deg(f) = \det(\mathbf{A}) = \pm 1$.*

PROOF. Recall that $O(n+1)$ has exactly two path-components, which can be labeled according to whether their elements have determinant $+1$ or -1 . A given $\mathbf{A} \in O(n+1)$ thus admits a continuous path in $O(n+1)$ to the identity matrix $\mathbf{1}$ 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 $f : S^n \rightarrow S^n$ 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 $S^1 \rightarrow S^1$ with winding number -1 , so the claim follows from Proposition 35.6. Now if we assume the claim is true for reflections $f : S^n \rightarrow S^n$, it suffices to observe that $\Sigma f : \Sigma S^n \rightarrow \Sigma S^n$ is also a reflection under a suitable identification $\Sigma S^n \cong S^{n+1}$, so by induction, the result follows from Proposition 35.7. \square

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 35.9. *Every map $f : S^n \rightarrow S^n$ with $\deg(f) \neq (-1)^{n+1}$ has a fixed point.*

PROOF. It is easy to think of a specific map $f : S^n \rightarrow S^n$ that has no fixed point: the antipodal map $x \mapsto -x$ has this property, and its degree according to Proposition 35.8 is $(-1)^{n+1}$. The theorem will follow from the claim that, in fact, every map $f : S^n \rightarrow S^n$ with no fixed point is homotopic to the antipodal map, and therefore also has degree $(-1)^{n+1}$.

Indeed, if $f : S^n \rightarrow S^n$ 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 $g_t : S^n \rightarrow \mathbb{R}^{n+1} \setminus \{0\}$ with $g_0 = f$ and $g_1(x) = -x$. Since $g_t(x)$ is never zero, we can then define a homotopy $h : I \times S^n \rightarrow S^n$ in S^n from

f to the antipodal map g_1 by

$$h(t, x) = \frac{g_t(x)}{|g_t(x)|}.$$

□

THEOREM 35.10 (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 $V : S^n \rightarrow \mathbb{R}^{n+1} \setminus \{0\}$ 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 $f_0 = \text{Id}$ 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. □

35.3. Exercises.

EXERCISE 35.1. The mapping torus of $f : S^1 \rightarrow S^1 : e^{i\theta} \mapsto e^{-i\theta}$ is homeomorphic to the Klein bottle K^2 . Use Theorem 35.1 to compute $H_*(K^2; \mathbb{Z})$ and $H_*(K^2; \mathbb{Z}_2)$.

EXERCISE 35.2. The goal of this exercise is to gain a more concrete picture of the connecting homomorphism $\Phi : H_1(X_f; \mathbb{Z}) \rightarrow H_0(X; \mathbb{Z})$ that appears in the long exact sequence of the mapping torus of a *homeomorphism* $f : X \rightarrow X$,

$$\dots \longrightarrow H_{k+1}(X_f; \mathbb{Z}) \xrightarrow{\Phi} H_k(X; \mathbb{Z}) \xrightarrow{1_* - f_*} H_k(X; \mathbb{Z}) \xrightarrow{i_*} H_k(X_f; \mathbb{Z}) \xrightarrow{\Phi} H_{k-1}(X; \mathbb{Z}) \longrightarrow \dots$$

in singular homology with integer coefficients. It will be useful to observe first that if $f : X \rightarrow X$ is a homeomorphism, then its mapping torus admits an alternative description as the quotient

$$X_f = (X \times \mathbb{R}) / (x, t) \sim (f(x), t + 1),$$

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 $\tilde{X} := X \times \mathbb{R}$ as a covering space for X_f , with the quotient projection defining a covering map $\tilde{X} \rightarrow X_f$ of infinite degree. Writing $S^1 := \mathbb{R}/\mathbb{Z}$, we also see a natural continuous surjective map $\pi : X_f \rightarrow S^1 : [(x, t)] \mapsto [t]$, 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 $\gamma : [0, 1] \rightarrow X_f$ with $\gamma(0) = \gamma(1) = x$. Now let $\tilde{\gamma} : [0, 1] \rightarrow \tilde{X}$ denote the unique lift of γ to the cover $\tilde{X} = X \times \mathbb{R}$ 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 $\Phi : H_1(X_f; \mathbb{Z}) \rightarrow \mathbb{Z}$ can be chosen⁵⁴ such that

$$\Phi([\gamma]) = m,$$

so in particular, $[\gamma] \in \ker \Phi$ if and only if the lift of γ to the cover \tilde{X} is a loop.

- (b) Prove directly from the characterization in part (a) that $\Phi : H_1(X_f; \mathbb{Z}) \rightarrow H_0(X; \mathbb{Z})$ is surjective.

Remark: Of course this can also be deduced less directly from the exact sequence.

EXERCISE 35.3 (*). There are only four possible maps $f : S^0 \rightarrow S^0$. What are their degrees?

EXERCISE 35.4. Fix $n \in \mathbb{N}$.

- (a) Prove that if n is even, every continuous map $f : S^n \rightarrow S^n$ has at least one point $x \in S^n$ where either $f(x) = x$ or $f(x) = -x$. Deduce that every continuous map $\mathbb{R}P^n \rightarrow \mathbb{R}P^n$ has a fixed point if n is even.
- (b) Construct counterexamples to the statement in part (a) for every odd n .

Hint: Consider linear transformations with no real eigenvalues.

36. Local and global mapping degree

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 $f : X \rightarrow Y$ 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 $f : \mathbb{R} \rightarrow \mathbb{R}$ 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 $f : X \rightarrow Y$ 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 with later in the semester, when we construct fundamental classes without triangulations.

Local orientations and degree. We begin by defining a “local” version of homology that is interesting for manifolds in particular, but not necessarily for more general spaces.

DEFINITION 36.1. Suppose M is a topological manifold of dimension $n \in \mathbb{N}$. For a homology theory h_* and any interior point $x \in M \setminus \partial M$, the **local homology** of M at x is the group $h_n(M, M \setminus \{x\})$.

It turns out that local homology groups are always isomorphic to the coefficient group $G = h_0(\{*\})$, though the isomorphism itself may depend on various choices. To see this, recall that every interior point x in an n -manifold M admits a so-called **Euclidean neighborhood**; we shall

⁵⁴There 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 $\text{im } \Phi$ would not change.

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

$$\varphi : \mathcal{U} \xrightarrow{\cong} \mathbb{R}^n \text{ or } \mathbb{D}^n, \quad \text{such that} \quad \varphi(x) = 0,$$

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. A Euclidean neighborhood $\mathcal{U} \subset M$ of x with coordinate chart $\varphi : \mathcal{U} \rightarrow \mathbb{D}^n$ determines a string of isomorphisms

$$(36.1) \quad h_n(M, M \setminus \{x\}) \leftarrow h_n(\mathcal{U}, \mathcal{U} \setminus \{x\}) \xrightarrow{\varphi_*} h_n(\mathbb{D}^n, \mathbb{D}^n \setminus \{0\}) \leftarrow h_n(\mathbb{D}^n, S^{n-1}) \xrightarrow{\partial_*} \tilde{h}_{n-1}(S^{n-1}) \cong G,$$

in which every arrow without a label is induced by a continuous inclusion map. Here, the first map is an isomorphism due to excision, the connecting homomorphism ∂_* is an isomorphism due to the long exact sequence of the pair (\mathbb{D}^n, S^{n-1}) in reduced homology, and the map induced by the inclusion $(\mathbb{D}^n, S^{n-1}) \hookrightarrow (\mathbb{D}^n, \mathbb{D}^n \setminus \{0\})$ is an isomorphism due to a combination of the homotopy and exactness axioms with the five-lemma: indeed, since $j : S^{n-1} \hookrightarrow \mathbb{D}^n \setminus \{0\}$ is a homotopy equivalence, one can apply the five-lemma to the diagram

$$\begin{array}{ccccccccc} \dots & \longrightarrow & h_n(S^{n-1}) & \longrightarrow & h_n(\mathbb{D}^n) & \longrightarrow & h_n(\mathbb{D}^n, S^{n-1}) & \longrightarrow & h_{n-1}(S^{n-1}) & \longrightarrow & h_{n-1}(\mathbb{D}^n) & \longrightarrow & \dots \\ & & \cong \downarrow j_* & & \cong \downarrow \mathbb{1} & & \downarrow j_* & & \cong \downarrow j_* & & \cong \downarrow \mathbb{1} & & \\ \dots & \longrightarrow & h_n(\mathbb{D}^n \setminus \{0\}) & \longrightarrow & h_n(\mathbb{D}^n) & \longrightarrow & h_n(\mathbb{D}^n, \mathbb{D}^n \setminus \{0\}) & \longrightarrow & h_{n-1}(\mathbb{D}^n \setminus \{0\}) & \longrightarrow & h_{n-1}(\mathbb{D}^n) & \longrightarrow & \dots \end{array}$$

If preferred, the same trick can be done with an open Euclidean neighborhood and chart $\varphi : \mathcal{U} \rightarrow \mathbb{R}^n$ after replacing the middle term of (36.1) with $h_n(\mathbb{R}^n, \mathbb{R}^n \setminus \{0\})$.

Specializing 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 36.2. A **local orientation** of an n -manifold M at an interior point $x \in M \setminus \partial M$ is a choice of generator $[M]_x$ for the group $H_n(M, M \setminus \{x\}; \mathbb{Z}) \cong \mathbb{Z}$.

Note that in light of the excision isomorphism

$$H_n(\mathcal{U}, \mathcal{U} \setminus \{x\}; \mathbb{Z}) \xrightarrow{\cong} H_n(M, M \setminus \{x\}; \mathbb{Z})$$

defined for any Euclidean neighborhood $\mathcal{U} \subset M$ of x , a local orientation can equivalently be regarded as a generator of $H_n(\mathcal{U}, \mathcal{U} \setminus \{x\}; \mathbb{Z}) \cong \mathbb{Z}$.

EXAMPLE 36.3. If M is a surface without boundary and $x \in M$ is a point, then a specific relative 2-cycle generating $H_2(M, M \setminus \{x\}; \mathbb{Z})$ can be defined via a single singular 2-simplex $\sigma : \Delta^2 \rightarrow M$ that embeds the triangle Δ^2 onto a neighborhood of x . Indeed, $\sigma \in C_2(M)$ is clearly a relative cycle in $(M, M \setminus \{x\})$ since σ maps $\partial \Delta^2$ to $M \setminus \{x\}$, and to see that it generates $H_2(M, M \setminus \{x\}; \mathbb{Z})$, one can follow the string of isomorphisms (36.1): 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 $\pi_1(\partial \Delta^2) \cong H_1(\partial \Delta^2; \mathbb{Z})$. 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 36.4. Suppose M and N are manifolds of dimension $n \in \mathbb{N}$, $f : M \rightarrow N$ is a map, and $x \in M \setminus \partial M$ and $y = f(x) \in N \setminus \partial N$ 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 $\mathcal{U} \subset M \setminus \partial M$ is a Euclidean neighborhood. Given local orientations $[M]_x \in H_n(\mathcal{U}, \mathcal{U} \setminus \{x\}; \mathbb{Z})$ and $[N]_y \in H_n(N, N \setminus \{y\}; \mathbb{Z})$, 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 $H_n(\mathcal{U}, \mathcal{U} \setminus \{x\}; \mathbb{Z}) \rightarrow H_n(N, N \setminus \{y\}; \mathbb{Z})$ induced by $f : (\mathcal{U}, \mathcal{U} \setminus \{x\}) \rightarrow (N, N \setminus \{y\})$ 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 $f_* : H_n(\mathcal{U}, \mathcal{U} \setminus \{x\}; \mathbb{Z}_2) \rightarrow H_n(N, N \setminus \{y\}; \mathbb{Z}_2)$ sends $[M]_x$ to $k[N]_y$, where $[M]_x$ and $[N]_y$ are now taken to be the unique nontrivial elements of $H_n(\mathcal{U}, \mathcal{U} \setminus \{x\}; \mathbb{Z}_2) \cong \mathbb{Z}_2$ and $H_n(N, N \setminus \{y\}; \mathbb{Z}_2) \cong \mathbb{Z}_2$ 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 (36.1), any choice of Euclidean neighborhood $\mathcal{U} \subset M$ of x gives rise to an isomorphism of $H_n(\mathcal{U}, \mathcal{U} \setminus \{x\}; \mathbb{Z})$ with $\tilde{H}_{n-1}(S^{n-1}; \mathbb{Z}) \cong \mathbb{Z}$. We can use this isomorphism to transform the definition above into a condition about maps between spheres:

PROPOSITION 36.5. *In the setting of Definition 36.4, fix a generator $[S^{n-1}] \in \tilde{H}_{n-1}(S^{n-1}; \mathbb{Z})$ and Euclidean neighborhoods $\mathcal{U} \subset M$ of x and $\mathcal{V} \subset N$ of y such that the resulting isomorphisms of $H_n(\mathcal{U}, \mathcal{U} \setminus \{x\}; \mathbb{Z})$ and $H_n(\mathcal{V}, \mathcal{V} \setminus \{y\}; \mathbb{Z})$ to $\tilde{H}_{n-1}(S^{n-1}; \mathbb{Z})$ 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*

$$\deg(f; x) = \deg \left(\left. \frac{\hat{f}}{|\hat{f}|} \right|_{\partial \mathbb{D}_\epsilon^n} : \partial \mathbb{D}_\epsilon^n \rightarrow S^{n-1} \right)$$

for all $\epsilon > 0$ sufficiently small, where \mathbb{D}_ϵ^n 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 $S^{n-1} \rightarrow S^{n-1}$. Similarly, $\deg_2(f; x)$ is related in the same way to the mod 2 degree of the same map $\partial \mathbb{D}_\epsilon^n \rightarrow S^{n-1}$. □

COROLLARY 36.6. *Suppose $\{f_t : M \rightarrow N\}_{t \in [0,1]}$ is a continuous family of maps between two manifolds of dimension $n \in \mathbb{N}$, with interior points $x \in M \setminus \partial M$ and $y \in N \setminus \partial N$ 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 36.5 as degrees of maps $S^{n-1} \rightarrow S^{n-1}$, and the assumption about the family f_t implies that these two maps between spheres are homotopic. □

EXAMPLE 36.7. Continuing the discussion of Example 36.3, suppose $f : M \rightarrow N$ is a map between surfaces such that x is an isolated point in $f^{-1}(y)$, and suppose we have fixed local orientations $[M]_x \in H_2(M, M \setminus \{x\}; \mathbb{Z})$ and $[N]_y \in H_2(N, N \setminus \{y\}; \mathbb{Z})$. 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 $\alpha : S^1 \hookrightarrow \mathcal{U} \setminus \{x\}$ winding once around x , so that $f \circ \alpha : S^1 \rightarrow \mathcal{V} \setminus \{y\}$ is also uniquely determined up to homotopy. The winding number of $f \circ \alpha$ 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 $S^1 \hookrightarrow \mathcal{V} \setminus \{y\}$ 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 $[\mathbb{R}^n]_x \in H_n(\mathbb{R}^n, \mathbb{R}^n \setminus \{x\}; \mathbb{Z})$ at every point $x \in \mathbb{R}^n$: if $\mathbb{D}_x^n \subset \mathbb{R}^n$ 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 (36.1) a string of natural isomorphisms

$$H_n(\mathbb{R}^n, \mathbb{R}^n \setminus \{x\}; \mathbb{Z}) \cong H_n(\mathbb{D}_x^n, \partial \mathbb{D}_x^n; \mathbb{Z}) \cong \tilde{H}_{n-1}(S^{n-1}; \mathbb{Z}),$$

so that any choice of generator $[S^{n-1}] \in \tilde{H}_{n-1}(S^{n-1}; \mathbb{Z})$ determines local orientations $[\mathbb{R}^n]_x \in H_n(\mathbb{R}^n, \mathbb{R}^n \setminus \{x\}; \mathbb{Z})$ for all $x \in \mathbb{R}^n$ simultaneously. With this choice in place, any continuous map $f : \mathcal{U} \rightarrow \mathbb{R}^n$ 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 36.8. *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 $f : \mathcal{U} \rightarrow \mathbb{R}^n$ is a map that is differentiable at a point $x \in \mathcal{U}$ such that its derivative $df(x) : \mathbb{R}^n \rightarrow \mathbb{R}^n$ 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 $f : \mathcal{U} \rightarrow \mathbb{R}^n$ 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 $\lim_{h \rightarrow 0} \eta(h) = 0$. If $df(x) : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is invertible, then there exists a constant $c > 0$ such that $|df(x)h| \geq c|h|$ for all $h \in \mathbb{R}^n$, so

$$|f(x+h) - y| = |df(x)h + |h|\eta(h)| \geq |df(x)h| - |h||\eta(h)| \geq (c - |\eta(h)|)|h|,$$

and the right hand side is positive for all $|h|$ sufficiently small since $\eta(h) \rightarrow 0$. 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 36.6, we can now assume without loss of generality that the remainder term vanishes completely, i.e. $f(x+h) = y + df(x)h$. Now observe that if we modify f by a further homotopy of the form

$$f_t(x+h) = y + A_t h,$$

where $A_t : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is a family of invertible linear transformations, then the local degree still will not change due to Corollary 36.6, thus we are free to assume without loss of generality that $df(x)$ is an *orthogonal* transformation. The corresponding map $S^{n-1} \rightarrow S^{n-1}$ is then of the type considered in Proposition 35.8, 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. \square

Many applications of the local degree are based on the following result, 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 . Its proof is Exercise 36.2.

PROPOSITION 36.9. *If $\mathcal{U} \subset \mathbb{R}^n$ is open and $f : \mathcal{U} \rightarrow \mathbb{R}^n$ 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 $\hat{f} : \mathcal{U} \rightarrow \mathbb{R}^n$ satisfying $|\hat{f} - f| < \epsilon$ maps some point in \mathcal{U}_x to y .* \square

Mapping degree for closed manifolds. We will now extend the global mapping degree previously defined for maps $f : S^n \rightarrow S^n$ to maps $f : M \rightarrow N$ between more general closed connected n -manifolds. Most of what follows can also be extended to maps $f : (M, \partial M) \rightarrow (N, \partial N)$ 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 $f : M \rightarrow N$ 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,⁵⁵ depending which coefficients we want to use—though it will be a while before we are in a position to fully prove this.

DEFINITION 36.10. Given an axiomatic homology theory h_* , a topological manifold M of dimension $n \in \mathbb{N}$ will be called h_* -**admissible**⁵⁶ if M is closed and the obvious inclusion $i^x : (M, \emptyset) \hookrightarrow (M, M \setminus \{x\})$ induces an isomorphism

$$i_*^x : h_n(M) \xrightarrow{\cong} h_n(M, M \setminus \{x\})$$

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:

PROPOSITION 36.11. *For each $n \in \mathbb{N}$, the sphere S^n is h_* -admissible for every axiomatic homology theory h_* .*

PROOF. See Exercise 36.1. □

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 semester. 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 36.12. *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 homology, and for the other manifolds mentioned above, it can be proved using Mayer-Vietoris sequences (see Lecture 34, especially Exercises 34.5 and 34.6), or in certain cases also the exact sequence of a mapping torus (Lecture 35). These two conditions in tandem give us the following. In Lecture 30, we saw that a triangulation of a closed n -manifold M gives rise to a fundamental class in simplicial homology, which we can feed into the natural map from ordered simplicial to singular homology to obtain a singular fundamental class

$$[M] \in H_n(M; \mathbb{Z}_2).$$

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. The homomorphism $i_*^x : H_n(M; \mathbb{Z}_2) \rightarrow$

⁵⁵We 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 36.2. A notion of orientations for topological manifolds of dimension $n \leq 2$ was discussed last semester in Lecture 20, and Example 36.3 hints at the relationship between that notion and local orientations in the case $n = 2$.

⁵⁶This is not a universally standard term, but it is convenient for our purposes at the moment.

$H_n(M, M \setminus \{x\}; \mathbb{Z}_2)$ then sends $[M]$ to the class in $H_n(M, M \setminus \{x\}; \mathbb{Z}_2)$ represented by a single singular n -simplex $\sigma : \Delta^n \rightarrow M$ whose image is the n -simplex of the triangulation containing x ; this n -chain is a relative cycle in $(M, M \setminus \{x\})$ since $\sigma(\partial\Delta^n) \subset M$ does not touch x . We can choose local coordinates in which x is identified with $0 \in \mathbb{R}^n$ and $\sigma(\Delta^n)$ with the unit disk \mathbb{D}^n ; following the chain of isomorphisms (36.1) from $H_n(M, M \setminus \{x\}; \mathbb{Z}_2)$ to $\tilde{H}_{n-1}(S^{n-1}; \mathbb{Z}_2)$ then identifies $i_*^x[M] = [\sigma]$ with a fundamental class determined by the obvious triangulation of $\partial\Delta^n \cong S^{n-1}$. It is not hard to prove by induction on n that this fundamental class generates $\tilde{H}_{n-1}(S^{n-1}; \mathbb{Z}_2) \cong \mathbb{Z}_2$, and it follows that $i_*^x[M]$ generates $H_n(M, M \setminus \{x\}; \mathbb{Z}_2) \cong \mathbb{Z}_2$, meaning it is the unique nontrivial element. If we also know $H_n(M; \mathbb{Z}_2) \cong \mathbb{Z}_2$, this implies that the map $i_*^x : H_n(M; \mathbb{Z}_2) \rightarrow H_n(M, M \setminus \{x\}; \mathbb{Z}_2)$ 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 can similarly deduce the stronger statement that $i_*^x : H_n(M; \mathbb{Z}) \rightarrow H_n(M, M \setminus \{x\}; \mathbb{Z}) \cong \mathbb{Z}$ 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 $H_n(M; \mathbb{Z}) \rightarrow H_n(M, M \setminus \{x\}; \mathbb{Z})$ 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 $H_n^\Delta(K; G)$ when K is the simplicial complex that triangulates a closed connected n -manifold (with or without orientation), see Exercise 30.1. 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 later from the general construction of fundamental classes in singular homology. \square

DEFINITION 36.13. 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 $f : M \rightarrow N$ 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 $S^n \rightarrow S^n$. 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 35.6 are satisfied for this new definition.

We can now state the main result relating global and local degrees.

THEOREM 36.14. *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 $f : M \rightarrow N$ and any point $y \in N$ such that $f^{-1}(y)$ is a finite set,*

$$(36.2) \quad \deg(f) = \sum_{x \in f^{-1}(y)} \deg(f; x).$$

Similarly, if M and N are \mathbb{Z}_2 -admissible and $f : M \rightarrow N$ 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 (36.2) as the *algebraic count* of points in $f^{-1}(y)$. One can check that if $f : M \rightarrow N$ 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 36.15. For a given map $f : M \rightarrow N$, 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 36.14 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 $f : M \rightarrow N$ 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 smoothed 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 36.8, and will always be ± 1 .

Theorem 36.14 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 $f : M \rightarrow N$, 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 36.15 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 (36.2) 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 36.14 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 $\{f_t : M \rightarrow N\}_{t \in [0,1]}$ 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].⁵⁷ 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 (36.2) matches the left hand side. The proof of that requires the formal properties of homology theories.

PROOF OF THEOREM 36.14. 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.

⁵⁷The differentiable approach to the mapping degree was also sketched in Exercise 19.14 last semester.

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

$$f(\mathcal{U}_k) \subset \mathcal{V} \quad \text{and} \quad \mathcal{U}_k \cap \mathcal{U}_j = \emptyset \text{ for } j \neq k.$$

These assumptions guarantee that $f(\mathcal{U}_k \setminus \{x_k\}) \subset \mathcal{V} \setminus \{y\}$, hence f also defines a map of pairs $(\mathcal{U}_k, \mathcal{U}_k \setminus \{x_k\}) \rightarrow (\mathcal{V}, \mathcal{V} \setminus \{y\})$ for every $k = 1, \dots, \ell$. Now consider the diagram

$$(36.3) \quad \begin{array}{ccccc} & & h_n(\mathcal{U}_k, \mathcal{U}_k \setminus \{x_k\}) & \xrightarrow{f_*} & h_n(\mathcal{V}, \mathcal{V} \setminus \{y\}) \\ & \swarrow & \downarrow \gamma_*^k & & \downarrow \beta \\ h_n(M, M \setminus \{x_k\}) & \xleftarrow{\alpha_*^k} & h_n(M, M \setminus f^{-1}(y)) & \xrightarrow{f_*} & h_n(N, N \setminus \{y\}) \\ & \swarrow p_*^k & \uparrow j_* & & \uparrow i_*^y \\ & & h_n(M) & \xrightarrow{f_*} & h_n(N) \\ & \nwarrow i_*^{x_k} & & & \end{array}$$

where the maps $\alpha^k, p^k, \gamma^k, j$ 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

$$p := (p_*^1, \dots, p_*^\ell) : h_n(M, M \setminus f^{-1}(y)) \rightarrow \bigoplus_{k=1}^\ell h_n(M, M \setminus \{x_k\}),$$

which fits into the following diagram:

$$\begin{array}{ccc} h_n(M, M \setminus f^{-1}(y)) & \xrightarrow{p} & \bigoplus_{k=1}^\ell h_n(M, M \setminus \{x_k\}) \\ \cong \uparrow & & \cong \uparrow \\ h_n\left(\prod_{k=1}^\ell (\mathcal{U}_k, \mathcal{U}_k \setminus \{x_k\})\right) & \xleftarrow{\cong} & \bigoplus_{k=1}^\ell h_n(\mathcal{U}_k, \mathcal{U}_k \setminus \{x_k\}) \end{array}$$

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 28.8), thus p is also an isomorphism. If we use this to replace $h_n(M, M \setminus f^{-1}(y))$ in (36.3) by $\bigoplus_{k=1}^\ell h_n(M, M \setminus \{x_k\})$, then the map p_*^k becomes simply the projection of $\bigoplus_{k=1}^\ell h_n(M, M \setminus \{x_k\})$ to the factor $h_n(M, M \setminus \{x_k\})$. With this replacement understood, we have

$$j_* = (j_*^{x_1}, \dots, j_*^{x_\ell}) : h_n(M) \rightarrow \bigoplus_{k=1}^\ell h_n(M, M \setminus \{x_k\}),$$

and the commutativity of the bottom right square in (36.3) then gives the formula

$$(36.4) \quad f_* j_* = \sum_{k=1}^\ell f_* i_*^{x_k} = i_*^y f_* : h_n(M) \rightarrow h_n(N, N \setminus \{y\}).$$

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

$$\sum_{k=1}^\ell f_* [M]_{x_k} = \sum_{k=1}^\ell \deg(f; x_k) [N]_y = i_*^y f_* [M] = \deg(f) i_*^y [M] = \deg(f) [N]_y,$$

from which the formula for the integer-valued degree follows. The formula for the mod 2 degree follows in the same way using $h_* = H_*(\cdot; \mathbb{Z}_2)$. \square

It is easy to see that a non-surjective map $f : S^n \rightarrow S^n$ must have degree 0, because its image then lies in the contractible space $S^n \setminus \{*\} \cong \mathbb{R}^n$, making f homotopic to a constant. The same argument does not work for maps $f : M \rightarrow N$ when N is a more general closed n -manifold, but Theorem 36.14 nonetheless gives us an easy proof of the same statement:

COROLLARY 36.16. *If M and N are \mathbb{Z} -admissible n -manifolds with $n \geq 1$ and $f : M \rightarrow N$ 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 36.14 to identify $\deg(f)$ or $\deg_2(f)$ with a suitable count of points in $f^{-1}(y)$ where $y \notin f(M)$. \square

In the axiomatic setting, the proof of Theorem 36.14 also gives rise to the following result, which will be of some theoretical importance when we develop cellular homology. Note that by Exercise 36.1, we already know S^n to be h_* -admissible for every h_* and every $n \in \mathbb{N}$.

THEOREM 36.17. *For any map $f : S^n \rightarrow S^n$ with $n \in \mathbb{N}$ and any axiomatic homology theory h_* , the induced homomorphism $f_* : h_n(S^n) \rightarrow h_n(S^n)$ takes the form $c \mapsto \deg(f)c$.*

PROOF. One can verify explicitly that the corresponding statement about reduced homology holds for all maps $f : S^0 \rightarrow S^0$; 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 (cf. Exercise 35.3). We now argue by induction on the dimension, assuming for a given n that homomorphisms $f_* : \tilde{h}_{n-1}(S^{n-1}) \rightarrow \tilde{h}_{n-1}(S^{n-1})$ are always given by multiplication with the integer-valued degree of maps $f : S^{n-1} \rightarrow S^{n-1}$. Using perturbation results from differential topology as mentioned in Remark 36.15, we can assume after a small perturbation of any given map $f : S^n \rightarrow S^n$ within its homotopy class that $f^{-1}(y)$ is a finite set for some $y \in S^n$. Now write $f^{-1}(y) = \{x_1, \dots, x_\ell\}$ and, given $c \in h_n(S^n)$, use (36.4) to write

$$i_*^y f_* c = \sum_{k=1}^{\ell} f_* i_*^{x_k} c \in h_n(S^n, S^n \setminus \{y\}),$$

where the individual terms on the right hand side involve the homomorphisms

$$f_* : h_n(S^n, S^n \setminus \{x_k\}) \rightarrow h_n(S^n, S^n \setminus \{y\}).$$

Using excision and connecting homomorphisms as in Proposition 36.5, one can identify both the domain and target of this map with $\tilde{h}_{n-1}(S^{n-1})$ so that f_* is equivalent to the homomorphism $\tilde{h}_{n-1}(S^{n-1}) \rightarrow \tilde{h}_{n-1}(S^{n-1})$ induced by a map $S^{n-1} \rightarrow S^{n-1}$, 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

$$\begin{array}{ccc} h_n(S^n, S^n \setminus \{x_k\}) & \xrightarrow{f_*} & h_n(S^n, S^n \setminus \{y\}) \\ \downarrow \cong & & \downarrow \cong \\ h_{n-1}(S^{n-1}) & \xrightarrow{\cdot \deg(f; x_k)} & h_{n-1}(S^{n-1}) \end{array}$$

Adding up these contributions for every $x_k \in f^{-1}(y)$ produces multiplication by $\deg(f)$ according to Theorem 36.14. \square

One consequence of this result is that the definition of $\deg(f)$ for maps $f : S^n \rightarrow S^n$ does not actually depend on the choice to use *singular* homology in particular—we could have replaced $H_*(\cdot; \mathbb{Z})$ with any other axiomatic 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 36.18. *If M and N are both \mathbb{Z} -admissible and \mathbb{Z}_2 -admissible, then for every map $f : M \rightarrow N$, $\deg_2(f)$ is the image of $\deg(f)$ under the natural projection $\mathbb{Z} \rightarrow \mathbb{Z}_2$. \square*

36.1. Exercises.

EXERCISE 36.1 (*). Prove Proposition 36.11 on the h_* -admissibility of S^n .

Hint: You can choose a neighborhood $\mathcal{U} \subset S^n$ of any $x \in S^n$ homeomorphic to a disk, and use $h_n(S^n, S^n \setminus \mathcal{U})$ as a substitute for $h_n(S^n, S^n \setminus \{x\})$ (why?). What kind of space is $S^n \setminus \mathcal{U}$?

EXERCISE 36.2. Prove Proposition 36.9 on the stability of solutions to the equation $f(x) = y$ when $\deg(f; x)$ or $\deg_2(f; x)$ is nonzero.

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 $\mathbb{R}^n \setminus \{y\}$?

EXERCISE 36.3. Viewing S^1 as the unit circle in \mathbb{C} , fix a generator $[S^1] \in H_1(S^1; \mathbb{Z}) \cong \mathbb{Z}$ and use it to determine local orientations $[\mathbb{C}]_z \in H_n(\mathbb{C}, \mathbb{C} \setminus \{z\}; \mathbb{Z})$ for every point $z \in \mathbb{C}$ via the natural isomorphisms $H_2(\mathbb{C}, \mathbb{C} \setminus \{z\}; \mathbb{Z}) \cong H_2(\mathbb{D}_z, \partial\mathbb{D}_z; \mathbb{Z}) \cong H_1(\partial\mathbb{D}_z; \mathbb{Z})$, where $\mathbb{D}_z \subset \mathbb{C}$ denotes the closed unit disk centered at z , whose boundary is canonically identified with S^1 . This choice will be used in the following for the definition of local degrees of maps $f : \mathcal{U} \rightarrow \mathbb{C}$ defined on open subsets $\mathcal{U} \subset \mathbb{C}$; note that changing the generator $[S^1] \in H_1(S^1; \mathbb{Z})$ does not change the definition of $\deg(f; z)$ since it changes *both* $[\mathbb{C}]_z$ and $[\mathbb{C}]_{f(z)}$ by a sign.

- Show that if $f : \mathbb{C} \rightarrow \mathbb{C}$ is of the form $f(z) = (z - z_0)^k g(z)$ for some $z_0 \in \mathbb{C}$, $k \in \mathbb{N}$ and g a continuous map with $g(z_0) \neq 0$, then $\deg(f; z_0) = k$.
- Can you modify the example in part (a) to produce one with $\deg(f; z_0) = -k$ for $k \in \mathbb{N}$?
- Let $f : S^2 \rightarrow S^2$ denote the natural continuous extension to $S^2 := \mathbb{C} \cup \{\infty\}$ of a complex polynomial $\mathbb{C} \rightarrow \mathbb{C}$ of degree n . What is $\deg(f)$?
- Pick a constant $t_0 \in S^1$ and let $A \cong S^1 \vee S^1$ denote the subset $\{(x, y) \mid x = t_0 \text{ or } y = t_0\} \subset S^1 \times S^1 = \mathbb{T}^2$. Show that $\mathbb{T}^2/A \cong S^2$, and that the quotient map $\mathbb{T}^2 \rightarrow \mathbb{T}^2/A$ has degree ± 1 (depending on choices of generators for $H_2(\mathbb{T}^2; \mathbb{Z})$ and $H_2(S^2; \mathbb{Z})$).

EXERCISE 36.4. Find an example of a smooth map $f : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ that has an isolated zero at the origin with $\deg(f; 0) = 0$ and admits arbitrarily small perturbations that are nowhere zero.

EXERCISE 36.5. Suppose $f : S^n \rightarrow S^n$ is any continuous map, and $p_+ \in \Sigma S^n = C_+ S^n \cup_{S^n} C_- S^n$ is the vertex of the top cone in the suspension $\Sigma S^n \cong S^{n+1}$. What is $\deg(\Sigma f; p_+)$? Use this to give a new proof (different from that of Proposition 35.7) that $\deg(\Sigma f) = \deg(f)$.

EXERCISE 36.6. Show that every map $S^n \rightarrow \mathbb{T}^n$ has degree 0 if $n \geq 2$.

Hint: Lift $S^n \rightarrow \mathbb{T}^n$ to the universal cover of \mathbb{T}^n .

EXERCISE 36.7. Show that for every $d \in \mathbb{Z}$ and every \mathbb{Z} -admissible n -dimensional manifold M with $n \geq 1$, there exists a map $M \rightarrow S^n$ of degree d .

Hint: Try a map that is interesting only on some n -ball in M and constant everywhere else.

EXERCISE 36.8. Using the mod 2 degree, describe examples of maps from $\mathbb{R}\mathbb{P}^2$ or the Klein bottle to S^2 that cannot be homotopic to a constant.

37. CW-complexes

37.1. Definitions and first examples. 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 37.1 (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

$$X^0 \subset X^1 \subset X^2 \subset \dots \subset \bigcup_{n \geq 0} X^n = X$$

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 $\{\varphi_\alpha : S^{n-1} \rightarrow X^{n-1}\}_{\alpha \in \mathcal{K}^n}$ 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.

$$(37.1) \quad X^n = X^{n-1} \cup_{\varphi^n} \coprod_{\alpha \in \mathcal{K}^n} \mathbb{D}^n, \quad \text{where} \quad \varphi^n := \coprod_{\alpha \in \mathcal{K}^n} \varphi_\alpha : \coprod_{\alpha \in \mathcal{K}^n} \partial \mathbb{D}^n \rightarrow X^{n-1}.$$

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

$$e_\alpha^n \subset X^n,$$

which is called an n -**cell** (*n-Zelle*) of the complex, and the associated map $\varphi_\alpha : S^{n-1} \rightarrow X^{n-1}$ is called its **attaching map** (*Anklebeabbildung*). The map

$$\Phi_\alpha^n : \mathbb{D}^n \rightarrow X$$

that satisfies $\Phi_\alpha^n|_{\partial \mathbb{D}^n} = \varphi_\alpha$ and restricts to the interior of the disk as the inclusion $e_\alpha^n \hookrightarrow X^n$ is called the **characteristic map** (*charakteristische Abbildung*) of the cell e_α^n . The complex is called n -**dimensional** if n is the largest number for which it contains an n -cell, i.e. $\mathcal{K}^m = \emptyset$ for all $m > n$ but $\mathcal{K}^n \neq \emptyset$.

Let us recall quickly what the notation in (37.1) means: we are defining X^n as a quotient of a disjoint union,

$$X^n = X^{n-1} \amalg \left(\coprod_{\alpha \in \mathcal{K}^n} \mathbb{D}^n \right) / \sim,$$

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 $n \geq 0$ 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 37.2 (CW-complexes, part 2 of 2). The topology of a CW-complex $X = X^0 \cup X^1 \cup X^2 \cup \dots$ 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 $n \geq 0$.

The next two results regarding the topology of CW-complexes are important but straightforward exercises in point-set topology.

PROPOSITION 37.3. *A subset $\mathcal{U} \subset X$ in a CW-complex is open if and only if for every $n \geq 0$ and every n -cell e_α^n , $\Phi_\alpha^{-1}(\mathcal{U})$ 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.* \square

PROPOSITION 37.4. *For any CW-complex X and any space Y , a map $f : X \rightarrow Y$ is continuous if and only if its restriction to the n -skeleton of X is continuous for every $n \geq 0$, or equivalently, if $f \circ \Phi_\alpha : \mathbb{D}^n \rightarrow Y$ is continuous for every $n \geq 0$ and $\alpha \in \mathcal{K}^n$.* \square

REMARK 37.5. 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 37.2 is not weak at all, but is the *strongest* with a given property. This discrepancy is apparently the fault of J.H.C. Whitehead,⁵⁸ 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. Proposition 37.4 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 37.6. A **cell decomposition** (*Zellenzerlegung*) of a space X is a choice of homeomorphism from X to a CW-complex.

EXAMPLE 37.7. 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 $k \leq n$. The attaching map $\partial\Delta^n \cong S^{n-1} \rightarrow X^{n-1}$ of each n -cell $e_\alpha^n \subset X$ 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 37.8. Recall that $S^n \cong \mathbb{D}^n / \partial\mathbb{D}^n$ for $n \geq 1$. 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 $\Phi : \mathbb{D}^n \rightarrow \mathbb{D}^n / \partial\mathbb{D}^n$. 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 37.8, 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 37.7. 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.

⁵⁸This terminological awkwardness is consistent with certain popular jokes among mathematicians, one of which is that J.H.C. Whitehead's initials actually stand for “Jesus, he's confused!”

EXAMPLE 37.9. 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 $S^{k-1} \rightarrow X^{k-1}$ are homeomorphisms, and all characteristic maps are inclusions.

EXAMPLE 37.10. 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 $S^n \rightarrow S^n$ sends \mathbb{D}_\pm^n to \mathbb{D}_\mp^n 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 37.9 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

$$X^k = \{(x_0, \dots, x_n) \in \mathbb{RP}^n = S^n/\mathbb{Z}_2 \mid x_{k+1} = \dots = x_n = 0\} \cong \mathbb{RP}^k.$$

In contrast to Example 37.9, the characteristic maps $\mathbb{D}^k \rightarrow \mathbb{RP}^n$ for this cell decomposition are not injective: indeed, the k -cells in Example 37.9 are attached to the $(k-1)$ -skeleton S^{k-1} via a homeomorphism $S^{k-1} \rightarrow S^{k-1}$, 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 $S^{k-1} \rightarrow S^{k-1}$ from Example 37.9 gets composed with the quotient projection $S^{k-1} \rightarrow \mathbb{RP}^{k-1}$ and becomes a covering map of degree 2.

EXAMPLE 37.11. This will be harder to picture, but one can adjust Example 37.9 by following the same procedure of attaching two k -cells along homeomorphisms $S^{k-1} \rightarrow X^{k-1}$ 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 $n \geq 1$ 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}^∞ that also happens to be contained in S^{k+1} , and S^∞ is the union of the nested sequence of spaces

$$S^0 \subset S^1 \subset S^2 \subset S^3 \subset \dots \subset \bigcup_{k \geq 0} S^k = S^\infty.$$

More concretely, S^∞ is just the subset of \mathbb{R}^∞ 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. The following observation (see Exercise 37.2) reveals that there is something a bit subtle about the topology of S^∞ : for any convergent sequence $x_k \in S^\infty$, there exists $n \in \mathbb{N}$ such that $x_k \in S^n$ for every k .

REMARK 37.12. The observation about sequences mentioned above demonstrates 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

$$\ell^2 := \left\{ \mathbf{x} = (x_1, x_2, \dots) \in \prod_{i=1}^\infty \mathbb{R} \mid \sum_{i=1}^\infty x_i^2 < \infty \right\}$$

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 37.13. Combining Examples 37.10 and 37.11 in the obvious way produces another infinite-dimensional CW-complex called $\mathbb{R}\mathbb{P}^\infty$, which has exactly one k -cell for every $k \geq 0$. 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 $X \rightarrow \mathbb{R}\mathbb{P}^\infty$. 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 37.14. Recall that the closed oriented surface Σ_g of genus $g \geq 0$ 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 $S^0 \rightarrow X^0$, and a single 2-cell attached via a map $S^1 \rightarrow X^1$ that defines the concatenation of loops indicated by the above word.

DEFINITION 37.15. 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 $n \geq 0$, such that every cell in A is also a cell in X with the same characteristic map.

37.2. Compact subsets. 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 to know a basic fact about the point-set topology of CW-complexes that is vacuous in the case of finite complexes, but nontrivial for infinite complexes:

PROPOSITION 37.16. *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 37.17. *The closure of each cell in a CW-complex intersects only finitely many other cells.* \square

PROOF OF PROPOSITION 37.16. *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 $n \geq 0$. 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 $\varphi_\alpha : S^{n-1} \rightarrow X^{n-1}$ and characteristic map $\Phi_\alpha : \mathbb{D}^n \rightarrow X$, $\varphi_\alpha^{-1}(A)$ 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, Proposition 37.3 now implies that $A \cap X^n \subset X^n$ is closed. By induction, this is true for every $n \geq 0$, 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 $n \geq 0$, 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 $\varphi_\alpha : S^{n-1} \rightarrow X^{n-1}$ of any n -cell e_α^n , providing a finite subcomplex $A \subset X^{n-1}$ whose union with e_α^n is a finite subcomplex of X^n containing e_α^n . 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_α^n of the complex, the compact subset $\overline{e_\alpha^n} = \Phi_\alpha(\mathbb{D}^n) \subset X$ is contained in a finite subcomplex, and combining this with the claim in step 2 proves the result. \square

37.3. Cellular homology. 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 R -module G to use for coefficients, and omit G from the notation whenever the choice of coefficients is unimportant. For $n \in \mathbb{Z}$, define $C_n^{\text{CW}}(X)$ to be the trivial module if $n < 0$, and otherwise

$$C_n^{\text{CW}}(X) = C_n^{\text{CW}}(X; G) := \bigoplus_{\alpha \in \mathcal{K}^n} G,$$

so e.g. $C_n^{\text{CW}}(X; \mathbb{Z})$ is the free abelian group generated by the set of n -cells e_α^n in our given cell decomposition of X . We shall regard the cells e_α^n as generators of $C_n^{\text{CW}}(X)$, thus writing elements of $C_n^{\text{CW}}(X)$ as sums

$$\sum_{\alpha \in \mathcal{K}^n} a_\alpha e_\alpha^n \in C_n^{\text{CW}}(X)$$

for coefficients $a_\alpha \in G$, with the understanding that the sum is always finite because at most finitely-many of the coefficients a_α are nonzero. The direct sum of all these groups produces a \mathbb{Z} -graded R -module

$$C_*^{\text{CW}}(X) = C_*^{\text{CW}}(X; G) := \bigoplus_{n \in \mathbb{Z}} C_n^{\text{CW}}(X),$$

which we shall now turn into a chain complex by defining a suitable boundary operator $\partial : C_*^{\text{CW}}(X) \rightarrow C_{*-1}^{\text{CW}}(X)$. There is a geometric motivation for the definition: for each generator e_α^n of $C_n^{\text{CW}}(X)$, 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 $n \geq 1$ the map $p_\alpha : X^n \rightarrow S^n$ determined by the following diagram:

$$(37.2) \quad \begin{array}{ccc} X^n & \xrightarrow{p_\alpha} & \mathbb{D}^n / \partial \mathbb{D}^n \equiv S^n \\ \downarrow \text{pr} & \swarrow \Phi_\alpha & \\ X^n / (X^n \setminus e_\alpha^n) & & \end{array}$$

Here pr denotes the quotient projection, and the fact that φ_α maps $\partial \mathbb{D}^n$ into $X^{n-1} \subset X^n \setminus e_\alpha^n$ implies that the characteristic map $\Phi_\alpha : \mathbb{D}^n \rightarrow X^n$ descends to a map of the quotients $\mathbb{D}^n / \partial \mathbb{D}^n \rightarrow X^n / (X^n \setminus e_\alpha^n)$. The key point is that the latter is a homeomorphism, thus we can invert it to define $p_\alpha = \Phi_\alpha^{-1} \circ \text{pr}$ as a map from $X^n \rightarrow S^n$ 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 $X^0 / (X^0 \setminus e_\alpha^0)$ 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 $X^0 \cong \{*\}$, i.e. when there is only one 0-cell e_α^0 and $X^0 \setminus e_\alpha^0 = \emptyset$, but in this case the map $X^0 / (X^0 \setminus e_\alpha^0) \rightarrow S^0$ is well

defined nonetheless, so we will adopt the convention of using it to define

$$(37.3) \quad \begin{array}{ccc} X^0 & \xrightarrow{p_\alpha} & S^0 \\ \downarrow \text{pr} & \nearrow & \\ X^0/(X^0 \setminus e_\alpha^0) & & \end{array}$$

as the analogue of (37.2) for $n = 0$.

DEFINITION 37.18. Given an n -cell e_α^n and an $(n-1)$ -cell e_β^{n-1} in a CW-complex X , we define the **incidence number**

$$[e_\beta^{n-1} : e_\alpha^n] \in \mathbb{Z}$$

as the degree of the map

$$S^{n-1} \xrightarrow{p_\beta \circ \varphi_\alpha} S^{n-1}$$

defined by composing the attaching map $\varphi_\alpha : S^{n-1} \rightarrow X^{n-1}$ for e_α^n with the map $p_\beta : X^{n-1} \rightarrow S^{n-1}$ defined by replacing e_α^n with e_β^{n-1} in the diagram (37.2) or (37.3).

Observe that whenever $\bar{e}_\alpha^n \cap e_\beta^{n-1} = \emptyset$, it follows that the image of $\varphi_\alpha : S^{n-1} \rightarrow X^{n-1}$ is disjoint from e_β^{n-1} and is thus mapped to a constant by $p_\beta : X^{n-1} \rightarrow S^{n-1}$, hence $p_\beta \circ \varphi_\alpha$ is a constant map and $[e_\beta^{n-1} : e_\alpha^n] = 0$. In light of Corollary 37.17, this implies that the sum in the following definition makes sense, because it can only have finitely-many nonzero terms.

DEFINITION 37.19. For each $n \in \mathbb{N}$, the boundary operator $C_n^{\text{CW}}(X) \xrightarrow{\partial} C_{n-1}^{\text{CW}}(X)$ is determined by linearity and the formula

$$\partial e_\alpha^n := \sum_{\beta \in \mathcal{K}^{n-1}} [e_\beta^{n-1} : e_\alpha^n] e_\beta^{n-1}$$

for each $\alpha \in \mathcal{K}^n$.

We define $\partial : C_n^{\text{CW}}(X) \rightarrow C_{n-1}^{\text{CW}}(X)$ to be the trivial map for every $n \leq 0$, as it must be since its target is then the trivial module.

Let us work out a more useful formula for $\partial : C_1^{\text{CW}}(X) \rightarrow C_0^{\text{CW}}(X)$. If $X^0 \cong \{*\}$, then $p_\beta \circ \varphi_\alpha : S^0 \rightarrow S^0$ always factors through a one-point space and is therefore a constant map, implying $[e_\beta^0 : e_\alpha^1] = 0$ 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 $p_\beta : X^0 \rightarrow S^0$ is the map that sends e_β^0 to $1 \in S^1$ and every other 0-cell to $-1 \in S^1$, so composing it with the attaching map $\varphi_\alpha : \partial \mathbb{D}^1 \rightarrow X^0$ produces the following possibilities:

- If $\varphi_\alpha(1) = e_\beta^0$ and $\varphi_\alpha(-1) \neq e_\beta^0$, then $p_\beta \circ \varphi_\alpha : S^0 \rightarrow S^0$ is the identity map and thus $[e_\beta^0 : e_\alpha^1] = 1$.
- If $\varphi_\alpha(1) \neq e_\beta^0$ but $\varphi_\alpha(-1) = e_\beta^0$, then $p_\beta \circ \varphi_\alpha(\pm 1) = \mp 1$ and thus $[e_\beta^0 : e_\alpha^1] = -1$.
- In all other cases, $p_\beta \circ \varphi_\alpha$ is constant and thus $[e_\beta^0 : e_\alpha^1] = 0$.

Since each point of X^0 is a 0-cell, we can identify it with a generator of $C_0^{\text{CW}}(X)$ and thus deduce from the remarks above the following:

PROPOSITION 37.20. *The map $\partial : C_1^{\text{CW}}(X) \rightarrow C_0^{\text{CW}}(X)$ is determined by the formula*

$$\partial e_\alpha^1 = \varphi_\alpha(1) - \varphi_\alpha(-1).$$

□

We shall now state two important theorems whose proofs will take up most of the next two lectures: the first states simply that $(C_*^{\text{CW}}(X), \partial)$ is a chain complex.

THEOREM 37.21. *The map $\partial : C_*^{\text{CW}}(X) \rightarrow C_*^{\text{CW}}(X)$ satisfies $\partial^2 = 0$.*

The **cellular homology** (*zelluläre Homologie*) of the CW-complex X with coefficient group G can now be defined as

$$H_*^{\text{CW}}(X) = H_*^{\text{CW}}(X; G) := H_*(C_*^{\text{CW}}(X; G), \partial).$$

The notation $C_*^{\text{CW}}(X)$ and $H_*^{\text{CW}}(X)$ 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 37.22. *For any CW-complex X and any axiomatic homology theory h_* with coefficient group G , there is a natural isomorphism $H_*^{\text{CW}}(X; G) \cong h_*(X)$.*

We'll say more in the next lecture about the precise meaning of the word “natural” in this statement. Theorem 37.22 has several remarkable consequences that can be recognized immediately: one is that $H_*^{\text{CW}}(X)$ 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 37.7, this also gives us a new explanation for why the *simplicial* homology of a polyhedron depends only on its topology: the oriented simplicial chain complex of a polyhedron is the same as its cellular chain complex!

Before trying to explain why all this is true, let's look at a couple of examples that will make Theorem 37.22 look more plausible.

EXAMPLE 37.23. We saw in Example 37.8 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 $X^0 = X^1 = \dots = X^{n-1} \cong \{*\}$ 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 $n \geq 2$, then the claim holds trivially because for every $k \in \mathbb{Z}$, either the domain or the target of the map $\partial : C_k^{\text{CW}}(S^n) \rightarrow C_{k-1}^{\text{CW}}(S^n)$ is trivial. When $n = 1$ there is still something to check: $\partial : C_1^{\text{CW}}(S^1) \rightarrow C_0^{\text{CW}}(S^1)$ 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 $p \circ \varphi : S^0 \rightarrow S^0$, where $\varphi : S^0 \rightarrow X^0 \cong \{*\}$ is the attaching map for e^1 and $p : X^0 \rightarrow S^0$ 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 37.24. We consider S^2 with the alternative cell decomposition described in Example 37.9, 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 $C_k^{\text{CW}}(S^2) = 0$ for $k < 0$ or $k > 2$, while $C_k^{\text{CW}}(S^2) = G \oplus G$ for each $k = 0, 1, 2$, with the two factors of the coefficient group G corresponding to the two generators $e_+^k, e_-^k \in C_k^{\text{CW}}(S^2; \mathbb{Z})$. Denote the attaching map for e_{\pm}^k by $\varphi_{\pm}^k : S^{k-1} \rightarrow X^{k-1}$, and denote the projection map as defined in (37.2) by $p_{\pm}^k : X^k \rightarrow S^k$, so $\partial : C_k^{\text{CW}}(S^2) \rightarrow C_{k-1}^{\text{CW}}(S^2)$ is now determined by

$$(37.4) \quad \begin{aligned} \partial e_+^k &= \deg(p_+^{k-1} \circ \varphi_+^k) e_+^{k-1} + \deg(p_-^{k-1} \circ \varphi_+^k) e_-^{k-1}, \\ \partial e_-^k &= \deg(p_+^{k-1} \circ \varphi_-^k) e_+^{k-1} + \deg(p_-^{k-1} \circ \varphi_-^k) e_-^{k-1}. \end{aligned}$$

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 $\Phi_{\pm}^1 : \mathbb{D}^1 \rightarrow S^2$ of the two 1-cells e_{\pm}^1 via the x coordinate, giving

$$\Phi_{\pm}^1 : \mathbb{D}^1 \rightarrow S^2 : x \mapsto (x, \pm\sqrt{1-x^2}, 0),$$

so the attaching maps $\varphi_{\pm}^1 : S^0 \rightarrow S^0$ are the restrictions of these to $\partial\mathbb{D}^1$ and are thus both the identity map $S^0 \rightarrow S^0$. Each of the maps $p_{\pm}^0 : X^0 \rightarrow S^0$ 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 $S^0 \rightarrow S^0$ 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 (37.4) 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

$$\Phi_{\pm}^2 : \mathbb{D}^2 \rightarrow S^2 : (x, y) \mapsto (x, y, \pm\sqrt{1-x^2-y^2}),$$

and the attaching maps $\varphi_{\pm}^2 : S^1 \rightarrow X^1$ thus become once again the identity map $S^1 \rightarrow S^1$. To understand the maps $p_{\pm}^1 : X^1 \rightarrow S^1$, let us first agree that the identification of $\mathbb{D}^1/\partial\mathbb{D}^1$ with S^1 should be defined via a path $\gamma : \mathbb{D}^1 \rightarrow S^1$ 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, $p_+^1 : S^1 \rightarrow \mathbb{D}^1/\partial\mathbb{D}^1$ 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 $\deg(p_+^1 \circ \varphi_{\pm}^2) = -1$. Meanwhile, $p_-^1 : S^1 \rightarrow \mathbb{D}^1/\partial\mathbb{D}^1$ 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 $\deg(p_-^1 \circ \varphi_{\pm}^2) = 1$. We thus have

$$\partial e_+^2 = \partial e_-^2 = -e_+^1 + e_-^1.$$

With these formulas in place, we can compute the homology of $C_*^{\text{CW}}(S^2)$ 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 $C_2^{\text{CW}}(S^2)$ with $G \oplus G$, the group of 2-cycles takes the form

$$\ker \partial_2 = \{(g, -g) \in G \oplus G \mid g \in G\} \subset C_2^{\text{CW}}(S^2),$$

which is isomorphic to G . Since $C_3^{\text{CW}}(S^2) = 0$, we conclude $H_2^{\text{CW}}(S^2) \cong G$. 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 $H_1^{\text{CW}}(S^2) = 0$. Finally, all 0-chains $ge_+^0 + he_-^0$ are cycles since $C_{-1}^{\text{CW}}(S^2) = 0$, but under the obvious isomorphism $C_0^{\text{CW}}(S^2) = G \oplus G$ we have

$$\text{im } \partial_1 = \{(g, -g) \in G \oplus G \mid g \in G\} \subset C_0^{\text{CW}}(S^2),$$

so $H_0^{\text{CW}}(S^2)$ 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 37.23.

It is not too hard to extend Example 37.24 to a computation of $H_*^{\text{CW}}(S^n)$ for every $n \in \mathbb{N}$ in terms of the cell decomposition $S^n = e_+^0 \cup e_-^0 \cup \dots \cup e_+^n \cup e_-^n$. 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 $k = 1, \dots, n-1$, but the groups $\ker \partial_n$ and $C_0^{\text{CW}}(S^n)/\text{im } \partial_1$ 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 37.23, so we will invest considerable effort during the next lecture into proving that this is allowed, because the isomorphism class of $H_*^{\text{CW}}(X)$ depends in general only on the topology of X and not on its cell decomposition.

Let's do one more easy example.

EXAMPLE 37.25. We saw in Example 37.14 that the closed oriented surface Σ_g of genus $g \geq 0$ has a cell decomposition with one 0-cell e^0 , $2g$ cells of dimension one which we can label

$$e_{a_1}^1, e_{b_1}^1, \dots, e_{a_g}^1, e_{b_g}^1,$$

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 $C_1^{\text{CW}}(\Sigma_g)$:

$$\partial e_{a_j}^1 = \partial e_{b_j}^1 = 0 \quad \text{for } j = 1, \dots, g.$$

The attaching map $\varphi : S^1 \rightarrow X^1$ 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 $p_{a_1} : X^1 \rightarrow S^1$ that collapses $X^1 \setminus e_{a_1}^1$ 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 $S^1 \rightarrow S^1$ 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

$$H_k^{\text{CW}}(\Sigma_g; G) = C_k^{\text{CW}}(\Sigma_g; G) \cong \begin{cases} G & \text{for } k = 0, 2, \\ G^{2g} & \text{for } k = 1, \\ 0 & \text{for } k < 0 \text{ and } k > 2. \end{cases}$$

37.4. Exercises.

EXERCISE 37.1 (*). Prove Propositions 37.3 and 37.4 concerning the definition of the topology of a CW-complex.

EXERCISE 37.2. 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 . (Note that this statement follows immediately from Proposition 37.16, but it is worth trying to do the exercise independently of this, in order to develop some intuition as to why Proposition 37.16 is true.)

Hint: Given $x \in S^\infty \subset S^\infty$ 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 .

EXERCISE 37.3. Convince yourself that if K is a simplicial complex and its polyhedron $X = |K|$ is viewed as a CW-complex as in Example 37.7, then its cellular chain complex $C_*^{\text{CW}}(X)$ is isomorphic to its oriented simplicial chain complex $C_*^\Delta(K)$.




FIGURE 18. The two cell decompositions of the Klein bottle that are considered in Exercise 37.4.

EXERCISE 37.4. Figure 18 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.

- Compute $H_*^{\text{CW}}(X; \mathbb{Z})$, $H_*^{\text{CW}}(X; \mathbb{Z}_2)$ and $H_*^{\text{CW}}(X; \mathbb{Q})$ for both complexes. (You'll know you've done something wrong if the answers you get from the two complexes are not isomorphic!)
- Recall that the **rank** (*Rang*) of a finitely generated abelian group G is the unique integer $k \geq 0$ 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

$$\sum_k (-1)^k \text{rank } H_k^{\text{CW}}(X; \mathbb{Z}) = \sum_k (-1)^k \dim_{\mathbb{Z}_2} H_k^{\text{CW}}(X; \mathbb{Z}_2) = \sum_k (-1)^k \dim_{\mathbb{Q}} H_k^{\text{CW}}(X; \mathbb{Q}) = 0.$$

(Congratulations, you've just computed the Euler characteristic of the Klein bottle! A comprehensive discussion of this invariant is coming up in Lecture 40.)

EXERCISE 37.5. The **complex projective n -space** $\mathbb{C}\mathbb{P}^n$ is a compact $2n$ -manifold defined as the set of all complex lines through the origin in \mathbb{C}^{n+1} , or equivalently,

$$\mathbb{C}\mathbb{P}^n = (\mathbb{C}^{n+1} \setminus \{0\}) / \sim$$

where two points $z, z' \in \mathbb{C}^{n+1} \setminus \{0\}$ are equivalent if and only if $z' = \lambda z$ for some $\lambda \in \mathbb{C}$. It is conventional to write elements of $\mathbb{C}\mathbb{P}^n$ in so-called *homogeneous coordinates*, meaning the equivalence class represented by $(z_0, \dots, z_n) \in \mathbb{C}^{n+1}$ is written as $[z_0 : \dots : z_n]$. Notice that $\mathbb{C}\mathbb{P}^n$ can be partitioned into two disjoint subsets

$$\mathbb{C}^n \cong \{[1 : z_1 : \dots : z_n] \in \mathbb{C}\mathbb{P}^n\} \quad \text{and} \quad \mathbb{C}\mathbb{P}^{n-1} \cong \{[0 : z_1 : \dots : z_n] \in \mathbb{C}\mathbb{P}^n\}.$$

- Show that the partition $\mathbb{C}\mathbb{P}^n = \mathbb{C}^n \cup \mathbb{C}\mathbb{P}^{n-1}$ gives rise to a cell decomposition of $\mathbb{C}\mathbb{P}^n$ with one $2k$ -cell for every $k = 0, \dots, n$.
- Compute $H_*(\mathbb{C}\mathbb{P}^n)$ and $H^*(\mathbb{C}\mathbb{P}^n)$ for an arbitrary coefficient group.
Hint: This is easy.

38. Invariance of cellular homology

Our goal in this lecture is to prove Theorem 37.22 for finite-dimensional CW-complexes: more precisely, given any axiomatic homology theory h_* with coefficients $G = h_0(\{*\})$ and a CW-complex X that is equal to its N -skeleton X^N for some $N \in \mathbb{N}$, we will establish an isomorphism $h_n(X) \cong H_n^{\text{CW}}(X; G)$ for every $n \geq 0$. Dropping the condition $\dim X < \infty$ will then be the main objective of the next lecture.

38.1. The isomorphism $H_*^{CW}(X) \cong h_*(X)$. The key idea behind the proof of Theorem 37.22 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) . It will turn out that the latter chain complex can be identified naturally with $C_*^{CW}(X)$.

LEMMA 38.1. *For all $n \in \mathbb{N}$, (X^n, X^{n-1}) is a good pair in the sense of Definition 34.3.*

PROOF. Since $X^n = X^{n-1} \cup_{\varphi^n} \coprod_{\alpha \in \mathcal{K}^n} \mathbb{D}^n$, it suffices to set $V := X^{n-1} \cup_{\varphi^n} \coprod_{\alpha \in \mathcal{K}^n} (\mathbb{D}^n \setminus \{0\})$. □

By Theorem 34.7, we now have a natural isomorphism $h_*(X^n, X^{n-1}) \cong \tilde{h}_*(X^n/X^{n-1})$ for each $n \geq 1$. Observe next that the disjoint union of the characteristic maps of n -cells defines a map of pairs

$$\Phi^n := \coprod_{\alpha \in \mathcal{K}^n} \Phi_\alpha : \coprod_{\alpha \in \mathcal{K}^n} (\mathbb{D}^n, \partial\mathbb{D}^n) \rightarrow (X^n, X^{n-1}).$$

We claim that this map descends to a *homeomorphism* between the quotients

$$\Phi^n : \coprod_{\alpha \in \mathcal{K}^n} \mathbb{D}^n / \coprod_{\alpha \in \mathcal{K}^n} \partial\mathbb{D}^n \xrightarrow{\cong} X^n/X^{n-1}.$$

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: $X^n \setminus X^{n-1}$ is the union of all the n -cells, which Φ^n identifies with copies of \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:

$$\begin{array}{ccc} h_*\left(\coprod_{\alpha \in \mathcal{K}^n} (\mathbb{D}^n, \partial\mathbb{D}^n)\right) & \xrightarrow{\Phi_*^n} & h_*(X^n, X^{n-1}) \\ \downarrow \cong & & \downarrow \cong \\ \tilde{h}_*\left(\coprod_{\alpha \in \mathcal{K}^n} \mathbb{D}^n / \coprod_{\alpha \in \mathcal{K}^n} \partial\mathbb{D}^n\right) & \xrightarrow{\Phi_*^n} & \tilde{h}_*(X^n/X^{n-1}). \end{array}$$

Applying the additivity axiom (in conjunction with the five-lemma as in Exercise 28.8) to identify $h_*\left(\coprod_{\alpha \in \mathcal{K}^n} (\mathbb{D}^n, \partial\mathbb{D}^n)\right)$ with $\bigoplus_{\alpha \in \mathcal{K}^n} h_*(\mathbb{D}^n, \partial\mathbb{D}^n)$, this proves:

LEMMA 38.2. *The characteristic maps $\Phi_\alpha : (\mathbb{D}^n, \partial\mathbb{D}^n) \rightarrow (X^n, X^{n-1})$ 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

$$h_k(\mathbb{D}^n, \partial\mathbb{D}^n) \xrightarrow{\ell_*} \tilde{h}_{k-1}(S^{n-1}) \cong \begin{cases} G & \text{if } k = n, \\ 0 & \text{if } k \neq n \end{cases}$$

are isomorphisms for all k and n , thus we've proved

$$(38.1) \quad h_k(X^n, X^{n-1}) \cong \begin{cases} C_n^{CW}(X; G) = \bigoplus_{\alpha \in \mathcal{K}^n} G & \text{if } k = n, \\ 0 & \text{if } k \neq n. \end{cases}$$

We've been assuming $n \geq 1$ so far, but it is not hard to incorporate $n = 0$ into this discussion: if we set

$$X^{-1} := \emptyset,$$

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

$$X^0 = \coprod_{\alpha \in \mathcal{K}^0} \{*\},$$

so that (38.1) 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 $C_n^{\text{CW}}(X; G)$ in our proof of Theorem 38.10.

The plan going forward is to use the Eilenberg-Steenrod axioms to construct a boundary operator 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 $\partial : C_*^{\text{CW}}(X) \rightarrow C_{*-1}^{\text{CW}}(X)$ 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

$$(38.2) \quad 0 = h_{k+1}(X^n, X^{n-1}) \rightarrow h_k(X^{n-1}) \rightarrow h_k(X^n) \rightarrow h_k(X^n, X^{n-1}) = 0,$$

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:

$$h_k(X^n) \cong h_k(X^{n-1}) \cong \dots \cong h_k(X^0) \cong \bigoplus_{\alpha \in \mathcal{K}^0} h_k(\{*\}) = 0,$$

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 38.3. *For every $k > n$, $h_k(X^n) = 0$.* □

Similarly, starting with $k < n$ and applying (38.2) repeatedly to increase n gives:

LEMMA 38.4. *For every $k < n$, the inclusions $X^n \hookrightarrow X^{n+1} \hookrightarrow X^{n+2} \hookrightarrow \dots$ induce isomorphisms $h_k(X^n) \cong h_k(X^{n+1}) \cong h_k(X^{n+2}) \cong \dots$* □

We can now proceed to the heart of the proof of Theorem 37.22. We define for each $n \geq 1$ a map

$$\beta_n : h_n(X^n, X^{n-1}) \rightarrow h_{n-1}(X^{n-1}, X^{n-2})$$

by combining the long exact sequences of the pairs (X^n, X^{n-1}) and (X^{n+1}, X^n) in the following diagram:

$$(38.3) \quad \begin{array}{ccccccc} & & & & h_n(X^{n-1}) & \overset{=}{=} & 0 \\ & & & & \downarrow & & \\ h_{n+1}(X^n) & \rightarrow & h_{n+1}(X^{n+1}) & \xrightarrow{j_{n+1}} & h_{n+1}(X^{n+1}, X^n) & \xrightarrow{\partial_{n+1}} & h_n(X^n) & \xrightarrow{i_n} & h_n(X^{n+1}) & \rightarrow & h_n(X^{n+1}, X^n) \\ & & & \searrow \beta_{n+1} & \downarrow j_n & & \downarrow & & & & \parallel \\ & & & & h_n(X^n, X^{n-1}) & & & & & & 0 \\ & & & & \downarrow \partial_n & & & & & & \\ & & & & h_{n-1}(X^{n-1}) & & & & & & \end{array}$$

In other words, we define $\beta_{n+1} := j_n \circ \partial_{n+1}$ for each $n \geq 0$, and of course $\beta_0 := 0$. (We can use the convention $X^{-1} := \emptyset$ 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 $n \geq 1$, 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

$$(38.4) \quad \dots \rightarrow h_n(X^n, X^{n-1}) \xrightarrow{\beta_n} h_{n-1}(X^{n-1}, X^{n-2}) \rightarrow \dots \rightarrow h_1(X^1, X^0) \xrightarrow{\beta_1} h_0(X^0) \xrightarrow{\beta_0} 0 \rightarrow \dots$$

as a chain complex whose individual chain groups are canonically isomorphic to the chain groups in $C_*^{CW}(X)$. 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

$$(38.5) \quad h_n(X^n) / \ker i_n \xrightarrow{\cong} h_n(X^{n+1}).$$

Second, j_{n-1} is injective, thus

$$\ker \beta_n = \ker(j_{n-1} \circ \partial_n) = \ker \partial_n = \text{im } j_n,$$

and since j_n is also injective, it maps $h_n(X^n)$ isomorphically to $\ker \beta_n$. Moreover, it maps the subgroup $\ker i_n = \text{im } \partial_{n+1}$ isomorphically to $\text{im } \beta_{n+1}$, implying that j_n descends to an isomorphism

$$(38.6) \quad h_n(X^n) / \ker i_n \xrightarrow{\cong} \ker \beta_n / \text{im } \beta_{n+1}.$$

The latter is of course the n th 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 $n \geq 0$ we can then take $N \geq n + 1$ without loss of generality, and use Lemma 38.4 to conclude via (38.5) and (38.6) that

$$\ker \beta_n / \text{im } \beta_{n+1} \cong h_n(X^{n+1}) \cong h_n(X^{n+2}) \cong \dots \cong h_n(X^N) = h_n(X).$$

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 $h_*(X) \cong H_*^{CW}(X)$ 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

$$\begin{array}{ccc} C_n^{\text{CW}}(X) & \xrightarrow{\cong} & h_n(X^n, X^{n-1}) \\ \downarrow \partial & & \downarrow \beta_n \\ C_{n-1}^{\text{CW}}(X) & \xrightarrow{\cong} & h_{n-1}(X^{n-1}, X^{n-2}) \end{array}$$

commutes for every n , where the horizontal maps are the canonical isomorphisms that we discussed above. The result that $\partial^2 = 0$ on $C_*^{\text{CW}}(X)$ will also follow from this, since we already know $\beta_{n-1} \circ \beta_n = 0$.

Here is a useful observation: the characteristic maps $\Phi_\alpha : (\mathbb{D}^n, \partial\mathbb{D}^n) \rightarrow (X^n, X^{n-1})$ also induce maps of quotients $\mathbb{D}^n/\partial\mathbb{D}^n \rightarrow X^n/X^{n-1}$ such that the direct sum of the induced map on reduced homology

$$(38.7) \quad \bigoplus_{\alpha \in \mathcal{K}^n} (\Phi_\alpha)_* : \bigoplus_{\alpha \in \mathcal{K}^n} \tilde{h}_n(\mathbb{D}^n/\partial\mathbb{D}^n) \rightarrow \tilde{h}_n(X^n/X^{n-1})$$

is an isomorphism. Indeed, under the natural isomorphisms between relative homology for good pairs and reduced homology of quotients, this is equivalent to Lemma 38.2. The advantage of rewriting this map in terms of quotients is, however, that we can explicitly write down its inverse. We recall the projections $p_\alpha : X^n \rightarrow X^n/(X^n \setminus e_\alpha^n) = \mathbb{D}^n/\partial\mathbb{D}^n$ 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

$$p_\alpha : X^n/X^{n-1} \rightarrow \mathbb{D}^n/\partial\mathbb{D}^n.$$

LEMMA 38.5. *The inverse of the map (38.7) is*

$$\prod_{\alpha \in \mathcal{K}^n} (p_\alpha)_* : \tilde{h}_n(X^n/X^{n-1}) \rightarrow \bigoplus_{\alpha \in \mathcal{K}^n} \tilde{h}_n(\mathbb{D}^n/\partial\mathbb{D}^n).$$

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 $\bigoplus_\alpha \tilde{h}_n(\mathbb{D}^n/\partial\mathbb{D}^n)$. This follows from the fact that $p_\alpha \circ \Phi_\alpha : \mathbb{D}^n/\partial\mathbb{D}^n \rightarrow \mathbb{D}^n/\partial\mathbb{D}^n$ is the identity map and thus induces the identity on $\tilde{h}_n(\mathbb{D}^n/\partial\mathbb{D}^n)$, 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 $\tilde{h}_n(\mathbb{D}^n/\partial\mathbb{D}^n)$ is trivial. \square

Here's a diagram to help us understand what β_n has to do with the cellular boundary map:

$$(38.8) \quad \begin{array}{ccccc} & & C_n^{\text{CW}}(X) & & C_{n-1}^{\text{CW}}(X) \\ & & \parallel & & \parallel \\ \bigoplus_{\alpha \in \mathcal{K}^n} h_n(\mathbb{D}^n, \partial\mathbb{D}^n) & \xrightarrow{\cong} & \bigoplus_{\alpha \in \mathcal{K}^n} \tilde{h}_{n-1}(\partial\mathbb{D}^n) & \xrightarrow{\partial^{\text{CW}}} & \bigoplus_{\beta \in \mathcal{K}^{n-1}} \tilde{h}_{n-1}(\mathbb{D}^{n-1}/\partial\mathbb{D}^{n-1}) \\ \cong \downarrow \bigoplus_\alpha (\Phi_\alpha)_* & & \downarrow \bigoplus_\alpha (p_\alpha)_* & & \uparrow \prod_\beta (p_\beta)_* \cong \\ h_n(X^n, X^{n-1}) & \xrightarrow{\partial_n} & \tilde{h}_{n-1}(X^{n-1}) & \xrightarrow{q_*} & \tilde{h}_{n-1}(X^{n-1}/X^{n-2}) \\ & \searrow \beta_n & \downarrow j_{n-1} & & \downarrow \cong \\ & & h_{n-1}(X^{n-1}, X^{n-2}) & \xrightarrow{\cong} & h_{n-1}(X^{n-1}/X^{n-2}, X^{n-2}/X^{n-2}) \end{array}$$

The following details deserve clarification:

- The map labeled q_* is induced by the quotient projection $q : X^{n-1} \rightarrow X^{n-1}/X^{n-2}$.

- Regarding the same quotient projection as a map of pairs produces the horizontal map at the bottom, which we proved in Theorem 34.7 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 $h_{n-1}(X^{n-1}, X^{n-2}) \xrightarrow{\cong} \tilde{h}_{n-1}(X^{n-1}/X^{n-2})$.
- We have replaced $h_{n-1}(X^{n-1})$ with $\tilde{h}_{n-1}(X^{n-1})$ 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 reduced homology anyway.
- The diagram is intended to serve as a **definition** of the map $\partial^{\text{CW}} : C_n^{\text{CW}}(X) \rightarrow C_{n-1}^{\text{CW}}(X)$, i.e. it is what $\beta_n : h_n(X^n, X^{n-1}) \rightarrow h_{n-1}(X^{n-1}, X^{n-2})$ 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 $h_{n-1}(X^{n-1}, X^{n-2})$ of β_n with $\tilde{h}_{n-1}(X^{n-1}/X^{n-2})$ so that we can then use Lemma 38.5 to identify the latter with $C_{n-1}^{\text{CW}}(X)$ via an explicit formula. The resulting formula for ∂^{CW} is

$$\prod_{\beta \in \mathcal{K}^{n-1}} (p_\beta)_* \circ \bigoplus_{\alpha \in \mathcal{K}^n} (\varphi_\alpha)_* : \bigoplus_{\alpha \in \mathcal{K}^n} \tilde{h}_{n-1}(S^{n-1}) \rightarrow \bigoplus_{\beta \in \mathcal{K}^{n-1}} \tilde{h}_{n-1}(S^{n-1}).$$

This is determined by the collection of endomorphisms of $\tilde{h}_{n-1}(S^{n-1})$ induced by $p_\beta \circ \varphi_\alpha$ for all $\alpha \in \mathcal{K}^n$ and $\beta \in \mathcal{K}^{n-1}$, and by Theorem 36.17, 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 $\partial^2 = 0$.

The proof of Theorem 37.22 in the case $\dim X < \infty$ is now complete up to one minor quibble: not every detail in the presentation above made sense in the case $n = 0$, in particular where the quotient $\mathbb{D}^n/\partial\mathbb{D}^n$ was mentioned. But these discrepancies are easy to fix; see Exercise 38.2.

38.2. The relative case. 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 $C_*^{\text{CW}}(A)$ is a subcomplex of $C_*^{\text{CW}}(X)$, i.e. it is a subgroup preserved by the boundary map, giving rise to a quotient chain complex

$$C_*^{\text{CW}}(X, A) = C_*^{\text{CW}}(X, A; G) := C_*^{\text{CW}}(X)/C_*^{\text{CW}}(A).$$

The homology of this complex is the **relative cellular homology**

$$H_*^{\text{CW}}(X, A) = H_*^{\text{CW}}(X, A; G) := H_*(C_*^{\text{CW}}(X, A; G)).$$

Theorem 37.22 generalizes to this setting as follows.

THEOREM 38.6. *For any CW-pair (X, A) and any axiomatic homology theory h_* with coefficient group G , there is a natural isomorphism $H_*^{\text{CW}}(X, A; G) \cong h_*(X, A)$.*

The proof of this theorem in the finite-dimensional case is Exercise 38.1. It is a somewhat lengthy exercise, but it is not fundamentally difficult—every step is simply a minor generalization of something that we saw in the proof of Theorem 37.22, and working through it is one of the best ways to achieve a deeper understanding of the isomorphism $H_*^{\text{CW}}(X) \cong h_*(X)$.

38.3. Cellular maps and naturality. By this point you should not be surprised to learn that one can define a category CW^{rel} whose objects are CW-pairs, such that $C_*^{\text{CW}} : \text{CW}^{\text{rel}} \rightarrow \text{Ch}(R\text{-Mod})$ and $H_n^{\text{CW}} : \text{CW}^{\text{rel}} \rightarrow R\text{-Mod}$ become functors. But I still need to tell you what the morphisms in CW^{rel} are.

DEFINITION 38.7. A continuous map $f : X \rightarrow Y$ between CW-complexes is called a **cellular map** (*zelluläre Abbildung*) if $f(X^n) \subset Y^n$ for every $n \geq 0$. More generally, if (X, A) and (Y, B) are CW-pairs, a **map of CW-pairs** is a cellular map $f : X \rightarrow Y$ such that $f(A) \subset B$. (Its restriction $f|_A : A \rightarrow B$ is then automatically a cellular map.)

EXAMPLE 38.8. If X and Y are polyhedra (and therefore also CW-complexes as explained in Example 37.7), then any simplicial map $f : X \rightarrow Y$ is also a cellular map.

In contrast to simplicial maps, a cellular map $f : X \rightarrow Y$ need not generally map cells of X to cells of Y . Instead, the image of an individual cell $e_\alpha^n \subset X$ may cover many n -cells $e_\beta^n \subset Y$, 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_α^0 is just a single point, $f(e_\alpha^0)$ is always a specific 0-cell in Y , so that for each 0-cell e_β^0 of Y we can define

$$[e_\beta^0 : e_\alpha^0] := \begin{cases} 1 & \text{if } f(e_\alpha^0) = e_\beta^0, \\ 0 & \text{otherwise.} \end{cases}$$

For $n \geq 1$, 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, $X^n/X^{n-1} \rightarrow Y^n/Y^{n-1}$ and we can therefore consider the composition

$$(38.9) \quad S^n \cong \mathbb{D}^n/\partial\mathbb{D}^n \xrightarrow{\Phi_\alpha} X^n/X^{n-1} \xrightarrow{f} Y^n/Y^{n-1} \xrightarrow{\text{pr}} Y^n/(Y^n \setminus e_\beta^n) \xrightarrow{\Phi_\beta^{-1}} \mathbb{D}^n/\partial\mathbb{D}^n \cong S^n,$$

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

$$[e_\beta^n : e_\alpha^n] \in \mathbb{Z}.$$

This incidence number vanishes whenever $e_\beta^n \cap f(\overline{e_\alpha^n}) = \emptyset$ since the map in (38.9) is in this case constant, so Proposition 37.16 implies that for each individual $e_\alpha^n \subset X$, there are at most finitely many $e_\beta^n \subset Y$ with $[e_\beta^n : e_\alpha^n] \neq 0$. This allows us to define a homomorphism

$$f_* : C_*^{\text{CW}}(X) \rightarrow C_*^{\text{CW}}(Y)$$

acting on the generators $e_\alpha^n \in C_n^{\text{CW}}(X)$ as

$$(38.10) \quad f_* e_\alpha^n = \sum_{e_\beta^n} [e_\beta^n : e_\alpha^n] e_\beta^n,$$

where the sum ranges over all n -cells $e_\beta^n \subset Y$ and has only finitely many nonzero terms.

As a sanity check, it is a simple but worthwhile exercise to show that if X and Y are the same CW-complex and $f : X \rightarrow Y$ is the identity map, the incidence number $[e_\beta^n : e_\alpha^n]$ is 1 for $\alpha = \beta$ and 0 otherwise, so in particular, $f_* : C_*^{\text{CW}}(X) \rightarrow C_*^{\text{CW}}(Y)$ is then the identity homomorphism. Another worthwhile exercise is to show that if $f : X \rightarrow Y$ and $g : Y \rightarrow Z$ are cellular maps, then

$$(g \circ f)_* = g_* \circ f_* : C_*^{\text{CW}}(X) \rightarrow C_*^{\text{CW}}(Z).$$

This discussion of induced maps extends in an obvious way to the relative case: if $f : (X, A) \rightarrow (Y, B)$ is a map of CW-pairs, then f_* maps $C_*^{\text{CW}}(A)$ into $C_*^{\text{CW}}(B)$ and thus descends to a homomorphism

$$f_* : C_*^{\text{CW}}(X, A) \rightarrow C_*^{\text{CW}}(Y, B).$$

The proof of the next theorem will arise naturally from the proof of the much bigger theorem that follows it.

THEOREM 38.9. *For any map of CW-pairs $f : (X, A) \rightarrow (Y, B)$, $f_* : C_*^{\text{CW}}(X, A) \rightarrow C_*^{\text{CW}}(Y, B)$ is a chain map and thus induces homomorphisms $f_* : H_n^{\text{CW}}(X, A) \rightarrow H_n^{\text{CW}}(Y, B)$ for every n . In particular, the cellular chain complex and cellular homology with coefficients in any given R -module G define functors*

$$C_*^{\text{CW}} = C_*^{\text{CW}}(\cdot; G) : \text{CW}^{\text{rel}} \rightarrow \text{Ch}(R\text{-Mod}), \quad \text{and} \quad H_n^{\text{CW}} = H_n^{\text{CW}}(\cdot; G) : \text{CW}^{\text{rel}} \rightarrow R\text{-Mod},$$

where CW^{rel} denotes the category of CW-pairs, with morphisms defined as maps of CW-pairs.

We can now clarify the meaning of the word “natural” in Theorems 37.22 and 38.6.

THEOREM 38.10. *Suppose h_* is an axiomatic homology theory with coefficient group G . Then the isomorphisms $H_n^{\text{CW}}(X, A; G) \cong h_n(X, A)$ for CW-pairs $(X, A) \in \text{CW}^{\text{rel}}$ are natural in the sense that for any map of CW-pairs $f : (X, A) \rightarrow (Y, B)$, the diagram*

$$\begin{array}{ccc} H_n^{\text{CW}}(X, A; G) & \xrightarrow[\cong]{\Psi_{(X,A)}} & h_n(X, A) \\ \downarrow f_* & & \downarrow f_* \\ H_n^{\text{CW}}(Y, B; G) & \xrightarrow[\cong]{\Psi_{(Y,B)}} & h_n(Y, B) \end{array}$$


commutes.

In the language of category theory, this theorem says the following. There is a functor $\text{CW}^{\text{rel}} \rightarrow \text{Top}^{\text{rel}}$ 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_n with this functor for any $n \in \mathbb{Z}$ produces a functor $\text{CW}^{\text{rel}} \rightarrow R\text{-Mod}$. Theorems 38.6 and 38.10 together define a natural transformation from $H_n^{\text{CW}}(\cdot; G)$ to the latter functor, associating to every CW-pair (X, A) the isomorphism $\Psi_{(X,A)}$.

For the proof of Theorem 38.10, we shall focus on the absolute case and leave the relative case as an exercise.

If $f : X \rightarrow Y$ is a cellular map, then it defines a map of pairs $(X^n, X^{n-1}) \rightarrow (Y^n, Y^{n-1})$ 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

$$\begin{array}{ccc} h_{n+1}(X^{n+1}, X^n) & \xrightarrow{\beta_{n+1}} & h_n(X^n, X^{n-1}) \\ \downarrow f_* & & \downarrow f_* \\ h_{n+1}(Y^{n+1}, Y^n) & \xrightarrow{\beta_{n+1}} & h_n(Y^n, Y^{n-1}), \end{array}$$

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 $H_*^{CW}(X) \rightarrow H_*^{CW}(Y)$. To relate this to the map $f_* : h_*(X) \rightarrow h_*(Y)$, recall that the isomorphism $H_n^{CW}(X) = \ker \beta_n / \text{im } \beta_{n+1} \cong h_n(X)$ 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

$$\begin{array}{ccc} H_n^{CW}(X) & \xrightarrow{\cong} & h_n(X) \\ \downarrow f_* & & \downarrow f_* \\ H_n^{CW}(Y) & \xrightarrow{\cong} & h_n(Y). \end{array}$$

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 $C_n^{CW}(X)$ and $C_n^{CW}(Y)$ respectively, the map $f_* : h_n(X^n, X^{n-1}) \rightarrow h_n(Y^n, Y^{n-1})$ matches the formula we gave in (38.10) for maps $C_n^{CW}(X) \rightarrow C_n^{CW}(Y)$ induced by cellular maps. This will prove simultaneously the statement that the homomorphism in (38.10)

is a chain map. Here is the analogue of the diagram (38.8) for the situation at hand:

$$(38.11) \quad \begin{array}{ccc} C_n^{\text{CW}}(X) & & C_n^{\text{CW}}(Y) \\ \parallel & & \parallel \\ \bigoplus_{e_\alpha^n \subset X} \tilde{h}_n(\mathbb{D}^n / \partial \mathbb{D}^n) & \xrightarrow{f_*^{\text{CW}}} & \bigoplus_{e_\beta^n \subset Y} \tilde{h}_n(\mathbb{D}^n / \partial \mathbb{D}^n) \\ \cong \downarrow \bigoplus_\alpha (\Phi_\alpha)_* & & \uparrow \prod_\beta (p_\beta)_* \cong \\ \tilde{h}_n(X^n / X^{n-1}) & \xrightarrow{f_*} & \tilde{h}_n(Y^n / Y^{n-1}) \end{array}$$

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 $f_*^{\text{CW}} : C_n^{\text{CW}}(X) \rightarrow C_n^{\text{CW}}(Y)$, which is equivalent to $f_* : h_n(X^n, X^{n-1}) \rightarrow h_n(Y^n, Y^{n-1})$ under the canonical isomorphisms. It produces the formula

$$f_*^{\text{CW}} = \prod_{e_\beta^n \subset Y} (p_\beta)_* \circ f_* \circ \bigoplus_{e_\alpha^n \subset X} (\Phi_\alpha)_* : \bigoplus_{e_\alpha^n \subset X} \tilde{h}_n(S^n) \rightarrow \bigoplus_{e_\beta^n \subset Y} \tilde{h}_n(S^n),$$

and this map is determined by the set of all its “matrix elements”

$$(p_\beta)_* \circ f_* \circ (\Phi_\alpha)_* = (p_\beta \circ f \circ \Phi_\alpha)_* : \tilde{h}_n(S^n) \rightarrow \tilde{h}_n(S^n)$$

for each individual $e_\alpha^n \subset X$ and $e_\beta^n \subset Y$. Applying Theorem 36.17 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 (38.10) for $f_* : C_n^{\text{CW}}(X) \rightarrow C_n^{\text{CW}}(Y)$.

The proof of Theorem 38.10 is now complete except for one detail that is outsourced to Exercise 38.2.

The major unresolved issue we still have is the simplifying assumption $\dim X < \infty$ that was 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.

38.4. Exercises.

EXERCISE 38.1 (*). Prove Theorem 38.6 on the isomorphism $h_*(X, A) \cong H_*^{\text{CW}}(X, A; G)$ for finite-dimensional CW-pairs (X, A) and any axiomatic homology theory h_* with coefficients G .

Hint: Start by showing that $C_n^{\text{CW}}(X, A; G)$ 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 $(X^n \cup A, X^{n-1} \cup A, A)$.

EXERCISE 38.2. Some portions of the proofs of Theorems 37.22 and 38.10 given above did not make sense for $n = 0$, especially when $\mathbb{D}^n / \partial \mathbb{D}^n$ was mentioned. Adapt the discussion as needed for that particular case.

39. Direct limits and infinite-dimensional cell complexes

If X is an infinite-dimensional CW-complex, then the arguments of the previous lecture do not suffice to prove $H_*^{\text{CW}}(X; G) \cong h_*(X)$ for every axiomatic homology theory h_* with coefficient group G . What they do prove is that for every integer $n \geq 0$, there are isomorphisms

$$H_n^{\text{CW}}(X; G) \cong h_n(X^{n+1}) \cong h_n(X^{n+2}) \cong h_n(X^{n+3}) \cong \dots,$$

where the maps $h_n(X^{n+k}) \rightarrow h_n(X^{n+k+1})$ 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 $f : X \rightarrow Y$, the

induced homomorphism $f_* : H_n^{\text{CW}}(X; G) \rightarrow H_n^{\text{CW}}(Y; G)$ fits into a commutative diagram

$$\begin{array}{ccccccc} H_n^{\text{CW}}(X; G) & \xrightarrow{\cong} & h_n(X^{n+1}) & \xrightarrow{\cong} & h_n(X^{n+2}) & \xrightarrow{\cong} & h_n(X^{n+3}) \xrightarrow{\cong} \dots \\ \downarrow f_* & & \downarrow f_* & & \downarrow f_* & & \downarrow f_* \\ H_n^{\text{CW}}(Y; G) & \xrightarrow{\cong} & h_n(Y^{n+1}) & \xrightarrow{\cong} & h_n(Y^{n+2}) & \xrightarrow{\cong} & h_n(Y^{n+3}) \xrightarrow{\cong} \dots \end{array}$$

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 $X^0, X^1, X^2, \dots, X^n, \dots$ as $n \rightarrow \infty$. If the functor h_n could be shown to be “continuous” with respect to such limits, we would conclude

$$(39.1) \quad h_n(X) = h_n\left(\lim_{k \rightarrow \infty} X^k\right) = \lim_{k \rightarrow \infty} h_n(X^k),$$

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 $H_n^{\text{CW}}(X; G)$. 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 $\{X^k\}_{k=0}^\infty$, and similarly for a sequence of abelian groups or R -modules such as $\{h_n(X^k)\}_{k=0}^\infty$. We will then see that the continuity relation (39.1) really does hold for any sequence X^k consisting of the skeleta of a CW-complex.

39.1. Direct systems, targets and limits. Suppose J is a set with a pre-order $<$, i.e. $<$ is reflexive ($\alpha < \alpha$) and transitive ($\alpha < \beta$ and $\beta < \gamma$ implies $\alpha < \gamma$), but the relations $\alpha < \beta$ and $\beta < \alpha$ need not imply $\alpha = \beta$, so $<$ need not be a partial order. Recall that $(J, <)$ is called a **directed set** (*gerichtete Menge*) if for every pair $\alpha, \beta \in J$, there exists $\gamma \in J$ with $\gamma > \alpha$ and $\gamma > \beta$. The most common directed set in our examples will be (\mathbb{N}, \leq) , or sometimes (\mathbb{N}_0, \leq) 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 39.5 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 \mathcal{C} , a **direct system** (*induktives System*) $\{X_\alpha, \varphi_{\beta\alpha}\}$ in \mathcal{C} over the directed set $(J, <)$ associates to each $\alpha \in J$ an object X_α of \mathcal{C} , along with morphisms

$$\varphi_{\beta\alpha} \in \text{Hom}(X_\alpha, X_\beta) \quad \text{for each } \alpha < \beta$$

such that

$$\varphi_{\alpha\alpha} = \text{Id}_{X_\alpha}$$

and the diagram

$$\begin{array}{ccccc} X_\alpha & \xrightarrow{\varphi_{\beta\alpha}} & X_\beta & \xrightarrow{\varphi_{\gamma\beta}} & X_\gamma \\ & & & \searrow \varphi_{\gamma\alpha} & \nearrow \\ & & & & \end{array}$$

commutes for every triple $\alpha, \beta, \gamma \in J$ with $\alpha < \beta < \gamma$.

REMARK 39.2. Exercise 27.3 shows that a pre-order $<$ on a set J can be encoded by calling J the collection of objects in a category \mathcal{J} , such that for each pair $x, y \in J$, the set of morphisms $\text{Hom}(x, y)$ contains exactly one element whenever $x < y$ and is otherwise empty. A direct system in \mathcal{C} over $(J, <)$ is then nothing other than a (covariant) functor $\mathcal{J} \rightarrow \mathcal{C}$.

EXAMPLE 39.3. For any CW-complex X , its collection of skeleta $\{X^n\}_{n=0}^\infty$ forms a direct system in \mathbf{Top} over (\mathbb{N}_0, \leq) , with the maps φ_{mn} for each $m \geq n$ defined as the inclusions $X^n \hookrightarrow X^m$. Similarly, the skeleta of a CW-pair define a direct system in $\mathbf{Top}^{\text{rel}}$.

EXAMPLE 39.4. For any axiomatic homology theory h_* taking values in the category of R -modules, the homologies of the skeleta of a CW-complex form direct systems in $R\text{-Mod}$ over (\mathbb{N}_0, \leq) : indeed, for each $k \in \mathbb{Z}$, there is a direct system consisting of the modules $\{h_k(X^n)\}_{n=0}^\infty$ and for every $m \geq n$ the map $h_k(X^n) \rightarrow h_k(X^m)$ induced by the inclusion $X^n \hookrightarrow X^m$.

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 \mathcal{A} over $(J, <)$, and $\mathcal{F} : \mathcal{A} \rightarrow \mathcal{B}$ is a covariant functor, then $\{\mathcal{F}(X_\alpha), \mathcal{F}(\varphi_{\beta\alpha})\}$ forms a direct system in \mathcal{B} over $(J, <)$. \square*

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 $\alpha \in J$ 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 \mathcal{C} over $(J, <)$, a **target** $\{Y, f_\alpha\}$ of the system consists of an object Y of \mathcal{C} together with associated morphisms $f_\alpha \in \text{Hom}(X_\alpha, Y)$ for each $\alpha \in J$ such that the diagram

$$\begin{array}{ccc} X_\alpha & \xrightarrow{\varphi_{\beta\alpha}} & X_\beta \\ & \searrow f_\alpha & \swarrow f_\beta \\ & & Y \end{array}$$

commutes for every pair $\alpha, \beta \in J$ with $\alpha < \beta$.

DEFINITION 39.7. A target $\{X_\infty, \varphi_\alpha\}$ of the direct system $\{X_\alpha, \varphi_{\beta\alpha}\}$ is called a **direct limit**⁵⁹ (*induktiver Limes*) of the system and written as

$$X_\infty = \varinjlim \{X_\alpha\}$$

if it satisfies the following “universal” property: for all targets $\{Y, f_\alpha\}$ of $\{X_\alpha, \varphi_{\beta\alpha}\}$, there exists a unique morphism $f_\infty \in \text{Hom}(X_\infty, Y)$ such that the diagram

$$\begin{array}{ccc} X_\alpha & \xrightarrow{\varphi_\alpha} & X_\infty \\ & \searrow f_\alpha & \downarrow f_\infty \\ & & Y \end{array}$$

commutes for every $\alpha \in J$.

The essential meaning of a direct limit can be encoded in the diagram

$$\begin{array}{ccccccc} X_\alpha & \xrightarrow{\varphi_{\beta\alpha}} & X_\beta & \xrightarrow{\varphi_{\gamma\beta}} & X_\gamma & \longrightarrow & \dots \longrightarrow \varinjlim \{X_\alpha\} \\ & & & & & & \downarrow \text{---} \\ & & & & & & Y \end{array}$$

⁵⁹Direct 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 39.8.

where we assume $\alpha < \beta < \gamma < \dots \in J$. The key feature of the object $\varinjlim\{X_\alpha\}$ is that whenever an object Y and morphisms $X_\alpha \rightarrow Y$ 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, the proof of the following is an easy exercise:

PROPOSITION 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 $\psi \in \text{Hom}(X, Y)$, then $\{Y, \psi \circ f_\alpha\}$ is also a direct limit of $\{X_\alpha, \varphi_{\beta\alpha}\}$. \square*

The non-uniqueness exhibited by the proposition 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 $g_\infty \in \text{Hom}(X, Y)$ and $f_\infty \in \text{Hom}(Y, X)$ satisfying $g_\infty \circ f_\alpha = g_\alpha$ and $f_\infty \circ g_\alpha = f_\alpha$ for every $\alpha \in J$. 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 $\alpha \in J$, which implies $f_\infty \circ g_\infty = \text{Id}_X$. A similar argument shows $g_\infty \circ f_\infty = \text{Id}_Y$, 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 result 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 $h_k(X^0) \rightarrow h_k(X^1) \rightarrow \dots \rightarrow h_k(X^n) \rightarrow h_k(X^{n+1}) \rightarrow \dots$ stabilizes as $n \rightarrow \infty$, 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. The proof is Exercise 39.2.

PROPOSITION 39.9. *Suppose $\{X_\alpha, \varphi_{\beta\alpha}\}$ is a direct system in \mathcal{C} over $(J, <)$ with the property that for some $\alpha_0 \in J$, $\varphi_{\gamma\beta} \in \text{Hom}(X_\beta, X_\gamma)$ is an isomorphism for every $\beta, \gamma \in J$ with $\beta > \alpha_0$ and $\gamma > \alpha_0$. For each $\alpha \in J$, choose $\gamma \in J$ such that $\gamma > \alpha$ and $\gamma > \alpha_0$, and define*

$$\varphi_\alpha := \varphi_{\gamma\alpha_0}^{-1} \circ \varphi_{\gamma\alpha} \in \text{Hom}(X_\alpha, X_{\alpha_0}).$$

Then the morphism φ_α does not depend on the choice of the element $\gamma \in J$, and $\{X_{\alpha_0}, \varphi_\alpha\}$ defines a direct limit of the system, i.e. $X_{\alpha_0} = \varinjlim\{X_\alpha\}$. \square

39.2. Constructions of direct limits. 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.

LEMMA 39.10. *Suppose $\{X_\alpha, \varphi_{\beta\alpha}\}$ is a direct system in \mathcal{C} over $(J, <)$, where \mathcal{C} is any category in which objects are sets (possibly with extra structure) and morphisms are maps between them. For any $\alpha, \beta \in J$, $x \in X_\alpha$ and $y \in X_\beta$, define the relation $x \sim y$ to mean*

$$x \sim y \iff \varphi_{\gamma\alpha}(x) = \varphi_{\gamma\beta}(y) \text{ for some } \gamma \in J \text{ with } \gamma > \alpha \text{ and } \gamma > \beta.$$

Then \sim is an equivalence relation on the set-theoretic disjoint union $\coprod_{\alpha \in J} X_\alpha$.⁶⁰

PROOF. See Exercise 39.3. \square

PROPOSITION 39.11. *If $\{X_\alpha, \varphi_{\beta\alpha}\}$ is a direct system in Top over $(J, <)$, then its direct limit is the space*

$$\varinjlim\{X_\alpha\} = \coprod_{\alpha \in J} X_\alpha / \sim,$$

⁶⁰The **set-theoretic disjoint union** of a collection of sets $\{X_\alpha\}_{\alpha \in J}$ can be defined in general as the set $\{(\alpha, x) \mid \alpha \in J, x \in X_\alpha\}$, 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 $\coprod_{\gamma \in J} X_\gamma$. 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 39.3 is not needed.

where the equivalence relation is defined as in Lemma 39.10, and the associated morphisms $\varphi_\alpha : X_\alpha \rightarrow \varinjlim\{X_\alpha\}$ are the compositions of the inclusions $X_\alpha \hookrightarrow \coprod_{\beta \in J} X_\beta$ with the quotient projection. Moreover, the topology on $\varinjlim\{X_\alpha\}$ is the strongest topology for which the maps $\varphi_\alpha : X_\alpha \rightarrow \varinjlim\{X_\alpha\}$ are continuous for all $\alpha \in J$.

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 $q : \coprod_\beta X_\beta \rightarrow X_\infty$ 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 $\alpha \in J$. 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 $\varphi_\alpha : X_\alpha \rightarrow X_\infty$ is continuous. An easy corollary of this observation is that for any other space Y , a map $f : X_\infty \rightarrow Y$ is continuous if and only if the maps $f \circ \varphi_\alpha : X_\alpha \rightarrow Y$ are continuous for all $\alpha \in J$ (cf. Prop. 37.4).

It is clear that $\{X_\infty, \varphi_\alpha\}$ is a target since for any $\alpha, \beta \in J$ with $\alpha < \beta$, the relation

$$\varphi_\beta \circ \varphi_{\beta\alpha}(x) = \varphi_\alpha(x) \quad \text{for all } x \in X_\alpha$$

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 $f_\infty : X_\infty \rightarrow Y$ satisfying the condition $f_\infty \circ \varphi_\alpha = f_\alpha$ for every $\alpha \in J$. To write down $f_\infty(x)$ for an arbitrary element $x \in X_\infty$, observe that since the quotient projection $q : \coprod_\beta X_\beta \rightarrow X_\infty$ is surjective, we have $x = q(x_\alpha) = \varphi_\alpha(x_\alpha)$ for some $\alpha \in J$ 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 $\beta \in J$ 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 $\gamma \in J$ satisfying $\gamma > \alpha$ and $\gamma > \beta$,

$$\varphi_{\gamma\alpha}(x_\alpha) = \varphi_{\gamma\beta}(x_\beta) =: x_\gamma \in X_\gamma,$$

and thus $f_\gamma(x_\gamma) = f_\alpha(x_\alpha) = f_\beta(x_\beta)$. This proves that a map $f_\infty : X_\infty \rightarrow Y$ with the desired properties is well defined and uniquely determined, though a remark is still required on why f_∞ is continuous: this follows from the previous paragraph since $f_\infty \circ \varphi_\alpha = f_\alpha : X_\alpha \rightarrow Y$ is continuous for every $\alpha \in J$. \square

REMARK 39.12. Proposition 39.11 extends in an obvious way to give a concrete description of any direct limit in the category $\mathbf{Top}^{\text{rel}}$ 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 $n \geq m$ defined by inclusion. Considering the quotient $X^\infty := \coprod_{n=0}^\infty X^n / \sim$ as in Proposition 39.11 along with the natural maps $\varphi_n : X^n \rightarrow X^\infty$, the disjoint union of the inclusion maps $i_n : X^n \hookrightarrow X$ descends to the quotient as a bijection

$$\coprod_{n=0}^\infty i_n : \coprod_{n=0}^\infty X^n / \sim \xrightarrow{\cong} X$$

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 $X^n \rightarrow \varinjlim\{X^n\}$ defined as the inclusions $X^n \hookrightarrow X$. □

We next consider the analogue of Proposition 39.11 in the category $R\text{-Mod}$ of R -modules.

PROPOSITION 39.14. *If $\{G_\alpha, \varphi_{\beta\alpha}\}$ is a direct system in $R\text{-Mod}$ over $(J, <)$, then its direct limit is the module*

$$\varinjlim\{G_\alpha\} = \bigoplus_{\alpha \in J} G_\alpha / H,$$

where $H \subset \bigoplus_{\alpha} G_\alpha$ is the submodule generated by all elements of the form $g - \varphi_{\beta\alpha}(g)$ for $g \in G_\alpha$ and $\beta > \alpha$, and the associated homomorphisms $\varphi_\alpha : G_\alpha \rightarrow \varinjlim\{G_\alpha\}$ are the compositions of the natural inclusions $G_\alpha \hookrightarrow \bigoplus_{\beta} G_\beta$ with the quotient projection.

PROOF. Abbreviating $G_\infty = \bigoplus_{\alpha} G_\alpha / H$, 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 $\beta > \alpha$ implies that the homomorphism

$$\bigoplus_{\alpha \in J} \psi_\alpha : \bigoplus_{\alpha \in J} G_\alpha \rightarrow A$$

vanishes on the submodule H and thus descends to a homomorphism $\psi_\infty : G_\infty \rightarrow A$ that satisfies $\psi_\infty \circ \varphi_\alpha = \psi_\alpha$ for all α . □

REMARK 39.15. A minor variation on Proposition 39.14 gives a similar description of direct limits in the category of chain complexes; see Exercise 39.4.

39.3. The mapping telescope. We have seen above that any CW-complex X can be identified with the direct limit of its skeleta. Combining Proposition 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 R -modules $\{h_k(X^n)\}_{n=0}^\infty$ stabilizes as $n \rightarrow \infty$ and thus has direct limit $h_k(X^n)$ for any n sufficiently large, which matches $H_k^{\text{CW}}(X^n) = H_k^{\text{CW}}(X)$. This gives an isomorphism

$$H_k^{\text{CW}}(X) \cong \varinjlim\{h_k(X^n)\}_{n=0}^\infty.$$

The isomorphism $H_*^{\text{CW}}(X) \cong h_*(X)$ will therefore follow if we can prove that the functors h_k behave “continuously” under this direct limit, i.e. the question becomes

$$\varinjlim\{h_k(X^n)\} \cong h_k(\varinjlim\{X^n\})?$$

It is time to insert a word of caution: the next example shows that singular homology does *not* always behave as nicely as one would hope with respect to direct limits.

EXAMPLE 39.16. Define $\{X_\alpha\}_{\alpha \in J}$ to be the collection of all countable subspaces of S^1 , with a partial order assigned to the index set such that

$$\alpha < \beta \iff X_\alpha \subset X_\beta.$$

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. One can show that $\varinjlim\{X_\alpha\}$ is homeomorphic to S^1 , but $\varinjlim\{H_1(X_\alpha; \mathbb{Z})\}$ is not isomorphic to $H_1(S^1; \mathbb{Z})$; see Exercise 39.5.

To see nonetheless why it might *sometimes* be true that $\varinjlim\{h_k(X_\alpha)\} \cong h_k(\varinjlim\{X_\alpha\})$, 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 \mathcal{A} over $(J, <)$, and $\mathcal{F} : \mathcal{A} \rightarrow \mathcal{B}$ is a covariant functor, thus producing a direct system $\{\mathcal{F}(X_\alpha), \mathcal{F}(\varphi_{\beta\alpha})\}$ in \mathcal{B} . If $\varinjlim\{X_\alpha\}$ exists, then the natural morphisms

$$X_\alpha \xrightarrow{\varphi_\alpha} \varinjlim\{X_\alpha\}$$

called a *homotopy colimit*. For present purposes, it is not necessary to know what this means, but the following easy observation will be important: by the homotopy axiom, applying the functor h_k to the entirety of the homotopy-commutative diagram (39.4) produces a genuine commutative diagram

$$\begin{array}{ccccccc}
 h_k(X^0) & \xrightarrow{f_*^0} & h_k(X^1) & \xrightarrow{f_*^1} & h_k(X^2) & \longrightarrow & \dots \\
 & & & & & \searrow^{i_*^2} & \\
 & & & & & & h_k(T) \\
 & & & & & \swarrow_{i_*^1} & \\
 & & & & & \swarrow_{i_*^0} & \\
 & & & & & &
 \end{array}$$

thus making $h_k(T)$ a target of the direct system of R -modules $\{h_k(X^n)\}_{n=0}^\infty$, so that there is a unique limiting homomorphism

(39.5)
$$\varinjlim \{h_k(X^n)\}_{n=0}^\infty \rightarrow h_k(T)$$

LEMMA 39.18. *The map (39.5) is always an isomorphism.*

PROOF. By Proposition 39.14, we can identify $\varinjlim \{h_k(X^n)\}_{n=0}^\infty$ with the quotient of the direct sum $\bigoplus_{n=0}^\infty h_k(X^n)$ by the submodule generated by all elements of the form $c - f_*^n c \in h_k(X^n) \oplus h_k(X^{n+1})$ for $n \geq 0$ and $c \in h_k(X^n)$. As it turns out, that is precisely the same description that we get for $h_k(T)$ if we compute it using the long exact sequence of the mapping torus established in Lecture 35. Indeed, T can be identified with the mapping torus of the map

$$\prod_{n=0}^\infty X^n \xrightarrow{f} \prod_{n=0}^\infty X^n$$

whose restriction to $X^m \subset \prod_{n=0}^\infty X^n$ for each $m \geq 0$ is $X^m \xrightarrow{f^m} X^{m+1}$. Theorem 35.1 thus gives a long exact sequence

$$\dots \rightarrow h_k\left(\prod_{n \geq 0} X^n\right) \xrightarrow{\Phi_k} h_k\left(\prod_{n \geq 0} X^n\right) \rightarrow h_k(T) \rightarrow h_{k-1}\left(\prod_{n \geq 0} X^n\right) \xrightarrow{\Phi_{k-1}} h_{k-1}\left(\prod_{n \geq 0} X^n\right) \rightarrow \dots,$$

where $\Phi_k := \mathbb{1} - f_*$. One can always turn a long exact sequence into a short exact sequence with a desired term in the middle: in the case at hand, one obtains from this trick a short exact sequence

$$0 \rightarrow \text{coker } \Phi_k \rightarrow h_k(T) \rightarrow \ker \Phi_{k-1} \rightarrow 0.$$

Now, using the additivity axiom, we can identify Φ_k with the map

$$\begin{pmatrix}
 \mathbb{1} & 0 & 0 & 0 & \dots \\
 -f_*^0 & \mathbb{1} & 0 & 0 & \dots \\
 0 & -f_*^1 & \mathbb{1} & 0 & \dots \\
 0 & 0 & -f_*^2 & \mathbb{1} & \dots \\
 0 & 0 & 0 & -f_*^3 & \dots \\
 \vdots & \vdots & \vdots & \vdots & \ddots
 \end{pmatrix} : \bigoplus_{n=0}^\infty h_k(X^n) \rightarrow \bigoplus_{n=0}^\infty h_k(X^n),$$

which is easily seen to be injective for every $k \in \mathbb{Z}$, turning the exact sequence above into an isomorphism

$$\text{coker } \Phi_k \xrightarrow{\cong} h_k(T).$$

But in fact, the image of Φ_k in $\bigoplus_{n=0}^\infty h_k(X^n)$ is precisely the submodule that gets quotiented out in order to construct the direct limit, and this isomorphism is precisely the map (39.5). \square

The mapping telescope T of the sequence $\{X^n, f^n\}$ is a different space from the direct limit $X := \varinjlim \{X^n\}$, and in general $h_k(T)$ and $h_k(X)$ also are not isomorphic. The situation is nicer, however, if we impose a few extra assumptions that hold in the case where $\{X^n\}$ is the sequence

of skeleta of a CW-complex. A simplifying feature in that situation is that the spaces X^n are all subsets of the direct limit X , and the maps $f^n : X^n \rightarrow X^{n+1}$ are all inclusions. The mapping telescope can thus be understood as a subset of $X \times [0, \infty)$,

$$T = (X^0 \times [0, 1]) \cup (X^1 \times [1, 2]) \cup (X^2 \times [2, 3]) \cup \dots \subset X \times [0, \infty).$$

PROPOSITION 39.19. *Assume X is a CW-complex and*

$$A_0 \subset A_1 \subset A_2 \subset \dots \subset \bigcup_{n=0}^{\infty} A_n = X$$

is a nested sequence of subcomplexes such that for every $n \geq 0$, the n -skeleton X^n is contained in A_m for some $m \geq 0$. Then for the mapping telescope $T \subset X \times [0, \infty)$ of the direct system consisting of the spaces $\{A_n\}_{n=0}^{\infty}$ and their inclusions, there exists a deformation retraction of $X \times [0, \infty)$ to T .

The proof of this proposition uses the following lemma.

LEMMA 39.20. *For any CW-pair (X, A) , there exists a deformation retraction of $X \times [0, 1]$ to the subset $(A \times [0, 1]) \cup (X \times \{1\})$.*

PROOF. We start by constructing a deformation retraction of $(X^0 \cup A) \times [0, 1]$ to $(A \times [0, 1]) \cup (X^0 \times \{1\})$. This is easy: one only has to push $e_{\alpha}^0 \times [0, 1]$ to $e_{\alpha}^0 \times \{1\}$ for each isolated 0-cell e_{α}^0 that does not belong to A . One then proceeds inductively: Assuming for some $n \in \mathbb{N}$ that a suitable deformation retraction of $(X^{n-1} \cup A) \times [0, 1]$ to $(A \times [0, 1]) \cup (X^{n-1} \times \{1\})$ has already been constructed, one must extend it so that X^{n-1} is replaced by X^n . This simply requires extending it to each individual n -cell e_{α}^n , which amounts to extending a given deformation retraction of $S^{n-1} \times [0, 1]$ to a deformation retraction of $\mathbb{D}^n \times [0, 1]$, and you will easily convince yourself that this can be done, roughly for the same reasons that $\mathbb{D}^n \times [0, 1]$ admits a deformation retraction to the subset $(\mathbb{D}^n \times \{1\}) \cup (S^{n-1} \times [0, 1])$. The topology of a CW-complex is defined such that any map constructed in this manner by induction on the dimensions of the skeleta is automatically continuous. \square

PROOF OF PROPOSITION 39.19. We define a homotopy $H : I \times X \times [0, \infty) \rightarrow X \times [0, \infty)$ as follows. For $t \in [1/2, 1]$, $H(t, \cdot) : X \times [0, \infty) \rightarrow X \times [0, \infty)$ is a map that fixes $X \times [1, \infty)$, is the identity for $t = 1$, and for $t = 1/2$ is a retraction to $(A_0 \times [0, 1]) \cup (X \times [1, \infty))$ constructed via Lemma 39.20. For $t \in [1/4, 1/2]$, we apply Lemma 39.20 again to construct a further deformation fixing $X \times [2, \infty)$ so that at $t = 1/4$, we have a retraction to $(A_0 \times [0, 1]) \cup (A_1 \times [1, 2]) \cup (X \times [2, \infty))$. Continuing in this manner with an infinite sequence of deformations, one obtains a deformation defined for $t \in (0, 1]$, but it has the property that for each $n \geq 0$, the deformation does not change anything on $A_n \times [0, \infty)$ for $t > 0$ sufficiently close to 0. Since each finite-dimensional skeleton of X is contained in one of the subsets A_n , it follows that the deformation has a unique continuous extension to $t = 0$, and this extension produces a retraction of $X \times [0, \infty)$ to T . \square

PROOF OF THEOREM 39.17. For each $n \geq 0$, we can put the inclusions $X^n \hookrightarrow X$, $i^n : X^n \hookrightarrow T$, $X = X \times \{0\} \hookrightarrow X \times [0, \infty)$ and $T \hookrightarrow X \times [0, \infty)$ together in a diagram

$$\begin{array}{ccc} X^n & \hookrightarrow & X \\ \downarrow i^n & & \downarrow \\ T & \hookrightarrow & X \times [0, \infty) \end{array},$$

which is not strictly commutative but commutes up to homotopy, and Proposition 39.19 tells us that the bottom arrow is a homotopy equivalence, which is also obviously true of the arrow at

the right. By the homotopy axiom, this gives rise to a strictly commutative diagram of homology groups

$$\begin{array}{ccc} h_k(X^n) & \longrightarrow & h_k(X) \\ \downarrow i_*^n & & \downarrow \cong \\ h_k(T) & \xrightarrow{\cong} & h_k(X \times [0, \infty)) \end{array},$$

in which two of the arrows are already known to be isomorphisms, and therefore also a diagram

$$\begin{array}{ccc} \bigoplus_{n=0}^{\infty} h_k(X^n) & \longrightarrow & h_k(X) \\ \bigoplus_n i_*^n \downarrow & & \downarrow \cong \\ h_k(T) & \xrightarrow{\cong} & h_k(X \times [0, \infty)) \end{array}.$$

Both of the maps defined on $\bigoplus_n h_k(X^n)$ in this diagram descend to its quotient by the image of the map Φ_k that appeared in the proof of Lemma 39.18, and this quotient is the direct limit of the homology groups, giving rise to another commutative diagram

$$\begin{array}{ccc} \varinjlim \{h_k(X^n)\}_{n=0}^{\infty} & \longrightarrow & h_k(X) \\ \cong \downarrow & & \downarrow \cong \\ h_k(T) & \xrightarrow{\cong} & h_k(X) \times [0, \infty) \end{array},$$

in which the top arrow is the natural map in the statement of the theorem, while the arrow at the left is the isomorphism in Lemma 39.18. \square

We will not carry out the details here, but this entire discussion can be extended to the setting of a CW-pair (X, A) . Using $h_n(X^n \cup A, X^{n-1} \cup A)$ as substitutes for the relative cellular chain groups $C_n^{\text{CW}}(X, A)$ as sketched in Exercise 38.1, one obtains isomorphisms

$$H_k^{\text{CW}}(X, A) \xrightarrow{\cong} h_k(X^{k+1} \cup A, A) \xrightarrow{\cong} h_k(X^{k+2} \cup A, A) \xrightarrow{\cong} \dots \xrightarrow{\cong} \varinjlim \{h_k(X^n \cup A, A)\}_{n=0}^{\infty} \rightarrow h_k(X, A),$$

in which the homomorphism at the end is determined via the universal property of the direct limit by the maps induced by the inclusions of pairs $(X^n \cup A, A) \hookrightarrow (X, A)$. It is possible to adapt the entire mapping telescope discussion to accommodate CW-pairs and thus show that this last map is an isomorphism. An alternative approach is to write down the obvious long exact sequence in cellular homology that results from the short exact sequence $0 \rightarrow C_*^{\text{CW}}(A) \rightarrow C_*^{\text{CW}}(X) \rightarrow C_*^{\text{CW}}(X, A) \rightarrow 0$, and show that its connecting homomorphisms fit into commutative diagrams together with $\partial_* : h_k(X, A) \rightarrow h_{k-1}(A)$ and the natural maps $H_k^{\text{CW}}(X, A) \rightarrow h_k(X, A)$ and $H_{k-1}^{\text{CW}}(A) \rightarrow h_{k-1}(A)$. This produces a diagram relating the long exact sequences of (X, A) in H_*^{CW} and h_* , together with the natural maps from the former to the latter, and the result in the absolute case then combines with the five-lemma to prove that $H_*^{\text{CW}}(X, A) \rightarrow h_*(X, A)$ is an isomorphism.

Finally, one should also check naturality, i.e. that the homomorphisms $f_* : H_*^{\text{CW}}(X, A) \rightarrow H_*^{\text{CW}}(Y, B)$ induced by cellular maps of pairs $f : (X, A) \rightarrow (Y, B)$ fit together into commutative diagrams with the maps $f_* : h_*(X, A) \rightarrow h_*(Y, B)$ and isomorphisms $H_*^{\text{CW}}(X, A) \cong h_*(X, A)$ and $H_*^{\text{CW}}(Y, B) \cong h_*(Y, B)$. In the absolute case, for a cellular map $f : X \rightarrow Y$, this amounts to

verifying that all squares in the diagram

$$\begin{array}{ccccccc}
 H_k^{\text{CW}}(X) & \xrightarrow{\cong} & h_k(X^{k+1}) & \xrightarrow{\cong} & h_k(X^{k+2}) & \xrightarrow{\cong} & \dots & \xrightarrow{\cong} & \varinjlim \{h_k(X^n)\}_{n=0}^\infty & \xrightarrow{\cong} & h_k(X) \\
 \downarrow f_* & & \downarrow f_* & & \downarrow f_* & & & & \downarrow & & \downarrow f_* \\
 H_k^{\text{CW}}(Y) & \xrightarrow{\cong} & h_k(Y^{k+1}) & \xrightarrow{\cong} & h_k(Y^{k+2}) & \xrightarrow{\cong} & \dots & \xrightarrow{\cong} & \varinjlim \{h_k(Y^n)\}_{n=0}^\infty & \xrightarrow{\cong} & h_k(Y)
 \end{array}$$

commute, where the map between the two direct limits is uniquely determined by the universal property. That the squares on the left commute was established in the proof of Theorem 38.10, so the only real question here involves the rightmost square, but here also, commutativity can be deduced from the uniqueness condition in the universal property.

We shall leave the further details as exercises, and thus regard the proofs of Theorems 37.22 and 38.10 as complete.

39.4. Singular homology of direct limits. The following discussion is not strictly necessary, but you may be interested to know that in the particular case of singular homology, there is also a more hands-on way to prove that the natural map $\varinjlim \{H_k(X^n)\}_{n=0}^\infty \rightarrow H_k(X)$ is an isomorphism for every CW-complex X .

Recall that for each $k \in \mathbb{Z}$, $H_k : \mathbf{Top} \rightarrow R\text{-Mod}$ is in fact the composition of two functors: the first is $C_* : \mathbf{Top} \rightarrow \mathbf{Ch}$, which sends each space X to its singular chain complex with coefficients in a given module G , and the second is $H_k : \mathbf{Ch} \rightarrow R\text{-Mod}$, sending a chain complex to its homology in degree k . The second of these two functors turns out to be extremely well behaved with respect to direct limits.

PROPOSITION 39.21. *Suppose $(J, <)$ is a directed set, with a chain complex C_*^α associated to each $\alpha \in J$ and a chain map $\varphi_{\beta\alpha} : C_*^\alpha \rightarrow C_*^\beta$ associated to each pair $\alpha < \beta \in J$ such that $\{C_*^\alpha, \varphi_{\beta\alpha}\}$ is a direct system in \mathbf{Ch} over $(J, <)$. Then choosing \mathcal{F} to be the functor $H_k : \mathbf{Ch} \rightarrow R\text{-Mod}$ for some $k \in \mathbb{Z}$, the map*

$$\Phi_{\mathcal{F}} : \varinjlim \{H_k(C_*^\alpha)\} \rightarrow H_k(\varinjlim \{C_*^\alpha\})$$

defined as in (39.2) is an isomorphism of R -modules.

The proof of the proposition uses the following consequence of Proposition 39.14, which makes proving things about direct limits of abelian groups or R -modules (or chain complexes thereof) considerably easier.

LEMMA 39.22. *The following statements hold for any direct system $\{G_\alpha, \varphi_{\beta\alpha}\}$ in $R\text{-Mod}$ or $\mathbf{Ch}(R\text{-Mod})$ over a directed set $(J, <)$:*

- (i) *For every $x \in \varinjlim \{G_\alpha\}$, there exists $\beta \in J$ and $x_\beta \in G_\beta$ such that $x = \varphi_\beta(x_\beta)$.*
- (ii) *For every $\beta \in J$ and $x_\beta \in G_\beta$ satisfying $\varphi_\beta(x_\beta) = 0 \in \varinjlim \{G_\alpha\}$, there exists $\gamma > \beta$ such that $\varphi_\gamma(x_\beta) = 0 \in G_\gamma$.*

PROOF. Writing $\varinjlim \{G_\alpha\} = \bigoplus_\alpha G_\alpha / H$, any given element $x \in \varinjlim \{G_\alpha\}$ is an equivalence class represented by an element

$$\sum_{\alpha \in J_0} g_\alpha \in \bigoplus_{\alpha \in J} G_\alpha$$

for some finite subset $J_0 \subset J$. Since $(J, <)$ is a directed set, we can then find an element $\beta \in J$ satisfying $\beta > \alpha$ for every $\alpha \in J_0$, so

$$\sum_{\alpha \in J_0} g_\alpha - \sum_{\alpha \in J_0} \varphi_{\beta\alpha}(g_\alpha) \in H,$$

implying that $x_\beta := \sum_{\alpha \in J_0} \varphi_{\beta\alpha}(g_\alpha) \in G_\beta$ 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 submodule H , meaning

$$(39.6) \quad x_\beta = \sum_{i=1}^N (g_i - \varphi_{\beta_i \alpha_i}(g_i))$$

for some finite collection of elements $\beta_i > \alpha_i \in J$ and $g_i \in G_{\alpha_i}$, $i = 1, \dots, N$. Choose a finite subset $J_0 \subset J$ that contains all the α_i, β_i for $i = 1, \dots, N$, along with an element $\gamma \in J$ such that $\gamma > \alpha$ for all $\alpha \in J_0$. Applying the homomorphism $\bigoplus_{\alpha \in J_0} \varphi_{\gamma \alpha}$ to both sides of (39.6) 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 ,

$$\left(\bigoplus_{\alpha \in J_0} \varphi_{\gamma \alpha} \right) (g_i - \varphi_{\beta_i \alpha_i}(g_i)) = \varphi_{\gamma \alpha_i}(g_i) - \varphi_{\gamma \beta_i} \circ \varphi_{\beta_i \alpha_i}(g_i) = \varphi_{\gamma \alpha_i}(g_i) - \varphi_{\gamma \alpha_i}(g_i) = 0.$$

We have thus proved $\varphi_{\gamma \beta}(x_\beta) = 0$. □

PROOF OF PROPOSITION 39.21. We prove first that Φ_∞ is surjective. Given a homology class $[c] \in H_k(\varinjlim \{C_*^\alpha\})$ represented by a k -cycle $c \in \varinjlim \{C_*^\alpha\}$, Lemma 39.22 implies $c = \varphi_\beta(c_\beta)$ for some $\beta \in J$ and $c_\beta \in C_k^\beta$, where $\varphi_\beta : C_*^\beta \rightarrow \varinjlim \{C_*^\alpha\}$ 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 Lemma 39.22, we can find some $\gamma > \beta$ and replace c_β with $c_\gamma := \varphi_{\gamma \beta}(c_\beta) \in C_k^\gamma$ 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 $[c_\gamma] \in H_k(C_*^\gamma)$. Now let

$$\Psi_\gamma : H_k(C_*^\gamma) \rightarrow \varinjlim \{H_k(C_*^\alpha)\}$$

denote the natural morphism associated to the direct limit of the system $\{H_k(C_*^\alpha), \Phi_{\gamma \alpha}\}$, where $\Phi_{\gamma \alpha} := (\varphi_{\gamma \alpha})_* : H_k(C_*^\alpha) \rightarrow H_k(C_*^\gamma)$ for $\gamma > \alpha$. Writing $\Phi_\gamma := (\varphi_\gamma)_* : H_k(C_*^\gamma) \rightarrow H_k(\varinjlim \{C_*^\alpha\})$, the diagram

$$\begin{array}{ccc} H_k(C_*^\gamma) & \xrightarrow{\Psi_\gamma} & \varinjlim \{H_k(C_*^\alpha)\} \\ & \searrow \Phi_\gamma & \downarrow \Phi_\infty \\ & & H_k(\varinjlim \{C_*^\alpha\}) \end{array}$$

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. □

An essential role in the proof above was played by Lemma 39.22, which is a tool for replacing statements about direct limits with corresponding statements about individual objects in the direct system. We saw in Exercise 39.5 that the singular homology functors $H_k : \mathbf{Top} \rightarrow R\text{-Mod}$ are not always continuous with respect to direct limits; since H_k is the composition of two functors $H_k : \mathbf{Ch} \rightarrow R\text{-Mod}$ and $C_* : \mathbf{Top} \rightarrow \mathbf{Ch}$, Proposition 39.21 implies that something must go wrong in general with the continuity of $C_* : \mathbf{Top} \rightarrow \mathbf{Ch}$. The following result therefore contains an extra hypothesis that is not satisfied in pathological examples such as Exercise 39.5, but certainly is satisfied (due to Proposition 37.16) by the direct system formed by the skeleta of any CW-complex. The key point is that since every singular n -simplex $\sigma : \Delta^n \rightarrow \varinjlim \{X_\alpha\}$ has image contained in a compact set, this extra hypothesis will allow us to write it as $\varphi_\beta \circ \sigma_\beta$ for some $\beta \in J$ and a singular n -simplex $\sigma_\beta : \Delta^n \rightarrow X_\beta$, producing an analogue of Lemma 39.22 for the situation at hand.

PROPOSITION 39.23. *Suppose $\{X_\alpha, \varphi_{\beta \alpha}\}$ is a direct system of topological spaces over $(J, <)$ satisfying the following conditions:*

- (1) For every $\alpha \in J$, X_α is a subspace of $X := \varinjlim\{X_\alpha\}$ and the maps $\varphi_{\beta\alpha} : X_\alpha \rightarrow X_\beta$ and $\varphi_\alpha : X_\alpha \rightarrow X$ are the natural inclusions;
- (2) Every compact subset $K \subset X$ is contained in X_α for some $\alpha \in J$.

Then choosing \mathcal{F} to be the singular chain complex functor $C_* : \mathbf{Top} \rightarrow \mathbf{Ch}$ with an arbitrary coefficient group G , the chain map

$$\Phi_\infty : \varinjlim\{C_*(X_\alpha)\} \rightarrow C_*(\varinjlim\{X_\alpha\})$$

defined as in (39.2) is an isomorphism of chain complexes.

PROOF. For surjectivity, given $c = \sum_i g_i \sigma_i \in C_n(\varinjlim\{X_\alpha\})$, the finitely many singular n -simplices $\sigma_i : \Delta^n \rightarrow \varinjlim\{X_\alpha\}$ can each be written as $\sigma_i = \varphi_{\alpha_i} \circ \sigma'_i$ for some $\alpha_i \in J$ and $\sigma'_i : \Delta^n \rightarrow X_{\alpha_i}$ since Δ^n is compact. We can then find $\beta \in J$ with $\beta > \alpha_i$ for all i and define $\sigma''_i := \varphi_{\beta\alpha_i} \circ \sigma'_i : \Delta^n \rightarrow X_\beta$, 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 $\Psi_\beta : C_*(X_\beta) \rightarrow \varinjlim\{C_*(X_\alpha)\}$ for the natural map associated to the direct limit, the diagram

$$\begin{array}{ccc} C_*(X_\beta) & \xrightarrow{\Psi_\beta} & \varinjlim\{C_*(X_\alpha)\} \\ & \searrow (\varphi_\beta)_* & \downarrow \Phi_\infty \\ & & C_*(\varinjlim\{X_\alpha\}) \end{array}$$

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. □

Applying Propositions 39.21 and 39.23 together, we've proved:

THEOREM 39.24. *Under the same hypotheses as in Proposition 39.23, there is a natural isomorphism*

$$\varinjlim\{H_k(X_\alpha)\} \xrightarrow{\cong} H_k(\varinjlim\{X_\alpha\})$$

for every $k \geq 0$ and every choice of coefficient group. □

39.5. Exercises.

EXERCISE 39.1. Prove Proposition 39.8 on the non-uniqueness of direct limits.

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.

EXERCISE 39.2 (*). Prove Proposition 39.9, on direct limits of systems that are eventually constant.

EXERCISE 39.3. Prove Lemma 39.10 on the equivalence relation defined on the set-theoretic disjoint union for any direct system of sets and maps.

EXERCISE 39.4. Prove the obvious analogue of Proposition 39.14 for direct systems in the category $\mathbf{Ch} = \mathbf{Ch}(R\text{-Mod})$ of chain complexes of R -modules.

EXERCISE 39.5. For the direct system $\{X_\alpha\}_{\alpha \in J}$ described in Example 39.16, prove the claim that $\varinjlim\{X_\alpha\} \cong S^1$ but $\varinjlim\{H_1(X_\alpha; \mathbb{Z})\}$ is not isomorphic to $H_1(S^1; \mathbb{Z})$.

Hint: Describing $\varinjlim\{X_\alpha\}$ as in Proposition 39.11, it is not hard to find a natural bijection between $\varinjlim\{X_\alpha\}$ and $\bigcup_{\alpha \in J} X_\alpha = S^1$, but you need to check that the topology of this direct limit matches the standard topology of S^1 .

EXERCISE 39.6. 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 $n \geq m$, 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 $\mathbb{R}\mathbb{P}^m \hookrightarrow \mathbb{R}\mathbb{P}^n$, and the complex version gives $\mathbb{C}\mathbb{P}^m \hookrightarrow \mathbb{C}\mathbb{P}^n$. Use cell decompositions to compute the homology with integer coefficients for each space:

- (a) $S^\infty = \varinjlim \{S^n\}_{n \in \mathbb{N}}$
- (b) $\mathbb{R}\mathbb{P}^\infty = \varinjlim \{\mathbb{R}\mathbb{P}^n\}_{n \in \mathbb{N}}$
- (c) $\mathbb{C}\mathbb{P}^\infty = \varinjlim \{\mathbb{C}\mathbb{P}^n\}_{n \in \mathbb{N}}$

EXERCISE 39.7. 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 $\varphi_{\beta\alpha} : X_\alpha \rightarrow X_\beta$ in this case are the natural inclusions. Let us use Proposition 39.11 to identify $\varinjlim \{X_\alpha\}$ with $\coprod_\alpha X_\alpha / \sim$, in terms of the equivalence relation

$$X_\alpha \ni x \sim y \in X_\beta \quad \Leftrightarrow \quad \varphi_{\gamma\alpha}(x) = \varphi_{\gamma\beta}(y) \text{ for some } \gamma \in J \text{ with } \gamma > \alpha, \gamma > \beta.$$

The disjoint union of the inclusions $X_\alpha \hookrightarrow \bigcup_{\beta \in J} X_\beta$ then descends to the quotient as a bijection

$$\varinjlim \{X_\alpha\} \rightarrow \bigcup_{\alpha \in J} X_\alpha,$$

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 $J = (0, 1)$ and consider the family of sets $X_t = \{0\} \cup (t, 1] \subset \mathbb{R}$ for $t \in J$, ordered by inclusion. The union of these sets is $[0, 1]$, but show that the topological space $\varinjlim \{X_t\}$ is not connected.

EXERCISE 39.8. 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 $(J, <)$ can be encoded as a category \mathcal{J} in which relations $\alpha < \beta$ are viewed as morphisms $\alpha \rightarrow \beta$, and from this perspective, a direct system $\{X_\alpha, \varphi_{\beta\alpha}\}$ over $(J, <)$ in the category \mathcal{C} is the same thing as a (covariant) functor $\mathcal{F} : \mathcal{J} \rightarrow \mathcal{C}$. To define targets and direct limits in this language, one can identify each object Y of \mathcal{C} with the “constant” functor $\mathcal{Y} : \mathcal{J} \rightarrow \mathcal{C}$ that sends every $\alpha \in J$ to Y and every morphism $\alpha \rightarrow \beta$ of \mathcal{J} to the identity morphism $Y \rightarrow Y$. Note that if $\mathcal{X}, \mathcal{Y} : \mathcal{J} \rightarrow \mathcal{C}$ are two such constant functors associated to objects X, Y respectively in \mathcal{C} , then a natural transformation from \mathcal{X} to \mathcal{Y} must associate to every $\alpha \in J$ the same morphism $X \rightarrow Y$, and conversely, every morphism $X \rightarrow Y$ determines a natural transformation $\mathcal{X} \rightarrow \mathcal{Y}$. A target $\{Y, f_\alpha\}$ of the system $\{X_\alpha, \varphi_{\beta\alpha}\}$ is now the same thing as a natural transformation $T_Y : \mathcal{F} \rightarrow \mathcal{Y}$, assigning to each object α of \mathcal{J} the morphism $T_Y(\alpha) := f_\alpha : X_\alpha \rightarrow Y$, and in this language, a target $T_X : \mathcal{F} \rightarrow \mathcal{X}$ is called universal (and is thus a direct limit of the system) if for every target $T_Y : \mathcal{F} \rightarrow \mathcal{Y}$, there is a unique natural transformation $\Phi : \mathcal{X} \rightarrow \mathcal{Y}$ 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 \mathcal{J} associated with the directed set $(J, <)$ by an arbitrary⁶¹ category \mathcal{A} . For reasons that may become clearer when you look at the examples below, a functor $\mathcal{F} : \mathcal{A} \rightarrow \mathcal{C}$ is then often referred to as a *diagram* in \mathcal{C} over \mathcal{A} . A target is again simply a natural transformation $T_X : \mathcal{F} \rightarrow \mathcal{X}$ to the constant functor $\mathcal{X} : \mathcal{A} \rightarrow \mathcal{C}$ determined by some object X

⁶¹I say “arbitrary,” but in practice, \mathcal{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, \mathcal{A} contains only finitely many objects, and there are already interesting examples (as in Exercise 39.8(c)) in which it has only two.

of \mathcal{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 $T_X(\alpha) : \mathcal{F}(\alpha) \rightarrow X$ associated to each object α of \mathcal{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 \mathcal{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 $\mathcal{A} \rightarrow \mathcal{C}$ is simply a collection $\{X_\alpha\}_{\alpha \in J}$ of objects in \mathcal{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 \mathbf{Top} and $\mathbf{Top}^{\text{rel}}$ are disjoint unions, and coproducts in \mathbf{Ab} , $\mathbf{R-Mod}$ and \mathbf{Ch} are direct sums.
- (b) Give a concrete description of coproducts in the categories \mathbf{Top}_* (pointed spaces) and \mathbf{Grp} (not necessarily abelian groups).
Hint: Both answers are constructions that were introduced in last semester's Topologie I course.
- (c) If \mathcal{A} contains only two objects α, β and its morphisms consist only of the identity morphisms on α, β plus exactly two morphisms $\alpha \rightarrow \beta$, then a diagram $\mathcal{F} : \mathcal{A} \rightarrow \mathcal{C}$ can be described as a pair of objects in \mathcal{C} with a pair of morphisms

$$X \begin{array}{c} \xrightarrow{f} \\ \xrightarrow{g} \end{array} Y,$$

and a colimit of such a diagram is called a **coequalizer**. Give an explicit description of coequalizers in the categories \mathbf{Top} and $\mathbf{R-Mod}$.

Hint: Use quotients.

- (d) Prove: If \mathcal{C} is a category in which coproducts and coequalizers always exist, then every direct system in \mathcal{C} has a direct limit.
Hint: Special cases of this yield the explicit descriptions of direct limits in \mathbf{Top} and $\mathbf{R-Mod}$ that appear in Propositions 39.11 and 39.14.

40. Euler characteristic and fixed points

We would now like to discuss a few applications of the isomorphism

$$H_*^{\text{CW}}(X, A; G) \cong H_*(X, A; G).$$

40.1. Finitely-generated homology. One of the advantages of cellular homology is that for compact spaces, cell decompositions are always finite, so that the cellular chain complex has only finitely-many generators. Working with coefficients in the ring R , this means in particular that the entire chain complex is a finitely-generated R -module, and therefore so is the homology:

COROLLARY 40.1. *If (X, A) is a compact CW-pair, then $H_*(X, A; R)$ is a finitely-generated R -module. In particular, $H_*(X, A; \mathbb{Z})$ is a finitely-generated abelian group, and for any field \mathbb{K} , $H_*(X, A; \mathbb{K})$ is a finite-dimensional vector space over \mathbb{K} . \square*

Note that the corollary is actually two statements in one, as it says on the one hand that $H_n(X, A; R)$ is finitely generated for each $n \in \mathbb{Z}$, but a second implication is that $H_n(X, A; R)$ is *trivial* for all but finitely many values of n .

It is similarly obvious that $C_k^{\text{CW}}(X)$ and therefore also $H_k^{\text{CW}}(X)$ must vanish for any CW-pair that has no k -cells:

COROLLARY 40.2. *If (X, A) is an n -dimensional CW-pair, then $H_k(X, A) = 0$ for all $k > n$ and every choice of coefficients. \square*

REMARK 40.3. 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 40.2 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 40.2 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 semester, 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 40.4. 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.

40.2. The Hopf trace formula. There are situations in which interesting topologically invariant information can be extracted from the cellular chain complex without even computing its homology. The following algebraic result makes this possible.

THEOREM 40.5 (Hopf trace formula). *Suppose C_* is a finite-dimensional chain complex of vector spaces over a field \mathbb{K} , and $f : C_* \rightarrow C_*$ 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 $\partial_n : C_n \rightarrow C_{n-1}$ for each $n \in \mathbb{Z}$, abbreviate

$$Z_n := \ker \partial_n \subset C_n, \quad B_n := \operatorname{im} \partial_{n+1} \subset C_n, \quad H_n := Z_n / B_n.$$

Denote $f_{C_n} := f|_{C_n} : C_n \rightarrow C_n$, and note that since f is a chain map, it restricts to these subspaces and the quotient as linear maps

$$Z_n \xrightarrow{f_{Z_n}} Z_n, \quad B_n \xrightarrow{f_{B_n}} B_n, \quad H_n \xrightarrow{f_{H_n}} H_n,$$

such that the following diagram commutes

$$\begin{array}{ccccccccc} 0 & \longrightarrow & Z_n & \hookrightarrow & C_n & \xrightarrow{\partial_n} & B_{n-1} & \longrightarrow & 0 \\ & & \downarrow f_{Z_n} & & \downarrow f_{C_n} & & \downarrow f_{B_{n-1}} & & \\ 0 & \longrightarrow & Z_n & \hookrightarrow & C_n & \xrightarrow{\partial_n} & B_{n-1} & \longrightarrow & 0 \end{array}$$

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 or modules—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 $f_{C_n} : C_n \rightarrow C_n$ 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 $g : B_{n-1} \rightarrow Z_n$. 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 $B_{n-1} \rightarrow B_{n-1}$ induced by the same chain map f , i.e. it is the lower-right term $f_{B_{n-1}}$. This formula proves

$$(40.1) \quad \operatorname{tr}(f_{C_n}) = \operatorname{tr}(f_{Z_n}) + \operatorname{tr}(f_{B_{n-1}}).$$

Now apply the same argument to the diagram

$$\begin{array}{ccccccc} 0 & \longrightarrow & B_n & \hookrightarrow & Z_n & \longrightarrow & H_n \longrightarrow 0 \\ & & \downarrow f_{B_n} & & \downarrow f_{Z_n} & & \downarrow f_{H_n} \\ 0 & \longrightarrow & B_n & \hookrightarrow & Z_n & \longrightarrow & H_n \longrightarrow 0 \end{array}$$

where the maps $Z_n \rightarrow H_n$ 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 (40.1) gives

$$\sum_{n \in \mathbb{Z}} (-1)^n [\operatorname{tr}(f_{C_n}) - \operatorname{tr}(f_{B_{n-1}})] = \sum_{n \in \mathbb{Z}} (-1)^n \operatorname{tr}(f_{Z_n}) = \sum_{n \in \mathbb{Z}} (-1)^n [\operatorname{tr}(f_{B_n}) + \operatorname{tr}(f_{H_n})],$$

which implies the desired result after dropping the extraneous terms $\operatorname{tr}(f_{B_n})$ from both sides. \square

Plugging the identity map $C_* \rightarrow C_*$ into the Hopf trace formula gives:

COROLLARY 40.6. *For C_* a finite-dimensional chain complex of vector spaces over a field \mathbb{K} ,*

$$\sum_{n \in \mathbb{Z}} (-1)^n \dim_{\mathbb{K}} C_n = \sum_{n \in \mathbb{Z}} (-1)^n \dim_{\mathbb{K}} H_n(C_*).$$

\square

40.3. Betti numbers and the Euler characteristic. Associating a sequence of abelian groups or R -modules 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.

DEFINITION 40.7. For any space X and integer $k \geq 0$, the *k th Betti number* of X is the nonnegative (or possibly infinite) integer

$$b_k(X) := \dim_{\mathbb{Q}} H_k(X; \mathbb{Q}).$$

REMARK 40.8. Several alternative definitions of $b_k(X)$ will become possible once we've proved the universal coefficient theorem, e.g. it will turn out that the coefficient field \mathbb{Q} in Definition 40.7 can freely be replaced by any other field of characteristic zero. One can also use \mathbb{Z} coefficients: 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 $n \geq 0$ 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 $n \geq 0$ that appears in the isomorphism above, and is called the **rank** (*Rang*) of G ,

$$\text{rank } G := n \geq 0, \quad \text{where} \quad G \cong \mathbb{Z}^n \oplus \text{torsion}.$$

We will see as a corollary of the universal coefficient theorem that $H_k(X; \mathbb{K}) \cong H_k(X; \mathbb{Z}) \otimes \mathbb{K}$ whenever \mathbb{K} is a field of characteristic zero. If $H_k(X; \mathbb{Z}) \cong \mathbb{Z}^n \oplus T$ for a torsion group T , it follows in this situation that

$$H_k(X; \mathbb{K}) \cong (\mathbb{Z}^n \otimes \mathbb{K}) \oplus (T \otimes \mathbb{K}) \cong \mathbb{K}^n$$

since $\mathbb{Z} \otimes \mathbb{K} \cong \mathbb{K}$ and $T \otimes \mathbb{K}$ is trivial, and another way of defining the Betti numbers is therefore

$$b_k(X) = \text{rank } H_k(X; \mathbb{Z}).$$

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 a 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 $H_{\text{dR}}^k(M)$, which is a vector space over \mathbb{R} . This matches our definition above due to de Rham's theorem, which provides an isomorphism between $H_{\text{dR}}^k(M)$ and the singular cohomology group $H^k(M; \mathbb{R})$ with coefficients in \mathbb{R} .

DEFINITION 40.9. For any space X with finitely-generated singular homology over \mathbb{Z} , the **Euler characteristic** (*Eulercharakteristik*) of X is the integer⁶²

$$\chi(X) = \sum_{k=0}^{\infty} (-1)^k b_k(X) \in \mathbb{Z}.$$

The usefulness of $\chi(X)$ as an invariant derives from Corollary 40.6, which gives us a very easy way to compute $\chi(X)$ whenever X is a finite cell complex, without even needing to compute its homology! The following result is a direct consequence of Corollary 40.6 and the isomorphism $H_*^{\text{CW}}(X; \mathbb{Q}) \cong H_*(X; \mathbb{Q})$.

THEOREM 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 $|\mathcal{K}^n| \geq 0$ denotes the number of n -cells in the decomposition. □

EXAMPLE 40.11. For $n \geq 0$, 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 $k = 0, \dots, n$.

⁶²Let's be clear about this notational detail: χ is the Greek letter "chi," not a variety of the letter "X" in a strange font. The χ of course stands for "characteristic".

EXAMPLE 40.12. For the closed oriented surface Σ_g of genus $g \geq 0$, we computed $H_*(\Sigma_g)$ in Example 37.25: taking rational coefficients, the nontrivial homology groups are $H_0(\Sigma_g; \mathbb{Q}) \cong \mathbb{Q}$, $H_1(\Sigma_g; \mathbb{Q}) \cong \mathbb{Q}^{2g}$ and $H_2(\Sigma_g; \mathbb{Q}) \cong \mathbb{Q}$, thus

$$\chi(\Sigma_g) = 1 - 2g + 1 = 2 - 2g.$$

But one can also compute $\chi(\Sigma_g)$ without computing $H_*(\Sigma_g)$ at all, 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 37.25, 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 $H_0(|\Gamma|; \mathbb{Q}) \cong \mathbb{Q}$. Since there are no 2-cells, $H_1(|\Gamma|; \mathbb{Q})$ is isomorphic to the subgroup of 1-cycles in $C_1^{\text{CW}}(|\Gamma|; \mathbb{Q})$, 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 $H_1(|\Gamma|; \mathbb{Q}) = 0$. This proves $\chi(|\Gamma|) = 1$, and combining it with Theorem 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 $v - e = 1$. \square*

Here is an application to covering spaces. The proof of the following lemma is Exercise 40.1.

LEMMA 40.14. *Suppose X is a compact cell complex and $\pi : Y \rightarrow X$ is a covering map of finite degree $d \in \mathbb{N}$. Then 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 $\mathbb{D}^n \rightarrow Y$ are lifts of the characteristic map $\mathbb{D}^n \rightarrow X$ for e_α^n . \square*

In conjunction with Theorem 40.10, the lemma implies:

THEOREM 40.15. *If X is a finite cell complex and $\pi : Y \rightarrow X$ is a covering map of finite degree $d \in \mathbb{N}$, then $\chi(Y) = d\chi(X)$. \square*

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 $\pi : S^n \rightarrow X$ is a d -fold covering map, n is even and X is a CW-complex, then d is either 1 or 2. \square*

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 $S^n \rightarrow \mathbb{R}P^n$.

40.4. The Lefschetz fixed point theorem. As another application of cellular homology and the Hopf trace formula, I'd like to address the following general question:

QUESTION 40.18. *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 36 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.

DEFINITION 40.19. 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 $f : X \rightarrow X$ is defined by

$$L_{\mathbb{K}}(f) := \sum_{n \in \mathbb{Z}} (-1)^n \operatorname{tr} (H_n(X; \mathbb{K}) \xrightarrow{f_*} H_n(X; \mathbb{K})) \in \mathbb{K}.$$

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 40.20. 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 $f_* : C_n^{\text{CW}}(X; \mathbb{Q}) \rightarrow C_n^{\text{CW}}(X; \mathbb{Q})$ that are represented in the canonical basis by matrices with integer entries. Without these assumptions, it follows more generally from the universal coefficient theorem, which will give us a natural isomorphism $H_*(X; \mathbb{Q}) \cong H_*(X; \mathbb{Z}) \otimes \mathbb{Q}$, so that the maps $f_* : H_*(X; \mathbb{Q}) \rightarrow H_*(X; \mathbb{Q})$ can also be presented as matrices with integer entries. More precisely, every endomorphism $H_n(X; \mathbb{Z}) \rightarrow H_n(X; \mathbb{Z})$ preserves the torsion subgroup $T_n \subset H_n(X; \mathbb{Z})$ and thus descends to an endomorphism of the *free part* of $H_n(X; \mathbb{Z})$,

$$H_n(X; \mathbb{Z})/T_n \xrightarrow{f_*} H_n(X; \mathbb{Z})/T_n,$$

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 $L(f)$.

EXAMPLE 40.21. If X has finitely-generated homology and $f : X \rightarrow X$ is homotopic to the identity map, then $L(f) = \chi(X)$.

THEOREM 40.22 (Lefschetz-Hopf). *If X is a compact polyhedron and \mathbb{K} is a field, then every map $f : X \rightarrow X$ 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 35.10), and its proof requires some knowledge of the *flow* of a smooth vector field from differential geometry.

COROLLARY 40.23. *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 $\varphi_X^t : M \rightarrow M$ with no fixed points, but is clearly also homotopic to the identity, thus $L(\varphi_X^t) = \chi(M) = 0$. \square

REMARK 40.24. 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 $r : \mathcal{U} \rightarrow X$ then makes X a retract of a compact polyhedron $K \subset \mathcal{U}$ containing X . Now if $f : X \rightarrow X$ has $L_{\mathbb{K}}(f) \neq 0$, one can consider the map

$$i \circ f \circ r : K \rightarrow K$$

where $i : X \hookrightarrow K$ is the inclusion, and use Exercise 40.2 to prove $L_{\mathbb{K}}(i \circ f \circ r) = L_{\mathbb{K}}(f)$, so that Theorem 40.22 guarantees a fixed point for $i \circ f \circ r$. But $i \circ f \circ r(x) = x$ implies $x \in X$ and $f(x) = x$.

REMARK 40.25. 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 40.22 lacks. For a map $f : M \rightarrow M$ 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 $f : M \rightarrow N$ 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 40.23 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 40.26. It is easy to see that the compactness of X in Theorem 40.22 is essential: for instance, \mathbb{R} has finitely-generated homology and $f : \mathbb{R} \rightarrow \mathbb{R} : x \mapsto x + 1$ is homotopic to the identity, hence $L(f) = \chi(\mathbb{R}) = 1$, even though f has no fixed points.

REMARK 40.27. Figure 19 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

$$H_*(X) \cong H_*(S^1) \oplus H_*(S^1) \oplus H_*(\mathbb{R}),$$

implying $\chi(X) = \chi(S^1) + \chi(S^1) + \chi(\mathbb{R}) = 0 + 0 + 1 = 1$. But it is easy to visualize a map $f : X \rightarrow X$ 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.)

The idea behind the proof of Theorem 40.22 is that a map $f : X \rightarrow X$ with no fixed points can be modified to a *cellular* map whose induced chain map has no diagonal terms, and must therefore have Lefschetz number zero. The main tool needed for this is the simplicial approximation theorem (see Theorem 31.9).

PROOF OF THEOREM 40.22. Assume X is a compact polyhedron, \mathbb{K} is a field and $f : X \rightarrow X$ 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

$$d(x, f(x)) \geq \epsilon > 0 \quad \text{for all } x \in X.$$

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




FIGURE 19. A compact space X with $\chi(X) = 1$ admitting maps homotopic to the identity that have no fixed point.

triangulation further subdivided so that the simplicial approximation theorem applies, giving a simplicial map

$$g : X' \rightarrow X$$

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 $n \geq 0$, one can also regard g as a *cellular* (though not simplicial) map

$$g : X' \rightarrow X'.$$

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,

$$d(x, g(x)) \geq d(x, f(x)) - d(f(x), g(x)) > \epsilon - \frac{\epsilon}{2} = \frac{\epsilon}{2},$$

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

$$C_*^{\text{CW}}(X'; \mathbb{K}) \xrightarrow{g_*} C_*^{\text{CW}}(X'; \mathbb{K})$$

has only zeroes along the diagonal, and its trace in every dimension is therefore 0. By the Hopf trace formula, it follows that $L_{\mathbb{K}}(g) = L_{\mathbb{K}}(f) = 0$. \square

40.5. Exercises.

EXERCISE 40.1. Prove Lemma 40.14 on covering spaces of CW-complexes.

Hint: The key point here is that characteristic maps $\mathbb{D}^n \rightarrow X$ will always lift to the cover since \mathbb{D}^n is simply connected. It's probably easiest if you argue by induction on n .

EXERCISE 40.2. Suppose $A \subset X$ is a subspace with inclusion $i : A \hookrightarrow X$ and a retraction $r : X \rightarrow A$, 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 $f : A \rightarrow A$, the induced maps $f_* : H_n(A; \mathbb{K}) \rightarrow H_n(A; \mathbb{K})$ and $(i \circ f \circ r)_* : H_n(X; \mathbb{K}) \rightarrow H_n(X; \mathbb{K})$ 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 $f_* r_* : H_n(X; \mathbb{K}) \rightarrow H_n(A; \mathbb{K})$ and $i_* : H_n(A; \mathbb{K}) \rightarrow H_n(X; \mathbb{K})$, and recall the formula $\mathrm{tr}(\mathbf{AB}) = \mathrm{tr}(\mathbf{BA})$.

41. Singular cohomology

Motivation. Singular cohomology assigns to each topological space X and each R -module G a sequence of R -modules $H^n(X; G)$ whose direct sum we denote by

$$H^*(X; G) = \bigoplus_{n \in \mathbb{Z}} H^n(X; G), \quad \text{or more succinctly,} \quad H^*(X) = \bigoplus_{n \in \mathbb{Z}} H^n(X).$$

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 collection of functors from **Top** to $R\text{-Mod}$, cohomology is *contravariant*, meaning that continuous maps $f : X \rightarrow Y$ induce homomorphisms

$$f^* : H^n(Y) \rightarrow H^n(X),$$

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 take coefficients in the ring $G := R$, then $H^*(X; R)$ has a natural product structure, called the **cup product**

$$H^k(X; R) \otimes H^\ell(X; R) \xrightarrow{\cup} H^{k+\ell}(X; R).$$

This structure 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**

$$H_{n-k}(X; \mathbb{Z}) \otimes H_{n-\ell}(X; \mathbb{Z}) \rightarrow H_{n-k-\ell}(X; \mathbb{Z}) : [M] \otimes [N] \mapsto [M] \cdot [N] := [M \cap N].$$

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 37.5 the complex projective space $\mathbb{C}\mathbb{P}^n$, defined as the space of all complex lines through the origin in \mathbb{C}^{n+1} , meaning literally the quotient space

$$\mathbb{C}\mathbb{P}^n = (\mathbb{C}^{n+1} \setminus \{0\}) / \mathbb{C}^*,$$

where the multiplicative group $\mathbb{C}^* := \mathbb{C} \setminus \{0\}$ is understood to act on $\mathbb{C}^{n+1} \setminus \{0\}$ by scalar multiplication.

THEOREM 41.1. *For every even $n \geq 0$, every continuous map $f : \mathbb{C}\mathbb{P}^n \rightarrow \mathbb{C}\mathbb{P}^n$ has a fixed point.*

PROOF (MODULO TECHNICAL DETAILS). We saw in Exercise 37.5 that $\mathbb{C}\mathbb{P}^n$ has a cell decomposition of the form $e^0 \cup e^2 \cup \dots \cup e^{2n}$, 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:

$$H^k(\mathbb{C}\mathbb{P}^n; \mathbb{Z}) \cong H_k(\mathbb{C}\mathbb{P}^n; \mathbb{Z}) \cong \begin{cases} \mathbb{Z} & \text{for } k = 0, 2, 4, \dots, 2n, \\ 0 & \text{for all other } k. \end{cases}$$

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 one deduces from this theorem two things: first, that the Lefschetz number $L(f) \in \mathbb{Z}$ of a map $f : X \rightarrow X$ can be computed equally well using homology or cohomology, and second, that we are free to use $H^*(X; \mathbb{Z})$ modulo torsion in place of $H^*(X; \mathbb{Q})$ for this computation. Thus for a map $f : \mathbb{C}\mathbb{P}^n \rightarrow \mathbb{C}\mathbb{P}^n$, we can write

$$L(f) = \sum_{k \in \mathbb{Z}} (-1)^k \operatorname{tr} (H^k(\mathbb{C}\mathbb{P}^n; \mathbb{Z}) \xrightarrow{f^*} H^k(\mathbb{C}\mathbb{P}^n; \mathbb{Z})) = \sum_{k=0}^n \operatorname{tr} (H^{2k}(\mathbb{C}\mathbb{P}^n; \mathbb{Z}) \xrightarrow{f^*} H^{2k}(\mathbb{C}\mathbb{P}^n; \mathbb{Z})).$$

Now we take advantage of the cup product on $H^*(\mathbb{C}\mathbb{P}^n; \mathbb{Z})$, which has the following properties:

- It is *natural*, i.e. for all $\alpha, \beta \in H^*(\mathbb{C}\mathbb{P}^n; \mathbb{Z})$, $f^*(\alpha \cup \beta) = f^*\alpha \cup f^*\beta$. (This is a general property of the cup product with respect to continuous maps between arbitrary spaces.)
- If $\alpha \in H^2(\mathbb{C}\mathbb{P}^n; \mathbb{Z}) \cong \mathbb{Z}$ is a generator, then for each $k = 0, 1, \dots, n$,

$$\alpha^k := \underbrace{\alpha \cup \dots \cup \alpha}_k \in H^{2k}(\mathbb{C}\mathbb{P}^n; \mathbb{Z}) \cong \mathbb{Z}$$

is also a generator. We will prove this as a corollary of Poincaré duality, which holds since $\mathbb{C}\mathbb{P}^n$ is a closed and oriented manifold.

Now fixing a generator $\alpha \in H^2(\mathbb{C}\mathbb{P}^n; \mathbb{Z})$, every continuous map $f : \mathbb{C}\mathbb{P}^n \rightarrow \mathbb{C}\mathbb{P}^n$ gives rise to a unique integer $m \in \mathbb{Z}$ such that

$$f^*\alpha = m\alpha$$

since $H^2(\mathbb{C}\mathbb{P}^n; \mathbb{Z}) \cong \mathbb{Z}$. It follows via the two properties above that for each $k = 0, \dots, n$, the generator $\alpha^k \in H^{2k}(\mathbb{C}\mathbb{P}^n; \mathbb{Z})$ satisfies

$$f^*(\alpha^k) = f^*(\alpha \cup \dots \cup \alpha) = f^*\alpha \cup \dots \cup f^*\alpha = m^k \alpha^k,$$

and the Lefschetz number of f is therefore

$$L(f) = 1 + m + \dots + m^n \in \mathbb{Z}.$$

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

41.1. The functor $\text{Hom}(\cdot, G)$ and cochains. Let's talk about algebra. Given a chain complex (C_*, ∂) of R -modules, we obtain its homology in degree $n \in \mathbb{Z}$ by applying the functor $H_n : \text{Ch} \rightarrow R\text{-Mod}$, which discards some of the information in (C_*, ∂) in the hope of obtaining something more computable. The algebraic idea behind cohomology is to pre-process the chain complex via a *dualization* functor before passing it to the functor H_n .

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 ,

$$\text{Hom}(A, R) = \text{Hom}_R(A, R) = \{R\text{-module homomorphisms } A \rightarrow R\},$$

which reproduces the definition familiar from linear algebra if R is a field. In particular, the dual of an abelian group A is the abelian group of group homomorphisms $A \rightarrow \mathbb{Z}$.

More generally, we can fix an arbitrary R -module G and consider the functor

$$R\text{-Mod} \rightarrow R\text{-Mod} : A \mapsto \text{Hom}(A, G).$$

This is perhaps the simplest example of a **contravariant** functor, as one can naturally associate to each homomorphism $\Phi : A \rightarrow B$ a homomorphism in the other direction

$$\Phi^* : \text{Hom}(B, G) \rightarrow \text{Hom}(A, G)$$

defined by

$$\Phi^*(\lambda) := \lambda \circ \Phi \in \text{Hom}(A, G) \quad \text{for } \lambda \in \text{Hom}(B, G).$$

You should take a moment to convince yourself that this satisfies the relations characteristic of a contravariant functor (see Definition 27.15): the identity map $\mathbb{1} : A \rightarrow A$ induces the identity map $\mathbb{1}^* : \text{Hom}(A, G) \rightarrow \text{Hom}(A, G)$, and $(\Phi\Psi)^* = \Psi^*\Phi^*$ whenever Φ and Ψ can be composed.

If A is a free R -module with basis \mathcal{B} , then homomorphisms $\varphi : A \rightarrow G$ are uniquely determined by the elements $\varphi(b) \in G$ for every $b \in \mathcal{B}$, and $\text{Hom}(A, G)$ thus has a natural bijective correspondence with the set of all functions $\mathcal{B} \rightarrow G$,

$$\text{Hom}(A, G) \cong \{\text{functions } \varphi : \mathcal{B} \rightarrow G\}, \quad \text{assuming} \quad A \cong \bigoplus_{b \in \mathcal{B}} R.$$

This special case is important, because the modules that we feed into $\text{Hom}(\cdot, G)$ in order to define singular cohomology will always be *free* R -modules.

We next define what $\text{Hom}(C_*, G)$ should mean when C_* is a chain complex with boundary operator $\partial : C_* \rightarrow C_{*-1}$. Since C_* is a \mathbb{Z} -graded R -module, we would like $\text{Hom}(C_*, G)$ to be another \mathbb{Z} -graded R -module: the obvious definition is then

$$\text{Hom}(C_*, G) := \bigoplus_{n \in \mathbb{Z}} \text{Hom}(C_n, G),$$

so that $\text{Hom}(C_n, G)$ is the submodule of elements with degree n in $\text{Hom}(C_*, G)$.⁶³ Now we can dualize the map $\partial : C_* \rightarrow C_*$ to obtain a map

$$\partial^* : \text{Hom}(C_*, G) \rightarrow \text{Hom}(C_*, G) : \alpha \mapsto \alpha \circ \partial,$$

which sends $\text{Hom}(C_n, G)$ to $\text{Hom}(C_{n+1}, G)$ 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 41.3), we're going to introduce

⁶³If 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, $\text{Hom}(C_*, G)$ as we've defined it is not literally the module of all homomorphisms $C_* \rightarrow G$. That would be $\prod_{n \in \mathbb{Z}} \text{Hom}(C_n, G)$, 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 $\text{Hom}(C_*, G)$ is not to be interpreted too literally.

an extra sign and define

$$(41.1) \quad \delta : \text{Hom}(C_*, G) \rightarrow \text{Hom}(C_*, G) : \alpha \mapsto (-1)^{|\alpha|+1} \partial^* \alpha,$$

where $\alpha \in \text{Hom}(C_*, G)$ here is assumed to be a homogeneous element of degree $|\alpha|$, i.e. it belongs to $\text{Hom}(C_n, G)$ for $n = |\alpha|$. This clearly also satisfies the relation

$$\delta^2 = 0,$$

and it is a map of degree +1, meaning it sends $\text{Hom}(C_n, G)$ to $\text{Hom}(C_{n+1}, G)$ for every $n \in \mathbb{Z}$.

We shall refer to any \mathbb{Z} -graded R -module $A^* = \bigoplus_{n \in \mathbb{Z}} A^n$ endowed with a homomorphism $\delta : A^* \rightarrow A^*$ 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 $\varphi, \psi : A^* \rightarrow B^*$ of cochain complexes (A^*, δ_A) and (B^*, δ_B) is a homomorphism $h : A^* \rightarrow B^*$ of degree -1 that satisfies the usual chain homotopy relation

$$\varphi - \psi = h\delta_A + \delta_B h.$$

The **cohomology** of a cochain complex (A^*, δ) is the \mathbb{Z} -graded R -module

$$H^*(A^*, \delta) = \bigoplus_{n \in \mathbb{Z}} H^n(A^*, \delta) := \ker \delta / \text{im } \delta,$$

so in other words $H^n(A^*, \delta) = \ker \delta_n / \text{im } \delta_{n-1}$ 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 collection of R -modules

$$H^n(C_*, \partial; G) := H^n(\text{Hom}(C_*, G), \delta), \quad n \in \mathbb{Z}.$$

The functor that replaces a chain complex with its cohomology in any given degree $n \in \mathbb{Z}$ can be expressed as the composition of two functors:

$$\text{Ch} \xrightarrow{\text{Hom}(\cdot, G)} \text{CoCh} \xrightarrow{H^n} R\text{-Mod}.$$

Here CoCh denotes the category whose objects are cochain complexes, with morphisms defined as chain maps, and $H^n : \text{CoCh} \rightarrow R\text{-Mod}$ is a covariant functor, in fact exactly the same functor as $H_n : \text{Ch} \rightarrow R\text{-Mod}$, but with the cosmetic difference that it is fed with cochain complexes instead of chain complexes. The functor

$$\text{Hom}(\cdot, G) : \text{Ch} \rightarrow \text{CoCh}$$

replaces a chain complex (C_*, ∂) with the cochain complex $(\text{Hom}(C_*, G), \delta)$ as defined above, and it is contravariant: it associates to each chain map $\varphi : (A_*, \partial_A) \rightarrow (B_*, \partial_B)$ the dual map

$$\varphi^* : (\text{Hom}(B_*, G), \delta_A) \rightarrow (\text{Hom}(A_*, G), \delta_B),$$

which is a chain map since for $\beta \in \text{Hom}(B_n, G)$,

$$\begin{aligned} \varphi^* \delta_B \beta &= \varphi^* ((-1)^{n+1} \partial_B^* \beta) = (-1)^{n+1} \varphi^* \partial_B^* \beta = (-1)^{n+1} (\partial_B \varphi)^* \beta = (-1)^{n+1} (\varphi \partial_A)^* \beta \\ &= (-1)^{n+1} \partial_A^* \varphi^* \beta = \delta_A \varphi^* \beta. \end{aligned}$$

As a consequence, the composition functors $H^n(\cdot; G) : \text{Ch} \rightarrow R\text{-Mod}$ are also contravariant: they associate to each chain map $\varphi : (A_*, \partial_A) \rightarrow (B_*, \partial_B)$ the homomorphisms $H^n(\text{Hom}(B_*, G), \delta_B) \rightarrow H^n(\text{Hom}(A_*, G), \delta_A)$ induced by the chain map φ^* , and we shall also denote the induced morphism of \mathbb{Z} -graded R -modules by

$$\varphi^* : H^*(B_*, \partial_B; G) \rightarrow H^*(A_*, \partial_A; G).$$

Two further algebraic observations are worth recording before we go back to topology.

PROPOSITION 41.2. *If $\varphi, \psi : (A_*, \partial_A) \rightarrow (B_*, \partial_B)$ are chain maps between chain complexes and $h : A_* \rightarrow B_*$ is a chain homotopy between φ and ψ , then the map $\eta : \text{Hom}(B_*, G) \rightarrow \text{Hom}(A_*, G)$ defined for each $n \in \mathbb{Z}$ by*

$$\text{Hom}(B_n, G) \xrightarrow{\eta} \text{Hom}(A_{n-1}, G) : \beta \mapsto (-1)^n h^* \beta$$

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 $\beta \in \text{Hom}(B_n, G)$,

$$(\delta_A \eta + \eta \delta_B) \beta = \partial_A^* h^* \beta + h^* \partial_B^* \beta = (\varphi^* - \psi^*) \beta.$$

□

In category-theoretic terms, the proposition means that $\text{Hom}(\cdot, G)$ descends to a well-defined functor

$$\text{Hom}(\cdot, G) : \mathbf{hCh} \rightarrow \mathbf{hCoCh},$$

where \mathbf{hCoCh} is the category with cochain complexes as objects and chain homotopy classes of chain maps as morphisms. As a consequence, $H^n(\cdot; G) : \mathbf{Ch} \rightarrow R\text{-Mod}$ for each $n \in \mathbb{Z}$ likewise descends to a functor

$$H^n(\cdot; G) : \mathbf{hCh} \rightarrow R\text{-Mod}.$$

The second observation is that for any chain complex (C_*, ∂) , the canonical pairing

$$(41.2) \quad \text{Hom}(C_n, G) \times C_n \rightarrow G : (\alpha, c) \mapsto \alpha(c)$$

descends to homology to give a well-defined pairing

$$(41.3) \quad H^n(C_*, \partial; G) \times H_n(C_*, \partial) \rightarrow G : ([\alpha], [c]) \mapsto \langle [\alpha], [c] \rangle := \alpha(c).$$

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 $\beta \in \text{Hom}(C_{n-1}, G)$,

$$(\delta\beta)(c) = \pm(\partial^* \beta)(c) = \pm\beta(\partial c) = 0.$$

We will often refer to (41.3) as the **evaluation** of cohomology classes on homology classes.

REMARK 41.3. The reason for the sign in (41.1) can be understood in terms of the “chain-level” evaluation map (41.2). Since it is bilinear, it can be expressed as a homomorphism

$$\text{Hom}(C_n, G) \otimes C_n \rightarrow G,$$

which extends in a trivial way to all degrees as a homomorphism

$$(41.4) \quad \text{Hom}(C_*, G) \otimes C_* \rightarrow G$$

if we define $\alpha(c) := 0$ whenever $\alpha \in \text{Hom}(C_k, G)$ and $c \in C_\ell$ for $k \neq \ell$. With a little care, we can then rephrase the fact that (41.3) is well defined as a corollary of the fact that (41.4) is a chain map. For this, we need to make sense of $\text{Hom}(C_*, G) \otimes C_*$ as a chain complex. Given two chain complexes (A_*, ∂^A) and (B_*, ∂^B) , there is a natural way to make $A_* \otimes B_*$ into a chain complex with

$$(A_* \otimes B_*)_n := \bigoplus_{k+\ell=n} A_k \otimes B_\ell$$

by defining $\partial : A_* \otimes B_* \rightarrow A_* \otimes B_*$ on tensor products of homogeneous elements $a \otimes b \in A_k \otimes B_\ell \subset A_* \otimes B_*$ by

$$\partial(a \otimes b) := \partial a \otimes b + (-1)^{|a|} a \otimes \partial b.$$

The sign in this definition is in keeping with the Koszul sign convention (cf. Remark 31.4), and we will see more motivation for it when we later study product structures on homology; one can easily check in any case that the presence of the sign $(-1)^{|\alpha|}$ ensures the relation $\partial^2 = 0$, thus making $(A_* \otimes B_*, \partial)$ a chain complex. Now, before we can regard $\text{Hom}(C_*, G) \otimes C_*$ as a chain complex, we must also face the fact that $\text{Hom}(C_*, G)$ 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

$$\text{Hom}(C_*, G)_n := \text{Hom}(C_{-n}, G)$$

and think of δ as a homomorphism that sends $\text{Hom}(C_*, G)_n$ to $\text{Hom}(C_*, G)_{n-1}$. The fact that $\alpha(c) = 0$ whenever $\alpha \in \text{Hom}(C_k, G)$ and $c \in C_\ell$ with $k \neq \ell$ then means that the map (41.4) 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 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 $\alpha \in \text{Hom}(C_k, G)$ 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 (41.1) is therefore necessary in order to make the evaluation $\text{Hom}(C_*, G) \otimes C_* \rightarrow G$ 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 $\text{im } \delta$. 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].

41.2. The singular cochain complex. The **singular cohomology** of a pair of spaces (X, A) with coefficients in an R -module G is now defined by applying the algebraic processing described above to the singular chain complex with coefficients in the ring R : that is,

$$H^*(X, A) = H^*(X, A; G) := H^*(C_*(X, A; R); G) = H^*(\text{Hom}(C_*(X, A; R), G)).$$

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 $\text{Hom}(C_*(X, A; R), G)$ by

$$C^*(X, A) = C^*(X, A; G) := \text{Hom}(C_*(X, A; R), G)$$

and refer to elements of $C^*(X, A)$ as **singular cochains** with coefficients in G . Elements of $\ker \delta \subset C^*(X, A)$ and $\text{im } \delta \subset C^*(X, A)$ are likewise called (singular) **cocycles** and **coboundaries** respectively. The reason to use $C_*(X, A; R)$ in the definition rather than chain groups with a different choice of coefficients is that for each $n \geq 0$, $C_n(X, A; R)$ is a *free* R -module. In particular, homomorphisms $\varphi \in C^n(X)$ from $C_n(X; R)$ to G are uniquely determined by their values on the generators of $C_n(X; R)$, i.e. the singular n -simplices $\sigma \in \mathcal{K}_n(X)$, and we therefore have a canonical identification

$$C^n(X) = G^{\mathcal{K}_n(X)} = \prod_{\sigma \in \mathcal{K}_n(X)} G = \{\text{functions } \varphi : \mathcal{K}_n(X) \rightarrow G\}.$$

We will often use this identification to regard cochains $\varphi \in C^n(X)$ simply as functions $\varphi : \mathcal{K}_n(X) \rightarrow G$. With this understood, we plug in (41.1) and the usual formula for the boundary operator

$\partial : C_{n+1}(X; R) \rightarrow C_n(X; R)$ to find a corresponding formula for the **coboundary operator** $\delta : C^n(X) \rightarrow C^{n+1}(X)$, in the form

$$(41.5) \quad (\delta\varphi)(\sigma) = (-1)^{n+1} \sum_{k=0}^{n+1} (-1)^k \varphi(\sigma|_{\partial_{(k)}\Delta^{n+1}}) \quad \text{for } \varphi : \mathcal{K}_n(X) \rightarrow G, \quad \sigma \in \mathcal{K}_{n+1}(X).$$

In the relative case, we can think of a homomorphism $\varphi : C_n(X, A; R) = C_n(X; R)/C_n(A; R) \rightarrow G$ as equivalent to a homomorphism $\varphi : C_n(X; R) \rightarrow G$ that vanishes on the submodule $C_n(A; R) \subset C_n(X; R)$, so this is the same thing as a function $\mathcal{K}_n(X) \rightarrow G$ that vanishes on the subset $\mathcal{K}_n(A) \subset \mathcal{K}_n(X)$:

$$C^n(X, A) = \left\{ \varphi : \mathcal{K}_n(X) \rightarrow G \mid \varphi|_{\mathcal{K}_n(A)} = 0 \right\}.$$

The formula (41.5) then gives the correct homomorphism $\delta : C^n(X, A) \rightarrow C^{n+1}(X, A)$ by restriction.

As a functor, $H^n = H^n(\cdot; G) : \mathbf{Top}^{\text{rel}} \rightarrow R\text{-Mod}$ is the composition of three functors,

$$\mathbf{Top}^{\text{rel}} \xrightarrow{C_*^*(\cdot; R)} \mathbf{Ch} \xrightarrow{\text{Hom}(\cdot, G)} \mathbf{CoCh} \xrightarrow{H^n} R\text{-Mod},$$

one of which is contravariant, thus $H^n(\cdot; G)$ is also contravariant. Concretely, this means that continuous maps of pairs $f : (X, A) \rightarrow (Y, B)$ induce “pullback” homomorphisms

$$f^* : H^n(Y, B) \rightarrow H^n(X, A)$$

for every $n \in \mathbb{Z}$. These maps are induced by the chain map $f^* : C^*(Y, B) \rightarrow C^*(X, A)$ defined by

$$(f^*\varphi)(c) := \varphi(f_*c) \quad \text{for } \varphi \in C^n(Y, B), \quad c \in C_n(X, A; R).$$

By the previous algebraic discussion, there is a natural pairing

$$H^*(X, A; G) \otimes H_*(X, A; R) \rightarrow G : [\varphi] \otimes [c] \mapsto \langle [\varphi], [c] \rangle := \varphi(c),$$

which we call the **evaluation** of the cohomology class $[\varphi]$ on the homology class $[c]$, and it satisfies

$$(41.6) \quad \langle f^*[\varphi], [c] \rangle = \langle [\varphi], f_*[c] \rangle \quad \text{for } [\varphi] \in H^*(Y, B; G), \quad [c] \in H_*(X, A; R), \quad (X, A) \xrightarrow{f} (Y, B).$$

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 $\gamma : I \rightarrow X$, (41.5) gives

$$(\delta\varphi)(\gamma) = \pm [\varphi(\gamma(1)) - \varphi(\gamma(0))] \quad \text{for } \varphi : X \rightarrow G, \quad \gamma : I \rightarrow X,$$

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 $\varphi : X \rightarrow G$ is therefore a 0-cocycle if and only if it is constant on path-components, meaning it is equivalent to a function $\pi_0(X) \rightarrow G$. We’ve proved:

THEOREM 41.4. *For any space X and R -module G , there is a canonical isomorphism*

$$H^0(X; G) \cong \prod_{\pi_0(X)} G.$$

□

REMARK 41.5. 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 modules $\{G_\alpha\}_{\alpha \in J}$, the direct sum $\bigoplus_{\alpha \in J} G_\alpha$ can be identified with the submodule of the direct product $\prod_{\alpha \in J} G_\alpha$ consisting of tuples $\{g_\alpha\}_{\alpha \in J}$ that have at most finitely-many nonzero coordinates. For

example, if the index set J is \mathbb{N} and $G_\alpha = \mathbb{Z}_2$ for every $\alpha \in J$, then $\bigoplus_{\alpha \in J} G_\alpha$ is countably infinite but $\prod_{\alpha \in J} G_\alpha$ 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 $\{\gamma : I \rightarrow X\}$ to G .

THEOREM 41.6. *For any path-connected pointed space (X, x) and any R -module G , there exists an isomorphism*

$$\Psi : H^1(X; G) \rightarrow \text{Hom}(\pi_1(X, x), G) : [\varphi] \mapsto \Psi_\varphi,$$

defined via the formula

$$\Psi_\varphi([\gamma]) := \varphi(\gamma) \quad \text{for paths } \gamma : I \rightarrow X \text{ with } \gamma(0) = \gamma(1) = x.$$

The proof of the theorem is divided up into Exercises 41.1, 41.2 and 41.3.

41.3. Exercises.

EXERCISE 41.1. 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 $\gamma : I \rightarrow X$, $\varphi(\gamma) \in G$ depends only on the homotopy class of γ with fixed end points;
- (ii) For every pair of paths $\alpha, \beta : I \rightarrow X$ with $\alpha(1) = \beta(0)$, $\varphi(\alpha \cdot \beta) = \varphi(\alpha) + \varphi(\beta)$.

Hint: If $\sigma : \Delta^2 \rightarrow X$ is a singular 2-simplex, one can identify its three boundary faces with paths $\alpha, \beta, \gamma : I \rightarrow X$ such that $\alpha \cdot \beta$ is homotopic to γ with fixed end points.

EXERCISE 41.2. Show that a singular 1-cochain $\varphi \in C^1(X; G)$ is a coboundary if and only if there exists a function⁶⁴ $\psi : X \rightarrow G$ such that for all paths $\gamma : I \rightarrow X$, $\varphi(\gamma) = \psi(\gamma(1)) - \psi(\gamma(0))$.

EXERCISE 41.3. Prove that for any $x \in X$, there is a well-defined homomorphism

$$\Psi : H^1(X; G) \rightarrow \text{Hom}(\pi_1(X, x), G) : [\varphi] \mapsto \Psi_\varphi$$

such that for each 1-cocycle $\varphi \in C^1(X; G)$, $\Psi_\varphi : \pi_1(X, x) \rightarrow G$ is given by

$$\Psi_\varphi([\gamma]) = \varphi(\gamma) \quad \text{for } x \stackrel{\gamma}{\sim} x.$$

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 41.2. 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, $\text{Hom}(\pi_1(X, x), G) = \text{Hom}(H_1(X; \mathbb{Z}), G)$, so the map $\Psi : H^1(X; G) \rightarrow \text{Hom}(\pi_1(X, x), G)$ can then be identified with

$$H^1(X; G) \rightarrow \text{Hom}(H_1(X; \mathbb{Z}), G) : [\varphi] \mapsto \langle [\varphi], \cdot \rangle.$$

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 $C_1(X; \mathbb{Z}) \rightarrow G$. Use the fact that $0 \rightarrow Z_1 \hookrightarrow C_1(X; \mathbb{Z}) \xrightarrow{\partial} B_0 \rightarrow 0$ is a split exact sequence. (Why?)

⁶⁴Note 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 $\psi : X \rightarrow G$.

42. Axioms for cohomology

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 42.1. An **axiomatic cohomology theory** h^* valued in the category of R -modules is a collection $\{h^n\}_{n \in \mathbb{Z}}$ of contravariant functors

$$\mathbf{Top}^{\text{rel}} \xrightarrow{h^n} R\text{-Mod} : (X, A) \mapsto h^n(X, A),$$

which also determine functors $h^n : \mathbf{Top} \rightarrow R\text{-Mod}$ by defining

$$h^n(X) := h^n(X, \emptyset),$$

with maps of pairs $f : (X, A) \rightarrow (Y, B)$ inducing homomorphisms

$$h^n(Y, B) \xrightarrow{f^*} h^n(X, A).$$

The data of the theory also includes natural transformations δ^* from the functor $\mathbf{Top}^{\text{rel}} \rightarrow R\text{-Mod} : (X, A) \mapsto h^n(X, A)$ to the functor $\mathbf{Top}^{\text{rel}} \rightarrow R\text{-Mod} : (X, A) \mapsto h^{n+1}(X, A)$ for each $n \in \mathbb{Z}$, such that the following axioms are satisfied:

- (HOMOTOPY) For any two homotopic maps of pairs $f, g : (X, A) \rightarrow (Y, B)$, the induced morphisms $f^*, g^* : h^*(Y, B) \rightarrow h^*(X, A)$ are identical.
- (EXACTNESS) For all pairs (X, A) with inclusion maps $i : A \hookrightarrow X$ and $j : (X, \emptyset) \hookrightarrow (X, A)$, the sequence

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

is exact.

- (EXCISION) For any pair (X, A) and any subset $B \subset X$ such that there exists a continuous function $u : X \rightarrow I$ equal to 0 on B and 1 on $X \setminus A$, the map induced by the inclusion $(X \setminus B, A \setminus B) \hookrightarrow (X, A)$ is an isomorphism

$$h^n(X, A) \xrightarrow{\cong} h^n(X \setminus B, A \setminus B) \quad \text{for every } n \in \mathbb{Z}.$$

- (DIMENSION) For any one-point space $\{*\}$, $h^n(\{*\}) = 0$ for all $n \neq 0$. The group $h^0(\{*\})$ is then called the **coefficient group** of the cohomology theory.
- (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 $i_\alpha^* : h^* \left(\coprod_{\beta \in J} X_\beta \right) \rightarrow h^*(X_\alpha)$ determine an isomorphism

$$\prod_{\alpha \in J} i_\alpha^* : h^* \left(\coprod_{\beta \in J} X_\beta \right) \xrightarrow{\cong} \prod_{\alpha \in J} h^*(X_\alpha).$$

THEOREM 42.2. For any R -module G , the singular cohomology $H^*(\cdot; G)$ is an axiomatic cohomology theory with coefficient group G .

PROOF. The main reason for the homotopy axiom is Proposition 41.2 in the previous lecture, which implies that if the two chain maps $f_*, g_* : C_*(X, A; R) \rightarrow C_*(Y, B; R)$ are chain homotopic, then so are the two chain maps $f^*, g^* : C^*(Y, B; G) \rightarrow C^*(X, A; G)$.

Exactness follows from the fact that if we dualize the usual short exact sequence of singular chain complexes $0 \rightarrow C_*(A; R) \xrightarrow{i_*} C_*(X; R) \xrightarrow{j_*} C_*(X, A; R) \rightarrow 0$, then the resulting sequence of chain maps

$$(42.1) \quad 0 \leftarrow C^*(A; G) \xleftarrow{i^*} C^*(X; G) \xleftarrow{j^*} C^*(X, A; G) \leftarrow 0$$

is also exact. Indeed, under the canonical identifications of these groups with sets of functions $\mathcal{K}_n(X) \rightarrow G$ or $\mathcal{K}_n(A) \rightarrow G$, j^* becomes the obvious inclusion

$$j^* : \left\{ \varphi : \mathcal{K}_n(X) \rightarrow G \mid \varphi|_{\mathcal{K}_n(A)} = 0 \right\} \hookrightarrow \left\{ \varphi : \mathcal{K}_n(X) \rightarrow G \right\},$$

and i^* becomes the restriction map

$$i^* : \left\{ \varphi : \mathcal{K}_n(X) \rightarrow G \right\} \rightarrow \left\{ \varphi : \mathcal{K}_n(A) \rightarrow G \right\} : \varphi \mapsto \varphi|_{\mathcal{K}_n(A)},$$

which is manifestly surjective and has kernel equal to $\text{im } j^*$. I should caution you against thinking that the exactness of this dualized sequence follows automatically from abstract nonsense—we will see when we study the universal coefficient theorem that not every short exact sequence remains exact after it is dualized. But this one does. As a result, (42.1) 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 excision property is where we need to make use of the *chain-level* excision result established in Theorem 33.9. Indeed, if $B \subset \bar{B} \subset \dot{A} \subset A \subset X$, then the inclusion $i : (X \setminus B, A \setminus B) \hookrightarrow (X, A)$ induces a chain homotopy equivalence $i_* : C_*(X \setminus B, A \setminus B; R) \rightarrow C_*(X, A; R)$, meaning in particular that there is a chain map $\rho_* : C_*(X, A; R) \rightarrow C_*(X \setminus B, A \setminus B; R)$ such that $\rho_* i_*$ and $i_* \rho_*$ are each chain homotopic to the identity. Dualizing both i_* and ρ_* then produces chain maps $i^* : C^*(X, A; G) \rightarrow C^*(X \setminus B, A \setminus B; G)$ and $\rho^* : C^*(X \setminus B, A \setminus B; G) \rightarrow C^*(X, A; G)$ such that by Proposition 41.2, $i^* \rho^*$ and $\rho^* i^*$ are also chain homotopic to the identity, hence

$$i^* : C^*(X, A; G) \rightarrow C^*(X \setminus B, A \setminus B; G)$$

is a chain homotopy equivalence and induces an isomorphism $H^*(X, A; G) \rightarrow H^*(X \setminus B, A \setminus B; G)$.

The dimension axiom and the computation of the coefficient group are straightforward since there is only one singular n -simplex $\sigma_n \in \mathcal{K}_n(\{*\})$ for each $n \geq 0$, giving canonical isomorphisms

$$C^n(\{*\}; G) \xrightarrow{\cong} G : \varphi \mapsto \varphi(\sigma_n).$$

The map $\delta : C^n(\{*\}; G) \rightarrow C^{n+1}(\{*\}; G)$ then becomes

$$\delta_n : G \rightarrow G : g \mapsto (-1)^{n+1} \sum_{k=0}^{n+1} (-1)^k g = \begin{cases} 0 & \text{if } n \text{ is even,} \\ (-1)^{n+1} g & \text{if } n \text{ is odd.} \end{cases}$$

For $n > 0$ even, this means $\ker \delta_n = \text{im } \delta_{n-1}$ and thus $H^n(\{*\}; G) = 0$. For $n > 0$ odd, we instead have $\ker \delta_n = 0$ and thus $H^n(\{*\}; G) = 0$. The only special case is $n = 0$, for which $H^0(\{*\}; G) = \ker \delta_0 = G$.

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 $C_*\left(\coprod_{\beta} X_{\beta}; R\right)$ splits naturally into a direct sum of chain complexes $\bigoplus_{\beta} C_*(X_{\beta}; R)$. 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

We mention in passing that there is also a cohomological version of the exact sequence of triples (34.3); see Exercise 42.1.

42.1. Reduced cohomology. Every cohomology theory h^* also has a **reduced** version, which is again defined in terms of the unique map

$$\epsilon : X \rightarrow \{*\}.$$

Choosing any embedding $i : \{*\} \rightarrow X$, the fact that $\epsilon \circ i$ is the identity map implies that

$$(\epsilon \circ i)^* = i^* \epsilon^* : h^*(\{*\}) \rightarrow h^*(\{*\})$$

is also the identity, so $\epsilon^* : h^*(\{*\}) \rightarrow h^*(X)$ is injective and has i^* as a left-inverse. We then define

$$\tilde{h}^*(X) := \text{coker } \epsilon^* = h^*(X) / \text{im } \epsilon^*,$$

so that the quotient projection $h^*(X) \rightarrow \tilde{h}^*(X)$ fits into a split exact sequence

$$0 \longrightarrow h^*(\{*\}) \xrightarrow{\epsilon^*} h^*(X) \longrightarrow \tilde{h}^*(X) \longrightarrow 0,$$

implying via the dimension axiom that if h^* has coefficient group G ,

$$h^n(X) \cong \begin{cases} \tilde{h}^n(X) \oplus G & \text{for } n = 0, \\ \tilde{h}^n(X) & \text{for } n \neq 0. \end{cases}$$

If X is contractible, then ϵ is a homotopy equivalence and $\epsilon^* : h^*(\{*\}) \rightarrow h^*(X)$ is thus an isomorphism, so its cokernel is trivial:

THEOREM 42.3. *For any axiomatic cohomology theory h^* , if X is contractible, $\tilde{h}^*(X) = 0$. \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 28 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 $A \neq \emptyset$ is defined to match the ordinary homology. We can do the same thing here: if we define

$$\tilde{h}^*(X, A) := h^*(X, A) \quad \text{if } A \neq \emptyset,$$

then repeating the arguments of Lecture 28 with reversed arrows gives:

THEOREM 42.4. *For any pair (X, A) and any axiomatic cohomology theory, the sequence*

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

is also well defined and exact. \square

It is a straightforward exercise to verify:

PROPOSITION 42.5. *For any space X and R -module G , the reduced singular cohomology $\tilde{H}^*(X; G)$ is also the cohomology (with coefficients in G) of the **augmented** chain complex*

$$\dots \longrightarrow C_2(X; R) \xrightarrow{\partial} C_1(X; R) \xrightarrow{\partial} C_0(X; R) \xrightarrow{\epsilon} \tilde{C}_{-1}(X; R) := R \longrightarrow 0 \longrightarrow 0 \longrightarrow \dots,$$

described in Exercise 33.4. \square

Other useful features of axiomatic homology that carry over to axiomatic cohomology with minimal effort include the computation

$$h^k(S^n) \cong \begin{cases} G & \text{if } k = 0 \text{ or } k = n, \\ 0 & \text{otherwise,} \end{cases}$$

the suspension isomorphisms

$$\tilde{h}^n(X) \cong \tilde{h}^{n+1}(\Sigma X),$$

and also the isomorphisms

$$h^n(X, A) \cong \tilde{h}^n(X/A)$$

for good pairs (X, A) . The details are worked out in the exercises at the end of this lecture.

REMARK 42.6. You may by now be getting the impression that cohomology is always isomorphic to homology, especially in light of the computation of $h^*(S^n)$ quoted above. There is a grain of truth in this, but the whole story is more complicated: e.g. we will see from the universal coefficient theorem that $H^*(X; G)$ is fully determined up to isomorphism by $H_*(X; R)$ and G , but it is not always isomorphic to $H_*(X; G)$, especially e.g. if we work over the ring $R = \mathbb{Z}$ and $H_*(X; \mathbb{Z})$ 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 $H_*(\cdot; R)$, not every axiomatic cohomology theory h^* has a corresponding axiomatic homology theory h_* (cf. Remark 42.10 at the end of this lecture).

42.2. The Mayer-Vietoris sequence. Under suitable assumptions on a space $X = A \cup B$ that is the union of two subsets, one can use a diagram chase as in Theorem 34.21 to derive from the axioms a Mayer-Vietoris sequence for any axiomatic cohomology theory. This is essentially just a matter of redoing all the arguments of §34.5 with reversed arrows, and we shall leave the details as an exercise.

For singular cohomology, the Mayer-Vietoris sequence can be also be seen more directly, and under the same assumptions as its homological counterpart. Suppose $A, B \subset X$ are subsets with inclusion maps

$$(42.2) \quad \begin{array}{ccc} & A & \\ i_A \nearrow & & \searrow j_A \\ A \cap B & & X \\ i_B \searrow & & \nearrow j_B \\ & B & \end{array}$$

such that $X = A \cup B$ and (A, B) is an excisive couple for singular homology, e.g. because the interiors of A and B cover X . Let

$$C_*(A; R) \xleftarrow{J_A} C_*(A + B; R) \xleftarrow{J_B} C_*(B; R)$$

denote the obvious inclusions of subcomplexes of $C_*(X; R)$, where we recall the notation

$$C_*(A + B; R) := C_*(A; R) + C_*(B; R) \subset C_*(X; R).$$

The Mayer-Vietoris sequence in singular homology was derived in §34.4 from a short exact sequence of chain complexes in the form

$$0 \longrightarrow C_*(A \cap B) \xrightarrow{((i_A)_*, -(i_B)_*)} C_*(A) \oplus C_*(B) \xrightarrow{J_A \oplus J_B} C_*(A + B) \longrightarrow 0.$$

For the present discussion, we take the coefficient module in the singular chain complex to be the ring R , and apply the functor $\text{Hom}(\cdot; G)$ to this short exact sequence; in light of the natural isomorphism

$$\begin{aligned} \text{Hom}(C_*(A), G) \oplus \text{Hom}(C_*(B), G) &\cong \text{Hom}(C_*(A) \oplus C_*(B), G) \\ (\varphi, \psi) &\mapsto \varphi \oplus \psi \end{aligned}$$

the result can be written as

$$0 \longleftarrow C^*(A \cap B; G) \xleftarrow{(i_A^* \oplus (-i_B^*))} C^*(A; G) \oplus C^*(B; G) \xleftarrow{(J_A^*, J_B^*)} C^*(A + B; G) \longleftarrow 0,$$

where we are abbreviating

$$C^*(A + B; G) := \text{Hom}(C_*(A + B; R), G).$$

The dual maps

$$\begin{aligned} i_A^* : C^*(A; G) &\rightarrow C^*(A \cap B; G), & i_B^* : C^*(B; G) &\rightarrow C^*(A \cap B; G), \\ J_A^* : C^*(A + B; G) &\rightarrow C^*(A; G), & J_B^* : C^*(A + B; G) &\rightarrow C^*(B; G) \end{aligned}$$

are all canonical restriction maps, e.g. J_A^* replaces a homomorphism $\varphi : C_*(A + B; R) \rightarrow G$ with its restriction to the submodule $C_*(A; R)$. 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 cohomology of the cochain complex $C^*(A + B; G)$ with something more familiar. The assumption that (A, B) is an excisive couple means that the inclusion $C_*(A + B; \mathbb{Z}) \hookrightarrow C_*(X; \mathbb{Z})$ induces an isomorphism on homology, but as explained in Remark 34.13, this implies the seemingly stronger statement that the inclusion of chain complexes is a chain homotopy equivalence, and it follows in turn that the inclusion $J : C_*(A + B; R) \hookrightarrow C_*(X; R)$ is likewise a chain homotopy equivalence. Now by Proposition 41.2, so is its dualization

$$J^* : C^*(X; G) \rightarrow C^*(A + B; G),$$

which therefore induces an isomorphism

$$H^*(X; G) \xrightarrow{\cong} H^*(C^*(A + B; G)).$$

Combining this with the usual diagram-chasing result gives:

THEOREM 42.7 (Mayer-Vietoris sequence for cohomology). *If $A, B \subset X$ are subsets with inclusions written as in (42.2) such that $X = A \cup B$ and (A, B) is an excisive couple for singular homology, then there exist connecting homomorphisms $\delta^* : H^n(A \cap B; G) \rightarrow H^{n+1}(X; G)$ for every $n \in \mathbb{Z}$ such that the sequence*

$$\begin{aligned} \dots \leftarrow H^{n+1}(X; G) \xleftarrow{\delta^*} H^n(A \cap B; G) \xleftarrow{i_A^* \oplus (-i_B^*)} H^n(A; G) \oplus H^n(B; G) \\ \xleftarrow{(j_A^*, j_B^*)} H^n(X; G) \xleftarrow{\delta^*} H^{n-1}(A \cap B; G) \leftarrow \dots \end{aligned}$$

is exact, and this sequence is also natural with respect to maps $f : X \rightarrow X' = \mathring{A}' \cup \mathring{B}'$ satisfying $f(A) \subset A'$ and $f(B) \subset B'$. \square

42.3. 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,

$$H_{\text{CW}}^*(X, A; G) := H_*(C_{\text{CW}}^*(X, A; G)),$$

where we define the **cellular cochain complex**

$$C_{\text{CW}}^*(X, A; G) := \text{Hom}(C_{\text{CW}}^{\text{CW}}(X, A; R), G).$$

This gives a collection of contravariant functors $H_{\text{CW}}^n : \text{CW}^{\text{rel}} \rightarrow R\text{-Mod}$ that are typically not very hard to compute. The coboundary map $\delta : C_{\text{CW}}^n(X, A; G) \rightarrow C_{\text{CW}}^{n+1}(X, A; G)$ can be expressed in terms of the same incidence numbers that describe the cellular boundary map: indeed, for each n -cell $e_\alpha^n \subset X$, define its dual cochain

$$\varphi_\alpha^n \in C_{\text{CW}}^n(X; R), \quad \varphi_\alpha^n(e_\beta^n) := \begin{cases} 1 & \text{if } \beta = \alpha, \\ 0 & \text{otherwise.} \end{cases}$$

These generators form a basis of $C_{\text{CW}}^n(X; R)$ if there are only finitely many n -cells, and any element of $C_{\text{CW}}^n(X; G)$ can then similarly be described as a linear combination of the φ_α^n 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 $\sum_{e_\alpha^n \subset X} g_\alpha \varphi_\alpha^n$, which have well-defined evaluations on $C_n^{CW}(X; R)$ 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 $\text{Hom}(R, G) \cong G$ defined by evaluation on the generator $1 \in R$, there are canonical isomorphisms

$$C_{CW}^n(X; G) = \text{Hom}(C_n^{CW}(X; R), G) = \text{Hom}\left(\bigoplus_{e_\alpha^n \subset X} R, G\right) \cong \prod_{e_\alpha^n \subset X} \text{Hom}(R, G) \cong \prod_{e_\alpha^n \subset X} G,$$

and $\sum_{e_\alpha^n \subset X} g_\alpha \varphi_\alpha^n \in C_{CW}^n(X; G)$ is to be interpreted as the element that corresponds canonically with $\{g_\alpha\} \in \prod_{e_\alpha^n \subset X} G$. With this understood, we shall abuse terminology by calling the elements φ_α^n a “basis” of $C_{CW}^n(X; R)$, and observe that the coboundary operator $\delta : C_{CW}^n(X; G) \rightarrow C_{CW}^{n+1}(X; G)$ is uniquely determined if we write down a formula for $\delta\varphi_\alpha^n \in C^{n+1}(X; R)$ for each n -cell $e_\alpha^n \subset X$. For any $(n + 1)$ -cell $e_\beta^{n+1} \subset X$, we have

$$(\delta\varphi_\alpha^n)(e_\beta^{n+1}) = (-1)^{n+1} \varphi_\alpha^n(\partial e_\beta^{n+1}) = (-1)^{n+1} \sum_{e_\gamma^n \subset X} \varphi_\alpha^n([e_\gamma^n : e_\beta^{n+1}]e_\gamma^n) = (-1)^{n+1} [e_\alpha^n : e_\beta^{n+1}],$$

thus the required formula is

$$\delta\varphi_\alpha^n = (-1)^{n+1} \sum_{e_\beta^{n+1} \subset X} [e_\alpha^n : e_\beta^{n+1}] \varphi_\beta^{n+1}.$$

As with homology, cellular cohomology provides a powerful tool for computing arbitrary axiomatic cohomology theories on spaces that have cell decompositions:

THEOREM 42.8. *For any axiomatic cohomology theory h^* with coefficient group G and every CW-pair (X, A) , there exist isomorphisms $H_{CW}^n(X, A; G) \rightarrow h^n(X, A)$ for every $n \in \mathbb{Z}$, and they are natural in the sense that every cellular map $f : (X, A) \rightarrow (Y, B)$ gives rise to a commutative diagram*

$$\begin{array}{ccc} H_{CW}^n(X, A; G) & \xrightarrow{\cong} & h^n(X, A) \\ f^* \uparrow & & f^* \uparrow \\ H_{CW}^n(Y, B; G) & \xrightarrow{\cong} & h^n(Y, B) \end{array}$$

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 38. One starts by deriving from the axioms (in particular the correspondence $h^*(X, A) \cong \tilde{h}^*(X/A)$ for good pairs) a natural isomorphism

$$h^k(X^n, X^{n-1}) \cong \prod_{e_\alpha^n \subset X} h^k(\mathbb{D}^n, \partial\mathbb{D}^n)$$

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 $h^n(\mathbb{D}^n, \partial\mathbb{D}^n) \cong \tilde{h}^{n-1}(S^{n-1}) \cong G$) is identical to the cellular n -cochain group $C_{CW}^n(X; G) \cong \prod_{e_\alpha^n \subset X} G$ when $k = n$. Putting $h^n(X^n, X^{n-1})$ for each $n \geq 0$ in the role of $C_{CW}^n(X; G)$, one then assembles the long exact sequences of (X^{n+1}, X^n) and (X^n, X^{n-1}) into the

Instead, $\bar{C}^n(X)$ 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 modules $\bar{C}^n(X)$ for $n \geq 0$ into a cochain complex, we associate to each function $\varphi : X^{n+1} \rightarrow G$ the function $\delta\varphi : X^{n+2} \rightarrow G$ defined by

$$(\delta\varphi)(x_0, \dots, x_{n+1}) := \sum_{k=0}^{n+1} (-1)^k \varphi(x_0, \dots, x_{k-1}, x_{k+1}, \dots, x_{n+1}).$$

This defines a homomorphism from the module of $(n+1)$ -functions to the module of $(n+2)$ -functions such that $\delta^2 = 0$, and it preserves the submodule of functions that vanish near the diagonal, thus it descends to a coboundary homomorphism

$$\delta : \bar{C}^n(X) \rightarrow \bar{C}^{n+1}(X).$$

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 cohomology is the **Alexander-Spanier cohomology** of X with coefficients in G , denoted by

$$\bar{H}^n(X) = \bar{H}^n(X; G) := H^n(\bar{C}^*(X; G)), \quad \text{for } n \in \mathbb{Z}.$$

It is not hard to give \bar{H}^* the structure of a contravariant functor: given a continuous map $f : X \rightarrow Y$, one defines a chain map

$$f^* : \bar{C}^*(Y) \rightarrow \bar{C}^*(X) : \varphi \mapsto \varphi \circ (f \times \dots \times f),$$

thus inducing homomorphisms $f^* : \bar{H}^*(Y) \rightarrow \bar{H}^*(X)$. 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 $\varphi : X \rightarrow G$; 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 $X \times X \rightarrow G$ 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 $X \rightarrow G$, 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 $X \rightarrow G$, 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 41.4, but not always: $\bar{H}^0(X)$ and $H^0(X)$ differ on spaces whose connected components and path-components do not match.

Elements of $\bar{C}^1(X)$ are represented by functions $\varphi : X \times X \rightarrow G$, and such a function represents a coboundary if and only if there exists a function $\psi : X \rightarrow G$ 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 proof of the following result in order to gain some intuition on what this cocycle condition means, and what $\bar{H}^1(X)$ actually measures.

PROPOSITION 42.9. *For any path-connected space X and any choice of base point $p \in X$, there is a well-defined and injective homomorphism*

$$\bar{H}^1(X; G) \rightarrow \text{Hom}(\pi(X, p), G) : [\varphi] \mapsto \Psi_{[\varphi]}$$

such that for any representative cocycle $\varphi \in \bar{C}^1(X; G)$ and loop $\gamma : [0, 1] \rightarrow X$ from p to itself, one can choose a sufficiently fine partition $0 =: t_0 < t_1 < \dots < t_{N-1} < t_N := 1$ of $[0, 1]$ to compute

$$\Psi_{[\varphi]}([\gamma]) = \sum_{j=1}^N \varphi(\gamma(t_j), \gamma(t_{j-1})).$$

The Alexander-Spanier theory \bar{H}^* satisfies an “extra” axiom that singular cohomology does not, the so-called *continuity* axiom, which we will come back to in the next lecture since it involves inverse limits (cf. Theorem 43.25). For this reason, \bar{H}^* is sometimes useful in applications that involves spaces which cannot be assumed to be as nice as CW-complexes.

REMARK 42.10. 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 certain technical drawbacks that we will also see in the next lecture in the context of Čech homology, namely it fails in general to satisfy the exactness axiom.

42.5. Exercises.

EXERCISE 42.1. Describe a cohomological version of the “braid” diagram (34.2) 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

$$\dots \longleftarrow h^{n+1}(X, A) \xleftarrow{\delta^*} h^n(A, B) \xleftarrow{i^*} h^n(X, B) \xleftarrow{j^*} h^n(X, A) \xleftarrow{\delta^*} h^{n-1}(A, B) \longleftarrow \dots$$

Give also an alternative proof of this for singular cohomology using a short exact sequence of cochain complexes.

EXERCISE 42.2 (*). Adapt the proof of Theorem 28.24 to prove that for any axiomatic cohomology theory h^* and any space X , there is a natural isomorphism $\tilde{h}^n(X) \rightarrow \tilde{h}^{n+1}(\Sigma X)$ for every $n \in \mathbb{Z}$.

EXERCISE 42.3 (*). For any axiomatic cohomology theory h^* and two spaces X and Y with maps $\epsilon_X : X \rightarrow \{*\}$ and $\epsilon_Y : Y \rightarrow \{*\}$, show that the isomorphism $h^*(X \amalg Y) \cong h^*(X) \times h^*(Y)$ given by the additivity axiom identifies $\tilde{h}_*(X \amalg Y)$ with the cokernel of the map

$$(\epsilon_X^*, \epsilon_Y^*) : h^*(\{*\}) \rightarrow h^*(X) \times h^*(Y).$$

Then apply this in the case $X = Y = \{*\}$ to identify $\tilde{h}^0(\{*\} \amalg \{*\})$ with the cokernel of the diagonal map $G \rightarrow G \times G$, where $G = h^0(\{*\})$. Conclude in particular

$$\tilde{h}^n(S^0) \cong \begin{cases} G & \text{if } n = 0, \\ 0 & \text{if } n \neq 0. \end{cases}$$

EXERCISE 42.4 (*). 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 42.5 (*). Adapt the proof of Theorem 34.7 to prove that for any axiomatic cohomology theory h^* and any good pair (X, A) , there is a natural isomorphism

$$h^*(X, A) \cong \tilde{h}^*(X/A).$$

EXERCISE 42.6. Adapt the diagram-chasing arguments in Lecture 34 to show that every axiomatic cohomology theory h^* admits a Mayer-Vietoris sequence under a suitable hypothesis on $X = A \cup B$, and that it also works if h^* is replaced by \tilde{h}^* .

EXERCISE 42.7 (*). Work out the further details of the proof of Theorem 42.8 for finite-dimensional CW-pairs.

EXERCISE 42.8. Prove Proposition 42.9, and show moreover that the map $\bar{H}^1(X; G) \rightarrow \text{Hom}(\pi_1(X, p), G)$ is an isomorphism if X is S^1 or \mathbb{R} .

43. Inverse limits, CW-complexes and Čech theory

As an initial goal in this lecture, we'd like to do for cohomology what was done for homology in Lecture 39: extend the isomorphism $H_{\text{CW}}^*(X, A; G) \cong h^*(X, A)$ so that it also applies to infinite-dimensional CW-pairs (X, A) . This necessitates a quick introduction to inverse limits, and we will observe a peculiarity of inverse limits in comparison with direct limits that makes the discussion slightly more complicated than its homological counterpart. Having inverse limits in the picture also gives us an opportunity to sketch a first example of an axiomatic homology theory other than singular homology: we will give a brief overview of Čech homology, and see in the process that it fails for technical reasons to satisfy *all* of the Eilenberg-Steenrod axioms without substantial caveats. On the other hand, its cohomological counterpart does not have this drawback: Čech cohomology is a fully legitimate axiomatic cohomology theory, and is a popular tool in certain branches of mathematics, especially algebraic geometry.

43.1. Inverse systems, targets and limits. It is easy to see how inverse limits naturally arise in the context of cellular cohomology. For an infinite-dimensional CW-complex X , we saw in Lecture 39 that X is the direct limit of a direct system of topological spaces $\{X^n\}_{n=0}^\infty$, consisting of the finite-dimensional skeleta of X with their obvious inclusions. Applying any homology functor $h_k : \text{Top} \rightarrow R\text{-Mod}$ to this system gives a direct system of R -modules $\{h_k(X^n)\}_{n=0}^\infty$, but something different happens if we instead apply a cohomology functor h^k , due to contravariance: we obtain a sequence of R -modules $\{h^k(X^n)\}_{n=0}^\infty$ together with the homomorphisms

$$h^k(X^n) \longrightarrow h^k(X^m) \quad \text{induced by inclusions} \quad X^m \longleftarrow X^n .$$

The key detail here is that the maps $h^k(X^n) \rightarrow h^k(X^m)$ are defined whenever $n \geq m$ rather than $n \leq m$, thus $\{h^k(X^n)\}_{n=0}^\infty$ with these maps does not define a direct system over the directed set (\mathbb{N}_0, \leq) . It is instead an example of an inverse system, and therefore requires a different notion of limit.

DEFINITION 43.1. Given a category \mathcal{C} and a directed set $(J, <)$, an **inverse system (projektives System)** $\{X_\alpha, \varphi_{\alpha\beta}\}$ in \mathcal{C} over $(J, <)$ associates to each $\alpha \in J$ an object X_α of \mathcal{C} , along with morphisms

$$\varphi_{\alpha\beta} \in \text{Hom}(X_\beta, X_\alpha) \quad \text{for each} \quad \alpha < \beta$$

such that

$$\varphi_{\alpha\alpha} = \text{Id}_{X_\alpha}$$

and the diagram

$$\begin{array}{ccccc} X_\alpha & \xleftarrow{\varphi_{\alpha\beta}} & X_\beta & \xleftarrow{\varphi_{\beta\gamma}} & X_\gamma \\ & \searrow & \swarrow & & \\ & & \varphi_{\alpha\gamma} & & \end{array}$$

commutes for every triple $\alpha, \beta, \gamma \in J$ with $\alpha < \beta < \gamma$.

REMARK 43.2. In terms of the category \mathcal{J} corresponding to the directed set $(J, <)$ as in Remark 39.2, an inverse system in \mathcal{C} over $(J, <)$ is the same thing as a contravariant functor $\mathcal{J} \rightarrow \mathcal{C}$.

Convergence of inverse systems is defined analogously to direct systems, the main difference being that most arrows go the other way.

DEFINITION 43.3. For an inverse system $\{X_\alpha, \varphi_{\alpha\beta}\}$ in \mathcal{C} over $(J, <)$, a **target** $\{Y, f_\alpha\}$ of the system consists of an object Y of \mathcal{C} together with associated morphisms $f_\alpha \in \text{Hom}(Y, X_\alpha)$ for each $\alpha \in J$ such that the diagram

$$\begin{array}{ccc} X_\alpha & \xleftarrow{\varphi_{\alpha\beta}} & X_\beta \\ & \searrow f_\alpha & \swarrow f_\beta \\ & & Y \end{array}$$

commutes for every pair $\alpha, \beta \in J$ with $\alpha < \beta$.

DEFINITION 43.4. A target $\{X_\infty, \varphi_\alpha\}$ of the inverse system $\{X_\alpha, \varphi_{\alpha\beta}\}$ is called an **inverse limit**⁶⁵ (*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 $f_\infty \in \text{Hom}(Y, X_\infty)$ such that the diagram

$$\begin{array}{ccc} X_\alpha & \xleftarrow{\varphi_\alpha} & X_\infty \\ & \searrow f_\alpha & \swarrow f_\infty \\ & & Y \end{array}$$

commutes for every $\alpha \in J$.

The meaning of an inverse limit can be encoded in the diagram

$$\begin{array}{ccccccc} X_\alpha & \xleftarrow{\varphi_{\alpha\beta}} & X_\beta & \xleftarrow{\varphi_{\beta\gamma}} & X_\gamma & \xleftarrow{\dots} & \varprojlim \{X_\alpha\} \\ & & & & & & \uparrow \text{dashed} \\ & & & & & & Y \end{array}$$

where we assume $\alpha < \beta < \gamma < \dots \in J$, 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

⁶⁵Inverse 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 39.8).

claim that it is nonempty—the empty set is also a topological space and can appear as the limit of an inverse system in \mathbf{Top} (see Example 43.10 below).

The proofs of the following propositions are worthwhile exercises.

PROPOSITION 43.5. *If $\{X_\alpha, \varphi_{\alpha\beta}\}$ is an inverse system in \mathbf{Top} over $(J, <)$, then its inverse limit is the space*

$$\varprojlim\{X_\alpha\} = \left\{ \{x_\alpha\} \in \prod_{\alpha \in J} X_\alpha \mid x_\alpha = \varphi_{\alpha\beta}(x_\beta) \text{ for all } \alpha, \beta \in J \text{ with } \alpha < \beta \right\},$$

with the associated morphisms $\varphi_\alpha : \varprojlim\{X_\beta\} \rightarrow X_\alpha$ defined via the natural projections $\prod_{\beta \in J} X_\beta \rightarrow X_\alpha$ for each $\alpha \in J$. Moreover, the topology on $\varprojlim\{X_\alpha\}$ is the weakest for which the maps $\varphi_\alpha : \varprojlim\{X_\beta\} \rightarrow X_\alpha$ are all continuous. \square

REMARK 43.6. Proposition 43.5 extends in obvious ways to describe inverse limits in the categories \mathbf{Set} of sets and $\mathbf{Top}^{\text{rel}}$ of pairs of spaces.

PROPOSITION 43.7. *Consider an inverse system $\{X_\alpha, \varphi_{\alpha\beta}\}$ in \mathbf{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 $\varphi_{\alpha\beta} : X_\beta \rightarrow X_\alpha$ are all inclusions. Then $\varprojlim\{X_\alpha\} = \bigcap_{\alpha \in J} X_\alpha$, with the associated morphisms $\varphi_\alpha : \varprojlim\{X_\beta\} \rightarrow X_\alpha$ given by the obvious inclusions. \square*

REMARK 43.8. The obvious analogue of Proposition 43.7 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 39.7 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 \mathbf{Top} are somewhat better behaved than direct systems.

The following result about inverse limits of compact Hausdorff spaces is a consequence of Tychonoff's theorem (see Exercise 43.3):

PROPOSITION 43.9. *For any inverse system $\{X_\alpha, \varphi_{\alpha\beta}\}$ of topological spaces such that every X_α is nonempty, compact and Hausdorff, $\varprojlim\{X_\alpha\} \neq \emptyset$. \square*

EXAMPLE 43.10. Combining the previous two propositions 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 Proposition 43.7, the inverse limit is

$$\varprojlim\{(0, 1/n]\}_{n \in \mathbb{N}} = \bigcap_{n \in \mathbb{N}} (0, 1/n] = \emptyset.$$

PROPOSITION 43.11. *If $\{G_\alpha, \varphi_{\alpha\beta}\}$ is an inverse system in $R\text{-Mod}$ over $(J, <)$, its inverse limit is a module of the form*

$$\varprojlim\{G_\alpha\} = \left\{ \{g_\alpha\} \in \prod_{\alpha \in J} G_\alpha \mid g_\alpha = \varphi_{\alpha\beta}(g_\beta) \text{ for all } \alpha, \beta \in J \text{ with } \alpha < \beta \right\},$$

with the associated homomorphisms $\varphi_\alpha : \varprojlim\{G_\beta\} \rightarrow G_\alpha$ defined via the projections $\prod_{\beta \in J} G_\beta \rightarrow G_\alpha$ all $\alpha \in J$. \square

REMARK 43.12. There is an obvious analogue of Proposition 43.11 for inverse systems in the category $\mathbf{Ch} = \mathbf{Ch}(R\text{-Mod})$ of chain complexes of R -modules.

For the following result, we say that a subset $J_0 \subset J$ of a directed set $(J, <)$ is **cofinal** if for every $\alpha \in J$ there exists some $\beta \in J_0$ such that $\beta > \alpha$.

PROPOSITION 43.13. *Assume $\{X_\alpha, \varphi_{\alpha\beta}\}$ is an inverse system over $(J, <)$ in any category, and suppose $J_0 \subset J$ is a cofinal set with the property that for every $\alpha, \beta \in J_0$ with $\alpha < \beta$, $\varphi_{\alpha\beta} \in \text{Hom}(X_\beta, X_\alpha)$ is an isomorphism. Then $\varinjlim\{X_\alpha\}$ is isomorphic to X_γ for any $\gamma \in J_0$. \square*

43.2. Exactness and the derived inverse limit functor. I want to point out a technical issue that sometimes makes inverse limits in the category of R -modules a bit trickier to deal with than direct limits.

The following pleasant feature of direct limits in $R\text{-Mod}$ was observed in Proposition 39.21: for a direct system of chain complexes, the homology of the direct limit is naturally isomorphic to the direct limit of the homologies. This was used in §39.4 to give a direct proof that

$$H_*(X) \cong H_*(\varinjlim\{X^n\}) \cong \varinjlim\{H_*(X^n)\} \cong \varinjlim\{H_*^{\text{CW}}(X^n)\} \cong H_*^{\text{CW}}(X)$$

for the direct system $\{X^n\}$ consisting of the finite-dimensional skeleta of a CW-complex X . If you try to prove the analogous result about inverse limits of chain complexes, you'll find that you get stuck: homology and inverse limits are two functors that, in general, do not commute with each other.

The problem can be distilled to a special case involving chain complexes with trivial homology, i.e. exact sequences. Consider first a direct system $\{A_n\}_{n=0}^\infty$ of R -modules given in the form of a sequence

$$A_0 \xrightarrow{\phi_0} A_1 \xrightarrow{\phi_1} A_2 \xrightarrow{\phi_2} \dots$$

By Proposition 39.14, the direct limit of $\{A_n\}$ can be expressed as the cokernel of the homomorphism

$$(43.1) \quad \Phi := \begin{pmatrix} \mathbf{1} & 0 & 0 & 0 & \dots \\ -\phi_0 & \mathbf{1} & 0 & 0 & \dots \\ 0 & -\phi_1 & \mathbf{1} & 0 & \dots \\ 0 & 0 & -\phi_2 & \mathbf{1} & \dots \\ 0 & 0 & 0 & -\phi_3 & \dots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{pmatrix} : \bigoplus_{n=0}^\infty A_n \longrightarrow \bigoplus_{n=0}^\infty A_n.$$

In the mapping telescope discussion in §39.3, we made use of the fact that Φ is always injective, because its kernel appeared in an exact sequence, thus giving rise to an isomorphism between its cokernel (i.e. a direct limit) and the homology of the mapping torus.

Here is another way in which the injectivity of Φ is useful. Suppose we have a short exact sequence of sequences $0 \rightarrow \{A_n\} \xrightarrow{f} \{B_n\} \xrightarrow{g} \{C_n\} \rightarrow 0$, meaning a diagram of the form

$$\begin{array}{ccccccccc} 0 & \longrightarrow & A_0 & \xrightarrow{f_0} & B_0 & \xrightarrow{g_0} & C_0 & \longrightarrow & 0 \\ & & \downarrow \phi_0 & & \downarrow \psi_0 & & \downarrow \theta_0 & & \\ 0 & \longrightarrow & A_1 & \xrightarrow{f_1} & B_1 & \xrightarrow{g_1} & C_1 & \longrightarrow & 0 \\ & & \downarrow \phi_1 & & \downarrow \psi_1 & & \downarrow \theta_1 & & \\ 0 & \longrightarrow & A_2 & \xrightarrow{f_2} & B_2 & \xrightarrow{g_2} & C_2 & \longrightarrow & 0 \\ & & \downarrow \phi_2 & & \downarrow \psi_2 & & \downarrow \theta_2 & & \\ & & \vdots & & \vdots & & \vdots & & \\ & & \downarrow & & \downarrow & & \downarrow & & \\ & & \varinjlim\{A_n\} & \xrightarrow{f} & \varinjlim\{B_n\} & \xrightarrow{g} & \varinjlim\{C_n\} & & \end{array}$$

in which all the rows are exact, except possibly the bottom row, containing the direct limits. Note that the maps $f : \varinjlim\{A_n\} \rightarrow \varinjlim\{B_n\}$ and $g : \varinjlim\{B_n\} \rightarrow \varinjlim\{C_n\}$ are uniquely determined by the commutativity of this diagram, due to the universal property of direct limits, e.g. the diagram makes $\varinjlim\{B_n\}$ a target of $\{A_n\}$ and thus determines uniquely a map $\varinjlim\{A_n\} \rightarrow \varinjlim\{B_n\}$. If you like Proposition 39.21, then you can apply it here—we have a special case involving chain complexes with trivial homology—and the conclusion will be that the bottom row of the this diagram must in fact form a short exact sequence

$$(43.2) \quad 0 \longrightarrow \varinjlim\{A_n\} \xrightarrow{f} \varinjlim\{B_n\} \xrightarrow{g} \varinjlim\{C_n\} \longrightarrow 0.$$

But without using Proposition 39.21, we can also see this as follows. The three direct limits can be expressed as cokernels of injective homomorphisms

$$\bigoplus_{n=0}^{\infty} A_n \xrightarrow{\Phi} \bigoplus_{n=0}^{\infty} A_n, \quad \bigoplus_{n=0}^{\infty} B_n \xrightarrow{\Psi} \bigoplus_{n=0}^{\infty} B_n, \quad \bigoplus_{n=0}^{\infty} C_n \xrightarrow{\Theta} \bigoplus_{n=0}^{\infty} C_n,$$

defined as in (43.1), and the commutativity of the diagram implies that the following diagram also commutes:

$$\begin{array}{ccccccc}
 & & \vdots & & \vdots & & \vdots \\
 & & \downarrow & & \downarrow & & \downarrow \\
 0 & \longrightarrow & 0 & \longrightarrow & 0 & \longrightarrow & 0 \longrightarrow 0 \\
 & & \downarrow & & \downarrow & & \downarrow \\
 0 & \longrightarrow & \bigoplus_n A_n & \xrightarrow{\bigoplus_n f_n} & \bigoplus_n B_n & \xrightarrow{\bigoplus_n g_n} & \bigoplus_n C_n \longrightarrow 0 \\
 & & \downarrow \Phi & & \downarrow \Psi & & \downarrow \Theta \\
 0 & \longrightarrow & \bigoplus_n A_n & \xrightarrow{\bigoplus_n f_n} & \bigoplus_n B_n & \xrightarrow{\bigoplus_n g_n} & \bigoplus_n C_n \longrightarrow 0 \\
 & & \downarrow & & \downarrow & & \downarrow \\
 0 & \longrightarrow & 0 & \longrightarrow & 0 & \longrightarrow & 0 \longrightarrow 0 \\
 & & \downarrow & & \downarrow & & \downarrow \\
 & & \vdots & & \vdots & & \vdots
 \end{array}$$

Here the rows are all short exact sequences, and extra rows of trivial modules have been added so that we can view each column in a trivial way as a chain complex, i.e. the diagram is a short exact sequence of chain complexes. Applying Proposition 32.13 then gives a long exact sequence that is not actually very long, since most of its terms are trivial: *a priori*, the only terms that can be nontrivial are

$$(43.3) \quad 0 \longrightarrow \ker \Phi \xrightarrow{f} \ker \Psi \xrightarrow{g} \ker \Theta \longrightarrow \operatorname{coker} \Phi \xrightarrow{f} \operatorname{coker} \Psi \xrightarrow{g} \operatorname{coker} \Theta \longrightarrow 0,$$

where the map $\ker \Theta \rightarrow \operatorname{coker} \Phi$ is a connecting homomorphism arising from Proposition 32.13, and all other maps are determined by f_n and g_n in obvious ways. Using the knowledge that Φ , Ψ and Θ are actually injective, and identifying their cokernels with direct limits, the conclusion is precisely the short exact sequence written in (43.2).

So what happens if we try to do the same trick with inverse limits? For an inverse system of the form

$$A_0 \xleftarrow{\phi_1} A_1 \xleftarrow{\phi_2} A_2 \xleftarrow{\phi_3} \dots,$$

the inverse limit of $\{A_n\}$ is identified via Proposition 43.11 with the kernel of the homomorphism

$$(43.4) \quad \Phi := \begin{pmatrix} \mathbb{1} & -\phi_1 & 0 & 0 & \cdots \\ 0 & \mathbb{1} & -\phi_2 & 0 & \cdots \\ 0 & 0 & \mathbb{1} & -\phi_3 & \cdots \\ 0 & 0 & 0 & \mathbb{1} & \cdots \\ 0 & 0 & 0 & 0 & \cdots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{pmatrix} : \prod_{n=0}^{\infty} A_n \longrightarrow \prod_{n=0}^{\infty} A_n.$$

You will find if you investigate it that the map Φ cannot be assumed surjective in general, and we shall therefore give its cokernel a name,

$$\varprojlim^1 \{A_n\} := \text{coker } \Phi.$$

Now suppose we have a diagram of the form

$$\begin{array}{ccccccccc} 0 & \longrightarrow & A_0 & \xrightarrow{f_0} & B_0 & \xrightarrow{g_0} & C_0 & \longrightarrow & 0 \\ & & \phi_1 \uparrow & & \psi_1 \uparrow & & \theta_1 \uparrow & & \\ 0 & \longrightarrow & A_1 & \xrightarrow{f_1} & B_1 & \xrightarrow{g_1} & C_1 & \longrightarrow & 0 \\ & & \phi_2 \uparrow & & \psi_2 \uparrow & & \theta_2 \uparrow & & \\ 0 & \longrightarrow & A_2 & \xrightarrow{f_2} & B_2 & \xrightarrow{g_2} & C_2 & \longrightarrow & 0, \\ & & \uparrow & & \uparrow & & \uparrow & & \\ & & \vdots & & \vdots & & \vdots & & \\ & & \uparrow & & \uparrow & & \uparrow & & \\ & & \varprojlim \{A_n\} & \xrightarrow{f} & \varprojlim \{B_n\} & \xrightarrow{g} & \varinjlim \{C_n\} & & \end{array}$$

in which the maps on the bottom row are again uniquely determined by the universal property, and every row except that one is assumed exact. Identifying the three inverse limits with the kernels of maps

$$\prod_{n=0}^{\infty} A_n \xrightarrow{\Phi} \prod_{n=0}^{\infty} A_n, \quad \prod_{n=0}^{\infty} B_n \xrightarrow{\Psi} \prod_{n=0}^{\infty} B_n, \quad \prod_{n=0}^{\infty} C_n \xrightarrow{\Theta} \prod_{n=0}^{\infty} C_n$$

defined as in (43.4), we can produce from this a short exact sequence of chain complexes

$$\begin{array}{ccccccc}
 & \vdots & & \vdots & & \vdots & \\
 & \downarrow & & \downarrow & & \downarrow & \\
 0 & \longrightarrow & 0 & \longrightarrow & 0 & \longrightarrow & 0 \longrightarrow 0 \\
 & & \downarrow & & \downarrow & & \downarrow \\
 0 & \longrightarrow & \prod_n A_n & \xrightarrow{\prod_n f_n} & \prod_n B_n & \xrightarrow{\prod_n g_n} & \prod_n C_n \longrightarrow 0 \\
 & & \downarrow \Phi & & \downarrow \Psi & & \downarrow \Theta \\
 0 & \longrightarrow & \prod_n A_n & \xrightarrow{\prod_n f_n} & \prod_n B_n & \xrightarrow{\prod_n g_n} & \prod_n C_n \longrightarrow 0 \\
 & & \downarrow & & \downarrow & & \downarrow \\
 0 & \longrightarrow & 0 & \longrightarrow & 0 & \longrightarrow & 0 \longrightarrow 0 \\
 & & \downarrow & & \downarrow & & \downarrow \\
 & & \vdots & & \vdots & & \vdots
 \end{array}$$

resulting in a long exact sequence of the same form as in (43.3), which can be rewritten as

$$0 \longrightarrow \varprojlim \{A_n\} \xrightarrow{f} \varprojlim \{B_n\} \xrightarrow{g} \varprojlim \{C_n\} \longrightarrow \varprojlim^1 \{A_n\} \xrightarrow{f} \varprojlim^1 \{B_n\} \xrightarrow{g} \varprojlim^1 \{C_n\} \longrightarrow 0.$$

The bad news we derive from this discussion is that taking the inverse limit of an inverse system of short exact sequences $0 \rightarrow \{A_n\} \rightarrow \{B_n\} \rightarrow \{C_n\} \rightarrow 0$ does not automatically produce a short exact sequence $0 \rightarrow \varprojlim \{A_n\} \rightarrow \varprojlim \{B_n\} \rightarrow \varprojlim \{C_n\} \rightarrow 0$. Exactness can fail in particular at the rightmost nontrivial term $\varprojlim \{C_n\}$, and its failure is measured by the term $\varprojlim^1 \{A_n\}$, which depends only on the first inverse system in the sequence. For this reason and others, it is useful to be able to recognize circumstances in which $\varprojlim^1 \{A_n\}$ vanishes; here is one that will be relevant for the discussion of CW-complexes:

LEMMA 43.14. *Suppose $A_0 \xleftarrow{\phi_1} A_1 \xleftarrow{\phi_2} A_2 \xleftarrow{\phi_3} \dots$ is an inverse system of R -modules such that for some $N \in \mathbb{N}$, the maps ϕ_n are isomorphisms for all $n \geq N$. Then $\varprojlim^1 \{A_n\} = 0$.*

PROOF. One needs to show that the map $\Phi : \prod_n A_n \rightarrow \prod_n A_n$ in (43.4) is surjective, i.e. that given any $\{b_n\} \in \prod_n A_n$, the equations $a_n - \phi_{n+1}(a_{n+1}) = b_n$ can be solved for $\{a_n\}_{n=0}^\infty$. In fact, there exists a unique solution satisfying $a_N = b_N$, which can be found by solving the equations for $n > N$ and $n < N$ recursively; the cases with $n > N$ require the assumption that ϕ_n is invertible. \square

We’ve touched in this section upon a few subjects that are large topics in homological algebra: the fact that direct limits of exact sequences are exact gets summarized by saying that \varinjlim is an *exact functor*, whereas the functor \varprojlim is only *left-exact*, and its failure to be exact is measured via the *derived functor* \varprojlim^1 . These notions will appear again in a slightly different context when we study the universal coefficient theorem.

43.3. Cohomology of the mapping telescope. With the subtle aspects of \varprojlim and \varprojlim^1 out of the way, we can now quickly finish the proof that $H_{CW}^*(X; G) \cong h^*(X)$ for infinite-dimensional CW-complexes and axiomatic cohomology theories with coefficient group G .

As in §39.3, we start by forming the mapping telescope T of the sequence

$$X^0 \hookrightarrow X^1 \hookrightarrow X^2 \hookrightarrow \dots,$$

which is also the mapping torus of a map $f : \coprod_{n=0}^{\infty} X^n \rightarrow \coprod_{n=0}^{\infty} X^n$. There is a cohomological version of the exact sequence of mapping tori discussed in Lecture 35, and in the present setting, it produces a long exact sequence of the form

$$h^{k-1}\left(\coprod_{n \geq 0} X^n\right) \xrightarrow{\Phi_{k-1}} h^{k-1}\left(\coprod_{n \geq 0} X^n\right) \longrightarrow h^k(T) \longrightarrow h^k\left(\coprod_{n \geq 0} X^n\right) \xrightarrow{\Phi_k} h^k\left(\coprod_{n \geq 0} X^n\right),$$

with $\Phi_k := 1 - f^*$. The usual algebraic trick turns this into a short exact sequence

$$0 \longrightarrow \operatorname{coker}(\Phi_{k-1}) \longrightarrow h^k(T) \longrightarrow \ker(\Phi_k) \longrightarrow 0,$$

and using the additivity axiom to identify $h^*\left(\coprod_n X^n\right)$ with $\prod_n h^*(X^n)$ puts Φ_k in precisely the form considered in the previous subsection, giving natural isomorphisms

$$\ker \Phi_k \cong \varprojlim \{h^k(X^n)\}_{n=0}^{\infty}, \quad \operatorname{coker} \Phi_k \cong \varprojlim^1 \{h^k(X^n)\}_{n=0}^{\infty}.$$

Our short exact sequence can thus be written as

$$0 \longrightarrow \varprojlim^1 \{h^{k-1}(X^n)\}_{n=0}^{\infty} \longrightarrow h^k(T) \longrightarrow \varprojlim \{h^k(X^n)\}_{n=0}^{\infty} \longrightarrow 0.$$

If we had no further information about the sequence $X^0 \hookrightarrow X^1 \hookrightarrow X^2 \hookrightarrow \dots$, then we would not be able to say now that we have computed the cohomology of its mapping telescope; a short exact sequence with $h^k(T)$ at the center is typically a weaker result than a full computation of $h^k(T)$. But we do have more information, because just as in the homological case, the vanishing of $h^k(X^n, X^{n-1})$ for $k \neq n$ implies that the maps $h^k(X^{n+1}) \rightarrow h^k(X^n)$ induced by inclusions $X^n \hookrightarrow X^{n+1}$ are isomorphisms for all $n > k$. Lemma 43.14 therefore implies

$$\varprojlim^1 \{h^{k-1}(X^n)\} = 0,$$

and we therefore have a natural isomorphism

$$h^k(T) \cong \varprojlim \{h^k(X^n)\}.$$

The rest of the story is the same as in §39.3: the mapping telescope T is a deformation retract of $X \times [0, \infty)$, giving rise to natural isomorphisms

$$h^k(X) \cong h^k(X \times [0, \infty)) \cong h^k(T) \cong \varprojlim \{h^k(X^n)\}_{n=0}^{\infty} \cong \varprojlim \{H_{\text{CW}}^k(X^n)\}_{n=0}^{\infty} \cong H_{\text{CW}}^k(X).$$

43.4. Čech homology and cohomology. Now that we have inverse limits, we can describe another axiomatic homology theory besides the singular theory. The caveat is that it doesn't quite work: the failure of \varprojlim to preserve exactness of sequences prevents the theory sketched below from satisfying the exactness axiom without some restrictions. On the other hand, its cohomological variant depends on direct limits instead of inverse limits, and consequently does not have this drawback.

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 43.15. 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

$$\bigcap_{\mathcal{U} \in \sigma} \mathcal{U} \neq \emptyset.$$

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

$$A \cap \bigcap_{\mathcal{U} \in \sigma} \mathcal{U} \neq \emptyset.$$

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 $\bigcap_{\mathcal{U} \in \sigma} \mathcal{U} \neq \emptyset$ 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 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 R -module G to be used for coefficients in simplicial homology, and it 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 20 shows some examples of open coverings \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 covering 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 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 covering $\mathfrak{U} \in \mathcal{O}(X)$ 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

$$F : \mathfrak{U}' \rightarrow \mathfrak{U}, \quad F(\mathcal{U}'_A) \subset \mathfrak{U}_A$$

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:

$$F : \mathcal{N}(\mathfrak{U}', \mathfrak{U}'_A) \rightarrow \mathcal{N}(\mathfrak{U}, \mathfrak{U}_A).$$




FIGURE 20. 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.

It therefore induces a chain map between the corresponding ordered simplicial chain complexes

$$(43.5) \quad F_* : C_*^o(\mathcal{N}(\mathfrak{U}', \mathfrak{U}'_A)) \rightarrow C_*^o(\mathcal{N}(\mathfrak{U}, \mathfrak{U}_A)).$$

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 43.16 ([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 (43.5) on ordered simplicial chain complexes is independent of choices.* \square

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

$$\varphi_{\alpha\beta} : H_*^o(\mathcal{N}(\beta)) \rightarrow H_*^o(\mathcal{N}(\alpha)).$$

One can view this as defining an inverse system: indeed, let us define a pre-order on $\mathcal{O}(X, A)$ by writing

$$\beta > \alpha$$

whenever β is a refinement of α . (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 $(\mathcal{O}(X, A), <)$ a directed set, and the result above associates homomorphisms $\varphi_{\alpha\beta}$ of homology groups in each degree to any pair $\alpha, \beta \in \mathcal{O}(X, A)$ with $\beta > \alpha$.

DEFINITION 43.17. The **Čech homology** of a pair of spaces (X, A) with coefficients in an R -module G is defined as the collection of R -modules

$$\check{H}_n(X, A) = \check{H}_n(X, A; G) := \varprojlim_{\{\mathfrak{U}, \mathfrak{U}_A\} \in \mathcal{O}(X, A)} \{H_n^o(\mathcal{N}(\mathfrak{U}, \mathfrak{U}_A); G)\}, \quad n \in \mathbb{Z}.$$

It is slightly harder than for singular homology to see why this should define a collection of functors $\text{Top}^{\text{rel}} \rightarrow R\text{-Mod}$, but still not so hard. The main point is that whenever $f : (X, A) \rightarrow (Y, B)$ 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 $f^{-1}(U)$ for $U \in \mathfrak{U}$, 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 between the open sets in $f^*\alpha$ and those in α then defines a simplicial map $\mathcal{N}(f^*\alpha) \rightarrow \mathcal{N}(\alpha)$, giving a homomorphism

$$f_* : H_*^o(\mathcal{N}(f^*\alpha)) \rightarrow H_*^o(\mathcal{N}(\alpha))$$

for every $\alpha \in \mathcal{O}(Y, B)$. Using the universal property of the inverse limit, one can derive from this a morphism

$$f_* : \check{H}_*(X, A) \rightarrow \check{H}_*(Y, B)$$

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 isomorphic to 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 31.5 the notion of the *open star* of a vertex v in K , which defines an open set

$$\text{st } v \subset X$$

containing all points that lie in simplices that have v as a vertex (see Figure 21). These sets define a distinguished open covering of X ,

$$\mathfrak{U}_K := \{\text{st } v \mid v \in V\},$$

and recall from Exercise 31.3 that for any finite collection of vertices $v_0, \dots, v_n \in V$, we have

$$\bigcap_{k=0}^n \text{st } v_k \neq \emptyset \iff \{v_0, \dots, v_n\} \in S.$$

In other words, the nerve of \mathfrak{U}_K is naturally isomorphic to 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, the induced map

$$H_*(\mathcal{N}(\mathfrak{U}_{K'})) \rightarrow H_*(\mathcal{N}(\mathfrak{U}_K))$$

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 Proposition 43.13, and thus provides enough information to compute the inverse limit. The result is:

THEOREM 43.18. *For any compact polyhedron $X = |K|$ with underlying simplicial complex K , $\check{H}_*(X; G) \cong H_*^o(K; G)$ for every choice of coefficients G . □*

Since $H_*^o(K; G) \cong H_*^\Delta(K; G) \cong H_*^{\text{CW}}(X; G) \cong H_*(X; G)$, it follows in particular that Čech homology is isomorphic to singular homology on compact polyhedra. Notice, by the way, that this implies yet another new proof that simplicial homology is independent of the triangulation of




FIGURE 21. The open stars of two neighboring vertices v_0 and v_1 in a simplicial complex.

a compact polyhedron, with no need to pass through any axiomatic homology theory—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 43.19. *If X is a connected space, then for every open covering \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. \square

THEOREM 43.20. *For any connected space X and any choice of coefficients G , $\check{H}_0(X; G) \cong G$.*

PROOF. Lemma 43.19 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 $H_0^o(\mathcal{N}(\mathfrak{U}'); G) \rightarrow H_0^o(\mathcal{N}(\mathfrak{U}); G)$ for any refinement \mathfrak{U}' of \mathfrak{U} is an isomorphism, and that the inverse limit is therefore isomorphic to G . \square

This result is different in general from singular homology in degree 0, which splits over a direct sum of the *path-components* (not connected components) of each space. So, for instance, Figure 19 in Lecture 40 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 43.18. 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 43.21. *For every choice of coefficient module G , the functors $\check{H}_n(\cdot; G) : \text{Top}^{\text{rel}} \rightarrow R\text{-Mod}$ satisfy all of the Eilenberg-Steenrod axioms except for exactness, but they do 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 discussion of Čech theory is only meant to be 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 43.21 is not the end of the story.

The problem with the exactness axiom results from the failure of inverse limits to preserve exactness of sequences. 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 short exact sequence of ordered simplicial chain complexes

$$0 \rightarrow C_*^o(\mathcal{N}(\mathfrak{U}_A)) \rightarrow C_*^o(\mathcal{N}(\mathfrak{U})) \rightarrow C_*^o(\mathcal{N}(\mathfrak{U}, \mathfrak{U}_A)) = C_*^o(\mathcal{N}(\mathfrak{U}), \mathcal{N}(\mathfrak{U}_A)) \rightarrow 0$$

gives rise to a long exact sequence of simplicial homology groups

$$(43.6) \quad \dots \rightarrow H_n^o(\mathcal{N}(\mathfrak{U}_A)) \rightarrow H_n^o(\mathcal{N}(\mathfrak{U})) \rightarrow H_n^o(\mathcal{N}(\mathfrak{U}, \mathfrak{U}_A)) \rightarrow H_{n-1}^o(\mathcal{N}(\mathfrak{U}_A)) \rightarrow \dots$$

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

$$\begin{array}{cccccccc} \dots & \rightarrow & H_n^o(\mathcal{N}(\mathfrak{U}_A)) & \rightarrow & H_n^o(\mathcal{N}(\mathfrak{U})) & \rightarrow & H_n^o(\mathcal{N}(\mathfrak{U}, \mathfrak{U}_A)) & \rightarrow & H_{n-1}^o(\mathcal{N}(\mathfrak{U}_A)) & \rightarrow & \dots \\ & & \uparrow & & \uparrow & & \uparrow & & \uparrow & & \\ \dots & \rightarrow & H_n^o(\mathcal{N}(\mathfrak{U}'_A)) & \rightarrow & H_n^o(\mathcal{N}(\mathfrak{U}')) & \rightarrow & H_n^o(\mathcal{N}(\mathfrak{U}', \mathfrak{U}'_A)) & \rightarrow & H_{n-1}^o(\mathcal{N}(\mathfrak{U}'_A)) & \rightarrow & \dots \end{array}$$

The universal property uniquely determines from this data a sequence

$$(43.7) \quad \dots \rightarrow \check{H}_n(A) \rightarrow \check{H}_n(X) \rightarrow \check{H}_n(X, A) \rightarrow \check{H}_{n-1}(A) \rightarrow \dots,$$

but it need not be exact in general. The following concrete example shows that, at least algebraically, this danger is real:

EXAMPLE 43.22. For every $n \in \mathbb{N}$, denote by $0 \rightarrow A_n \rightarrow B_n \rightarrow C_n \rightarrow 0$ the short exact sequence $0 \rightarrow \mathbb{Z} \xrightarrow{\cdot 2} \mathbb{Z} \xrightarrow{\text{pr}} \mathbb{Z}_2 \rightarrow 0$, and define homomorphisms $\varphi_{n-1, n}$ for each $n \geq 2$ by

$$A_n \xrightarrow{\cdot 3} A_{n-1}, \quad B_n \xrightarrow{\cdot 3} B_{n-1}, \quad C_n \xrightarrow{1} C_{n-1}.$$

Then the resulting diagram

$$\begin{array}{cccccccc} \dots & \longrightarrow & 0 & \longrightarrow & A_1 & \xrightarrow{\cdot 2} & B_1 & \xrightarrow{\text{pr}} & C_1 & \longrightarrow & 0 & \longrightarrow & \dots \\ & & \uparrow & & \cdot 3 \uparrow & & \cdot 3 \uparrow & & 1 \uparrow & & \uparrow & & \\ \dots & \longrightarrow & 0 & \longrightarrow & A_2 & \xrightarrow{\cdot 2} & B_2 & \xrightarrow{\text{pr}} & C_2 & \longrightarrow & 0 & \longrightarrow & \dots \\ & & \uparrow & & \cdot 3 \uparrow & & \cdot 3 \uparrow & & 1 \uparrow & & \uparrow & & \\ \dots & \longrightarrow & 0 & \longrightarrow & A_3 & \xrightarrow{\cdot 2} & B_3 & \xrightarrow{\text{pr}} & C_3 & \longrightarrow & 0 & \longrightarrow & \dots \\ & & \uparrow & & \uparrow & & \uparrow & & \uparrow & & \uparrow & & \\ & & \vdots & & \vdots & & \vdots & & \vdots & & \vdots & & \end{array}$$

commutes. By Proposition 43.11, the inverse limits of the individual columns are as follows: first,

$$\varprojlim \{A_n\} = \left\{ (a_1, a_2, a_3, \dots) \in \prod_{n \in \mathbb{N}} \mathbb{Z} \mid a_{n-1} = 3a_n \text{ for all } n \geq 2 \right\} = 0,$$

and $\varprojlim \{B_n\}$ similarly vanishes since no integer is divisible by arbitrarily large powers of 3. On the other hand,

$$\varprojlim \{C_n\} = \left\{ (c_1, c_2, c_3, \dots) \in \prod_{n \in \mathbb{N}} \mathbb{Z}_2 \mid c_{n-1} = c_n \text{ for all } n \geq 2 \right\} \cong \mathbb{Z}_2,$$

so the resulting sequence of inverse limits is of the form

$$\dots \longrightarrow 0 \longrightarrow 0 \longrightarrow 0 \longrightarrow \mathbb{Z}_2 \longrightarrow 0 \longrightarrow \dots,$$

which is not an exact sequence.

In general, suppose we are given an inverse system of chain complexes $\{C_*^\alpha, \varphi_{\alpha\beta}\}$ indexed by α in some directed set $(J, <)$, so we have a commuting diagram for every $\beta > \alpha$ in the form

$$(43.8) \quad \begin{array}{ccccccc} \dots & \longrightarrow & C_{n+1}^\alpha & \xrightarrow{\partial^\alpha} & C_n^\alpha & \xrightarrow{\partial^\alpha} & C_{n-1}^\alpha & \xrightarrow{\partial^\alpha} & \dots \\ & & \varphi_{\alpha\beta} \uparrow & & \varphi_{\alpha\beta} \uparrow & & \varphi_{\alpha\beta} \uparrow & & \\ \dots & \longrightarrow & C_{n+1}^\beta & \xrightarrow{\partial^\beta} & C_n^\beta & \xrightarrow{\partial^\beta} & C_{n-1}^\beta & \xrightarrow{\partial^\beta} & \dots \end{array}$$

and assume moreover that the rows of these diagrams are always exact. By Proposition 43.11 and Remark 43.12, the inverse limit is a chain complex

$$C_*^\infty := \varprojlim \{C_*^\alpha\} = \left\{ \{x_\alpha\} \in \prod_{\alpha \in J} C_*^\alpha \mid \varphi_{\alpha\beta}(x_\beta) = x_\alpha \text{ for all } \beta > \alpha \right\},$$

where the chain complex boundary map can be written as

$$\partial^\infty := \prod_{\alpha \in J} \partial^\alpha \Big|_{C_*^\infty} : C_*^\infty \rightarrow C_*^\infty,$$

the restriction to the submodule $C_*^\infty \subset \prod_\alpha C_*^\alpha$ being well defined since $\varphi_{\alpha\beta}(\partial^\beta x_\beta) = \partial^\alpha \varphi_{\alpha\beta}(x_\beta) = \partial^\alpha x_\alpha$ for all $\beta > \alpha$ and $x_\beta \in C_*^\beta$. Given $x = \{x_\alpha\}_{\alpha \in J} \in C_n^\infty$ with $\partial^\infty \{x_\alpha\} = 0$, we have $\partial^\alpha x_\alpha = 0$ for all $\alpha \in J$, and the exactness of the rows in (43.8) then implies $x_\alpha = \partial^\alpha y_\alpha$ for some $y_\alpha \in C_{n+1}^\alpha$. We can now observe a concrete reason why the exactness of $\dots \rightarrow C_{n+1}^\alpha \rightarrow C_n^\alpha \rightarrow C_{n-1}^\alpha \rightarrow \dots$ for every α might fail to imply the exactness of $\dots \rightarrow C_{n+1}^\infty \rightarrow C_n^\infty \rightarrow C_{n-1}^\infty \rightarrow \dots$: the trouble is that the elements y_α are not generally unique, and if they are chosen arbitrarily, then they need not satisfy

$$(43.9) \quad \varphi_{\alpha\beta}(y_\beta) = y_\alpha \quad \text{for all } \beta > \alpha,$$

without which $\{y_\alpha\}_{\alpha \in J}$ will not be an element of C_{n+1}^∞ .

To get a firmer handle on this problem, define for each $\alpha \in J$ 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

$$\varphi_{\alpha\beta}(K_\beta) \subset K_\alpha \quad \text{for all } \beta > \alpha,$$

which makes the collection of sets $\{K_\alpha\}_{\alpha \in J}$ with maps $K_\beta \xrightarrow{\varphi_{\alpha\beta}} K_\alpha$ into an inverse system in Set over $(J, <)$. By Proposition 43.5 and Remark 43.6, $\varprojlim \{K_\alpha\}$ is then the set of all elements $\{y_\alpha\} \in \prod_{\alpha \in J} K_\alpha$ such that (43.9) 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:

Is $\varprojlim \{K_\alpha\}$ nonempty?

Example 43.22 implies that the answer must sometimes be no, and indeed, we know from Example 43.10 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 $\alpha \in J$ 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 Hausdorff spaces. In this

case there is a positive result we can use: Proposition 43.9 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 Proposition 43.9 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 modules 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 (43.6) 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 43.23. *If G is either a finite abelian group or a finite-dimensional vector space over a field, then the restriction of Čech homology $\check{H}_*(\cdot; G)$ to the category Cpct^{rel} of compact pairs defines an axiomatic homology theory on Cpct^{rel} . \square*

Having defined Čech homology, you will have little trouble guessing how to define Čech cohomology: we must replace the simplicial homology of the nerve of an open covering $(\mathfrak{U}, \mathfrak{U}_A)$ of (X, A) with its (ordered) **simplicial cohomology**,

$$H_o^*(\mathcal{N}(\mathfrak{U}, \mathfrak{U}_A); G) := H^*(C_*^o(\mathfrak{U}, \mathfrak{U}_A; R); G).$$

For any refinement $(\mathfrak{U}', \mathfrak{U}'_A)$ of $(\mathfrak{U}, \mathfrak{U}_A)$, the resulting simplicial map $F : \mathcal{N}(\mathfrak{U}', \mathfrak{U}'_A) \rightarrow \mathcal{N}(\mathfrak{U}, \mathfrak{U}_A)$ induces a chain map $F_* : C_*^o(\mathcal{N}(\mathfrak{U}', \mathfrak{U}'_A); R) \rightarrow C_*^o(\mathcal{N}(\mathfrak{U}, \mathfrak{U}_A); R)$ uniquely up to chain homotopy, which therefore dualizes to a map of cochain complexes that is (by Proposition 41.2) likewise unique up to chain homotopy, producing a canonically defined morphism

$$F^* : H_o^*(\mathcal{N}(\mathfrak{U}, \mathfrak{U}_A); G) \rightarrow H_o^*(\mathcal{N}(\mathfrak{U}', \mathfrak{U}'_A); G).$$

Notice what has happened as a result of dualization: the collection of simplicial homology groups $\{H_*^o(\mathcal{N}(\mathfrak{U}, \mathfrak{U}_A); G)\}_{(\mathfrak{U}, \mathfrak{U}_A) \in \mathcal{O}(X, A)}$ 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: \varinjlim is an exact functor, and one can use this to prove that unlike \check{H}_* , the cohomology \check{H}^* satisfies the exactness axiom without any restrictions.

THEOREM 43.24 (see [ES52, Spa95]). *For any choice of R -module G , the Čech cohomology $\check{H}^* = \check{H}^*(\cdot; G)$ is an axiomatic cohomology theory with coefficient group G . \square*

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 43.19 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 42.2 then implies $\check{H}^1(\Sigma X; G) = 0$. But if X has more than one path-component, then $\check{H}^0(X; G)$ and $H^1(\Sigma X; G)$ are both nontrivial; the latter is isomorphic to $\text{Hom}(\pi_1(\Sigma X), G)$ since the suspension ΣX is always path-connected, thus ΣX is an example of a space for which $\check{H}^1(\Sigma X; G) \not\cong \text{Hom}(\pi_1(\Sigma X), G) \cong H^1(\Sigma X; G)$.

The Čech theory also has one nice property that the singular theory does not: it is *continuous* with respect to inverse limits of spaces. The condition 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 $(J, <)$. The associated morphisms $\varphi_\alpha : \varprojlim\{(X_\beta, A_\beta)\} \rightarrow (X_\alpha, A_\alpha)$ then induce homomorphisms

$$(\varphi_\alpha)_* : h_n(\varprojlim\{(X_\beta, A_\beta)\}) \rightarrow h_n(X_\alpha, A_\alpha),$$

for each $n \in \mathbb{Z}$, which make $\{h_n(\varprojlim\{(X_\beta, A_\beta)\}), (\varphi_\alpha)_*\}$ a target of the inverse system

$$\{h_n(X_\alpha, A_\alpha), (\varphi_{\alpha\beta})_*\}$$

in the category of R -modules. By the universal property of inverse limits, there is then a canonical limit morphism

$$h_n(\varprojlim\{(X_\alpha, A_\alpha)\}) \rightarrow \varprojlim\{h_n(X_\alpha, A_\alpha)\}.$$

For a cohomology theory h^* , an inverse system of pairs of spaces instead gives rise to a *direct* system of cohomologies $\{h^n(X_\alpha, A_\alpha; G), \varphi_{\alpha\beta}^*\}$, for which the maps $h^n(X_\alpha, A_\alpha) \rightarrow h^n(\varprojlim\{X_\beta, A_\beta\})$ induced by $\varprojlim\{X_\beta, A_\beta\} \rightarrow (X_\alpha, A_\alpha)$ make $h^n(\varprojlim\{X_\beta, A_\beta\})$ into a target and thus determine a canonical limit morphism

$$\varinjlim\{h^n(X_\alpha, A_\alpha)\} \rightarrow h^n(\varprojlim\{(X_\alpha, A_\alpha)\}).$$

The following result is often quoted as a selling point of the Čech theory in comparison with singular homology and cohomology. 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 43.25 (continuity in Čech theory; see [ES52, Chapter X]). *For any inverse system of compact pairs $\{(X_\alpha, A_\alpha), \varphi_{\alpha\beta}\}$ and any choice of coefficients G , the canonical maps*

$$\begin{aligned} \check{H}_n(\varprojlim\{(X_\alpha, A_\alpha)\}; G) &\longrightarrow \varprojlim\{\check{H}_n(X_\alpha, A_\alpha; G)\}, \\ \varinjlim\{\check{H}^n(X_\alpha, A_\alpha)\} &\longrightarrow \check{H}^n(\varprojlim\{(X_\alpha, A_\alpha)\}). \end{aligned}$$

are isomorphisms for all $n \in \mathbb{Z}$. □

For the failure of singular homology or cohomology to satisfy the continuity condition described above, see Exercise 43.5.

REMARK 43.26. In the previous lecture, we gave a brief sketch of yet another axiomatic cohomology theory \bar{H}^* , the Alexander-Spanier cohomology. One can show that \bar{H}^* also has the continuity property described in Theorem 43.25. Since all compact Hausdorff spaces are inverse limits of 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].)

43.5. Exercises.

EXERCISE 43.1 (*). Prove Propositions 43.5 and 43.11, giving explicit descriptions of inverse limits in the categories **Top** and **R-Mod**.

EXERCISE 43.2. Prove Proposition 43.7, on the intersection of a family of subspaces as an inverse limit.

EXERCISE 43.3. Prove Proposition 43.9 on nonempty inverse limits of compact Hausdorff spaces.

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 $x^{\beta} = \{x_{\alpha}^{\beta}\} \in \prod_{\alpha} X_{\alpha}$ whose coordinates satisfy $x_{\alpha}^{\beta} = \varphi_{\alpha\beta}(x_{\beta}^{\beta})$ for every $\alpha < \beta$ and are arbitrary for all other α . The collection $\{x^{\beta} \in \prod_{\alpha} X_{\alpha}\}_{\beta \in J}$ 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.)

EXERCISE 43.4 (*). Prove Proposition 43.13 on limits of inverse systems that are “eventually constant,” and describe the associated morphisms $\varprojlim \{X_{\beta}\} \xrightarrow{\varphi_{\alpha}} X_{\alpha}$ for every $\alpha \in J$.

*Advice: This problem becomes a bit easier if you work in any of the categories **Top**, **R-Mod** or **Ch** so that you can use the results of Propositions 43.5 or 43.11 (plus Remark 43.12). But it can also be done without that assumption, just by using the universal property and playing with commutative diagrams.*

EXERCISE 43.5. 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,

$$H_*(\varprojlim \{X_{\alpha}\}) \not\cong \varprojlim \{H_*(X_{\alpha})\}.$$

Hint: Use Proposition 43.7.

EXERCISE 43.6. 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.

44. Universal coefficient theorems and exact functors

Our goal in this and the next lecture is to clarify precisely how $H_*(X; G_1)$ and $H_*(X; G_2)$ or the corresponding cohomologies are related to each other for different choices of coefficient modules G_1 and G_2 . In fact, we are stating this question for singular (co)homology, but the answer will apply equally well to simplicial or cellular (co)homology. It is, in reality, an algebraic question, and the methods we need in order to answer it come from homological algebra.

There are two versions of the universal coefficient theorem: one for homology, and one for cohomology. Both posit the existence of a natural split exact sequence that reduces in favorable cases to the statement that a certain natural map is an isomorphism. In the case of singular cohomology with coefficients in an abelian group G , the natural map in question is

$$H^n(X; G) \xrightarrow{h} \text{Hom}(H_n(X; \mathbb{Z}), G) : [\alpha] \mapsto \langle [\alpha], \cdot \rangle,$$

defined in terms of the evaluation pairing $\langle [\alpha], [c] \rangle := \alpha(c)$, and the theorem gives a split exact sequence of the form

$$0 \longrightarrow \text{Ext}(H_{n-1}(X; \mathbb{Z}), G) \longrightarrow H^n(X; G) \xrightarrow{h} \text{Hom}(H_n(X; \mathbb{Z}), G) \longrightarrow 0,$$

where $\text{Ext} : \mathbf{Ab} \times \mathbf{Ab} \rightarrow \mathbf{Ab}$ is a functor that is guaranteed to vanish under various conditions. When it vanishes, the conclusion is that h is an isomorphism, and more generally, the theorem implies that h is always surjective, and in light of the splitting,

$$H^n(X; G) \cong \text{Hom}(H_n(X; \mathbb{Z}), G) \oplus \text{Ext}(H_{n-1}(X; \mathbb{Z}), G).$$

The version described above assumes only that G is an abelian group, but this is where it can become extremely useful to distinguish between abelian groups and modules over a commutative ring R other than \mathbb{Z} . If G is viewed as an R -module, then the natural map arising from the evaluation pairing takes the form

$$H^n(X; G) \xrightarrow{h} \text{Hom}(H_n(X; R), G),$$

and we will see that whenever R is a principal ideal domain (e.g. \mathbb{Z} or a field), the universal coefficient theorem gives a split exact sequence

$$0 \longrightarrow \text{Ext}(H_{n-1}(X; R), G) \longrightarrow H^n(X; G) \xrightarrow{h} \text{Hom}(H_n(X; R), G) \longrightarrow 0,$$

where Ext is in this case a functor $R\text{-Mod} \times R\text{-Mod} \rightarrow R\text{-Mod}$. The usefulness of this added generality is that one obtains a wider range of criteria for the vanishing of Ext , e.g. it turns out that it *always* vanishes if R is a field \mathbb{K} , and the universal coefficient theorem then reduces to the straightforward statement that for any vector space G over \mathbb{K} , the natural map

$$H^n(X; G) \xrightarrow{h} \text{Hom}(H_n(X; \mathbb{K}), G)$$

is a vector space isomorphism. That is one of the results that will be proved in the present lecture.

For singular homology with coefficients in an R -module G , the natural map to consider takes the form

$$H_n(X; R) \otimes G \xrightarrow{h} H_n(X; G) : \left[\sum_i a_i \sigma_i \right] \otimes g \mapsto \left[\sum_i a_i g \sigma_i \right],$$

and the homological version of the universal coefficient theorem gives a natural split exact sequence

$$0 \longrightarrow H_n(X; R) \otimes G \xrightarrow{h} H_n(X; G) \longrightarrow \text{Tor}(H_{n-1}(X; R), G) \longrightarrow 0,$$

where $\text{Tor} : R\text{-Mod} \times R\text{-Mod} \rightarrow R\text{-Mod}$ is another functor that is guaranteed to vanish under various conditions, e.g. whenever either of the R -modules $H_{n-1}(X; R)$ or G is free, implying in such cases that $h : H_n(X; R) \otimes G \rightarrow H_n(X; G)$ is an isomorphism.

Analogous results hold for simplicial or cellular (co)homology, as consequences of the algebraic versions of the universal coefficient theorem. In the present lecture, we will prove the special cases of these results in which the extra Tor or Ext terms vanish, implying that the natural maps labelled above as h are isomorphisms. The definitions of the functors Tor and Ext will be given in the next lecture, along with the general versions of the universal coefficient theorems.

REMARK 44.1. Readers who have not previously encountered tensor products in the category of abelian groups or R -modules may want to look at Exercise 44.1 before continuing.

44.1. Principal ideal domains. Up until this point in our development of homology and cohomology, there has always been a fixed commutative ring R in the background: most algebraic objects we've considered have been modules over R , and the choice of the ring R has typically been completely arbitrary. In the universal coefficient theorems, the choice of R is no longer arbitrary, because we will need to assume that R is a **principal ideal domain**. 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 44.2 (see [Lan02, §III.7]). *If R is a principal ideal domain, then every submodule of a free R -module is also free.* \square

In particular, this is true when R is \mathbb{Z} , hence subgroups of free abelian groups are also free, and it also holds whenever R is a field \mathbb{K} , since *all* vector spaces over \mathbb{K} are free \mathbb{K} -modules. Those are the main cases typically of interest.

44.2. The functors $\otimes G$ and $\text{Hom}(\cdot, G)$. For any space X and any R -module G , there is an obvious relationship between $H_*(X; R)$ and $H^*(X; G)$ at the level of their underlying chain/cochain complexes: the cochain complex $C^*(X; G)$ is obtained by feeding the chain complex $C_*(X; R)$ into the contravariant functor

$$\text{Hom}(\cdot, G) : \text{Ch}(R\text{-Mod}) \rightarrow \text{CoCh}(R\text{-Mod}).$$

In fact, $\text{Hom}(\cdot, G)$ is, in the first place, a contravariant functor

$$\text{Hom}(\cdot, G) : R\text{-Mod} \rightarrow R\text{-Mod},$$

and it has certain properties that we will formalize in the next subsection such that any chain complex

$$(44.1) \quad \dots \longrightarrow C_{n+1} \xrightarrow{\partial_{n+1}} C_n \xrightarrow{\partial_n} C_{n-1} \longrightarrow \dots$$

can be fed term-by-term into $\text{Hom}(\cdot, G)$ in order to produce a cochain complex

$$(44.2) \quad \dots \longleftarrow \text{Hom}(C_{n+1}, G) \xleftarrow{\partial_{n+1}^*} \text{Hom}(C_n, G) \xleftarrow{\partial_n^*} \text{Hom}(C_{n-1}, G) \longleftarrow \dots,$$

with each of the maps ∂_n^* defined by feeding ∂_n into the functor. The way that we actually defined the cochain complex $\text{Hom}(C_*, G)$ in Lecture 41 was slightly different: the coboundary operators were given some extra signs depending on the degree, but the cohomology is unchanged by this detail, so if our goal is to relate $H_*(C_*)$ and $H^*(C_*; G)$, we may as well regard $H^*(C_*; G)$ as the cohomology of the complex in (44.2).

CONVENTION. For this and the next lecture, we will use the notation $\text{Hom}(C_*, G)$ to denote the cochain complex in (44.2) rather than the variant with extra signs defined in Lecture 41. This is a harmless abuse of notation since the cohomology does not depend on the extra signs.

For homology, it is also possible to obtain the chain complex underlying $H_*(X; G)$ by feeding the chain complex underlying $H_*(X; R)$ into a functor, this time one that is covariant. For this purpose, we use the tensor product of R -modules and consider the functor

$$\otimes G : R\text{-Mod} \rightarrow R\text{-Mod},$$

which sends a module $A \in R\text{-Mod}$ to the module $A \otimes G \in R\text{-Mod}$, and sends a homomorphism $\phi : A \rightarrow B$ to the homomorphism

$$\phi \otimes 1 : A \otimes G \rightarrow B \otimes G : a \otimes g \mapsto \phi(a) \otimes g.$$

Feeding the chain complex C_* in (44.1) term-by-term into $\otimes G$ produces a new chain complex

$$\dots \longrightarrow C_{n+1} \otimes G \xrightarrow{\hat{c}_{n+1} \otimes \mathbb{1}} C_n \otimes G \xrightarrow{\hat{c}_n \otimes \mathbb{1}} C_{n-1} \otimes G \longrightarrow \dots,$$

which we will denote by

$$C_* \otimes G \in \text{Ch}(R\text{-Mod}).$$

We observe that for the singular chain complex $C_*(X; R)$ of a space X and any R -module G , there is a natural isomorphism of chain complexes

$$C_*(X; R) \otimes G \xrightarrow{\cong} C_*(X; G) : \left(\sum_i r_i \sigma_i \right) \otimes g \mapsto \sum_i (r_i g) \sigma_i.$$

For the relative singular homology of a pair of spaces (X, A) , there is a similar isomorphism $C_*(X, A; R) \otimes G \cong C_*(X, A; G)$, and there are also analogous isomorphisms for cellular or simplicial chain complexes.

In summary: chain or cochain complexes with coefficients in an arbitrary R -module G can be obtained by starting from $C_*(X; R)$ (or its cellular or simplicial counterparts) and applying the functors $\otimes G$ or $\text{Hom}(\cdot, G)$. The following detail will turn out to be important: the chain groups $C_n(X; R)$ are *free* modules over R , thus all versions of singular/cellular/simplicial homology or cohomology can be obtained algebraically from one that is built out of free R -modules.

44.3. Additive and exact functors. Both $\otimes G$ and $\text{Hom}(\cdot, G)$ are examples of *additive* functors $R\text{-Mod} \rightarrow R\text{-Mod}$, a notion that we shall now make precise.

REMARK 44.3. Several of the definitions in this section can be extended to the more general context of functors $\mathcal{F} : \mathcal{A} \rightarrow \mathcal{B}$ 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 unnecessarily extreme level of abstraction. We will therefore keep things more concrete by working only in the specific abelian category $R\text{-Mod}$, and considering only functors from $R\text{-Mod}$ to itself.

DEFINITION 44.4. A functor $\mathcal{F} : R\text{-Mod} \rightarrow R\text{-Mod}$ is called **additive** if for every pair of R -modules A and B , the map that \mathcal{F} defines from $\text{Hom}(A, B)$ to either $\text{Hom}(\mathcal{F}(A), \mathcal{F}(B))$ (in the covariant case) or $\text{Hom}(\mathcal{F}(B), \mathcal{F}(A))$ (in the contravariant case) 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 $g \circ f : A \rightarrow C$ vanishes, then in the covariant case, the fact that \mathcal{F} defines a group homomorphism $\text{Hom}(A, C) \rightarrow \text{Hom}(\mathcal{F}(A), \mathcal{F}(C))$ implies

$$\mathcal{F}(g) \circ \mathcal{F}(f) = \mathcal{F}(g \circ f) = 0 \in \text{Hom}(\mathcal{F}(A), \mathcal{F}(C)),$$

simply because group homomorphisms preserve the 0 element. It follows that covariant additive functors preserve the property of being a chain complex, i.e. any chain complex

$$\dots \longrightarrow A_{n+1} \xrightarrow{f_{n+1}} A_n \xrightarrow{f_n} A_{n-1} \longrightarrow \dots$$

gives rise to a chain complex

$$\dots \longrightarrow \mathcal{F}(A_{n+1}) \xrightarrow{\mathcal{F}(f_{n+1})} \mathcal{F}(A_n) \xrightarrow{\mathcal{F}(f_n)} \mathcal{F}(A_{n-1}) \longrightarrow \dots$$

If \mathcal{F} is instead a contravariant additive functor, then one obtains a cochain complex

$$\dots \longleftarrow \mathcal{F}(A_{n+1}) \xleftarrow{\mathcal{F}(f_{n+1})} \mathcal{F}(A_n) \xleftarrow{\mathcal{F}(f_n)} \mathcal{F}(A_{n-1}) \longleftarrow \dots$$

For a chain complex A_* , we will denote by $\mathcal{F}(A_*)$ the chain complex or cochain complex obtained by feeding A_* into an additive functor \mathcal{F} .

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

$$(44.3) \quad \begin{array}{ccccc} A & & & & A \\ & \searrow^{i_A} & & & \nearrow^{\pi_A} \\ & & A \oplus B & & \\ & \nearrow_{i_B} & & & \searrow_{\pi_B} \\ B & & & & B \end{array}$$

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

$$(44.4) \quad \pi_A i_A = \mathbb{1}_A, \quad \pi_B i_B = \mathbb{1}_B, \quad \pi_A i_B = 0, \quad \pi_B i_A = 0, \quad i_A \pi_A + i_B \pi_B = \mathbb{1}_C.$$

Conversely, suppose C is some other R -module, and that we are given four homomorphisms

$$\begin{array}{ccccc} A & & & & A \\ & \searrow^{i_A} & & & \nearrow^{\pi_A} \\ & & C & & \\ & \nearrow_{i_B} & & & \searrow_{\pi_B} \\ B & & & & B \end{array}$$

that likewise satisfy the five relations in (44.4). It then follows that the maps

$$A \oplus B \xrightarrow{i_A \oplus i_B} C \quad \text{and} \quad C \xrightarrow{(\pi_A, \pi_B)} A \oplus B$$

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 a covariant additive functor, then plugging the diagram (44.3) into \mathcal{F} gives us four maps

$$\begin{array}{ccccc} \mathcal{F}(A) & & & & \mathcal{F}(A) \\ & \searrow^{\mathcal{F}(i_A)} & & & \nearrow^{\mathcal{F}(\pi_A)} \\ & & \mathcal{F}(A \oplus B) & & \\ & \nearrow_{\mathcal{F}(i_B)} & & & \searrow_{\mathcal{F}(\pi_B)} \\ \mathcal{F}(B) & & & & \mathcal{F}(B) \end{array}$$

that similarly satisfy the five relations

$$\begin{aligned} \mathcal{F}(\pi_A)\mathcal{F}(i_A) &= \mathbb{1}_{\mathcal{F}(A)}, & \mathcal{F}(\pi_B)\mathcal{F}(i_B) &= \mathbb{1}_{\mathcal{F}(B)}, \\ \mathcal{F}(\pi_A)\mathcal{F}(i_B) &= 0, & \mathcal{F}(\pi_B)\mathcal{F}(i_A) &= 0, \\ \mathcal{F}(i_A)\mathcal{F}(\pi_A) &+ \mathcal{F}(i_B)\mathcal{F}(\pi_B) &= \mathbb{1}_C, \end{aligned}$$

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. If \mathcal{F} is instead a *contravariant* additive functor, then both the arrows and the order

of composition get reversed when plugging the maps from (44.3) into \mathcal{F} , thus producing four homomorphisms

$$\begin{array}{ccccc} \mathcal{F}(A) & & & & \mathcal{F}(A) \\ & \searrow^{\mathcal{F}(\pi_A)} & & \nearrow^{\mathcal{F}(i_A)} & \\ & & \mathcal{F}(A \oplus B) & & \\ & \nearrow_{\mathcal{F}(\pi_B)} & & \searrow_{\mathcal{F}(i_B)} & \\ \mathcal{F}(B) & & & & \mathcal{F}(B) \end{array}$$

that satisfy the relations

$$\begin{aligned} \mathcal{F}(i_A)\mathcal{F}(\pi_A) &= \mathbb{1}_{\mathcal{F}(A)}, & \mathcal{F}(i_B)\mathcal{F}(\pi_B) &= \mathbb{1}_{\mathcal{F}(B)}, \\ \mathcal{F}(i_B)\mathcal{F}(\pi_A) &= 0, & \mathcal{F}(i_A)\mathcal{F}(\pi_B) &= 0, \\ \mathcal{F}(\pi_A)\mathcal{F}(i_A) + \mathcal{F}(\pi_B)\mathcal{F}(i_B) &= \mathbb{1}_{\mathcal{F}(C)} \end{aligned}$$

for $C := \mathcal{F}(A \oplus B)$. These relations again 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 $\mathcal{F}(\pi_A) : \mathcal{F}(A) \rightarrow \mathcal{F}(A \oplus B)$ with the inclusion $\mathcal{F}(A) \rightarrow \mathcal{F}(A) \oplus \mathcal{F}(B)$, $\mathcal{F}(i_A) : \mathcal{F}(A \oplus B) \rightarrow \mathcal{F}(A)$ with the projection $\mathcal{F}(A) \oplus \mathcal{F}(B) \rightarrow \mathcal{F}(A)$, and so forth.

Recall that if $0 \rightarrow A \rightarrow C \rightarrow B \rightarrow 0$ is a split exact sequence, then C can be identified isomorphically with $A \oplus B$ so that the map $A \rightarrow C$ becomes the inclusion i_A and $C \rightarrow B$ becomes the projection π_B . One easy consequence of the natural isomorphisms explained above is then that split-exactness is preserved by additive functors:

PROPOSITION 44.5. *For any additive functor \mathcal{F} and any split exact sequence $0 \rightarrow A \xrightarrow{f} B \xrightarrow{g} C \rightarrow 0$, the sequence*

$$\begin{aligned} 0 \longrightarrow \mathcal{F}(A) \xrightarrow{\mathcal{F}(f)} \mathcal{F}(B) \xrightarrow{\mathcal{F}(g)} \mathcal{F}(C) \longrightarrow 0, & \quad \text{if } \mathcal{F} \text{ is covariant, or} \\ 0 \longrightarrow \mathcal{F}(C) \xrightarrow{\mathcal{F}(g)} \mathcal{F}(B) \xrightarrow{\mathcal{F}(f)} \mathcal{F}(A) \longrightarrow 0, & \quad \text{if } \mathcal{F} \text{ is contravariant} \end{aligned}$$

is also split exact. □

The main examples of additive functors you should have in mind at the moment are $\otimes G$ and $\text{Hom}(\cdot, G)$ for any fixed R -module G , and there are indeed natural isomorphisms

$$(A \oplus B) \otimes G \cong (A \otimes G) \oplus (B \otimes G), \quad \text{Hom}(A \oplus B, G) \cong \text{Hom}(A, G) \oplus \text{Hom}(B, G)$$

for any pair of R -modules A, B . 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.

If the sequence $0 \rightarrow A \rightarrow B \rightarrow C \rightarrow 0$ in Proposition 44.5 is assumed exact but not split exact, then there is no guarantee that the sequence obtained by feeding it into an additive functor will also be exact. We will see in particular that this is not true in general for either $\otimes G$ or $\text{Hom}(\cdot, G)$, but it is true sometimes, and this property is useful enough to deserve a name:

DEFINITION 44.6. An additive functor $\mathcal{F} : R\text{-Mod} \rightarrow R\text{-Mod}$ is called an **exact functor** if for every short exact sequence $0 \rightarrow A \xrightarrow{f} B \xrightarrow{g} C \rightarrow 0$, the sequence

$$\begin{aligned} 0 \longrightarrow \mathcal{F}(A) \xrightarrow{\mathcal{F}(f)} \mathcal{F}(B) \xrightarrow{\mathcal{F}(g)} \mathcal{F}(C) \longrightarrow 0, & \quad \text{if } \mathcal{F} \text{ is covariant, or} \\ 0 \longrightarrow \mathcal{F}(C) \xrightarrow{\mathcal{F}(g)} \mathcal{F}(B) \xrightarrow{\mathcal{F}(f)} \mathcal{F}(A) \longrightarrow 0, & \quad \text{if } \mathcal{F} \text{ is contravariant} \end{aligned}$$

is also exact.

Definition 44.6 appears to give a special role to *short* exact sequences, but the following result shows that this appearance is deceptive:

PROPOSITION 44.7. *An additive functor \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 $A \rightarrow B \rightarrow C$, the induced sequence $\mathcal{F}(A) \rightarrow \mathcal{F}(B) \rightarrow \mathcal{F}(C)$ is also exact.*

PROOF. See Exercise 44.2. □

We will see in the next lecture that the functors $\otimes G$ and $\text{Hom}(\cdot, G)$ are not exact in general, and this is the reason why both versions of the universal coefficient theorem assert the existence of natural split exact sequences rather than straightforward isomorphisms. On the other hand, it is useful to note that they are exact under certain circumstances, such as when the ring R is assumed to be a field:

PROPOSITION 44.8. *For any field \mathbb{K} , every additive functor $\mathcal{F} : \mathbb{K}\text{-Vect} \rightarrow \mathbb{K}\text{-Vect}$ is exact.*

PROOF. Every short exact sequence $0 \rightarrow A \rightarrow B \rightarrow C \rightarrow 0$ splits since C admits a basis, so Proposition 44.5 implies that the sequence remains exact after feeding it into \mathcal{F} . □

44.4. The main argument. We now explain the core of the argument behind both versions of the universal coefficient theorem.

We saw above that any additive functor $\mathcal{F} : R\text{-Mod} \rightarrow R\text{-Mod}$ turns a chain complex C_* of R -modules

$$\dots \rightarrow C_{n+1} \xrightarrow{\partial_{n+1}} C_n \xrightarrow{\partial_n} C_{n-1} \rightarrow \dots$$

into a new chain complex $\mathcal{F}(C_*)$ if \mathcal{F} is covariant, or a cochain complex if \mathcal{F} is contravariant. We now ask:

QUESTION 44.9. *For each $n \in \mathbb{Z}$, what relation is there between $\mathcal{F}(H_n(C_*))$ and $H_n(\mathcal{F}(C_*))$ or $H^n(\mathcal{F}(C_*))$ (in the covariant or contravariant case respectively)?*

We will be interested mainly in two specific examples: in the first, \mathcal{F} is $\otimes G$, and the problem is thus to find a relation between $H_n(C_*) \otimes G$ and $H_n(C_* \otimes G)$, so for instance, we would in this way find a relation between $H_n(X; R) \otimes G$ and $H_n(X; G)$. In the second example, \mathcal{F} is the contravariant functor $\text{Hom}(\cdot, G)$, so we are looking to relate $\text{Hom}(H_n(C_*), G)$ to $H^n(C_*; G) = H^n(\text{Hom}(C_*, G))$; in the setting of singular (co)chain complexes, this will mean a relation between $\text{Hom}(H_n(X; R), G)$ and $H^n(X; G)$. We shall impose two general assumptions that make the problem somewhat more tractable:

- ASSUMPTION 1: The chain groups $C_n \subset C_*$ are *free* R -modules;
- ASSUMPTION 2: The ring R is a principal ideal domain.

Both assumptions hold, for instance, whenever C_* is a singular/cellular/simplicial chain complex with coefficients in either an abelian group or a vector space over a field. They imply in particular that not only the modules C_n for each $n \in \mathbb{Z}$ but also their submodules

$$B_n := \text{im} \left(C_{n+1} \xrightarrow{\partial_{n+1}} C_n \right), \quad Z_n := \ker \left(C_n \xrightarrow{\partial_n} C_{n-1} \right)$$

consisting of boundaries and cycles respectively are free R -modules. Let

$$B_n \xleftarrow{i_n} Z_n \xleftarrow{j_n} C_n$$

denote the inclusion maps, and observe that since ∂_n vanishes on both B_n and Z_n , we can regard

$$B_* := \bigoplus_{n \in \mathbb{Z}} B_n \quad \text{and} \quad Z_* := \bigoplus_{n \in \mathbb{Z}} Z_n$$

as chain complexes with trivial boundary operators, so that the inclusions

$$B_* \xleftarrow{i} Z_* \xleftarrow{j} C_*$$

are chain maps.

Consider first the case where the additive functor \mathcal{F} is covariant. 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 $j_{n-1}i_{n-1} : B_{n-1} \hookrightarrow C_{n-1}$ produces a homomorphism

$$\mathcal{F}(B_{n-1}) \xrightarrow{\mathcal{F}(j_{n-1}i_{n-1})} \mathcal{F}(C_{n-1})$$

that need not be injective in general! What we can say instead is the following: the obvious commutative diagram

$$(44.5) \quad \begin{array}{ccc} C_n & \xrightarrow{\partial'_n} & B_{n-1} \\ \downarrow \partial_n & \swarrow j_{n-1}i_{n-1} & \\ C_{n-1} & & \end{array}$$

remains commutative after applying \mathcal{F} , thus producing

$$(44.6) \quad \mathcal{F}(\partial_n) = \mathcal{F}(j_{n-1}i_{n-1})\mathcal{F}(\partial'_n) = \mathcal{F}(j_{n-1})\mathcal{F}(i_{n-1})\mathcal{F}(\partial'_n)$$

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 bit of preparation out of the way, the next step in finding a relation between $\mathcal{F}(H_n(C_*))$ and $H_n(\mathcal{F}(C_*))$ is to observe that

$$(44.7) \quad 0 \longrightarrow Z_* \xrightarrow{j} C_* \xrightarrow{\partial'} B_{*-1} \longrightarrow 0$$

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. It follows via Proposition 44.5 that

$$0 \longrightarrow \mathcal{F}(Z_*) \xrightarrow{\mathcal{F}(j)} \mathcal{F}(C_*) \xrightarrow{\mathcal{F}(\partial')} \mathcal{F}(B_{*-1}) \longrightarrow 0$$

is also a split 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

$$\dots \longrightarrow \mathcal{F}(B_n) \xrightarrow{\Phi_n} \mathcal{F}(Z_n) \xrightarrow{\mathcal{F}(j)_*} H_n(\mathcal{F}(C_*)) \xrightarrow{\mathcal{F}(\partial')_*} \mathcal{F}(B_{n-1}) \xrightarrow{\Phi_{n-1}} \mathcal{F}(Z_{n-1}) \longrightarrow \dots,$$

where Φ_n and Φ_{n-1} are connecting homomorphisms. Let's take a closer look at what the map $\Phi_n : \mathcal{F}(B_n) \rightarrow \mathcal{F}(Z_n)$ actually is. If you recall how connecting homomorphisms were constructed

in Proposition 32.13, the diagram that we need to chase is:

$$\begin{array}{ccccccc}
 & & \vdots & & \vdots & & \vdots \\
 & & \downarrow & & \downarrow & & \downarrow \\
 0 & \longrightarrow & \mathcal{F}(Z_{n+1}) & \xrightarrow{\mathcal{F}(j_{n+1})} & \mathcal{F}(C_{n+1}) & \xrightarrow{\mathcal{F}(\partial'_{n+1})} & \mathcal{F}(B_n) \longrightarrow 0 \\
 & & \downarrow 0 & & \downarrow \mathcal{F}(\partial_{n+1}) & & \downarrow 0 \\
 0 & \longrightarrow & \mathcal{F}(Z_n) & \xrightarrow{\mathcal{F}(j_n)} & \mathcal{F}(C_n) & \xrightarrow{\mathcal{F}(\partial'_n)} & \mathcal{F}(B_{n-1}) \longrightarrow 0 \\
 & & \downarrow & & \downarrow \mathcal{F}(\partial_n) & & \downarrow \\
 & & \vdots & & \vdots & & \vdots
 \end{array}$$

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 (44.6), $\mathcal{F}(\partial_{n+1})c = \mathcal{F}(j_n)\mathcal{F}(i_n)\mathcal{F}(\partial'_{n+1})c = \mathcal{F}(j_n)\mathcal{F}(i_n)b$, thus the required condition is satisfied by $\Phi_n(b) := \mathcal{F}(i_n)b$, and our long exact sequence therefore becomes

$$\dots \longrightarrow \mathcal{F}(B_n) \xrightarrow{\mathcal{F}(i_n)} \mathcal{F}(Z_n) \xrightarrow{\mathcal{F}(j)_*} H_n(\mathcal{F}(C_*)) \xrightarrow{\mathcal{F}(\partial')_*} \mathcal{F}(B_{n-1}) \xrightarrow{\mathcal{F}(i_{n-1})} \mathcal{F}(Z_{n-1}) \longrightarrow \dots$$

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

$$(44.8) \quad 0 \longrightarrow \operatorname{coker} \mathcal{F}(i_n) \xrightarrow{\mathcal{F}(j)_*} H_n(\mathcal{F}(C_*)) \xrightarrow{\mathcal{F}(\partial')_*} \ker \mathcal{F}(i_{n-1}) \longrightarrow 0.$$

We will eventually deduce the general version of the universal coefficient theorem for homology from this short exact sequence.

In order to understand more precisely what $\operatorname{coker} \mathcal{F}(i_n)$ and $\ker \mathcal{F}(i_{n-1})$ are, it helps to examine a second short exact sequence

$$(44.9) \quad 0 \longrightarrow B_n \xleftarrow{i_n} Z_n \xrightarrow{q_n} H_n(C_*) \longrightarrow 0,$$

in which q_n denotes the quotient projection. Feeding this into the functor \mathcal{F} produces a chain complex

$$(44.10) \quad 0 \longrightarrow \mathcal{F}(B_n) \xrightarrow{\mathcal{F}(i_n)} \mathcal{F}(Z_n) \xrightarrow{\mathcal{F}(q_n)} \mathcal{F}(H_n(C_*)) \longrightarrow 0,$$

though since $H_n(C_*)$ need not be a free R -module, there is no guarantee that the sequence (44.9) splits, and therefore no guarantee that (44.10) will be exact. It will thus be helpful at this point to impose a third assumption on our setting, though we will put considerable effort into lifting this assumption in the next lecture:

- ASSUMPTION 3: The functor \mathcal{F} is exact.

The added assumption guarantees that (44.10) is indeed an exact sequence for every value of $n \in \mathbb{Z}$. One immediate consequence is that $\mathcal{F}(i_{n-1})$ is injective, thus

$$\ker \mathcal{F}(i_{n-1}) = 0,$$

and another is that $\mathcal{F}(q_n)$ is surjective and descends to an isomorphism

$$\operatorname{coker} \mathcal{F}(i_n) \xrightarrow[\cong]{\mathcal{F}(q_n)} \mathcal{F}(H_n(C_*)) .$$

Using these to replace the first and last nontrivial terms in the short exact sequence (44.8), the sequence now contains only two nontrivial terms, and we conclude that there is a natural isomorphism

$$\mathcal{F}(H_n(C_*)) \xrightarrow[\cong]{h} H_n(\mathcal{F}(C_*)).$$

At this level of generality, one cannot quite write down an explicit formula for the map h , but the argument above characterizes h as the unique map such that

$$(44.11) \quad h(x) = [\mathcal{F}(j_n)z] \in H_n(\mathcal{F}(C_*)) \quad \text{for any } z \in \mathcal{F}(Z_n) \text{ with } \mathcal{F}(q_n)z = x \in \mathcal{F}(H_n(C_*)).$$

Indeed, it will be useful to note for later that this characterization of h does not require \mathcal{F} to be an exact functor, but only requires the sequence (44.10) to be exact at its middle and last terms, which makes $\mathcal{F}(q_n) : \mathcal{F}(Z_n) \rightarrow \mathcal{F}(H_n(C_*))$ a surjective map that descends to an isomorphism $\text{coker } \mathcal{F}(i_n) \cong \mathcal{F}(H_n(C_*))$. This means in particular that 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 $\mathcal{F}(Z_n)/\text{im } \mathcal{F}(i_n) = \text{coker } \mathcal{F}(i_n)$ is unique. The map $\text{coker } \mathcal{F}(i_n) \rightarrow H_n(\mathcal{F}(C_*))$ in the exact sequence (44.8) then arises by letting the chain map $\mathcal{F}(j) : \mathcal{F}(Z_*) \rightarrow \mathcal{F}(C_*)$ descend to homology as the map $\mathcal{F}(Z_n) \rightarrow H_n(\mathcal{F}(C_*)) : z \mapsto [\mathcal{F}(j_n)z]$, and then letting the latter descend to the quotient of $\mathcal{F}(Z_n)$ by $\text{im } \mathcal{F}(i_n)$. The result is precisely the characterization in (44.11). For the special case $\mathcal{F} = \otimes G$, one easily deduces from this that h is the canonical map

$$(44.12) \quad H_n(C_*) \otimes G \xrightarrow{h} H_n(C_* \otimes G) : [c] \otimes g \mapsto [c \otimes g].$$

We've proved a first special case of the universal coefficient theorem for homology:

THEOREM 44.10. *Suppose C_* is a chain complex of free modules over a principal ideal domain R , and G is an R -module such that $\otimes G : R\text{-Mod} \rightarrow R\text{-Mod}$ is an exact functor. Then the natural map $h : H_n(C_*) \otimes G \rightarrow H_n(C_* \otimes G)$ described in (44.12) is an isomorphism for every $n \in \mathbb{Z}$. \square*

By Proposition 44.8, the extra hypothesis that $\otimes G$ is exact holds in particular whenever R is a field \mathbb{K} , so for any vector space G over \mathbb{K} , we obtain natural vector space isomorphisms

$$H_n(X, A; \mathbb{K}) \otimes G \cong H_n(X, A; G)$$

for all $n \geq 0$ and all pairs of spaces (X, A) , and similarly in cellular or simplicial homology.

Adapting the argument above for the case of a contravariant exact functor \mathcal{F} is straightforward. Instead of a chain complex, $\mathcal{F}(C_*)$ is now a cochain complex

$$\dots \longrightarrow \mathcal{F}(C_{n-1}) \xrightarrow{\mathcal{F}(\partial_n)} \mathcal{F}(C_n) \xrightarrow{\mathcal{F}(\partial_{n+1})} \mathcal{F}(C_{n+1}) \longrightarrow \dots,$$

and applying \mathcal{F} to the diagram (44.5) produces

$$(44.13) \quad \mathcal{F}(\partial_n) = \mathcal{F}(\partial'_n)\mathcal{F}(j_{n-1}i_{n-1}) = \mathcal{F}(\partial'_n)\mathcal{F}(i_{n-1})\mathcal{F}(j_{n-1})$$

as a relation between $\mathcal{F}(C_{n-1}) \xrightarrow{\mathcal{F}(\partial_n)} \mathcal{F}(C_n)$ and $\mathcal{F}(B_{n-1}) \xrightarrow{\mathcal{F}(\partial'_n)} \mathcal{F}(C_n)$. Plugging the split exact sequence (44.7) into \mathcal{F} produces a short exact sequence of cochain complexes

$$0 \longleftarrow \mathcal{F}(Z_*) \xleftarrow{\mathcal{F}(j)} \mathcal{F}(C_*) \xleftarrow{\mathcal{F}(\partial')} \mathcal{F}(B_{*-1}) \longleftarrow 0,$$

and therefore a long exact sequence

$$\dots \longleftarrow \mathcal{F}(B_n) \xleftarrow{\Phi_n} \mathcal{F}(Z_n) \xleftarrow{\mathcal{F}(j)^*} H^n(\mathcal{F}(C_*)) \xleftarrow{\mathcal{F}(\partial')^*} \mathcal{F}(B_{n-1}) \xleftarrow{\Phi_{n-1}} \mathcal{F}(Z_{n-1}) \longleftarrow \dots$$

One needs to do another 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

$$\Phi_n = \mathcal{F}(i_n) : \mathcal{F}(Z_n) \rightarrow \mathcal{F}(B_n),$$

and the proof depends on the relation (44.13). Turning the long exact sequence into a short exact sequence with $H^n(\mathcal{F}(C_*))$ as the middle term thus gives

$$0 \longrightarrow \operatorname{coker} \mathcal{F}(i_{n-1}) \longrightarrow H^n(\mathcal{F}(C_*)) \longrightarrow \ker \mathcal{F}(i_n) \longrightarrow 0.$$

If the functor \mathcal{F} is exact, then it also turns the short exact sequence $0 \rightarrow B_n \xrightarrow{i_n} Z_n \xrightarrow{q_n} H_n(C_*) \rightarrow 0$ into a short exact sequence

$$0 \longrightarrow \mathcal{F}(H_n(C_*)) \xrightarrow{\mathcal{F}(q_n)} \mathcal{F}(Z_n) \xrightarrow{\mathcal{F}(i_n)} \mathcal{F}(B_n) \longrightarrow 0,$$

implying that $\operatorname{coker} \mathcal{F}(i_n) = 0$ for all $n \in \mathbb{Z}$ and giving a natural isomorphism

$$\mathcal{F}(H_n(C_*)) \xrightarrow[\cong]{\mathcal{F}(q_n)} \ker \mathcal{F}(i_n) \subset \mathcal{F}(Z_n).$$

Our previous short exact sequence thus becomes an isomorphism

$$(44.14) \quad H^n(\mathcal{F}(C_*)) \xrightarrow[\cong]{h} \mathcal{F}(H_n(C_*)).$$

The map $h : H^n(\mathcal{F}(C_*)) \rightarrow \mathcal{F}(H_n(C_*))$ is characterized uniquely by the relation

$$(44.15) \quad \mathcal{F}(q_n)h([\varphi]) = \mathcal{F}(j_n)\varphi \quad \text{for all } \varphi \in \ker \mathcal{F}(\partial_{n+1}) \subset \mathcal{F}(C_n),$$

due to the fact that $\mathcal{F}(q_n)$ maps $\mathcal{F}(H_n(C_*))$ isomorphically to $\ker \mathcal{F}(i_n)$. In the special case $\mathcal{F} = \operatorname{Hom}(\cdot, G)$, this makes h the map $H^n(C_*; G) \xrightarrow{h} \operatorname{Hom}(H_n(C_*), G)$ defined via evaluation of cochains on chains,

$$(44.16) \quad h([\alpha])[c] = \alpha(c).$$

The resulting special case of the universal coefficient theorem for cohomology reads:

THEOREM 44.11. *Suppose C_* is a chain complex of free modules over a principal ideal domain R , and G is an R -module such that $\operatorname{Hom}(\cdot, G) : R\text{-Mod} \rightarrow R\text{-Mod}$ is an exact functor. Then the natural map $h : H^n(C_*; G) \rightarrow \operatorname{Hom}(H_n(C_*), G)$ described in (44.16) is an isomorphism for every $n \in \mathbb{Z}$. \square*

Once again, Proposition 44.8 provides the exactness hypothesis for free whenever R is a field \mathbb{K} , so for the singular cohomology of a pair of spaces (X, A) and any vector space G over \mathbb{K} , one obtains vector space isomorphisms

$$H^n(X, A; G) \cong \operatorname{Hom}(H_n(X, A; \mathbb{K}), G),$$

and similarly for cellular or simplicial homology. This result is already interesting when G is taken to be the field \mathbb{K} itself, because it then identifies the cohomology $H^n(X, A; \mathbb{K})$ with the dual vector space of the homology $H_n(X, A; \mathbb{K})$. One immediate application is that we can freely replace homology by cohomology in order to compute the Betti numbers of a space,

$$b_k(X) = \dim_{\mathbb{Q}} H_k(X; \mathbb{Q}) = \dim_{\mathbb{Q}} H^k(X; \mathbb{Q}),$$

and since every matrix has the same trace as its transpose, one can similarly use cohomology to compute Lefschetz numbers, as mentioned in the motivational introduction to Lecture 41; for details on the latter result, see Exercise 44.3.

44.5. Exercises.

EXERCISE 44.1. The following exercise is for readers who have not previously encountered tensor products in the categories of abelian groups or R -modules. We first recall the definition. Given a set S , let $F(S) := \bigoplus_{e \in S} R \in R\text{-Mod}$ denote the free R -module generated by this set; each of its elements can be written as $\sum_{e \in S} r_e e$ for a unique set of coefficients $r_e \in R$, assuming that only finitely-many of them are nonzero so that the sum is well defined. The **tensor product** of two R -modules A and B is then defined as a certain quotient of the free R -module generated by $A \times B$, namely

$$A \otimes B := F(A \times B)/N,$$

where $N \subset F(A \times B)$ 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 $(a, b) \in F(A \times B)$ in the quotient by

$$a \otimes b := [(a, b)] \in A \otimes B.$$

Recall that for R -modules A, B, C , a map $\Phi : A \oplus B \rightarrow C$ is called **bilinear** if the maps $A \rightarrow C : a \mapsto \Phi(a, b_0)$ and $B \rightarrow C : b \mapsto \Phi(a_0, b)$ are R -module homomorphisms for all choices of fixed elements $a_0 \in A$ and $b_0 \in B$.

- (a) Show that the map $G \oplus H \rightarrow G \otimes H : (g, h) \mapsto g \otimes h$ 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 $\Phi : G \oplus H \rightarrow K$ of R -modules, there exists a unique R -module homomorphism $\Psi : G \otimes H \rightarrow K$ such that $\Phi(g, h) = \Psi(g \otimes h)$ for all $(g, h) \in G \oplus H$.
- (c) Show that for any R -module G , the map $G \rightarrow G \otimes R : g \mapsto g \otimes 1$ is an 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 $F(S) \otimes F(T)$ to $F(S \times T)$.
- (f) For any R -modules A, B, C, D and homomorphisms $f : A \rightarrow B$, $g : C \rightarrow D$, show that there exists a homomorphism

$$f \otimes g : A \otimes C \rightarrow B \otimes D$$

defined uniquely by the condition $(f \otimes g)(a \otimes c) = f(a) \otimes g(c)$ for all $a \in A$ and $c \in C$.

- (g) In the case $R = \mathbb{Z}$, an element $a \in G$ of an abelian group G is said to be **torsion** if $ma = 0$ for some $m \in \mathbb{N}$. Show that if every element of G is torsion and \mathbb{K} is a field of characteristic zero (regarded as an abelian group with respect to addition), then $G \otimes \mathbb{K} = 0$.

EXERCISE 44.2 (*). Prove Proposition 44.7 on conditions equivalent to the exactness of an additive functor.

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 $0 \rightarrow A/\ker(f) \rightarrow B \rightarrow \text{im}(g) \rightarrow 0$.

EXERCISE 44.3 (*). Fix an R -module G and an integer $n \geq 0$.

- (a) Show that the canonical homomorphisms $h : H_n(X; R) \otimes G \rightarrow H_n(X; G)$ and $h : H^n(X; G) \rightarrow \text{Hom}(H_n(X; R), G)$ have the following naturality properties: for any continuous map $f : X \rightarrow Y$, the diagrams

$$\begin{array}{ccc} H_n(X; R) \otimes G & \xrightarrow{h} & H_n(X; G) \\ \downarrow f_* \otimes \mathbb{1} & & \downarrow f_* \\ H_n(Y; R) \otimes G & \xrightarrow{h} & H_n(Y; G) \end{array} \quad \text{and} \quad \begin{array}{ccc} H^n(X; G) & \xrightarrow{h} & \text{Hom}(H_n(X; R), G) \\ f^* \uparrow & & f^* \uparrow \\ H^n(Y; G) & \xrightarrow{h} & \text{Hom}(H_n(Y; R), G) \end{array}$$

both commute, where $f^* : \text{Hom}(H_n(Y; R), G) \rightarrow \text{Hom}(H_n(X; R), G)$ denotes the dualization of $f_* : H_n(X; R) \rightarrow H_n(Y; R)$.

- (b) Prove that for any field \mathbb{K} , the Lefschetz number $L_{\mathbb{K}}(f)$ of a map $f : X \rightarrow X$ defined in §40.4 satisfies

$$L_{\mathbb{K}}(f) = \sum_{n \in \mathbb{Z}} (-1)^n \text{tr} \left(H^n(X; \mathbb{K}) \xrightarrow{f^*} H^n(X; \mathbb{K}) \right).$$

Remark: For the computation of Lefschetz numbers in terms of homology or cohomology with coefficients in \mathbb{Z} , see Exercises 45.11 and 45.12 in the next lecture.

45. The derived functors Tor and Ext

Now comes the bad news: the natural maps

$$H_n(X; R) \otimes G \xrightarrow{h} H_n(X; G), \quad \text{and} \quad H^n(X; G) \xrightarrow{h} \text{Hom}(H_n(X; R), G)$$

are not always isomorphisms. We proved in the previous lecture that they are isomorphisms whenever R is a principal ideal domain and G is an R -module for which $\otimes G$ or $\text{Hom}(\cdot, G)$ respectively is an exact functor $R\text{-Mod} \rightarrow R\text{-Mod}$. We saw that both are exact when R is a field, but in the present lecture, we shall confront the reality that both can, in general, fail to be exact functors. On the other hand, this failure is not catastrophic, and it can be measured in precise terms via the nontriviality of certain auxiliary functors called

$$\text{Tor}(\cdot, G) : R\text{-Mod} \rightarrow R\text{-Mod}, \quad \text{and} \quad \text{Ext}(\cdot, G) : R\text{-Mod} \rightarrow R\text{-Mod}.$$

Once these have been defined and their main properties elucidated, they will furnish us with a third term in each of the short exact sequences that constitute the general versions of the universal coefficient theorems.

45.1. The properties of Tor and Ext. We now state two theorems whose proofs will occupy the bulk of this lecture.

THEOREM 45.1. *Assume R is a commutative ring with unit. There exists a sequence of functors*

$$\text{Tor}_n : R\text{-Mod} \times R\text{-Mod} \rightarrow R\text{-Mod}, \quad n \in \mathbb{N},$$

which are canonically defined up to natural isomorphisms and covariant in both variables, and have the following properties:

- (1) *For every short exact sequence of R -modules $0 \rightarrow A \xrightarrow{f} B \xrightarrow{g} C \rightarrow 0$ and every fixed R -module G , the maps $f_* : \text{Tor}_n(A, G) \rightarrow \text{Tor}_n(B, G)$ and $g_* : \text{Tor}_n(B, G) \rightarrow \text{Tor}_n(C, G)$ defined by feeding f and g respectively into the functors $\text{Tor}_n(\cdot, G) : R\text{-Mod} \rightarrow R\text{-Mod}$ fit into a long exact sequence*

$$\begin{aligned} \dots \rightarrow \text{Tor}_2(A, G) \xrightarrow{f_*} \text{Tor}_2(B, G) \xrightarrow{g_*} \text{Tor}_2(C, G) \rightarrow \text{Tor}_1(A, G) \xrightarrow{f_*} \text{Tor}_1(B, G) \xrightarrow{g_*} \text{Tor}_1(C, G) \\ \rightarrow A \otimes G \xrightarrow{f \otimes 1} B \otimes G \xrightarrow{g \otimes 1} C \otimes G \rightarrow 0. \end{aligned}$$

- (2) *For every fixed R -module G and each $n \in \mathbb{N}$, the functors $\text{Tor}_n(\cdot, G)$ and $\text{Tor}_n(G, \cdot)$ are additive; in particular, there are natural isomorphisms*

$$\begin{aligned} \text{Tor}_n(A \oplus B, G) &\cong \text{Tor}_n(A, G) \oplus \text{Tor}_n(B, G), \\ \text{Tor}_n(G, A \oplus B) &\cong \text{Tor}_n(G, A) \oplus \text{Tor}_n(G, B) \end{aligned}$$

for all pairs of R -modules A, B .

- (3) $\text{Tor}_n(A, G) = 0$ for all $n \geq 1$ whenever either of A or G is a free R -module.⁶⁶
 (4) $\text{Tor}_n(A, G) = 0$ for all $n \geq 2$ and all A, G whenever R is a principal ideal domain.
 (5) In the case $R = \mathbb{Z}$, $\text{Tor}_1(A, G) = 0$ whenever either A or G has no torsion.
 (6) If $k \in \mathbb{N}$ has the property that no nonzero element $x \in R$ satisfies $kx = 0$, then for every R -module G , one has

$$\text{Tor}_1(R/kR, G) \cong \ker \left(G \xrightarrow{\cdot k} G \right), \quad \text{and} \quad \text{Tor}_n(R/kR, G) = 0 \text{ for } n \geq 2.$$

- (7) There are natural isomorphisms $\text{Tor}_n(A, G) \cong \text{Tor}_n(G, A)$ for all $n \geq 1$ and all A, G .

A few comments on this statement before we continue. First, the words “canonically defined up to natural isomorphisms” may seem mysterious at first, but the proof of the theorem will make their meaning clear (see in particular Remark 45.27). The definitions of the functors Tor_n depend on choices, but the resulting functors are canonical in the sense that for any two such functors Tor_n and Tor'_n defined via different choices, there is a natural isomorphism $\text{Tor}_n \Rightarrow \text{Tor}'_n$, meaning a natural transformation such that the resulting morphisms $\text{Tor}_n(A, G) \rightarrow \text{Tor}'_n(A, G)$ are isomorphisms for every A, G .

Second, we will be concerned mainly with cases where R is a principal ideal domain, so that the “higher” Tor functors Tor_n all vanish for $n \geq 2$, leaving only the case $n = 1$, which gets a special name

$$\text{Tor}(A, G) := \text{Tor}_1(A, G).$$

The long exact sequence arising from any short exact sequence $0 \rightarrow A \rightarrow B \rightarrow C \rightarrow 0$ then becomes a sequence with at most six nontrivial terms,

$$0 \rightarrow \text{Tor}(A, G) \rightarrow \text{Tor}(B, G) \rightarrow \text{Tor}(C, G) \rightarrow A \otimes G \rightarrow B \otimes G \rightarrow C \otimes G \rightarrow 0.$$

Third, if you are familiar with the classification of finitely-generated abelian groups, then you will notice that the properties listed in the theorem make computations of $\text{Tor}(A, G)$ a straightforward matter in the case $R = \mathbb{Z}$, i.e. when A and G are just abelian groups, at least if they are finitely generated. In light of additivity and the vanishing of Tor on torsion-free groups, $\text{Tor}(A, G)$ can always be expressed as a direct sum of terms arising from the formula

$$\text{Tor}(\mathbb{Z}_k, \mathbb{Z}_\ell) \cong \ker \left(\mathbb{Z}_\ell \xrightarrow{\cdot k} \mathbb{Z}_\ell \right) \cong \mathbb{Z}_{\text{gcd}(k, \ell)}.$$

Lastly: just like the symbol $A \otimes G$, $\text{Tor}(A, G)$ and $\text{Tor}_n(A, G)$ have different meanings depending on whether A and G are regarded as mere abelian groups (i.e. \mathbb{Z} -modules) or as modules over some other ring $R \neq \mathbb{Z}$. In situations where this distinction is important, one sometimes sees notation such as

$$\text{Tor}^R(A, G) := \text{Tor}(A, G)$$

to specify that we are working with a functor on $R\text{-Mod}$ rather than Ab . I prefer to avoid such notational clutter and let the context determine the precise meaning of $\text{Tor}(A, G)$, but one should be aware that in some books, the convention is to reserve this notation for the case $\text{Tor}^{\mathbb{Z}}(A, G)$, and write Tor^R explicitly whenever any ring R other than \mathbb{Z} is to be used.

Here’s the contravariant cousin of Theorem 45.1.

THEOREM 45.2. *Assume R is a commutative ring with unit. There exists a sequence of functors*

$$\text{Ext}_n : R\text{-Mod} \times R\text{-Mod} \rightarrow R\text{-Mod}, \quad n \in \mathbb{N},$$

which are canonically defined up to natural isomorphisms and contravariant in the first variable but covariant in the second variable, and have the following properties:

⁶⁶We will see in §45.4 that this statement also holds under a somewhat weaker hypothesis, namely that A or G is a *projective* R -module, though in typical cases of interest, the two conditions are equivalent (cf. Remark 45.19).

- (1) For every short exact sequence of R -modules $0 \rightarrow A \xrightarrow{f} B \xrightarrow{g} C \rightarrow 0$ and every fixed R -module G , the maps $f^* : \text{Ext}_n(B, G) \rightarrow \text{Ext}_n(A, G)$ and $g^* : \text{Ext}_n(C, G) \rightarrow \text{Ext}_n(B, G)$ defined by feeding f and g respectively into the functors $\text{Ext}_n(\cdot, G) : R\text{-Mod} \rightarrow R\text{-Mod}$ fit into a long exact sequence

$$0 \rightarrow \text{Hom}(C, G) \xrightarrow{g^*} \text{Hom}(B, G) \xrightarrow{f^*} \text{Hom}(A, G) \rightarrow \text{Ext}_1(C, G) \xrightarrow{g^*} \text{Ext}_1(B, G) \xrightarrow{f^*} \text{Ext}_1(A, G) \\ \rightarrow \text{Ext}_2(C, G) \xrightarrow{g^*} \text{Ext}_2(B, G) \xrightarrow{f^*} \text{Ext}_2(A, G) \rightarrow \dots$$

- (2) For every fixed R -module G and each $n \in \mathbb{N}$, the functors $\text{Ext}_n(\cdot, G)$ and $\text{Ext}_n(G, \cdot)$ are additive; in particular, there are natural isomorphisms

$$\begin{aligned} \text{Ext}_n(A \oplus B, G) &\cong \text{Ext}_n(A, G) \oplus \text{Ext}_n(B, G), \\ \text{Ext}_n(G, A \oplus B) &\cong \text{Ext}_n(G, A) \oplus \text{Ext}_n(G, B) \end{aligned}$$

for all pairs of R -modules A, B .

- (3) $\text{Ext}_n(A, G) = 0$ for all $n \geq 1$ whenever A is a free R -module⁶⁷ or $\text{Hom}(\cdot, G)$ is an exact functor.
 (4) $\text{Ext}_n(A, G) = 0$ for all $n \geq 2$ and all A, G whenever R is a principal ideal domain.
 (5) If $k \in \mathbb{N}$ has the property that no nonzero element $x \in R$ satisfies $kx = 0$, then for every R -module G , one has

$$\text{Ext}_1(R/kR, G) \cong \text{coker} \left(\text{Hom}(R, G) \xrightarrow{\cdot k} \text{Hom}(R, G) \right), \quad \text{and} \quad \text{Ext}_n(R/kR, G) = 0 \text{ for } n \geq 2.$$

The same remarks made above for the Tor functors regarding their uniqueness up to natural isomorphism and their computability in the case $R = \mathbb{Z}$ apply equally well to the Ext functors. In light of the vanishing of Ext_n for $n \geq 2$ when R is a principal ideal domain, we similarly adopt the notation

$$\text{Ext}(A, G) := \text{Ext}_1(A, G),$$

thus turning the long exact sequence arising from any short exact sequence $0 \rightarrow A \rightarrow B \rightarrow C \rightarrow 0$ into another six-term sequence

$$0 \rightarrow \text{Hom}(C, G) \rightarrow \text{Hom}(B, G) \rightarrow \text{Hom}(A, G) \rightarrow \text{Ext}(C, G) \rightarrow \text{Ext}(B, G) \rightarrow \text{Ext}(A, G) \rightarrow 0.$$

In the case $R = \mathbb{Z}$, computing $\text{Ext}(A, G)$ when A is a finitely-generated abelian group reduces to taking direct sums of terms arising from the formula

$$\text{Ext}(\mathbb{Z}_k, G) \cong \text{coker} \left(\text{Hom}(\mathbb{Z}, G) \xrightarrow{\cdot k} \text{Hom}(\mathbb{Z}, G) \right) \cong G/kG.$$

The symmetry isomorphism $\text{Tor}_n(A, G) \cong \text{Tor}_n(G, A)$ also has an analogue for the Ext functors, but it is less straightforward to state; it involves an alternative definition of $\text{Ext}_n(A, G)$ that will be sketched in §45.6.

45.2. Left- and right-exact functors. The most important features of Theorems 45.1 and 45.2 are the long exact sequences, which quantify the failure of $\otimes G$ and $\text{Hom}(\cdot, G)$ to be exact functors, e.g. we see from Theorem 45.1 that for any given short exact sequence $0 \rightarrow A \rightarrow B \rightarrow C \rightarrow 0$ and any R -module G , the induced sequence $0 \rightarrow A \otimes G \rightarrow B \otimes G \rightarrow C \otimes G \rightarrow 0$ will be exact whenever $\text{Tor}(C, G) = 0$. Even when this is not the case, the theorem implies that the sequence

$$A \otimes G \rightarrow B \otimes G \rightarrow C \otimes G \rightarrow 0$$

⁶⁷As in Theorem 45.1, this statement also holds under the seemingly weaker (but typically equivalent) hypothesis that A is projective.

will be exact in any case, i.e. exactness always holds at the second and third nontrivial terms, and we see in Theorem 45.2 that $\text{Hom}(\cdot, G)$ has a similar property with the arrows reversed. There is terminology to describe additive functors with these properties.

DEFINITION 45.3. An additive covariant functor $\mathcal{F} : R\text{-Mod} \rightarrow R\text{-Mod}$ is called **right-exact** if for every exact sequence $A \xrightarrow{i} B \xrightarrow{j} C \rightarrow 0$, the sequence

$$\mathcal{F}(A) \xrightarrow{\mathcal{F}(i)} \mathcal{F}(B) \xrightarrow{\mathcal{F}(j)} \mathcal{F}(C) \longrightarrow 0$$

is also exact, and **left-exact** if for every exact sequence $0 \rightarrow A \xrightarrow{i} B \xrightarrow{j} C$, the sequence

$$0 \rightarrow \mathcal{F}(A) \xrightarrow{\mathcal{F}(i)} \mathcal{F}(B) \xrightarrow{\mathcal{F}(j)} \mathcal{F}(C)$$

is also exact. Similarly, if \mathcal{F} is additive and contravariant, it is called **right-exact** if for every exact sequence $0 \rightarrow A \xrightarrow{i} B \xrightarrow{j} C$, the sequence

$$\mathcal{F}(C) \xrightarrow{\mathcal{F}(j)} \mathcal{F}(B) \xrightarrow{\mathcal{F}(i)} \mathcal{F}(A) \longrightarrow 0$$

is also exact, and **left-exact** if for every exact sequence $A \xrightarrow{i} B \xrightarrow{j} C \rightarrow 0$, the sequence

$$0 \rightarrow \mathcal{F}(C) \xrightarrow{\mathcal{F}(j)} \mathcal{F}(B) \xrightarrow{\mathcal{F}(i)} \mathcal{F}(A)$$

is also exact.

Recalling Definition 44.6 from the previous lecture, we see that a functor is exact if and only if it is both right-exact and left-exact.

PROPOSITION 45.4. For any R -module G , the functors $\otimes G : R\text{-Mod} \rightarrow R\text{-Mod}$ and $\text{Hom}(\cdot, G) : R\text{-Mod} \rightarrow R\text{-Mod}$ are right-exact and left-exact respectively.

PROOF. Exactness of a sequence $A \xrightarrow{i} B \xrightarrow{j} C \rightarrow 0$ means that j vanishes on $\text{im}(i)$ and descends to an isomorphism

$$(45.1) \quad \text{coker}(i) = B/\text{im}(i) \xrightarrow[\cong]{j} C.$$

The goal is now to show that the induced sequences

$$A \otimes G \xrightarrow{i \otimes 1} B \otimes G \xrightarrow{j \otimes 1} C \otimes G \longrightarrow 0$$

and

$$0 \longrightarrow \text{Hom}(C, G) \xrightarrow{j^*} \text{Hom}(B, G) \xrightarrow{i^*} \text{Hom}(A, G)$$

are both exact. The first is equivalent to showing that the map

$$(45.2) \quad \text{coker}(i \otimes 1) = (B \otimes G)/\text{im}(i \otimes 1) \xrightarrow{j \otimes 1} C \otimes G$$

to which $j \otimes 1$ descends is an isomorphism, while the second is equivalent showing that j^* defines an isomorphism

$$(45.3) \quad \text{Hom}(C, G) \xrightarrow{j^*} \ker(i^*) \subset \text{Hom}(B, G).$$

We can establish both by writing down explicit inverses: first, one easily checks that the map

$$C \otimes G \rightarrow \text{coker}(i \otimes 1) : c \otimes g \mapsto [b \otimes g]$$

defined by choosing any $b \in j^{-1}(c)$ is a well-defined homomorphism inverse to the map in (45.2). To write down an inverse of (45.3), it is useful to note that $\ker(i^*) \subset \text{Hom}(B, G)$ consists of all homomorphisms $\varphi : B \rightarrow G$ that vanish on $\text{im}(i) \subset B$ and thus descend to maps $\text{coker}(i) \rightarrow G$,

giving a natural identification of $\ker(i^*)$ with $\text{Hom}(\text{coker}(i), G)$. Using this identification, the inverse of (45.3) is the map

$$\text{Hom}(\text{coker}(i), G) \rightarrow \text{Hom}(C, G) : \varphi \mapsto \varphi \circ j^{-1},$$

where j^{-1} in this context denotes the inverse of the isomorphism (45.1). □

The reason that $\otimes G$ and $\text{Hom}(\cdot, G)$ can each fail in general to be exact is that the injectivity of a map $i : A \hookrightarrow B$ does not guarantee the injectivity of $i \otimes \mathbb{1} : A \otimes G \rightarrow B \otimes G$, nor the surjectivity of $i^* : \text{Hom}(B, G) \rightarrow \text{Hom}(A, G)$. Both can be framed as conditions on the R -module G .

DEFINITION 45.5. An R -module G is **flat** if for every pair of R -modules A, B with an injective homomorphism $i : A \hookrightarrow B$, the homomorphism $i \otimes \mathbb{1} : A \otimes G \rightarrow B \otimes G$ is also injective.

DEFINITION 45.6. An R -module G is **injective** if for every pair of R -modules A, B with an injective homomorphism $i : A \rightarrow B$ and a homomorphism $\varphi : A \rightarrow G$, one can find a third homomorphism $\tilde{\varphi} : B \rightarrow G$ such that the diagram

$$\begin{array}{ccc} B & & \\ \uparrow i & \searrow \tilde{\varphi} & \\ A & \xrightarrow{\varphi} & G \end{array}$$

commutes. In other words: Every injection $A \xrightarrow{i} B$ has a surjective dual $\text{Hom}(B, G) \xrightarrow{i^*} \text{Hom}(A, G)$.

The right-exactness of $\otimes G$ and left-exactness of $\text{Hom}(\cdot, G)$ now imply:

COROLLARY 45.7. An R -module G is flat if and only if the functor $\otimes G$ is exact, and G is injective if and only if the functor $\text{Hom}(\cdot, G)$ is exact. □

EXAMPLE 45.8. Consider the exact sequence of abelian groups

$$(45.4) \quad 0 \longrightarrow 2\mathbb{Z} \xrightarrow{i} \mathbb{Z} \xrightarrow{\text{pr}} \mathbb{Z}_2 \longrightarrow 0$$

formed by the inclusion i and quotient projection. Feeding it into the functor $\otimes \mathbb{Z}_2 : \text{Ab} \rightarrow \text{Ab}$ gives

$$0 \longrightarrow 2\mathbb{Z} \otimes \mathbb{Z}_2 \cong \mathbb{Z}_2 \xrightarrow{\cdot 2} \mathbb{Z}_2 \cong \mathbb{Z} \otimes \mathbb{Z}_2 \xrightarrow{\text{pr} \otimes \mathbb{1}} \mathbb{Z}_2 \otimes \mathbb{Z}_2 \rightarrow 0,$$

which is not exact since multiplication by 2 is the trivial map on \mathbb{Z}_2 , hence the map $i \otimes \mathbb{1} : 2\mathbb{Z} \otimes \mathbb{Z}_2 \rightarrow \mathbb{Z} \otimes \mathbb{Z}_2$ fails to be injective. This shows that \mathbb{Z}_2 is not a flat \mathbb{Z} -module and $\otimes \mathbb{Z}_2 : \text{Ab} \rightarrow \text{Ab}$ is not an exact functor.

EXAMPLE 45.9. The isomorphism of \mathbb{Z} -modules $2\mathbb{Z} \rightarrow \mathbb{Z} : m \mapsto m/2$ cannot be extended from the subgroup $2\mathbb{Z} \subset \mathbb{Z}$ to a homomorphism $\mathbb{Z} \rightarrow \mathbb{Z}$. This shows that \mathbb{Z} is not an injective \mathbb{Z} -module, and in particular, feeding the exact sequence (45.4) into $\text{Hom}(\cdot, \mathbb{Z}) : \text{Ab} \rightarrow \text{Ab}$ produces a non-exact sequence

$$0 \longrightarrow \text{Hom}(\mathbb{Z}_2, \mathbb{Z}) \xrightarrow{\text{pr}^*} \text{Hom}(\mathbb{Z}, \mathbb{Z}) \xrightarrow{i^*} \text{Hom}(2\mathbb{Z}, \mathbb{Z}) \longrightarrow 0,$$

with exactness failing at the term $\text{Hom}(2\mathbb{Z}, \mathbb{Z})$ because the map $i^* : \text{Hom}(\mathbb{Z}, \mathbb{Z}) \rightarrow \text{Hom}(2\mathbb{Z}, \mathbb{Z})$ dual to the inclusion $i : 2\mathbb{Z} \hookrightarrow \mathbb{Z}$ is not surjective.

In the previous lecture, we made use of the fact that feeding a short exact sequence $0 \rightarrow A \rightarrow B \rightarrow C \rightarrow 0$ into any additive functor \mathcal{F} does produce an exact sequence $0 \rightarrow \mathcal{F}(A) \rightarrow \mathcal{F}(B) \rightarrow \mathcal{F}(C) \rightarrow 0$ whenever the original sequence splits, which is true for instance if C is a free R -module. It will be useful to note 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 following dualization of Definition 45.6 does something similar in algebraic settings.

DEFINITION 45.10. An R -module G is called **projective** if for every surjective homomorphism $\pi : B \rightarrow A$, every homomorphism $\varphi : G \rightarrow A$ can be lifted to a homomorphism $\tilde{\varphi} : G \rightarrow B$ so that the following diagram commutes:

$$\begin{array}{ccc} & & B \\ & \nearrow \tilde{\varphi} & \downarrow \pi \\ G & \xrightarrow{\varphi} & A \end{array}$$

The following result makes more precise the sense in which the notions of projectivity and injectivity are dual to each other; its proof is an exercise.

PROPOSITION 45.11. *For any R -module G , the covariant functor $\text{Hom}(G, \cdot) : R\text{-Mod} \rightarrow R\text{-Mod}$ is left-exact, and it is exact if and only if G is projective.* \square

EXAMPLE 45.12. Every free R -module is also projective. Indeed, if G has a basis $\mathcal{B} \subset G$, then the required lift $\tilde{\varphi} : G \rightarrow B$ 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 45.13. The abelian group \mathbb{Z}_2 is not a projective \mathbb{Z} -module. For instance, the lift in the diagram

$$\begin{array}{ccc} & & \mathbb{Z} \\ & \nearrow \tilde{\varphi} & \downarrow \text{pr} \\ \mathbb{Z}_2 & \xrightarrow{\mathbb{1}} & \mathbb{Z}_2 \end{array}$$

can never exist since $\text{Hom}(\mathbb{Z}_2, \mathbb{Z}) = 0$.

Now suppose $0 \rightarrow A \xrightarrow{i} B \xrightarrow{j} C \rightarrow 0$ is a short exact sequence of R -modules and, instead of assuming C is free, suppose it is projective. The identity map $C \rightarrow C$ then admits a lift $\varphi : C \rightarrow B$ satisfying $j \circ \varphi = \mathbb{1}$, so φ is a right-inverse of j . We conclude:

PROPOSITION 45.14. *If C is a projective R -module and $0 \rightarrow A \rightarrow B \rightarrow C \rightarrow 0$ is a short exact sequence, then the sequence splits, and the sequence induced by feeding it into any additive functor is therefore also split exact.* \square

Let us now collect some tools for recognizing the flatness of a module.

LEMMA 45.15. *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 $i : A \rightarrow B$, 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 $i \otimes \mathbb{1} : A \otimes G_\alpha \rightarrow B \otimes G_\alpha$ are injective. \square

LEMMA 45.16. *Every free R -module is flat.*

PROOF. Free modules are isomorphic to direct sums of copies of R , so thanks to Lemma 45.15, this follows from the trivial observation that R itself is a flat R -module. \square

LEMMA 45.17. *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 $\pi : F \rightarrow G$, 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

$$0 \rightarrow \ker \pi \hookrightarrow F \xrightarrow{\pi} G \longrightarrow 0,$$

and if G is projective, then Proposition 45.14 gives a splitting of this sequence, and therefore an isomorphism $F \cong G \oplus \ker \pi$. \square

COROLLARY 45.18. *All projective R -modules are flat.* \square

REMARK 45.19. If R is a principal ideal domain, then Lemma 45.17 and Proposition 44.2 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.

The following result lies in the background of the fact stated in Theorem 45.1 that $\text{Tor}(A, B) = 0$ whenever A or B is a torsion-free \mathbb{Z} -module; its proof is Exercise 45.3.

PROPOSITION 45.20. *Every torsion-free abelian group is a flat \mathbb{Z} -module.* \square

45.3. Projective resolutions. By now we have some motivation to believe that there is something special about R -modules that are projective: in particular, Proposition 45.14 and Corollary 45.18 imply that the results of the previous lecture produce a natural isomorphism

$$H_n(C_*) \otimes G \cong H_n(C_* \otimes G)$$

whenever either $H_{n-1}(C_*)$ or G happens to be projective. Unfortunately, plenty of interesting R -modules (such as the abelian group \mathbb{Z}_2) are not projective, and the natural map $h : H_n(C_*) \otimes G \rightarrow H_n(C_* \otimes G)$ will then sometimes fail to be an isomorphism; we still need a better understanding of what what happens 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

$$0 \rightarrow P \rightarrow H_{n-1}(C_*) \rightarrow 0$$

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 $P \rightarrow H_{n-1}(C_*)$, 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 45.21. A **projective resolution** $A_* \xrightarrow{\alpha} A$ of an R -module A consists of an exact sequence

$$\dots \longrightarrow A_2 \xrightarrow{\alpha_2} A_1 \xrightarrow{\alpha_1} A_0 \xrightarrow{\alpha} A \longrightarrow 0$$

such that all of the A_n for $n = 0, 1, 2, \dots$ 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

$$\dots \longrightarrow A_2 \xrightarrow{\alpha_2} A_1 \xrightarrow{\alpha_1} A_0 \longrightarrow A_{-1} := 0 \longrightarrow A_{-2} := 0 \longrightarrow \dots$$

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 $A_* \rightarrow A$, also known as an *augmentation* of the chain complex A_* , for which the resulting augmented chain complex

$$\tilde{A}_* := \bigoplus_{n=-1}^{\infty} \tilde{A}_n \quad \text{with} \quad \tilde{A}_{-1} := A \quad \text{and} \quad \tilde{A}_n := A_n \text{ for } n \geq 0$$

is the original exact sequence in the projective resolution. 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 $\alpha : A_* \rightarrow A$ for which the reduced homology $\tilde{H}_*(A_*) := H_*(\tilde{A}_*)$ is trivial.

PROPOSITION 45.22. *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 $n \geq 2$.*

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 consider the free R -module $A_0 := \bigoplus_{e \in S_0} R$ generated by the set S_0 , and define $\alpha : A_0 \rightarrow A$ 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 $A_1 := \bigoplus_{e \in S_1} R$ and $\alpha_1 : A_1 \rightarrow \ker \alpha$ analogously; this defines $\alpha_1 : A_1 \rightarrow A_0$ such that $\text{im } \alpha_1 = \ker \alpha$. 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 44.2 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 $n \geq 2$. □

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 45.23. *Given an R -module homomorphism $\varphi : A \rightarrow B$ 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 $\varphi_n : A_n \rightarrow B_n$ for $n \geq 0$ which, together with $\varphi : A \rightarrow B$, form a chain map $\varphi_* : \tilde{A}_* \rightarrow \tilde{B}_*$ between the corresponding augmented chain complexes, i.e. the diagram*

$$(45.5) \quad \begin{array}{ccccccc} \dots & \longrightarrow & A_2 & \xrightarrow{\alpha_2} & A_1 & \xrightarrow{\alpha_1} & A_0 & \xrightarrow{\alpha} & A \\ & & \downarrow \varphi_2 & & \downarrow \varphi_1 & & \downarrow \varphi_0 & & \downarrow \varphi \\ \dots & \longrightarrow & B_2 & \xrightarrow{\beta_2} & B_1 & \xrightarrow{\beta_1} & B_0 & \xrightarrow{\beta} & B \end{array}$$

commutes. Moreover, this chain map is unique up to chain homotopy.

COROLLARY 45.24. *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 $\varphi_* : A_* \rightarrow A'_*$ whose restriction $\varphi_0 : A_0 \rightarrow A'_0$ to degree 0 satisfies $\alpha' \circ \varphi_0 = \alpha$.*

PROOF. Proposition 45.23 can be applied with $\varphi : A \rightarrow A$ 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. □

PROOF OF PROPOSITION 45.23. 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 $k \geq 0$ that the maps $\varphi_{-1}, \dots, \varphi_{k-1}$ in (45.5) have already been constructed so that all the relevant squares commute. We must then find a map $\varphi_k : A_k \rightarrow B_k$ 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 $\text{im}(\varphi_{k-1} \alpha_k) \subset \ker \beta_{k-1} = \text{im} \beta_k$, and we can therefore define φ_k to be any solution to the lifting problem

$$\begin{array}{ccc} & & B_k \\ & \nearrow \varphi_k & \downarrow \beta_k \\ A_k & \xrightarrow{\varphi_{k-1} \alpha_k} & \ker \beta_{k-1} \end{array}$$

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 $h_k : A_k \rightarrow B_{k+1}$ for $k \geq 0$ 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 $h_{-1} : A \rightarrow B_0$, which we define as $h_{-1} := 0$. Assume for some $k \geq 0$ that h_{-1}, \dots, h_{k-1} have already been constructed, so we now need to find a map $h_k : A_k \rightarrow B_{k+1}$ 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 $\text{im}(\varphi_k - \psi_k - h_{k-1} \alpha_k) \subset \ker \beta_k = \text{im} \beta_{k+1}$, and h_k can now be defined as any solution to the lifting problem

$$\begin{array}{ccc} & & B_{k+1} \\ & \nearrow h_k & \downarrow \beta_{k+1} \\ A_k & \xrightarrow{\varphi_k - \psi_k - h_{k-1} \alpha_k} & \ker \beta_k \end{array}$$

The result now follows again by induction on k . □

The next result serves as the technical engine behind the long exact sequences in Theorems 45.1 and 45.2.

PROPOSITION 45.25 (horseshoe lemma). *Suppose $0 \rightarrow A \xrightarrow{f} B \xrightarrow{g} C \rightarrow 0$ 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 $n \geq 0$ such that the diagram in Figure 22 commutes, and the rows $0 \rightarrow A_n \rightarrow B_n \rightarrow C_n \rightarrow 0$ for all $n \geq 0$ are split exact sequences.*

PROOF. Let us label the given exact sequence $0 \rightarrow A_{-1} \xrightarrow{f_{-1}} B_{-1} \xrightarrow{g_{-1}} C_{-1} \rightarrow 0$ and, arguing by induction, suppose for some $n \geq 0$ 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




FIGURE 22. The diagram constructed in the horseshoe lemma. The modules B_n and the dashed arrows need to be constructed; everything else is given.

the trivial module. As a preliminary observation, we claim that the short exact sequence at row $n - 1$ restricts to a short exact sequence

$$(45.6) \quad 0 \longrightarrow \ker \alpha_{n-1} \xrightarrow{f_{n-1}} \ker \beta_{n-1} \xrightarrow{g_{n-1}} \ker \gamma_{n-1} \longrightarrow 0.$$

If $n = 0$ this just means that the original sequence $0 \rightarrow A \rightarrow B \rightarrow C \rightarrow 0$ is exact, and for $n \geq 1$, 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 (45.6) 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 $0 \rightarrow A_n \rightarrow B_n \rightarrow C_n \rightarrow 0$ 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. 45.14). 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 $g_n : A_n \oplus C_n \rightarrow C_n$ as the obvious inclusion and projection respectively. Note that by Exercise 45.2, $A_n \oplus C_n$ is projective. The remaining task is thus to find a map

$$A_n \oplus C_n \xrightarrow{\beta_n} B_{n-1}$$

that satisfies

$$(45.7) \quad \text{im } \beta_n = \ker \beta_{n-1}, \quad \gamma_n \circ g_n = g_{n-1} \circ \beta_n, \quad \text{and} \quad \beta_n \circ f_n = f_{n-1} \circ \alpha_n.$$

Denote the natural inclusions into the direct sum by

$$i_A : A_n \hookrightarrow A_n \oplus C_n, \quad i_C : C_n \hookrightarrow A_n \oplus C_n,$$

so we have

$$i_A = f_n \quad \text{and} \quad g_n \circ i_C = \mathbf{1}_{C_n}.$$

Any pair of maps $\beta_A : A_n \rightarrow B_{n-1}$ and $\beta_C : C_n \rightarrow B_{n-1}$ will determine a unique $\beta_n : A_n \oplus C_n \rightarrow B_{n-1}$ via the relations

$$\beta_A := \beta_n \circ i_A : A_n \rightarrow B_{n-1} \quad \beta_C := \beta_n \circ i_C : C_n \rightarrow B_{n-1},$$

and the relation $\beta_{n-1} \circ \beta_n = 0$ then holds trivially if $n = 0$, and for $n \geq 1$, it holds if and only if

$$\beta_{n-1} \circ \beta_A = 0 \quad \text{and} \quad \beta_{n-1} \circ \beta_C = 0.$$

Since $i_A = f_n$, (45.7) 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 $n \geq 1$, 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 (45.7) 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 $C_n \rightarrow \ker \beta_{n-1}$ satisfying

$$(45.8) \quad g_{n-1} \circ \beta_C = \gamma_n \circ g_n \circ i_C = \gamma_n.$$

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 $\beta_C : C_n \rightarrow \ker \beta_{n-1}$ that lifts $\gamma_n : C_n \rightarrow \ker \gamma_{n-1}$ via the surjection $g_{n-1} : \ker \beta_{n-1} \rightarrow \ker \gamma_{n-1}$ in precisely the sense of (45.8), 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 $\text{im } \beta_n = \ker \beta_{n-1}$, 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 row $n-1$ with $0 \rightarrow \ker \alpha_{n-1} \rightarrow \ker \beta_{n-1} \rightarrow \ker \gamma_{n-1} \rightarrow 0$, i.e. the diagram in Figure 23. The resulting long exact sequence of homologies⁶⁸ takes the form

$$\dots \rightarrow 0 \rightarrow 0 \rightarrow \ker \alpha_n \rightarrow \ker \beta_n \rightarrow \ker \gamma_n \rightarrow \frac{\ker \alpha_{n-1}}{\text{im } \alpha_n} \rightarrow \frac{\ker \beta_{n-1}}{\text{im } \beta_n} \rightarrow \frac{\ker \gamma_{n-1}}{\text{im } \gamma_n} \rightarrow 0 \rightarrow 0 \rightarrow \dots,$$

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. \square

45.4. Left and right derived functors. Your first instinct when you see a chain map like $\varphi_* : A_* \rightarrow B_*$ as in Proposition 45.23 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

$$\mathcal{F} : R\text{-Mod} \rightarrow R\text{-Mod}$$

that is covariant and right-exact, or contravariant and left-exact.

⁶⁸What 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.




FIGURE 23. The short exact sequence of chain complexes that implies $\text{im } \beta_n = \ker \beta_{n-1}$.

We consider first the covariant case. 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

$$\dots \longrightarrow \mathcal{F}(A_2) \xrightarrow{\mathcal{F}(\alpha_2)} \mathcal{F}(A_1) \xrightarrow{\mathcal{F}(\alpha_1)} \mathcal{F}(A_0) \longrightarrow 0 \longrightarrow 0 \longrightarrow \dots,$$

along with a homomorphism

$$\mathcal{F}(A_0) \xrightarrow{\mathcal{F}(\alpha)} \mathcal{F}(A).$$

If \mathcal{F} is an exact functor, then by Proposition 44.7, the sequence

$$\dots \longrightarrow \mathcal{F}(A_2) \xrightarrow{\mathcal{F}(\alpha_2)} \mathcal{F}(A_1) \xrightarrow{\mathcal{F}(\alpha_1)} \mathcal{F}(A_0) \xrightarrow{\mathcal{F}(\alpha)} \mathcal{F}(A) \longrightarrow 0$$

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 $\mathcal{F}(\alpha) : \mathcal{F}(A_0) \rightarrow \mathcal{F}(A)$ to be a surjective map that descends to an isomorphism

$$(45.9) \quad H_0(\mathcal{F}(A_*)) = \mathcal{F}(A_0) / \text{im}(\mathcal{F}(\alpha_1)) = \mathcal{F}(A_0) / \ker(\mathcal{F}(\alpha)) \xrightarrow{\cong} \mathcal{F}(A).$$

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 45.23 and Corollary 45.24 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 (45.5) to produce a diagram

$$\begin{array}{ccccccc}
 \dots & \longrightarrow & \mathcal{F}(A_2) & \xrightarrow{\mathcal{F}(\alpha_2)} & \mathcal{F}(A_1) & \xrightarrow{\mathcal{F}(\alpha_1)} & \mathcal{F}(A_0) \xrightarrow{\mathcal{F}(\alpha)} \mathcal{F}(A) \\
 & & \downarrow \mathcal{F}(\varphi_2) & & \downarrow \mathcal{F}(\varphi_1) & & \downarrow \mathcal{F}(\varphi_0) & & \downarrow \mathcal{F}(\varphi) \\
 \dots & \longrightarrow & \mathcal{F}(B_2) & \xrightarrow{\mathcal{F}(\beta_2)} & \mathcal{F}(B_1) & \xrightarrow{\mathcal{F}(\beta_1)} & \mathcal{F}(B_0) \xrightarrow{\mathcal{F}(\beta)} \mathcal{F}(B)
 \end{array}$$

that represents a chain map $\mathcal{F}(\varphi_*) : \mathcal{F}(A_*) \rightarrow \mathcal{F}(B_*)$ associated to any homomorphism $\varphi : A \rightarrow B$, 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 $\varphi : A \rightarrow B$ a sequence of R -module homomorphisms

$$H_n(\mathcal{F}(A_*)) \xrightarrow{\varphi_*} H_n(\mathcal{F}(B_*)) \quad \text{for each } n \geq 0,$$

which are defined independently of the choices of chain maps in Proposition 45.23, and at the degree 0 level, the relation $\varphi \circ \alpha = \beta \circ \varphi_0$ implies that they match $\mathcal{F}(\varphi) : \mathcal{F}(A) \rightarrow \mathcal{F}(B)$ 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

$$\begin{array}{ccc} H_0(\mathcal{F}(A_*)) & \xrightarrow{\cong} & \mathcal{F}(A) \\ \downarrow \varphi_* & & \downarrow \varphi \\ H_0(\mathcal{F}(B_*)) & \xrightarrow{\cong} & \mathcal{F}(B) \end{array}$$

commutes. If we now apply this construction with $\varphi : A \rightarrow B$ as the identity map $A \rightarrow A$ 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

$$H_n(\mathcal{F}(A_*)) \xrightarrow{\cong} H_n(\mathcal{F}(A'_*)) \quad \text{for each } n \geq 0,$$

which for $n = 0$ fit into the diagram

$$\begin{array}{ccc} H_0(\mathcal{F}(A_*)) & \xrightarrow{\cong} & \mathcal{F}(A) \\ \downarrow \cong & & \parallel \\ H_0(\mathcal{F}(A'_*)) & \xrightarrow{\cong} & \mathcal{F}(A) \end{array}$$

This discussion justifies the following definition.

DEFINITION 45.26. Fix a choice of projective resolution $A_* \xrightarrow{\alpha} A$ for each R -module A . Given a covariant right-exact functor $\mathcal{F} : R\text{-Mod} \rightarrow R\text{-Mod}$, the associated **left derived functors**

$$L_n \mathcal{F} : R\text{-Mod} \rightarrow R\text{-Mod}$$

are defined for each integer $n \geq 0$ by associating to each R -module A the R -module

$$L_n \mathcal{F}(A) := H_n(\mathcal{F}(A_*)),$$

and to each homomorphism $\varphi : A \rightarrow B$ 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 $A_* \rightarrow B_*$ provided by Proposition 45.23.

REMARK 45.27. You would be within your rights to find something slightly unsatisfactory about the way Definition 45.26 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 $L_n \mathcal{F} : R\text{-Mod} \rightarrow R\text{-Mod}$ rather far from unique or canonical. On the other hand, Corollary 45.24 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 $n \geq 0$, 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 45.28. For any R -module G and each integer $n \geq 0$, the functor

$$\mathrm{Tor}_n(\cdot, G) : R\text{-Mod} \rightarrow R\text{-Mod} : A \mapsto \mathrm{Tor}_n(A, G)$$

is defined as the left derived functor $\mathrm{Tor}_n(\cdot, G) := L_n\mathcal{F}$ associated to the right-exact functor $\mathcal{F} := \otimes G$. More explicitly,

$$\mathrm{Tor}_n(A, G) := H_n(A_* \otimes G)$$

for any R -module A with a choice of projective resolution $A_* \xrightarrow{\alpha} A$.

Aside from some details involving the functoriality of $\mathrm{Tor}_n(A, G)$ with respect to G , the following result about left derived functors implies most of the properties of Tor_n listed in Theorem 45.1:

THEOREM 45.29. For any covariant right-exact functor $\mathcal{F} : R\text{-Mod} \rightarrow R\text{-Mod}$, the left derived functors $L_n\mathcal{F}$ have the following properties:

- (1) Every short exact sequence $0 \rightarrow A \xrightarrow{f} B \xrightarrow{g} C \rightarrow 0$ gives rise to a long exact sequence
- $$\begin{aligned} \dots \rightarrow L_2\mathcal{F}(A) \xrightarrow{L_2\mathcal{F}(f)} L_2\mathcal{F}(B) \xrightarrow{L_2\mathcal{F}(g)} L_2\mathcal{F}(C) \rightarrow L_1\mathcal{F}(A) \xrightarrow{L_1\mathcal{F}(f)} L_1\mathcal{F}(B) \xrightarrow{L_1\mathcal{F}(g)} L_1\mathcal{F}(C) \\ \rightarrow \mathcal{F}(A) \xrightarrow{\mathcal{F}(f)} \mathcal{F}(B) \xrightarrow{\mathcal{F}(g)} \mathcal{F}(C) \rightarrow 0. \end{aligned}$$

- (2) The functors $L_n\mathcal{F}$ are additive; in particular, there are natural isomorphisms

$$L_n\mathcal{F}(A \oplus B) \cong L_n\mathcal{F}(A) \oplus L_n\mathcal{F}(B)$$

for all pairs of R -modules A, B .

- (3) $L_n\mathcal{F}(A) = 0$ for all $n \geq 1$ whenever A is projective or \mathcal{F} is exact.
 (4) $L_n\mathcal{F}(A) = 0$ for all $n \geq 2$ and all A whenever R is a principal ideal domain.
 (5) There are natural isomorphisms $L_0\mathcal{F}(A) \cong \mathcal{F}(A)$ for all R -modules A .
 (6) If $k \in \mathbb{N}$ has the property that no nonzero element $x \in R$ satisfies $kx = 0$, then one has

$$L_1\mathcal{F}(R/kR) \cong \ker \left(\mathcal{F}(R) \xrightarrow{\cdot k} \mathcal{F}(R) \right), \quad \text{and} \quad L_n\mathcal{F}(R/kR) = 0 \text{ for } n \geq 2.$$

PROOF. The natural isomorphisms $L_0\mathcal{F}(A) \cong \mathcal{F}(A)$ asserted in property (5) were already explained in the discussion surrounding (45.9).

For property (1), given a short exact sequence $0 \rightarrow A \xrightarrow{f} B \xrightarrow{g} C \rightarrow 0$, the horseshoe lemma (Proposition 45.25) allows us to choose suitable projective resolutions so that we get a corresponding short exact sequence of chain complexes

$$0 \rightarrow A_* \xrightarrow{f_*} B_* \xrightarrow{g_*} C_* \rightarrow 0.$$

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

$$0 \rightarrow \mathcal{F}(A_*) \xrightarrow{\mathcal{F}(f_*)} \mathcal{F}(B_*) \xrightarrow{\mathcal{F}(g_*)} \mathcal{F}(C_*) \rightarrow 0.$$

Passing to homology via Proposition 32.13 and replacing $L_0\mathcal{F}$ with \mathcal{F} via the natural isomorphisms now gives a long exact sequence of derived functors as claimed.

Property (2) requires proving that for any two R -module homomorphisms $f, g : A \rightarrow B$ and each $n \geq 0$, the homomorphisms $L_n\mathcal{F}(A) \rightarrow L_n\mathcal{F}(B)$ given by $L_n\mathcal{F}(f + g)$ and $L_n\mathcal{F}(f) + L_n\mathcal{F}(g)$ are identical. In fact, if $f_*, g_* : A_* \rightarrow B_*$ denote the corresponding chain maps as provided by Proposition 45.23, one can take $f_* + g_*$ as the chain map associated to $f + g : A \rightarrow B$, from which the result follows.

For property (3), note first that if A is projective, then it admits a projective resolution of the form $\dots \rightarrow 0 \rightarrow 0 \rightarrow A \xrightarrow{\mathbb{1}} A \rightarrow 0$, in which $A_n = 0$ for all $n \geq 1$, hence $L_n \mathcal{F}(A) = H_n(\mathcal{F}(A_n))$ also vanishes. If on the other hand \mathcal{F} is an exact functor, then $L_n \mathcal{F}(A) = H_n(\mathcal{F}(A_n))$ vanishes for all $n \geq 1$ because for any projective resolution $A_* \rightarrow A$, the induced sequence $\dots \rightarrow \mathcal{F}(A_2) \rightarrow \mathcal{F}(A_1) \rightarrow \mathcal{F}(A_0)$ is exact.

Property (4) follows by choosing a projective resolution $A_* \rightarrow A$ with $A_n = 0$ for all $n \geq 2$, as provided by Proposition 45.22.

Property (6) is an interesting exercise; specifically, it is Exercise 45.4. \square

Next, suppose that $\mathcal{F} : R\text{-Mod} \rightarrow R\text{-Mod}$ is contravariant and left-exact. Starting from a choice of projective resolution $A_* \xrightarrow{\alpha} A$ for each R -module A , one feeds the chain complex A_* into \mathcal{F} and obtains a cochain complex $\mathcal{F}(A_*)$, whose cohomologies define the **right derived functors**

$$R_n \mathcal{F}(A) := H_n(\mathcal{F}(A_*)), \quad n \geq 0.$$

These are contravariant functors due to Proposition 45.23, which associates to any homomorphism $\varphi : A \rightarrow B$ a chain map $\varphi_* : A_* \rightarrow B_*$, giving rise to a chain map $\mathcal{F}(\varphi_*) : \mathcal{F}(B_*) \rightarrow \mathcal{F}(A_*)$ that induces maps $R_n \mathcal{F}(B) \rightarrow R_n \mathcal{F}(A)$ for each $n \geq 0$. The chain map $\varphi_* : A_* \rightarrow B_*$ is of course not unique, but any other such map $\psi_* : A_* \rightarrow B_*$ is related to it by a chain homotopy $h_* : A_* \rightarrow B_{*+1}$, which can similarly be fed into \mathcal{F} to produce a chain homotopy $\mathcal{F}(h_*) : \mathcal{F}(B_*) \rightarrow \mathcal{F}(A_{*-1})$ between $\mathcal{F}(\varphi_*)$ and $\mathcal{F}(\psi_*)$, showing that the induced map $R_n \mathcal{F}(B) \rightarrow R_n \mathcal{F}(A)$ is independent of choices. Applying the same argument to $\mathbb{1} : A \rightarrow A$ 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 cohomology of that complex, and we therefore have a natural isomorphism

$$\mathcal{F}(A) \xrightarrow[\cong]{\mathcal{F}(\alpha)} R_0 \mathcal{F}(A).$$

The contravariant analogue of Theorem 45.29 is proved by variations on the same arguments, which we shall leave as an exercise:

THEOREM 45.30. *For any contravariant left-exact functor $\mathcal{F} : R\text{-Mod} \rightarrow R\text{-Mod}$, the right derived functors $R_n \mathcal{F} : R\text{-Mod} \rightarrow R\text{-Mod}$ have the following properties:*

(1) *Every short exact sequence $0 \rightarrow A \xrightarrow{f} B \xrightarrow{g} C \rightarrow 0$ gives rise to a long exact sequence*

$$\begin{aligned} 0 \longrightarrow \mathcal{F}(C) \xrightarrow{\mathcal{F}(g)} \mathcal{F}(B) \xrightarrow{\mathcal{F}(f)} \mathcal{F}(A) \longrightarrow R_1 \mathcal{F}(C) \xrightarrow{R_1 \mathcal{F}(g)} R_1 \mathcal{F}(B) \xrightarrow{R_1 \mathcal{F}(f)} R_1 \mathcal{F}(A) \\ \longrightarrow R_2 \mathcal{F}(C) \xrightarrow{R_2 \mathcal{F}(g)} R_2 \mathcal{F}(B) \xrightarrow{R_2 \mathcal{F}(f)} R_2 \mathcal{F}(A) \longrightarrow \dots \end{aligned}$$

(2) *The functors $R_n \mathcal{F}$ are additive; in particular, there are natural isomorphisms*

$$R_n \mathcal{F}(A \oplus B) \cong R_n \mathcal{F}(A) \oplus R_n \mathcal{F}(B)$$

for all pairs of R -modules A, B .

(3) *$R_n \mathcal{F}(A) = 0$ for all $n \geq 1$ whenever A is projective or \mathcal{F} is exact.*

(4) *$R_n \mathcal{F}(A) = 0$ for all $n \geq 2$ and all A whenever R is a principal ideal domain.*

(5) *There are natural isomorphisms $R_0 \mathcal{F}(A) \cong \mathcal{F}(A)$ for all R -modules A .*

(6) If $k \in \mathbb{N}$ has the property that no nonzero element $x \in R$ satisfies $kx = 0$, then one has

$$R_1\mathcal{F}(R/kR) \cong \operatorname{coker}\left(\mathcal{F}(R) \xrightarrow{\cdot k} \mathcal{F}(R)\right), \quad \text{and} \quad R_n\mathcal{F}(R/kR) = 0 \text{ for } n \geq 2.$$

□

45.5. Symmetry of Tor. We have not yet discussed in what sense $\operatorname{Tor}_n(A, \cdot) : R\text{-Mod} \rightarrow R\text{-Mod}$ is also a functor for each $n \geq 0$, 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 $A \otimes : R\text{-Mod} \rightarrow R\text{-Mod}$ 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

$$C_* \otimes : R\text{-Mod} \rightarrow \operatorname{Ch}(R\text{-Mod})$$

that sends any R -module G to the chain complex $C_* \otimes G$ and any homomorphism $\varphi : G \rightarrow H$ to a chain map of the form

$$\mathbb{1} \otimes \varphi : C_* \otimes G \rightarrow C_* \otimes H,$$

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 $\varphi : G \rightarrow H$, and one easily checks that for $n = 0$, the natural isomorphism identifies φ_* with $\mathbb{1} \otimes \varphi : A \otimes G \rightarrow A \otimes H$. We leave it as an exercise to show that the functors $\operatorname{Tor}_n(A, \cdot) : R\text{-Mod} \rightarrow R\text{-Mod}$ are also additive for every $n \geq 0$ and $A \in R\text{-Mod}$.

With that understood, one can now show that any short exact sequence of “coefficient” modules

$$0 \longrightarrow G \xrightarrow{\varphi} H \xrightarrow{\psi} K \longrightarrow 0$$

naturally gives rise to a long exact sequence

$$(45.10) \quad \begin{aligned} \dots \rightarrow \operatorname{Tor}_2(A, G) \rightarrow \operatorname{Tor}_2(A, H) \rightarrow \operatorname{Tor}_2(A, K) \rightarrow \operatorname{Tor}_1(A, G) \rightarrow \operatorname{Tor}_1(A, H) \rightarrow \operatorname{Tor}_1(A, K) \\ \rightarrow A \otimes G \rightarrow A \otimes H \rightarrow A \otimes K \rightarrow 0, \end{aligned}$$

much as in the first property listed in Theorem 45.1. The proof does not require the horseshoe lemma; see Exercise 45.6.

The last detail of Theorem 45.1 remaining to be addressed is the symmetry isomorphism

$$\operatorname{Tor}_n(A, B) \cong \operatorname{Tor}_n(B, A).$$

This symmetry is the reason why any criterion on B that implies the vanishing of $\operatorname{Tor}_n(A, B)$ can be applied equally well to A or vice versa, e.g. according to Proposition 45.20, B is a flat \mathbb{Z} -module whenever it has no torsion, hence $\otimes B$ is an exact functor, and it follows via symmetry that $\operatorname{Tor}_n(A, B)$ also vanishes whenever A is torsion free.

The isomorphism $\operatorname{Tor}_n(A, B) \cong \operatorname{Tor}_n(B, A)$ should not be surprising since there is already a natural isomorphism $A \otimes B \cong B \otimes A$ for every pair of R -modules, which is the special case $n = 0$. Here is a nice way to see the general case.

After choosing projective resolutions $A_* \xrightarrow{\alpha} A$ and $B_* \xrightarrow{\beta} B$, we can form the commutative diagram in Figure 24, 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 $m, n \geq 0$, Corollary 45.18 implies that the functors $A_m \otimes$ and $\otimes B_n$ are exact, which makes the m th row and n th column for

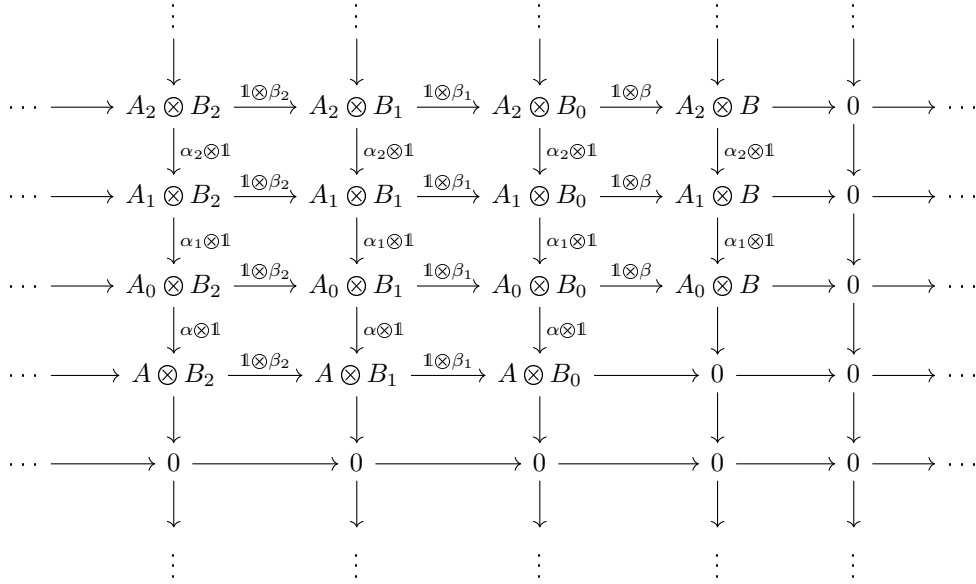


FIGURE 24. The double complex that implies $\text{Tor}_n(A, B) \cong \text{Tor}_n(B, A)$.

each $m, n \geq 0$ 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

$$H_m(A_* \otimes B) = \text{Tor}_m(A, B),$$

and similarly the complex $A \otimes B_*$ in row -1 , whose homology we can identify with

$$H_n(A \otimes B_*) \cong \text{Tor}_n(B, A)$$

using the natural isomorphism $A \otimes B_* \cong B_* \otimes A$. The rest is a diagram-chasing exercise:

PROPOSITION 45.31. *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 25, with the property that all other rows and columns are exact, and $C_j^i = 0$ whenever $i < -1$ or $j < -1$. Then there is a natural isomorphism*

$$H_n(C^*, \partial) \cong H_n(C_*, d)$$

for every $n \in \mathbb{Z}$. □

Applying Proposition 45.31 to the specific double complex in Figure 24 yields the promised isomorphisms $\text{Tor}_n(A, B) \cong \text{Tor}_n(B, A)$. The proof of Theorem 45.1 is now complete.

45.6. The other definition of Ext. We will not need the following detail in subsequent developments, but it would seem criminal not to mention the property of the functors $\text{Ext}_n : R\text{-Mod} \times R\text{-Mod} \rightarrow R\text{-Mod}$ that is analogous to the symmetry of Tor_n .

Since Ext_n is contravariant in one of its variables but covariant in the other, symmetry is out of the question. The key instead is to fix an R -module G and develop derived functors for the covariant functor $\text{Hom}(G, \cdot)$.

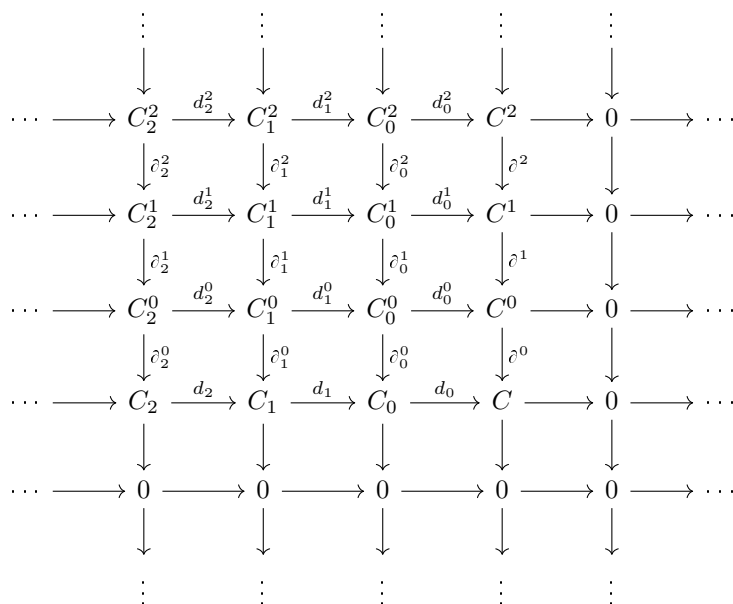


FIGURE 25. The abstract double complex in Exercise 45.8.

According to Proposition 45.11, $\text{Hom}(G, \cdot)$ is left-exact, and it is exact if and only if G is projective. The theory of left derived functors developed so far 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 $n \geq 0$ are injective. It is pretty easy to think of modules that are *not* injective, e.g. we saw in Example 45.9 that \mathbb{Z} is not an injective \mathbb{Z} -module, since the isomorphism $2\mathbb{Z} \rightarrow \mathbb{Z} : m \mapsto m/2$ cannot be extended from the subgroup $2\mathbb{Z} \subset \mathbb{Z}$ to a homomorphism $\mathbb{Z} \rightarrow \mathbb{Z}$. 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 divisible will fail to be injective due to examples such as $\varphi : k\mathbb{Z} \rightarrow G$ for $k \geq 2$ 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 45.32. *Every R -module is isomorphic to a submodule of one that is injective.* □

If you believe Lemma 45.32, then you will easily convince yourself that every R -module admits an injective resolution: the lemma can be used to define $\alpha : A \rightarrow A^0$ as the inclusion of A into

a larger injective module A^0 , whose quotient by $\text{im}(\alpha) \subset A^0$ 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 $n \geq 0$ by

$$R_n\mathcal{F}(A) := H^n(\mathcal{F}(A^*)).$$

It is straightforward to prove an injective analogue of Proposition 45.23, so that any homomorphism $\varphi : A \rightarrow B$ gives rise to a unique chain homotopy class of chain maps $\varphi_* : A^* \rightarrow B^*$ between the corresponding injective resolutions. Feeding such a chain map into \mathcal{F} gives a chain map $\mathcal{F}(A^*) \rightarrow \mathcal{F}(B^*)$, thus inducing a natural homomorphism $R_n\mathcal{F}(A) \rightarrow R_n\mathcal{F}(B)$ 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 45.25), turning any short exact sequence $0 \rightarrow A \rightarrow B \rightarrow C \rightarrow 0$ into a short exact sequence of cochain complexes $0 \rightarrow \mathcal{F}(A^*) \rightarrow \mathcal{F}(B^*) \rightarrow \mathcal{F}(C^*) \rightarrow 0$; the exactness of the latter follows from the observation that a short exact sequence $0 \rightarrow A^n \rightarrow B^n \rightarrow C^n \rightarrow 0$ splits whenever A^n is injective, because injectivity guarantees the existence of a left-inverse for the injective map $A^n \rightarrow B^n$. The result is that for a left-exact covariant functor \mathcal{F} , any short exact sequence $0 \rightarrow A \rightarrow B \rightarrow C \rightarrow 0$ gives rise to a long exact sequence

$$\begin{aligned} 0 \rightarrow \mathcal{F}(A) \rightarrow \mathcal{F}(B) \rightarrow \mathcal{F}(C) \rightarrow R_1\mathcal{F}(A) \rightarrow R_1\mathcal{F}(B) \rightarrow R_1\mathcal{F}(C) \\ \rightarrow R_2\mathcal{F}(A) \rightarrow R_2\mathcal{F}(B) \rightarrow R_2\mathcal{F}(C) \rightarrow \dots \end{aligned}$$

The functors $R_n\mathcal{F}$ are easily shown to be additive, and the left-exactness of \mathcal{F} gives rise to a natural isomorphism

$$\mathcal{F}(A) \cong R_0\mathcal{F}(A).$$

Moreover, $R_n\mathcal{F}(A) = 0$ for every $n \geq 1$ whenever \mathcal{F} is exact or A is injective.

Since $\text{Hom}(A, \cdot)$ is left-exact for any A , we can apply this machinery to define a second variant of the sequence of Ext functors, which we shall denote for now by

$$\text{Ext}^n(A, \cdot) := R_n(\text{Hom}(A, \cdot)),$$

so explicitly,

$$\text{Ext}^n(A, B) = H^n(\text{Hom}(A, B^*))$$

for the cochain complex $\text{Hom}(A, B^*)$ arising from any injective resolution $B \xrightarrow{\beta} B^*$ of B . We can then observe that Ext^n has several properties matching those of Ext_n , though seemingly for different reasons. Indeed,

$$\text{Ext}^n(A, B) = 0 \quad \text{for every } n \geq 1 \text{ if } A \text{ is projective or } B \text{ is injective,}$$

the reason being that $\text{Hom}(A, \cdot)$ is exact if A is projective, and $0 \rightarrow B \xrightarrow{1} B \rightarrow 0 \rightarrow \dots$ is an injective resolution if B is injective. Similarly, a short exact sequence $0 \rightarrow G \rightarrow H \rightarrow K \rightarrow 0$ produces a long exact sequence

$$\begin{aligned} 0 \rightarrow \text{Hom}(A, G) \rightarrow \text{Hom}(A, H) \rightarrow \text{Hom}(A, K) \rightarrow \text{Ext}^1(A, G) \rightarrow \text{Ext}^1(A, H) \\ \rightarrow \text{Ext}^1(A, K) \rightarrow \text{Ext}^2(A, G) \rightarrow \text{Ext}^2(A, H) \rightarrow \text{Ext}^2(A, K) \rightarrow \dots \end{aligned}$$

due to the general properties of right derived functors, and a similar sequence for the functors Ext_n can be obtained via more direct arguments (with no need of the horseshoe lemma), using the covariant functoriality of Ext_n in the second variable; see Exercise 45.6. The analogous result




FIGURE 26. The double complex that implies $\text{Ext}_n(A, B) \cong \text{Ext}^n(A, B)$.

for Ext^n shows that it is likewise a contravariant functor in its first variable, and that short exact sequences $0 \rightarrow A \rightarrow B \rightarrow C \rightarrow 0$ give rise to long exact sequences

$$\begin{aligned}
 0 \rightarrow \text{Hom}(C, G) \rightarrow \text{Hom}(B, G) \rightarrow \text{Hom}(A, G) \rightarrow \text{Ext}^1(C, G) \rightarrow \text{Ext}^1(B, G) \\
 \rightarrow \text{Ext}^1(A, G) \rightarrow \text{Ext}^2(C, G) \rightarrow \text{Ext}^2(B, G) \rightarrow \text{Ext}^2(A, G) \rightarrow \dots,
 \end{aligned}$$

thus matching the form of the sequence for Ext_n obtained from the general theory of contravariant right derived functors. All this provides strong evidence that Ext_n and Ext^n are, secretly, the same thing.

We will leave the details as an exercise, but an explicit (and natural) isomorphism $\text{Ext}_n(A, B) \cong \text{Ext}^n(A, B)$ can be obtained from the double complex in Figure 26. 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 n th cohomology of the cochain complex $\text{Hom}(A_*, B)$ in the leftmost nontrivial column is $\text{Ext}_n(A, B)$, the n th cohomology of the complex $\text{Hom}(A, B^*)$ in the lowest nontrivial row is $\text{Ext}^n(A, B)$, 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 $\text{Hom}(A_n, \cdot)$ and $\text{Hom}(\cdot, B_n)$ exact functors. The rest, as they say, is diagram chasing (cf. Proposition 45.31).

45.7. The universal coefficient theorems. We now have more than enough machinery in place to prove the general versions of the universal coefficient theorems.

The setting is again as follows: R is a principal ideal domain, C_* is a chain complex of free R -modules, G is a fixed R -module, and \mathcal{F} is either the covariant functor $\otimes G$ or the contravariant functor $\text{Hom}(\cdot, G)$. In contrast with the previous lecture, we will now avoid assuming that \mathcal{F} is an exact functor.

In the covariant case $\mathcal{F} = \otimes G$, the arguments in the previous lecture produced a short exact sequence

$$(45.11) \quad 0 \longrightarrow \text{coker } \mathcal{F}(i_n) \xrightarrow{\mathcal{F}(j)_*} H_n(C_* \otimes G) \xrightarrow{\mathcal{F}(\partial')_*} \ker \mathcal{F}(i_{n-1}) \longrightarrow 0,$$

where $i_n : B_n \hookrightarrow Z_n$ is the inclusion and ∂' denotes the chain map $C_* \rightarrow B_{*-1}$ defined via the boundary operator on C_* , which induces a chain map $\mathcal{F}(\partial') : C_* \otimes G \rightarrow B_{*-1} \otimes G$ and then

descends to homology as a homomorphism $\mathcal{F}(\partial')_* : H_n(C_* \otimes G) \rightarrow B_{n-1} \otimes G$ whose image is the kernel of

$$\mathcal{F}(i_{n-1}) = i_{n-1} \otimes \mathbb{1} : B_{n-1} \otimes G \rightarrow Z_{n-1} \otimes G.$$

The step where exactness was used was when the short exact sequence

$$(45.12) \quad 0 \longrightarrow B_n \xrightarrow{i_n} Z_n \xrightarrow{q_n} H_n(C_*) \longrightarrow 0$$

was fed into the functor \mathcal{F} . Without exactness, the resulting chain complex will not generally be an exact sequence, but for $\mathcal{F} = \otimes G$, we can instead extract from Theorem 45.1 a long (but not very long) exact sequence

$$0 \longrightarrow \text{Tor}(H_n(C_*), G) \longrightarrow B_n \otimes G \xrightarrow{i_n \otimes \mathbb{1}} Z_n \otimes G \xrightarrow{q_n \otimes \mathbb{1}} H_n(C_*) \otimes G \longrightarrow 0,$$

in which the initial term is actually $\text{Tor}(Z_n, G)$, which vanishes because Z_n is a free R -module. This sequence is only slightly more complicated than what we had in the exact case: we can still conclude from it that $\mathcal{F}(q_n) = q_n \otimes \mathbb{1}$ is surjective and descends to an isomorphism

$$\text{coker}(i_n \otimes \mathbb{1}) \xrightarrow[\cong]{q_n \otimes \mathbb{1}} H_n(C_*) \otimes G,$$

but the major difference is that while the kernel of $\mathcal{F}(i_n) = i_n \otimes \mathbb{1}$ was previously trivial, we now instead have an isomorphism

$$\text{Tor}(H_n(C_*), G) \xrightarrow{\cong} \ker(i_n \otimes \mathbb{1}).$$

Making these substitutions in (45.11) gives a short exact sequence

$$0 \longrightarrow H_n(C_*) \otimes G \xrightarrow{h} H_n(C_* \otimes G) \longrightarrow \text{Tor}(H_{n-1}(C_*), G) \longrightarrow 0,$$

in which the map h admits the same characterization as in (44.11) and is thus the usual canonical homomorphism $[c] \otimes g \mapsto [c \otimes g]$.

One further detail should be addressed before we state a theorem: the sequence obtained above splits. One sees this by choosing left-inverses $p_n : C_n \rightarrow Z_n$ 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 $h : \mathcal{F}(H_n(C_*)) \rightarrow H_n(\mathcal{F}(C_*))$. Indeed, the existence of left-inverses $p_n : C_n \rightarrow Z_n$ is guaranteed because the exact sequence $0 \rightarrow Z_n \hookrightarrow C_n \rightarrow B_{n-1} \rightarrow 0$ 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 $j : Z_* \rightarrow C_*$, denoted by $p : C_* \rightarrow Z_*$. Unfortunately, p is not typically a chain map, for fairly obvious reasons: the boundary operator on Z_* is trivial, but $p_{n-1}\partial_n : C_n \rightarrow Z_{n-1}$ 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 $q_n : Z_n \rightarrow H_n(C_*)$ 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 $q_n \circ p_n : C_n \rightarrow H_n(C_*)$. Feeding this into \mathcal{F} and regarding the \mathbb{Z} -graded module $\mathcal{F}(H_*(C_*)) = \bigoplus_{n \in \mathbb{Z}} \mathcal{F}(H_n(C_*))$ 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

$$\mathcal{F}(qp)_* : H_n(\mathcal{F}(C_*)) \rightarrow \mathcal{F}(H_n(C_*)) : [y] \mapsto \mathcal{F}(q_n p_n)y.$$

We claim that this map is a left-inverse of $h : \mathcal{F}(H_n(C_*)) \rightarrow H_n(\mathcal{F}(C_*))$. 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 (44.11), we find

$$\mathcal{F}(qp)_*h(x) = \mathcal{F}(qp)_*([\mathcal{F}(j_n)z]) = \mathcal{F}(q_np_n)\mathcal{F}(j_n)z = \mathcal{F}(q_n)\mathcal{F}(p_nj_n)z = \mathcal{F}(q_n)z = x,$$

due to the fact that p_n is a left-inverse of j_n .

I will leave it as an exercise to verify naturality in the following statement, but otherwise, we've proved:

THEOREM 45.33 (universal coefficient theorem for homology). *For any chain complex C_* of free modules over a principal ideal domain R , any fixed R -module G and any $n \in \mathbb{Z}$, there exists a split exact sequence*

$$0 \longrightarrow H_n(C_*) \otimes G \xrightarrow{h} H_n(C_* \otimes G) \longrightarrow \text{Tor}(H_{n-1}(C_*), G) \longrightarrow 0,$$

where h is the natural map $[c] \otimes g \mapsto [c \otimes g]$. Moreover, the sequence (but not its splitting) is natural, in the sense that for any chain map $\Phi : A_* \rightarrow B_*$ between two chain complexes of free R -modules, the diagram

$$\begin{array}{ccccccc} 0 & \longrightarrow & H_n(A_*) \otimes G & \xrightarrow{h} & H_n(A_* \otimes G) & \longrightarrow & \text{Tor}(H_{n-1}(A_*), G) \longrightarrow 0 \\ & & \downarrow \Phi_* \otimes \mathbb{1} & & \downarrow (\Phi \otimes \mathbb{1})_* & & \downarrow \\ 0 & \longrightarrow & H_n(B_*) \otimes G & \xrightarrow{h} & H_n(B_* \otimes G) & \longrightarrow & \text{Tor}(H_{n-1}(B_*), G) \longrightarrow 0 \end{array}$$

commutes, where $\text{Tor}(H_{n-1}(A_*), G) \rightarrow \text{Tor}(H_{n-1}(B_*), G)$ is the map induced by $\Phi_* : H_{n-1}(A_*) \rightarrow H_{n-1}(B_*)$ via the functoriality of Tor. \square

The cohomological variant of this result follows by analogous arguments, most of which already appeared in the previous lecture: for the left-exact functor $\mathcal{F} := \text{Hom}(\cdot, G)$, we obtain a short exact sequence

$$(45.13) \quad 0 \longrightarrow \text{coker } \mathcal{F}(i_{n-1}) \longrightarrow H^n(\mathcal{F}(C_*)) \longrightarrow \text{ker } \mathcal{F}(i_n) \longrightarrow 0,$$

while feeding (45.12) into \mathcal{F} gives the not-very-long exact sequence

$$0 \longrightarrow \text{Hom}(H_n(C_*), G) \xrightarrow{q_n^*} \text{Hom}(Z_n, G) \xrightarrow{i_n^*} \text{Hom}(B_n, G) \longrightarrow \text{Ext}(H_n(C_*), G) \longrightarrow 0,$$

in which the rightmost term is $\text{Ext}(Z_n, G)$, again vanishing because Z_n is free. This sequence identifies the kernel of $\mathcal{F}(i_n) = i_n^*$ with $\text{Hom}(H_n(C_*), G)$ and the cokernel of $\mathcal{F}(i_{n-1}) = i_{n-1}^*$ with $\text{Ext}(H_{n-1}(C_*), G)$, so that (45.13) becomes

$$0 \longrightarrow \text{Ext}(H_{n-1}(C_*), G) \longrightarrow H^n(C_*; G) \xrightarrow{h} \text{Hom}(H_n(C_*), G) \longrightarrow 0.$$

To see that this sequence splits, we can again make use of a left-inverse $p : C_* \rightarrow Z_*$ of the inclusion chain map $j : Z_* \rightarrow C_*$. Feeding $qp : C_* \rightarrow H_*(C_*)$ into \mathcal{F} produces a chain map

$$\mathcal{F}(qp) : \mathcal{F}(H_*(C_*)) \rightarrow \mathcal{F}(C_*)$$

between two cochain complexes, where $\mathcal{F}(H_*(C_*)) = \bigoplus_{n \in \mathbb{Z}} \mathcal{F}(H_n(C_*))$ is regarded as a cochain complex with trivial coboundary operator. The induced map on cohomology in degree n is then a homomorphism

$$\mathcal{F}(qp)_* : \mathcal{F}(H_n(C_*)) \rightarrow H^n(\mathcal{F}(C_*)),$$

which can be verified to be a right-inverse of the surjective map $h : H^n(\mathcal{F}(C_*)) \rightarrow \mathcal{F}(H_n(C_*))$.

We conclude:

THEOREM 45.34 (universal coefficient theorem for cohomology). *For any chain complex C_* of free modules over a principal ideal domain R , any fixed R -module G and any $n \in \mathbb{Z}$, there exists a split exact sequence*

$$0 \longrightarrow \text{Ext}(H_{n-1}(C_*), G) \longrightarrow H^n(C_*; G) \xrightarrow{h} \text{Hom}(H_n(C_*), G) \longrightarrow 0,$$

where h is the natural map $[\alpha] \mapsto \langle [\alpha], \cdot \rangle$. Moreover, the sequence (but not its splitting) is natural, in the sense that for any chain map $\Phi : A_* \rightarrow B_*$ between two chain complexes of free R -modules, the diagram

$$\begin{array}{ccccccc} 0 & \longrightarrow & \text{Ext}(H_{n-1}(A_*), G) & \longrightarrow & H^n(A_*; G) & \xrightarrow{h} & \text{Hom}(H_n(A_*), G) \longrightarrow 0 \\ & & \uparrow \Phi^* & & \uparrow \Phi^* & & \uparrow \\ 0 & \longrightarrow & \text{Ext}(H_{n-1}(B_*), G) & \longrightarrow & H^n(B_*; G) & \longrightarrow & \text{Hom}(H_n(B_*), G) \longrightarrow 0 \end{array}$$

commutes, where $\text{Ext}(H_{n-1}(B_*), G) \rightarrow \text{Ext}(H_{n-1}(A_*), G)$ is the map induced by $\Phi_* : H_{n-1}(A_*) \rightarrow H_{n-1}(B_*)$ via the contravariant functoriality of Ext in its first variable. \square

In both of Theorems 45.33 and 45.34, one of the most useful features is the splitting, which gives us isomorphisms

$$\begin{aligned} H_n(C_* \otimes G) &\cong (H_n(C_*) \otimes G) \oplus \text{Tor}(H_{n-1}(C_*), G), \\ H^n(C_*; G) &\cong \text{Hom}(H_n(C_*), G) \oplus \text{Ext}(H_{n-1}(C_*), G), \end{aligned}$$

or e.g. for the singular (co)homology of a space X ,

$$\begin{aligned} H_n(X; G) &\cong (H_n(X; R) \otimes G) \oplus \text{Tor}(H_{n-1}(X; R), G), \\ H^n(X; G) &\cong \text{Hom}(H_n(X; R), G) \oplus \text{Ext}(H_{n-1}(X; R), G). \end{aligned}$$

It should be emphasized however that these splittings are not *natural*: they depend on arbitrary choices and thus cannot be expected to be preserved by the homomorphisms induced by maps between different spaces.

45.8. Applications. We briefly sketch here a few useful applications of the universal coefficient theorems, leaving the proofs as exercises.

Recall that for any path-connected space X , there is a natural map $\pi_1(X) \rightarrow H_1(X; \mathbb{Z})$ that descends to the abelianization of $\pi_1(X)$ as an isomorphism; moreover, Theorem 41.6 then gives a natural isomorphism $H^1(X; G) \cong \text{Hom}(\pi_1(X), G)$ for any coefficient group G . Since G is necessarily abelian, every homomorphism $\pi_1(X) \rightarrow G$ descends to the abelianization of $\pi_1(X)$, thus identifying $\text{Hom}(\pi_1(X), G)$ with $\text{Hom}(H_1(X; \mathbb{Z}), G)$. But in fact, the isomorphism

$$H^1(X; G) \cong \text{Hom}(H_1(X; \mathbb{Z}), G)$$

can be obtained without mentioning $\pi_1(X)$, and without assuming X is path-connected; it follows directly from the universal coefficient theorem for cohomology, according to Exercise 45.9.

Another application concerns the definition of the Betti numbers $b_n(X)$ of a space: we defined $b_n(X)$ in Lecture 40 as the dimension of the rational vector space $H_n(X; \mathbb{Q})$, but in Exercises 45.11 and 45.12, one uses the universal coefficient theorems to deduce several alternative formulas for $b_n(X)$, namely

$$b_n(X) = \text{rank } H_n(X; \mathbb{Z}) = \text{rank } H^n(X; \mathbb{Z}) = \dim_{\mathbb{K}} H_n(X; \mathbb{K}) = \dim_{\mathbb{K}} H^n(X; \mathbb{K}),$$

where \mathbb{K} is allowed to be any field of characteristic zero. This formula is false in general if \mathbb{K} has finite characteristic, but remarkably, the usual formula for the Euler characteristic

$$\chi(X) = \sum_{n \in \mathbb{Z}} (-1)^n b_n(X) = \sum_{n \in \mathbb{Z}} (-1)^n \dim_{\mathbb{K}} H_n(X; \mathbb{K}) = \sum_{n \in \mathbb{Z}} (-1)^n \dim_{\mathbb{K}} H^n(X; \mathbb{K})$$

turns out to be valid nonetheless for coefficients in an arbitrary field \mathbb{K} , even in cases where $b_n(X)$ and $\dim_{\mathbb{K}} H_n(X; \mathbb{K})$ do not match for individual values of n . These exercises also reveal that the Lefschetz number $L(f)$ of a map $f : X \rightarrow X$, which was defined in §40.4 using homology with rational coefficients, can be computed in terms of the homomorphism induced by f on the *free part* of $H_*(X; \mathbb{Z})$ or $H^*(X; \mathbb{Z})$,

$$H_n^{\text{free}}(X) := H_n(X; \mathbb{Z})/\text{torsion}, \quad H_{\text{free}}^n(X) := H^n(X; \mathbb{Z})/\text{torsion},$$

namely

$$L(f) = \sum_{n=0}^{\infty} (-1)^n \text{tr} \left(H_n^{\text{free}}(X) \xrightarrow{f_*} H_n^{\text{free}}(X) \right) = \sum_{n=0}^{\infty} (-1)^n \text{tr} \left(H_{\text{free}}^n(X) \xrightarrow{f^*} H_{\text{free}}^n(X) \right).$$

This proves on the one hand that $L(f)$ is always an integer (not just a rational number), and it also provides the freedom to compute $L(f)$ in terms of cohomology, as in the application sketched at the beginning of Lecture 41.

Finally, Exercise 45.13 establishes a technical result that lies somewhere in the background of the famous Poincaré conjecture. When combined with some results about the homology of manifolds to be proved later in this semester, it will imply that for every closed topological n -manifold M ,

$$\text{torsion}(H_{n-1}(M; \mathbb{Z})) = 0.$$

This will serve as one step in a somewhat intricate argument proving that every simply-connected closed 3-manifold is homotopy equivalent to S^3 , thus motivating the most popular statement of the Poincaré conjecture in dimension three: every simply-connected closed 3-manifold is homeomorphic to S^3 . The obvious analogue of that statement in dimensions four and upward is easily seen to be false—the correct statement of the Poincaré conjecture for arbitrary dimensions is rather that every closed n -manifold *homotopy equivalent* to S^n is also homeomorphic to it. This is in fact a known theorem, and is highly nontrivial for every $n \geq 3$: the cases $n \geq 5$ were established by Smale in the 1960's, the case $n = 4$ followed by work of Friedman in the 1980's, and the hardest case is $n = 3$, proved by Perelman at the beginning of the current century.

45.9. Exercises.

EXERCISE 45.1. Prove Proposition 45.11 on the left-exactness of $\text{Hom}(G, \cdot)$ and exactness in the case that G is projective.

EXERCISE 45.2. 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 $R\text{-Mod}$ (cf. Exercise 39.8). Using this language essentially makes the result valid not just for R -modules but in arbitrary abelian categories.

EXERCISE 45.3 (*). Prove Proposition 45.20, stating that torsion-free \mathbb{Z} -modules are flat. *Hint:* Given abelian groups A, B, G and an injective homomorphism $i : A \hookrightarrow B$, show that any nontrivial element in the kernel of $i \otimes \mathbb{1} : A \otimes G \rightarrow B \otimes G$ is also in the kernel of the restriction of this map to $A \otimes G_0 \rightarrow B \otimes G_0$ 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 ?

EXERCISE 45.4 (*). Prove the formula for $L_n \mathcal{F}(R/kR)$ stated in Theorem 45.29. *Hint:* Every additive functor $\mathcal{F} : R\text{-Mod} \rightarrow R\text{-Mod}$ has the following property (why?). For every R -module A and every integer $k \in \mathbb{Z}$, \mathcal{F} sends the morphism $A \rightarrow A : x \mapsto kx$ to the morphism $\mathcal{F}(A) \rightarrow \mathcal{F}(A) : y \mapsto ky$.

Further hint: you already know this for $k = 1$, just because \mathcal{F} is a functor.

EXERCISE 45.5 (*). Prove Theorem 45.30, the contravariant analogue of Theorem 45.29 on the properties of derived functors.

EXERCISE 45.6. Suppose $0 \rightarrow G \xrightarrow{\varphi} H \xrightarrow{\psi} K \rightarrow 0$ is a short exact sequence of R -modules, and A is another R -module.

- (a) Derive the long exact sequence (45.10) that ends with $\dots \rightarrow \text{Tor}_1(A, K) \rightarrow A \otimes G \rightarrow A \otimes H \rightarrow A \otimes K \rightarrow 0$.

Hint: Show that for any projective resolution $A_ \xrightarrow{\alpha} A$ of A , there is a short exact sequence of chain complexes*

$$0 \rightarrow A_* \otimes G \xrightarrow{1 \otimes \varphi} A_* \otimes H \xrightarrow{1 \otimes \psi} A_* \otimes K \rightarrow 0.$$

The fact that the modules A_n are projective is relevant here; it implies that each of the functors $A_n \otimes : R\text{-Mod} \rightarrow R\text{-Mod}$ is exact.

- (b) Derive a similar long exact sequence of the form $\dots \rightarrow \text{Ext}_1(A, K) \rightarrow \text{Hom}(A, G) \rightarrow \text{Hom}(A, H) \rightarrow \text{Hom}(A, K) \rightarrow 0$.

EXERCISE 45.7. Set $R := \mathbb{Z}$, and show that for an abelian group G with torsion subgroup $T(G) \subset G$ and any other abelian group A , the map

$$\text{Tor}(A, T(G)) \rightarrow \text{Tor}(A, G)$$

induced by the inclusion $T(G) \hookrightarrow G$ is an isomorphism.

Hint: Use a long exact sequence.

EXERCISE 45.8. Prove Proposition 45.31 on the isomorphism $H_n(C^*, \partial) \cong H_n(C_*, d)$ arising from a double complex, and clarify what it means to call this isomorphism *natural*.

Hint: Chasing the diagram should feel a bit like climbing stairs.

EXERCISE 45.9. Show that for every space X and abelian group G , the natural map $h : H^1(X; G) \rightarrow \text{Hom}(H_1(X; \mathbb{Z}), G)$ is an isomorphism.

Hint: The computation of $H_0(X; \mathbb{Z})$ tells you something about $\text{Ext}(H_0(X; \mathbb{Z}), G)$.

EXERCISE 45.10 (*). In this exercise, \mathbb{K} is a field, but it will be regarded as merely an abelian group for the purposes of tensor products and Tor and Ext functors.

- (a) Show that for any $m \in \mathbb{N}$, the three abelian groups $\mathbb{Z}_m \otimes \mathbb{K}$, $\text{Tor}(\mathbb{Z}_m, \mathbb{K})$ and $\text{Ext}(\mathbb{Z}_m, \mathbb{K})$ are all isomorphic to each other, and are all either trivial or isomorphic to \mathbb{K} .

Hint: $\mathbb{K} \rightarrow \mathbb{K} : k \mapsto mk$ is a linear map between two 1-dimensional vector spaces.

- (b) Deduce that for any *finite* abelian group T , $T \otimes \mathbb{K}$, $\text{Tor}(T, \mathbb{K})$ and $\text{Ext}(T, \mathbb{K})$ are all isomorphic abelian groups, and they also all have natural \mathbb{K} -module structures that make them isomorphic vector spaces.

EXERCISE 45.11 (*). Assume in this exercise that \mathbb{K} is a field of characteristic zero.

- (a) Show that for any chain complex C_* of free abelian groups, 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. Note that on both sides of this map, \mathbb{K} should be regarded as an abelian group rather than a \mathbb{K} -module (since $H_*(C_*)$ and C_* are not assumed to be \mathbb{K} -modules), thus \otimes denotes the tensor product of abelian groups.

Remark: It follows in particular that for any space X and any field \mathbb{K} of characteristic zero, the canonical map $h : H_(X; \mathbb{Z}) \otimes \mathbb{K} \rightarrow H_*(X; \mathbb{K})$ is an isomorphism.*

- (b) Where does the proof in part (a) fail if \mathbb{K} is a field with finite characteristic? What happens, for instance, if $\mathbb{K} = \mathbb{Z}_2$ and $C_* = C_*(\mathbb{R}\mathbb{P}^2; \mathbb{Z})$?

- (c) Recall that the *rank* of a finitely-generated abelian group G with torsion subgroup $T \subset G$ is the unique integer $r \geq 0$ such that $G/T \cong \mathbb{Z}^r$. Deduce from part (a) that for any space X with finitely-generated homology groups over \mathbb{Z} , the Betti numbers $b_n(X) = \dim_{\mathbb{Q}} H_n(X; \mathbb{Q})$ satisfy

$$b_n(X) = \dim_{\mathbb{K}} H_n(X; \mathbb{K}) = \dim_{\mathbb{K}} H^n(X; \mathbb{K}) = \text{rank } H_n(X; \mathbb{Z}).$$

Find a counterexample showing that this need not hold if \mathbb{K} has finite characteristic.

- (d) Let $H_n^{\text{free}}(X)$ denote the quotient of $H_n(X; \mathbb{Z})$ by its torsion subgroup, and observe that for any map $f : X \rightarrow X$, the induced homomorphism $f_* : H_n(X; \mathbb{Z}) \rightarrow H_n(X; \mathbb{Z})$ descends to a map $f_* : H_n^{\text{free}}(X) \rightarrow H_n^{\text{free}}(X)$. If $H_n(X; \mathbb{Z})$ is finitely generated, then $H_n^{\text{free}}(X)$ is isomorphic to \mathbb{Z}^r for $r := \text{rank } H_n(X; \mathbb{Z})$, thus $f_* : H_n^{\text{free}}(X) \rightarrow H_n^{\text{free}}(X)$ can be represented by an r -by- r matrix with integer entries whose trace $\text{tr}(f_*) \in \mathbb{Z}$ is independent of choices. Using the naturality of the map $h : H_n(X; \mathbb{Z}) \otimes \mathbb{Q} \rightarrow H_n(X; \mathbb{Q})$, show that the Lefschetz number $L(f) \in \mathbb{Q}$ of a map $f : X \rightarrow X$ as defined in §40.4 satisfies

$$L(f) = \sum_{n=0}^{\infty} (-1)^n \text{tr} \left(H_n^{\text{free}}(X) \xrightarrow{f_*} H_n^{\text{free}}(X) \right).$$

This proves in particular that $L(f)$ is an integer.

EXERCISE 45.12 (*). Suppose $n \in \mathbb{Z}$ and C_* is a chain complex of free abelian groups such that the homology groups $H_n(C_*)$ and $H_{n-1}(C_*)$ are finitely generated. For $k \in \{n-1, n\}$, the classification of finitely-generated abelian groups allows us to write

$$H_k(C_*) = H_k^{\text{free}} \oplus T_k,$$

where $T_k \subset H_k(C_*)$ is a finite subgroup (the torsion) and $H_k^{\text{free}} := H_k(C_*)/T_k$ is a free abelian group of finite rank. Let H_{free}^k in turn denote the quotient of $H^k(C_*; \mathbb{Z})$ by its torsion subgroup.

- (a) Show that the map $h : H^n(C_*; \mathbb{Z}) \rightarrow \text{Hom}(H_n(C_*), \mathbb{Z})$ determines a natural isomorphism

$$H_{\text{free}}^n \xrightarrow[\cong]{h} \text{Hom}(H_n^{\text{free}}, \mathbb{Z}),$$

implying in particular that there is a (non-natural) isomorphism $H_{\text{free}}^n \cong H_n^{\text{free}}$. One obtains from this yet another new formula for the n th Betti number of a space X ,

$$b_n(X) = \text{rank } H^n(X; \mathbb{Z}).$$

- (b) Show that the torsion subgroup of $H^n(C_*; \mathbb{Z})$ is isomorphic to the torsion subgroup of $H_{n-1}(C_*)$.
 (c) Show that for any space X with finitely-generated homology, the formula

$$\chi(X) = \sum_{k \in \mathbb{Z}} (-1)^k \dim_{\mathbb{K}} H_k(X; \mathbb{K}) = \sum_{k \in \mathbb{Z}} (-1)^k \dim_{\mathbb{K}} H^k(X; \mathbb{K})$$

holds for arbitrary fields \mathbb{K} (not just with characteristic zero).

Hint: Exercise 45.10 will be helpful here.

- (d) Show that the Lefschetz number of a map $f : X \rightarrow X$ satisfies

$$L(f) = \sum_{n=0}^{\infty} (-1)^n \text{tr} \left(H_{\text{free}}^n(X) \xrightarrow{f_*} H_{\text{free}}^n(X) \right).$$

EXERCISE 45.13 (*). 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.

46. The cross product

The main question for this lecture is straightforward: if we understand $H_*(X)$ and $H_*(Y)$, can we use them to compute $H_*(X \times Y)$?

In the bordism theory sketched in Lecture 27, it is easy to see that there is a product operation

$$\begin{aligned} \Omega_m^\bullet(X) \otimes \Omega_n^\bullet(Y) &\xrightarrow{\times} \Omega_{m+n}^\bullet(X \times Y), \\ [(M, \varphi)] \otimes [(N, \psi)] &\longmapsto [(M, \varphi)] \times [(N, \psi)] := [(M \times N, \varphi \times \psi)]. \end{aligned}$$

The reason this is straightforward to define is that the product of two closed manifolds is also a closed manifold. We will see in this lecture that one can define a similarly straightforward product in cellular homology, using the fact that k -cells $e^k \cong \mathbb{D}^k$ can be identified homeomorphically with cubes, so that the product of a k -cell with an ℓ -cell becomes a $(k + \ell)$ -cell. Defining such a product in singular homology is less straightforward, because the product of two simplices is not a simplex in any obvious way, but here we can make use of the triangulation of $\Delta^m \times \Delta^n$ introduced in §31.2, or as a purely algebraic alternative, the method of acyclic models (see Lecture 32).

46.1. 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

$$(X \times Y)^n = \bigcup_{0 \leq k \leq n} X^k \times Y^{n-k}.$$

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 $\Phi_\alpha : I^k \rightarrow X$ and $\Phi_\beta : I^\ell \rightarrow Y$ respectively gives rise to a $(k + \ell)$ -cell

$$e_\alpha^k \times e_\beta^\ell \subset X \times Y$$

with characteristic map

$$\Phi_\alpha \times \Phi_\beta : I^{k+\ell} \rightarrow X \times Y : (s, t) \mapsto (\Phi_\alpha(s), \Phi_\beta(t)).$$

Using coefficients in a commutative ring R , the bilinear operation

$$C_k^{\text{CW}}(X; R) \times C_\ell^{\text{CW}}(Y; R) \xrightarrow{\times} C_{k+\ell}^{\text{CW}}(X \times Y; R)$$

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 46.3).⁶⁹

PROPOSITION 46.1. *For any pair of CW-complexes X and Y with a k -cell $e_\alpha^k \subset X$ and an ℓ -cell $e_\beta^\ell \subset Y$, the cellular chain complex of $X \times Y$ with its induced cell decomposition and coefficients in any commutative ring R satisfies*

$$\partial(e_\alpha^k \times e_\beta^\ell) = \partial e_\alpha^k \times e_\beta^\ell + (-1)^k e_\alpha^k \times \partial e_\beta^\ell \in C_{k+\ell-1}^{\text{CW}}(X \times Y; R).$$

□

⁶⁹An explicit proof of the formula in Prop. 46.1 can also be found in [Hat02, Prop. 3B.1].

Being a bilinear map of R -modules, the cross product is equivalent to an R -module homomorphism

$$C_k^{\text{CW}}(X; R) \otimes C_\ell^{\text{CW}}(Y; R) \xrightarrow{\times} C_{k+\ell}^{\text{CW}}(X \times Y; R),$$

namely the unique homomorphism that sends $e_\alpha^k \otimes e_\beta^\ell$ to $e_\alpha^k \times e_\beta^\ell$ for every pair of cells of the appropriate dimensions in X and Y . More generally, one can include two arbitrary coefficient modules $G, H \in R\text{-Mod}$ in the picture and define

$$C_k^{\text{CW}}(X; G) \otimes C_\ell^{\text{CW}}(Y; H) \xrightarrow{\times} C_{k+\ell}^{\text{CW}}(X \times Y; G \otimes H),$$

$$ge_\alpha^k \otimes he_\beta^\ell \mapsto (g \otimes h)(e_\alpha^k \times e_\beta^\ell)$$

for $g \in G$ and $h \in H$. The first version follows from this one by taking $G = H := R$ and using the canonical R -module isomorphism $R \otimes R \cong R : r \otimes s \mapsto rs$. Taking the direct sum of these maps over all pairs of integers $k, \ell \geq 0$, the chain-level cellular cross product now determines an R -module homomorphism

$$C_*^{\text{CW}}(X; G) \otimes C_*^{\text{CW}}(Y; H) \xrightarrow{\times} C_*^{\text{CW}}(X \times Y; G \otimes H).$$

In the case $G = H := R$, this map is in fact 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.

46.2. The Künneth formula. The following purely algebraic definition should now hopefully seem quite natural.

DEFINITION 46.2. Given chain complexes (A_*, ∂^A) and (B_*, ∂^B) of R -modules, the **tensor product chain complex** $(A_* \otimes B_*, \partial)$ is defined by

$$(46.1) \quad (A_* \otimes B_*)_n = \bigoplus_{k+\ell=n} A_k \otimes B_\ell,$$

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

$$(46.2) \quad \partial(a \otimes b) = \partial^A a \otimes b + (-1)^k a \otimes \partial^B b \quad \text{for } a \in A_k, b \in B_\ell.$$

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 $f : A_* \rightarrow A'_*$ and $g : B_* \rightarrow B'_*$, there is a chain map

$$(46.3) \quad f \otimes g : A_* \otimes B_* \rightarrow A'_* \otimes B'_* : a \otimes b \mapsto f(a) \otimes g(b).$$

We can now rephrase Proposition 46.1 as follows:

PROPOSITION 46.3. *For any choices of coefficient modules $G, H \in R\text{-Mod}$, the chain-level cellular cross product determines a chain map*

$$C_*^{\text{CW}}(X; G) \otimes C_*^{\text{CW}}(Y; H) \rightarrow C_*^{\text{CW}}(X \times Y; G \otimes H) : a \otimes b \rightarrow a \times b,$$

and it is an isomorphism of chain complexes in the case $G = H := R$. □

REMARK 46.4. The signs in formulas such as (46.2) 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 31.4). 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 $\partial : A_* \rightarrow A_*$ 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

$$H_k(A_*) \otimes H_\ell(B_*) \rightarrow H_{k+\ell}(A_* \otimes B_*) : [a] \otimes [b] \mapsto [a \otimes b].$$

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

$$(46.4) \quad \bigoplus_{k+\ell=n} H_k(A_*) \otimes H_\ell(B_*) \rightarrow H_n(A_* \otimes B_*)$$

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 46.5 (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 (46.4) for every $n \in \mathbb{Z}$ fits into a natural short exact sequence*

$$0 \rightarrow \bigoplus_{k+\ell=n} H_k(C_*) \otimes H_\ell(C'_*) \rightarrow H_n(C_* \otimes C'_*) \rightarrow \bigoplus_{k+\ell=n-1} \text{Tor}(H_k(C_*), H_\ell(C'_*)) \rightarrow 0,$$

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 $A_* \rightarrow B_*$ and $A'_* \rightarrow B'_*$.

The statement becomes a bit more concise if we define the operation \otimes and the functor Tor on pairs of \mathbb{Z} -graded R -modules $A_* = \bigoplus_{k \in \mathbb{Z}} A_k$ and $B_* = \bigoplus_{\ell \in \mathbb{Z}} B_\ell$, i.e. a grading on $A_* \otimes B_*$ is defined via (46.1), and similarly,

$$\text{Tor}(A_*, B_*) := \bigoplus_{n \in \mathbb{Z}} (\text{Tor}(A_*, B_*))_n, \quad \text{where} \quad (\text{Tor}(A_*, B_*))_n := \bigoplus_{k+\ell=n} \text{Tor}(A_k, B_\ell).$$

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 46.5 can now be written as

$$0 \rightarrow H_*(C_*) \otimes H_*(C'_*) \rightarrow H_*(C_* \otimes C'_*) \rightarrow (\text{Tor}(H_*(C_*), H_*(C'_*)))_{*-1} \rightarrow 0,$$

where the subscript “ $* - 1$ ” on the last term indicates the downward degree shift. The splitting gives rise to a (non-canonical) isomorphism

$$H_*(C_* \otimes C'_*) \cong (H_*(C_*) \otimes H_*(C'_*)) \oplus (\text{Tor}(H_*(C_*), H_*(C'_*)))_{*-1},$$

which can be used in practice to compute the cellular homology of products. We will see at the end of this 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 for homology 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 46.5 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 the category

$$R\text{-Mod}_{\mathbb{Z}}$$

of \mathbb{Z} -graded R -modules $A_* = \bigoplus_{n \in \mathbb{Z}} A_n$, whose morphisms consist of R -module homomorphisms $A_* \rightarrow B_*$ that map A_n to B_n for every $n \in \mathbb{Z}$. One can speak of additive, exact, left-exact and right-exact functors on $R\text{-Mod}_{\mathbb{Z}}$ in exactly the same way as for functors on $R\text{-Mod}$: 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

$$\otimes G_* : R\text{-Mod}_{\mathbb{Z}} \rightarrow R\text{-Mod}_{\mathbb{Z}} : A_* \mapsto A_* \otimes G_*$$

is right-exact for the same reasons that $\otimes G : R\text{-Mod} \rightarrow R\text{-Mod}$ is right-exact, and one can similarly define left derived functors $\text{Tor}_n(\cdot, G_*) := L_n(\otimes G_*) : R\text{-Mod}_{\mathbb{Z}} \rightarrow R\text{-Mod}_{\mathbb{Z}}$ 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

$$\dots \longrightarrow A_{*,2} \xrightarrow{\alpha_2} A_{*,1} \xrightarrow{\alpha_1} A_{*,0} \xrightarrow{\alpha} A_* \longrightarrow 0,$$

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

$$\begin{array}{ccc} & & H_* \\ & \nearrow \tilde{\varphi} & \downarrow \pi \\ A_{*,n} & \xrightarrow{\varphi} & G_* \end{array}$$

can be solved in the category $R\text{-Mod}_{\mathbb{Z}}$ whenever $\varphi : A_{*,n} \rightarrow G_*$ and $\pi : H_* \rightarrow G_*$ 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 $R\text{-Mod}_{\mathbb{Z}}$ is thus equivalent to a collection of projective resolutions $A_{k,*} \xrightarrow{\alpha} A_k$ in $R\text{-Mod}$, one for each $k \in \mathbb{Z}$. For this reason, defining $\text{Tor}_n(A_*, G_*)$ as a left derived functor in the category $R\text{-Mod}_{\mathbb{Z}}$ 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 $\text{Tor}_n(A_*, G_*)$ the \mathbb{Z} -grading such that

$$\text{Tor}_n(A_*, G_*)_m = \bigoplus_{k+\ell=m} \text{Tor}_n(A_k, G_{\ell})$$

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 46.5. 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 $0 \rightarrow Z_* \hookrightarrow C_* \xrightarrow{\partial} B_{*-1} \rightarrow 0$ splits, and so therefore does the sequence

$$0 \longrightarrow Z_* \otimes C'_* \longrightarrow C_* \otimes C'_* \rightarrow B_{*-1} \otimes C'_* \longrightarrow 0,$$

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

$$(Z_* \otimes C'_*)_n = \bigoplus_{k+\ell=n} Z_k \otimes C'_\ell,$$

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

$$(Z_* \otimes C'_*)_n \cong \bigoplus_{k+\ell=n} \bigoplus_{e \in \mathcal{B}_k} R \otimes C'_\ell \cong \bigoplus_{k+\ell=n} \bigoplus_{e \in \mathcal{B}_k} C'_\ell,$$

so that $(Z_* \otimes C'_*)_n \xrightarrow{\partial} (Z_* \otimes C'_*)_{n-1}$ becomes the corresponding direct sum of the boundary maps $C'_\ell \rightarrow C'_{\ell-1}$. The homology of this complex is thus

$$H_n(Z_* \otimes C'_*) \cong \bigoplus_{k+\ell=n} \bigoplus_{e \in \mathcal{B}_k} H_\ell(C'_*) \cong \bigoplus_{k+\ell=n} Z_k \otimes H_\ell(C'_*),$$

a formula that can be written concisely as a natural isomorphism of \mathbb{Z} -graded R -modules

$$H_*(Z_* \otimes C'_*) \cong Z_* \otimes H_*(C'_*).$$

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

$$(B_{*-1} \otimes C'_*)_n = \bigoplus_{k+\ell=n} B_{k-1} \otimes C'_\ell = \bigoplus_{k+\ell=n-1} B_k \otimes C'_\ell$$

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,

$$H_*(B_{*-1} \otimes C'_*) = H_{*-1}(B_* \otimes C'_*) \cong (B_* \otimes H_*(C'_*))_{*-1}.$$

The short exact sequence of chain complexes now gives rise as usual to a long exact sequence of homologies

$$\dots \rightarrow H_n(B_* \otimes C'_*) \xrightarrow{\Phi_n} H_n(Z_* \otimes C'_*) \rightarrow H_n(C_* \otimes C'_*) \rightarrow H_{n-1}(B_* \otimes C'_*) \xrightarrow{\Phi_{n-1}} H_{n-1}(Z_* \otimes C'_*) \rightarrow \dots,$$

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

$$0 \rightarrow \operatorname{coker} \Phi_n \rightarrow H_n(C_* \otimes C'_*) \rightarrow \ker \Phi_{n-1} \rightarrow 0,$$

or if we take the direct sum over all $n \in \mathbb{Z}$, a short exact sequence of \mathbb{Z} -graded R -modules

$$(46.5) \quad 0 \rightarrow \operatorname{coker} \Phi_* \rightarrow H_*(C_* \otimes C'_*) \rightarrow \ker \Phi_{*-1} \rightarrow 0.$$

Inspecting the diagram chase behind the long exact sequence reveals that the map $\Phi_* : H_*(B_* \otimes C'_*) \rightarrow H_*(Z_* \otimes C'_*)$ is the obvious thing: it is the map

$$H_*(B_* \otimes C'_*) \cong B_* \otimes H_*(C'_*) \xrightarrow{i_* \otimes \mathbb{1}} Z_* \otimes H_*(C'_*) \cong H_*(Z_* \otimes C'_*)$$

induced by the inclusions $B_* \xrightarrow{i_*} Z_*$. In order to understand the kernel and cokernel of $i_* \otimes \mathbb{1}$, one feeds the short exact sequence

$$0 \rightarrow B_* \xrightarrow{i_*} Z_* \xrightarrow{q_*} H_*(C_*) \rightarrow 0$$

into the functor $\otimes H_*(C'_*) : R\text{-Mod}_{\mathbb{Z}} \rightarrow R\text{-Mod}_{\mathbb{Z}}$ and obtains a long (but not very long) exact sequence

$$\begin{aligned} 0 \longrightarrow \text{Tor}(H_*(C_*), H_*(C'_*)) \longrightarrow B_* \otimes H_*(C'_*) \xrightarrow{i_* \otimes \mathbb{1}} Z_* \otimes H_*(C'_*) \\ \xrightarrow{q_* \otimes \mathbb{1}} H_*(C_*) \otimes H_*(C'_*) \longrightarrow 0, \end{aligned}$$

in which the leftmost term $\text{Tor}(Z_*, H_*(C'_*)) = 0$ vanishes because $Z_* \subset C_*$ is free. This leads to the isomorphisms

$$\text{coker}(i_* \otimes \mathbb{1}) \cong H_*(C_*) \otimes H_*(C'_*),$$

and

$$\text{ker}(i_* \otimes \mathbb{1}) \cong \text{Tor}(H_*(C_*), H_*(C'_*)).$$

The sequence we were looking for is now obtained by plugging these isomorphisms into (46.5), 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

$$\text{Tor}(A, B) = 0 \quad \text{for all vector spaces } A, B \text{ over } \mathbb{K},$$

and we therefore obtain:

COROLLARY 46.6. *For any field \mathbb{K} and any two chain complexes C_* and C'_* of vector spaces over \mathbb{K} , the canonical map*

$$\bigoplus_{k+\ell=n} H_k(C_*) \otimes H_\ell(C'_*) \rightarrow H_n(C_* \otimes C'_*)$$

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.

46.3. 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

$$(46.6) \quad C_*^{\text{CW}}(X; R) \otimes C_*^{\text{CW}}(Y; R) \xrightarrow{\cong} C_*^{\text{CW}}(X \times Y; R) : a \otimes b \mapsto a \times b,$$

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 $e_\alpha^k \times e_\beta^\ell \subset X \times Y$. Composing the induced R -module isomorphism on homology with the natural map

$$H_k^{\text{CW}}(X; R) \otimes H_\ell^{\text{CW}}(Y; R) \rightarrow H_{k+\ell}(C_*^{\text{CW}}(X; R) \otimes C_*^{\text{CW}}(Y; R)) : [x] \otimes [y] \mapsto [x \otimes y]$$

gives rise to a bilinear cross product on homology,

$$H_k^{\text{CW}}(X; R) \otimes H_\ell^{\text{CW}}(Y; R) \xrightarrow{\times} H_{k+\ell}^{\text{CW}}(X \times Y; R).$$

If R is additionally a principal ideal domain, then the Künneth formula also holds, producing for each integer $n \geq 0$ a natural (and non-naturally split) short exact sequence of R -modules

$$0 \longrightarrow \bigoplus_{k+\ell=n} H_k^{\text{CW}}(X; R) \otimes H_\ell^{\text{CW}}(Y; R) \xrightarrow{\times} H_n^{\text{CW}}(X \times Y; R) \\ \longrightarrow \bigoplus_{k+\ell=n-1} \text{Tor}(H_k^{\text{CW}}(X; R), H_\ell^{\text{CW}}(Y; R)) \longrightarrow 0,$$

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 section how to establish such an exact sequence directly in singular homology, without needing to assume that X and Y are CW-complexes.

REMARK 46.7. 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 $K \rightarrow X \times Y$ from some compact space K is continuous, then both topologies give the same answer (see Exercise 46.2). Applying this observation for maps $\Delta^n \rightarrow X \times Y$, 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 $H_*(X \times Y; G) \cong H_*^{\text{CW}}(X \times Y; G)$ holds as usual. With this in mind, we shall assume from now on that $X \times Y$ carries the product topology.

46.4. The singular cross product. Since cellular and singular homologies are isomorphic, the cellular cross product determines a homomorphism

$$(46.7) \quad H_k(X; R) \otimes H_\ell(Y; R) \xrightarrow{\times} H_{k+\ell}(X \times Y; R)$$

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 (46.7) 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 (46.7) be defined in general?

Since there is always a canonical homomorphism $H_*(X; R) \otimes H_*(Y; R) \rightarrow H_*(C_*(X; R) \otimes C_*(Y; R))$, we would obtain a map (46.7) if we had a chain map

$$C_*(X; R) \otimes C_*(Y; R) \xrightarrow{x} C_*(X \times Y; R)$$

to play the role in singular homology that (46.6) 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 $\sigma : \Delta^k \rightarrow X$ and $\tau : \Delta^\ell \rightarrow Y$. Unfortunately, $\Delta^k \times \Delta^\ell$ is not a simplex in any canonical way, so we cannot simply write down the continuous map

$$(46.8) \quad \sigma \times \tau : \Delta^k \times \Delta^\ell \rightarrow X \times Y : (s, t) \mapsto (\sigma(s), \tau(t))$$

and call it a generator of $C_{k+\ell}(X \times Y; R)$. 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 $k, \ell \geq 0$, 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 (46.8) 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 (46.8) 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 §31.2, and the formula (31.2) for $\partial c_{\Delta^k \times \Delta^\ell}$ guarantees that the unique R -module homomorphism $C_*(X) \otimes C_*(Y) \xrightarrow{x} C_*(X \times Y)$ 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 §31.2 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: the method of acyclic models, which was used in Lecture 32 to show that the natural chain map from the ordered to the oriented versions of the simplicial chain complex is a chain homotopy equivalence. In the setting of singular chain complexes, acyclic model arguments produce an abundance of natural chain maps that descend to product structures on homology and cohomology.

REMARK 46.8. 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 is difficult to digest, as it is typically expressed in heavily abstract category-theoretic language. We prefer in these notes to follow the approach of [Bre93] and demonstrate the method by example.

A preparatory comment is in order before we continue. If we can define a cross product on the singular chain complex with integer coefficients, then for the reasons explained in Remark 32.10, the definition and its important properties will almost immediately extend to coefficients in an arbitrary commutative ring R .

LEMMA 46.9. *One can assign to every tuple of topological spaces X, Y a chain map*

$$\Phi : C_*(X; \mathbb{Z}) \otimes C_*(Y; \mathbb{Z}) \rightarrow C_*(X \times Y; \mathbb{Z})$$

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 $f : X \rightarrow X'$ and $g : Y \rightarrow Y'$, the diagram

$$\begin{CD} C_*(X; \mathbb{Z}) \otimes C_*(Y; \mathbb{Z}) @>\Phi>> C_*(X \times Y; \mathbb{Z}) \\ @V f_* \otimes g_* VV @VV (f \times g)_* V \\ C_*(X'; \mathbb{Z}) \otimes C_*(Y'; \mathbb{Z}) @>\Phi>> C_*(X' \times Y'; \mathbb{Z}) \end{CD}$$

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 $\Phi : C_0(X) \otimes C_0(Y) \rightarrow C_0(X \times Y)$ 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 $f : X \rightarrow X'$, $g : Y \rightarrow Y'$ and points $x \in X$, $y \in Y$ (regarded as singular 0-simplices). We shall now argue by induction and assume that maps $\Phi : C_k(X) \otimes C_\ell(Y) \rightarrow C_{k+\ell}(X \times Y)$ have been defined for all spaces X, Y and all integers $k, \ell \geq 0$ with $k + \ell \leq n - 1$ for some $n \geq 1$, 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 $k \geq 0$, let $i_k : \Delta^k \rightarrow \Delta^k$ 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 $k, \ell \geq 0$ 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

$$(46.9) \quad \partial \Phi(i_k \otimes i_\ell) = \Phi(\partial(i_k \otimes i_\ell)) \in C_{n-1}(\Delta^k \times \Delta^\ell),$$

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

$$(46.10) \quad \partial \Phi \partial(i_k \otimes i_\ell) = \Phi \partial^2(i_k \otimes i_\ell) = 0,$$

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 $\epsilon : C_0(\Delta^k \times \Delta^\ell) \rightarrow \mathbb{Z}$ denote the augmentation in the augmented chain complex $\tilde{C}_*(\Delta^k \times \Delta^\ell)$, 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,

$$(46.11) \quad \epsilon \Phi \partial(i_k \otimes i_\ell) = 0 \quad \text{when} \quad n = 1.$$

Now comes the crucial point: $\Delta^k \times \Delta^\ell$ is contractible, so its reduced singular homology is trivial. In light of (46.10) and (46.11), this means

$$[\Phi \partial(i_k \otimes i_\ell)] = 0 \in \tilde{H}_{n-1}(\Delta^k \times \Delta^\ell),$$

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 (46.9) 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

$$(46.12) \quad \Phi(i_k \otimes i_\ell) \in \partial^{-1}(\Phi \partial(i_k \otimes i_\ell)).$$

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 $k, \ell \geq 0$ with $k + \ell = n$, we claim that the general extension of $\Phi : C_*(X) \otimes C_*(Y) \rightarrow C_*(X \times Y)$ 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 $\sigma : \Delta^k \rightarrow X$ and $\tau : \Delta^\ell \rightarrow Y$ with $k + \ell = n$, we have

$$\sigma = \sigma_* i_k \in C_k(X), \quad \tau = \tau_* i_\ell \in C_\ell(Y),$$

so naturality requires $\Phi : C_k(X) \otimes C_\ell(Y) \rightarrow C_n(X \times Y)$ 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 $\sigma_* : C_*(\Delta^k) \rightarrow C_*(X)$, $\tau_* : C_*(\Delta^\ell) \rightarrow C_*(Y)$ and $(\sigma \times \tau)_* : C_*(\Delta^k \times \Delta^\ell) \rightarrow C_*(X \times Y)$ 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 (46.3)). This establishes the chain map property. To see that naturality also holds, consider two continuous maps $f : X \rightarrow X'$ and $g : Y \rightarrow Y'$: then

$$\begin{aligned} \Phi(f_* \otimes g_*)(\sigma \otimes \tau) &= \Phi((f \circ \sigma)_* \otimes (g \circ \tau)_*)(i_k \otimes i_\ell) = ((f \circ \sigma) \times (g \circ \tau))_* \Phi(i_k \otimes i_\ell) \\ &= (f \times g)_*(\sigma \times \tau)_* \Phi(i_k \otimes i_\ell) = (f \times g)_* \Phi(\sigma \otimes \tau). \end{aligned}$$

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

$$h : C_k(X) \otimes C_\ell(Y) \rightarrow C_{n+1}(X \times Y)$$

for every pair of integers $k, \ell \geq 0$ 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

$$\begin{array}{ccc} C_k(X) \otimes C_\ell(Y) & \xrightarrow{h} & C_{n+1}(X \times Y) \\ \downarrow f_* \otimes g_* & & \downarrow (f \times g)_* \\ C_k(X') \otimes C_\ell(Y') & \xrightarrow{h} & C_{n+1}(X' \times Y') \end{array}$$

commutes for every pair of continuous maps $f : X \rightarrow X'$ and $g : Y \rightarrow Y'$.

Since Φ and Ψ match precisely on all 0-chains, we are free to define $h : C_0(X) \otimes C_0(Y) \rightarrow C_1(X \times Y)$ 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 $n \geq 1$. 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

$$h(i_k \otimes i_\ell) \in \partial^{-1}((-h\partial + \Phi - \Psi)(i_k \otimes i_\ell)).$$

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 $f : X \rightarrow X'$ and $g : Y \rightarrow Y'$, we have

$$\begin{aligned} h(f_* \otimes g_*)(\sigma \otimes \tau) &= h((f \circ \sigma)_* \otimes (g \circ \tau)_*)(i_k \otimes i_\ell) = ((f \circ \sigma) \times (g \circ \tau))_* h(i_k \otimes i_\ell) \\ &= (f \times g)_*(\sigma \times \tau)_* h(i_k \otimes i_\ell) = (f \times g)_* h(\sigma \otimes \tau). \end{aligned}$$

This completes the inductive step and finishes the proof. \square

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 last 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.

The chain map $\Phi : C_*(X; \mathbb{Z}) \otimes C_*(Y; \mathbb{Z}) \rightarrow C_*(X \times Y; \mathbb{Z})$ from Lemma 46.9 uniquely determines a chain map of R -modules

$$C_*(X; R) \otimes C_*(Y; R) \rightarrow C_*(X \times Y; R)$$

for any commutative ring R . For any two choices of the chain map Φ with coefficients in \mathbb{Z} , a chain homotopy between them similarly determines a chain homotopy between the resulting chain maps $C_*(X; R) \otimes C_*(Y; R) \rightarrow C_*(X \times Y; R)$. The induced map on homology is therefore independent

of choices, and it can be composed with the canonical map from $H_*(X; R) \otimes H_*(Y; R)$ to define what we will henceforth call the **singular cross product**

$$H_*(X; R) \otimes H_*(Y; R) \longrightarrow H_*(C_*(X; R) \otimes C_*(Y; R)) \xrightarrow{\Phi_*} H_*(X \times Y; R).$$

$\xrightarrow{\quad \times \quad}$

This definition is not only independent of choices, but is also natural in the sense that there is a commutative diagram

$$\begin{array}{ccc} H_*(X; R) \otimes H_*(Y; R) & \xrightarrow{\times} & H_*(X \times Y; R) \\ \downarrow f_* \otimes g_* & & \downarrow (f \times g)_* \\ H_*(X'; R) \otimes H_*(Y'; R) & \xrightarrow{\times} & H_*(X' \times Y'; R) \end{array}$$

for any pair of continuous maps $f : X \rightarrow X'$ and $g : Y \rightarrow Y'$.

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 $\Phi : C_*(X; R) \otimes C_*(Y; R) \rightarrow C_*(X \times Y; R)$ 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 46.10. *One can assign to every tuple of topological spaces X, Y chain maps*

$$\begin{aligned} C_*(X \times Y; \mathbb{Z}) &\xrightarrow{\theta} C_*(X; \mathbb{Z}) \otimes C_*(Y; \mathbb{Z}), \\ C_*(X; \mathbb{Z}) \otimes C_*(Y; \mathbb{Z}) &\xrightarrow{\alpha} C_*(X \times Y; \mathbb{Z}), \\ C_*(X \times Y; \mathbb{Z}) &\xrightarrow{\beta} C_*(X \times Y; \mathbb{Z}), \end{aligned}$$

which are uniquely determined up to chain homotopy by a naturality condition and their definitions on 0-chains,

$$\theta(x, y) = x \otimes y, \quad \alpha(x \otimes y) = x \otimes y, \quad \beta(x, y) = (x, y).$$

Here, naturality of θ means that there is a commutative diagram

$$\begin{array}{ccc} C_*(X \times Y; \mathbb{Z}) & \xrightarrow{\theta} & C_*(X; \mathbb{Z}) \otimes C_*(Y; \mathbb{Z}) \\ \downarrow (f \times g)_* & & \downarrow f_* \otimes g_* \\ C_*(X' \times Y'; \mathbb{Z}) & \xrightarrow{\theta} & C_*(X'; \mathbb{Z}) \otimes C_*(Y'; \mathbb{Z}) \end{array}$$

for any pair of continuous maps $f : X \rightarrow X'$ and $g : Y \rightarrow Y'$, 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 46.11 (Eilenberg-Zilber theorem). *The natural chain maps*

$$C_*(X; \mathbb{Z}) \otimes C_*(Y; \mathbb{Z}) \xrightleftharpoons[\theta]{\Phi} C_*(X \times Y; \mathbb{Z})$$

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 $C_*(X; R) \otimes C_*(Y; R)$ and $C_*(X \times Y; R)$ that these determine. \square

PROOF OF LEMMA 46.10. 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 $n \geq 1$ by induction on n . For $\theta : C_*(X \times Y) \rightarrow C_*(X) \otimes C_*(Y)$, 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

$$d_n : \Delta^n \rightarrow \Delta^n \times \Delta^n : t \mapsto (t, t),$$

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

$$H_m(C_*(\Delta^n) \otimes C_*(\Delta^n)) \cong \bigoplus_{k+\ell=m} H_k(\Delta^n) \otimes H_\ell(\Delta^n) \cong \begin{cases} \mathbb{Z} & \text{if } m = 0, \\ 0 & \text{otherwise,} \end{cases}$$

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 $n \geq 2$, and we can therefore choose $\theta(d_n)$ to satisfy

$$(46.13) \quad \theta(d_n) \in \partial^{-1}(\theta\partial d_n).$$

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

$$\theta\partial(d_1) = \theta((1, 1) - (0, 0)) = 1 \otimes 1 - 0 \otimes 0$$

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.⁷⁰ In either case, $\theta(d_n)$ can be defined so that (46.13) holds.

Now for an arbitrary singular n -simplex $\sigma : \Delta^n \rightarrow X \times Y$, we can use the projection maps $\pi_X : X \times Y \rightarrow X$ and $\pi_Y : X \times Y \rightarrow Y$ to write

$$\begin{array}{ccc} \Delta^n & \xrightarrow{d_n} & \Delta^n \times \Delta^n \xrightarrow{(\pi_X \circ \sigma) \times (\pi_Y \circ \sigma)} X \times Y, \\ & \searrow \sigma & \nearrow \end{array}$$

so naturality requires that we define

$$\theta(\sigma) = \theta((\pi_X \circ \sigma) \times (\pi_Y \circ \sigma))_* d_n := ((\pi_X \circ \sigma)_* \otimes (\pi_Y \circ \sigma)_*) \theta(d_n).$$

⁷⁰Equivalently, 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.

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. \square

REMARK 46.12. 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 46.10. It's relaxing.

Corollary 46.11 implies that for any commutative ring R , the natural map

$$\Phi_* : H_*(C_*(X; R) \otimes C_*(Y; R)) \rightarrow H_*(X \times Y; R)$$

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 46.13 (topological Künneth formula). *For any principal ideal domain R , any spaces X, Y and every integer $n \geq 0$, the singular cross product fits into a natural short exact sequence*

$$\begin{aligned} 0 \longrightarrow \bigoplus_{k+\ell=n} H_k(X; R) \otimes H_\ell(Y; R) \xrightarrow{\times} H_n(X \times Y; R) \\ \longrightarrow \bigoplus_{k+\ell=n-1} \text{Tor}(H_k(X; R), H_\ell(Y; R)) \longrightarrow 0, \end{aligned}$$

and the sequence splits (but not naturally). \square

In particular, we can always choose field coefficients to make the Tor terms vanish:

COROLLARY 46.14. *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*

$$\times : \bigoplus_{k+\ell=n} H_k(X; \mathbb{K}) \otimes H_\ell(Y; \mathbb{K}) \xrightarrow{\cong} H_n(X \times Y; \mathbb{K}).$$

for every integer $n \geq 0$. \square

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 next week, after introducing the cohomology cup product.

46.5. Exercises.

EXERCISE 46.1. 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 $0 \leq k \leq n$.

EXERCISE 46.2. Recall that the topology of a CW-complex X is defined normally as the strongest topology for which the characteristic maps of all cells $\Phi_\alpha : \mathbb{D}^k \rightarrow X$ 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 $Z' \rightarrow Z$ is continuous.

*Remark: In general, the identity map $Z' \rightarrow Z$ might not be a homeomorphism!*⁷¹

- (b) Prove that the identity map $Z' \rightarrow Z$ 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.

EXERCISE 46.3. This problem is intended to elucidate in differential-geometric terms the intuitive reason behind the formula $\partial(e_\alpha^k \times e_\beta^\ell) = \partial e_\alpha^k \times e_\beta^\ell + (-1)^k e_\alpha^k \times \partial e_\beta^\ell$ stated in Proposition 46.1 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 $\text{GL}(n, \mathbb{R})$ 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.⁷² On an oriented vector space, we call a basis **positive** whenever it belongs to the equivalence class determined by the orientation. A linear isomorphism $V \rightarrow W$ 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** $T_x M$ 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 $f : \mathbb{R}^k \rightarrow \mathbb{R}^{k-n}$ for some $k \in \mathbb{N}$; a famous theorem of Whitney says that every smooth n -manifold can be described in this way if $k \geq 2n$. The tangent space $T_x M$ 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 $T_x M$ 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 $T_x M$ such that all of them are positive. If M and N are smooth manifolds of the same dimension, then any smooth map $f : M \rightarrow N$ has a derivative $df(x) : T_x M \rightarrow T_{f(x)} N$ 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 $f : M \rightarrow N$ is called **orientation preserving/reversing** if $df(x) : T_x M \rightarrow T_{f(x)} N$ is orientation preserving/reversing for every $x \in M$.

- (a) Convince yourself that S^2 admits an orientation (i.e. it is **orientable**), but $\mathbb{R}P^2$ 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

⁷¹This 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].

⁷²Dimension 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.

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 $M \times N \rightarrow N \times M : (x, y) \mapsto (y, x)$ 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 $T_xM \setminus T_x(\partial M)$ 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.

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 $S^k \rightarrow S^k$ 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 $\partial(e_\alpha^k \times e_\beta^\ell) = \partial e_\alpha^k \times e_\beta^\ell + (-1)^k e_\alpha^k \times \partial e_\beta^\ell$.

47. Products in singular cohomology

The main goal of this lecture is to define the cup product on singular cohomology and establish its basic properties, so that for any space X and any choice of commutative coefficient ring R , $H^*(X; R)$ becomes a graded-commutative ring with unit. Before that, we also have some loose ends to tie up regarding the cross product on homology, including its extension to cohomology. In this lecture, we will get more mileage than ever before out of the method of acyclic models: the basic pattern is that in the singular chain and cochain complexes, acyclic model arguments give rise to certain natural chain maps that are canonical up to chain homotopy, and letting these chain maps descend to (co)homology tells us things about the cross and cup products.

CONVENTION. In this lecture, the coefficient module for singular homology and cohomology is always—unless otherwise noted—a fixed commutative ring R with unit. Having a ring structure on the coefficient group is important for several reasons, which I will endeavor to point out explicitly whenever they are relevant. It will usually not be necessary to assume that R is a principal ideal domain.

47.1. Properties of the cross product. Recall that the singular cross product $H_k(X) \otimes H_\ell(Y) \xrightarrow{\times} H_{k+\ell}(X \times Y)$ was defined in the previous lecture in terms of a natural chain map

$$(47.1) \quad C_*(X) \otimes C_*(Y) \xrightarrow{\times := \Phi} C_*(X \times Y) ,$$

which we will refer to henceforth as the *chain-level* singular cross product. It gives rise to the homological cross product via the simple formula

$$[a] \times [b] := [a \times b] \in H_{k+\ell}(X \times Y)$$

for cycles $a \in C_k(X)$ and $b \in C_\ell(Y)$, where the homology class $[a \times b]$ is independent of choices because the chain-level cross product is canonically defined up to chain homotopy, even though the chain map itself depends on a multitude of arbitrary choices. These definitions do not depend in any essential way on the choice to use ring coefficients: the acyclic model argument behind the construction of (47.1) was carried out with coefficients in \mathbb{Z} , but one can immediately deduce from it the existence of a canonical (up to chain homotopy) chain map

$$C_*(X; G) \otimes C_*(Y; H) \xrightarrow{\times} C_*(X \times Y; G \otimes H)$$

for any pair of coefficient modules G, H , thus giving rise to a more general homological product of the form

$$(47.2) \quad H_k(X; G) \otimes H_\ell(Y; H) \xrightarrow{\times} H_{k+\ell}(X \times Y; G \otimes H).$$

This reduces to the case with arbitrary ring coefficients if one sets $G = H := R$ and uses the canonical isomorphism $R \otimes R \cong R$.

A crucial detail for which ring coefficients are important is that the chain map (47.1) is a chain homotopy equivalence, because another acyclic model argument carried out in Lemma 46.10 produces a chain homotopy inverse

$$(47.3) \quad C_*(X \times Y) \xrightarrow{\theta} C_*(X) \otimes C_*(Y) .$$

The construction of θ was again carried out with integer coefficients, and it used the fact that \mathbb{Z} is a principal ideal domain, because the acyclicity of the complex $C_*(\Delta^n; \mathbb{Z}) \otimes C_*(\Delta^n; \mathbb{Z})$ in positive degrees was deduced from the algebraic Künneth formula. But defining the map (47.3) over \mathbb{Z} immediately determines a similar definition over any coefficient ring R , due to the fact that $C_*(X \times Y; R)$ and $C_*(X; R) \otimes C_*(Y; R)$ are both free R -modules with the same set of generators as their counterparts over \mathbb{Z} ; this observation does not require R to be a principal ideal domain.

To take the discussion further, we shall now state three more results involving natural chain maps that are canonical up to homotopy, each of which takes the form of a diagram that *commutes up to chain homotopy*. Such statements always have the following meaning: for any two paths between the same pair of terms in the diagram, the associated compositions of chain maps are chain homotopic.

LEMMA 47.1. *For any three spaces X, Y, Z , the diagram*

$$\begin{array}{ccc} C_*(X) \otimes C_*(Y) \otimes C_*(Z) & \xrightarrow{\times \otimes \mathbb{1}} & C_*(X \times Y) \otimes C_*(Z) \\ \downarrow \mathbb{1} \otimes \times & & \downarrow \times \\ C_*(X) \otimes C_*(Y \times Z) & \xrightarrow{\times} & C_*(X \times Y \times Z) \end{array}$$

commutes up to chain homotopy.

PROOF. We leave the details as an exercise, but the point is to show via an acyclic model argument that natural chain maps $C_*(X) \otimes C_*(Y) \otimes C_*(Z) \rightarrow C_*(X \times Y \times Z)$ that take the form $x \otimes y \otimes z \mapsto (x, y, z)$ at the degree 0 level are unique up to chain homotopy. The result follows because $\times \circ (\times \otimes \mathbb{1})$ and $\times \circ (\mathbb{1} \otimes \times)$ are both examples of chain maps with these properties. \square

The two ways of proceeding from $C_*(X) \otimes C_*(Y) \otimes C_*(Z)$ to $C_*(X \times Y \times Z)$ in the diagram of Lemma 47.1 represent two ways of computing the chain-level cross product of three chains $a \in C_k(X)$, $b \in C_\ell(Y)$ and $c \in C_m(Z)$, either as $(a \times b) \times c$ or as $a \times (b \times c)$. The result does not say that these two products are equal, because the two chain maps are only chain homotopic, not identical. But the chain homotopy implies that if we take the three chains a, b, c to be cycles and care only about the homology class of the threefold product, we get equality, hence:

COROLLARY 47.2. *The cross product on singular homology is associative, i.e. for any $[a] \in H_k(X)$, $[b] \in H_\ell(Y)$ and $[c] \in H_m(Z)$, one has*

$$([a] \times [b]) \times [c] = [a] \times ([b] \times [c]) \in H_{k+\ell+m}(X \times Y \times Z).$$

\square

So much for associativity; let's talk about commutativity. A straightforward relation such as $[a] \times [b] = [b] \times [a]$ would not make sense, since $X \times Y$ and $Y \times X$ are not technically the same space, but of course there is a canonical homeomorphism between them, which we shall denote by

$$X \times Y \xrightarrow{\tau} Y \times X : (x, y) \mapsto (y, x).$$

On top of this detail, the following must be added: the obvious map $C_*(X) \otimes C_*(Y) \rightarrow C_*(Y) \otimes C_*(X)$ defined by $a \otimes b \mapsto b \otimes a$ is not a chain map, but it is an easy exercise to verify that it becomes a chain map if we introduce an appropriate sign, an accordance with the Koszul sign convention: the chain map we need will be denoted by

$$C_*(X) \otimes C_*(Y) \xrightarrow{-\sigma} C_*(Y) \otimes C_*(X) : a \otimes b \mapsto (-1)^{|a| \cdot |b|} b \otimes a.$$

The proof of the following lemma is then a straightforward acyclic model argument, showing the uniqueness of certain natural chain maps $C_*(X) \otimes C_*(Y) \rightarrow C_*(Y \times X)$ up to chain homotopy:

LEMMA 47.3. *For any two spaces X, Y , the diagram*

$$\begin{array}{ccc} C_*(X) \otimes C_*(Y) & \xrightarrow{\times} & C_*(X \times Y) \\ \cong \downarrow \sigma & & \cong \downarrow \tau_* \\ C_*(Y) \otimes C_*(X) & \xrightarrow{\times} & C_*(Y \times X) \end{array}$$

commutes up to chain homotopy.

\square

COROLLARY 47.4. *The cross product on singular homology is graded commutative in the sense that for any $[a] \in H_k(X)$ and $[b] \in H_\ell(Y)$,*

$$\tau_*([a] \times [b]) = (-1)^{k\ell} [b] \times [a].$$

\square

The next result gives us something akin to a “unit” element for the singular cross product. Note that for a one-point space $\{*\}$, there is a canonical R -module isomorphism

$$C_0(\{*\}) \cong R,$$

which we shall use in the following to identify R with a submodule of the chain complex $C_*(\{*\})$. With that understood, one can use the unit $1 \in R$ to define chain maps

$$\begin{aligned} C_*(X) &\xrightarrow{u_L} C_*(\{*\}) \otimes C_*(X) : a \mapsto 1 \otimes a, \\ C_*(X) &\xrightarrow{u_R} C_*(X) \otimes C_*(\{*\}) : a \mapsto a \otimes 1. \end{aligned}$$

LEMMA 47.5. *For any space X , the diagrams*

$$\begin{array}{ccc} C_*(X) & \xrightarrow{u_L} & C_*(\{*\}) \otimes C_*(X) \\ & \searrow \cong & \downarrow \times \\ & & C_*(\{*\} \times X) \end{array} \quad \text{and} \quad \begin{array}{ccc} C_*(X) & \xrightarrow{u_R} & C_*(X) \otimes C_*(\{*\}) \\ & \searrow \cong & \downarrow \times \\ & & C_*(X \times \{*\}) \end{array}$$

commute up to chain homotopy, where the isomorphisms indicated by the diagonal arrows are induced by the obvious homeomorphisms $\{*\} \times X \cong X \cong X \times \{*\}$.

PROOF. Acyclic models blablabla. □

COROLLARY 47.6. *Under the obvious identifications $\{*\} \times X = X = X \times \{*\}$, the homology class $[1] \in H_0(\{*\})$ represented by the unit $1 \in R \subset C_0(\{*\})$ defines a unit for the cross product on singular homology, in the sense that*

$$[1] \times [a] = [a] = [a] \times [1]$$

for all $[a] \in H_*(X)$. □

REMARK 47.7. With the exception of the special role played by the unit element $1 \in R$ in Corollary 47.6, none of the three corollaries stated above really depends on the use of ring coefficients; all three can be extended to statements about the cross product (47.2) on homology with arbitrary coefficients. We do not plan to make use of such a generalization in this course, and will thus leave the details as an exercise.

47.2. The cohomological cross product. In cellular cohomology with coefficients in the ring R , there is a cross product

$$H_{CW}^k(X) \otimes H_{CW}^\ell(Y) \xrightarrow{\times} H_{CW}^{k+\ell}(X \times Y) : [\varphi] \otimes [\psi] \mapsto [\varphi] \times [\psi] := [\varphi \times \psi]$$

determined uniquely for any two CW-complexes X, Y by the condition that for all $\varphi \in C_{CW}^k(X)$, $\psi \in C_{CW}^\ell(Y)$, $a \in C_k^{CW}(X)$ and $b \in C_\ell^{CW}(Y)$,

$$(47.4) \quad (\varphi \times \psi)(a \times b) = (-1)^{|a| \cdot |\psi|} \varphi(a) \psi(b) \in R.$$

Here, the sign is motivated by the Koszul sign convention, and we will see below why it must appear in order for certain natural maps to satisfy the chain map condition. Defining a cochain-level cross product via (47.4) uses the ring structure of R in two ways: most obviously, that structure is the reason why the product $\varphi(a)\psi(b) \in R$ is well defined. More subtly, (47.4) uniquely determines $\varphi \times \psi \in C_{CW}^{k+\ell}(X \times Y)$ due to the fact that for ring coefficients, the chain-level cellular cross product $C_*^{CW}(X) \otimes C_*^{CW}(Y) \rightarrow C_*^{CW}(X \times Y)$ is an isomorphism, so that evaluating $\varphi \times \psi$ on $(k + \ell)$ -chains of the form $a \times b$ determines its evaluation on all other $(k + \ell)$ -chains in $X \times Y$. For a more direct formula, we could denote the inverse of the chain-level cross product by $\theta : C_*^{CW}(X \times Y) \rightarrow C_*^{CW}(X) \otimes C_*^{CW}(Y)$ and write

$$(\varphi \times \psi)(c) := (\varphi \otimes \psi) \circ \theta(c)$$

for $c \in C_*^{CW}(X \times Y)$, where the evaluation of $\varphi \otimes \psi \in C_{CW}^*(X) \otimes C_{CW}^*(Y)$ on elements of $C_*^{CW}(X) \otimes C_*^{CW}(Y)$ is determined by the formula

$$(47.5) \quad (\varphi \otimes \psi)(a \otimes b) := (-1)^{|a| \cdot |\psi|} \varphi(a) \psi(b) \in R,$$

with the understanding that evaluations such as $\varphi(a)$ vanish for any homogeneous elements $\varphi \in C_{CW}^*(X)$ and $a \in C_*^{CW}(X)$ whose degrees do not match. This notion of evaluation determines a natural homomorphism

$$(47.6) \quad C_{CW}^*(X) \otimes C_{CW}^*(Y) \xrightarrow{F} \text{Hom}(C_*^{CW}(X) \otimes C_*^{CW}(Y), R).$$

We can now explain the reason for the sign in (47.5), which gives rise to the sign in (47.4): this sign makes (47.6) into a chain map, which it would not otherwise be. The cochain-level cellular cross product is therefore a composition of two chain maps

$$C_{CW}^*(X) \otimes C_{CW}^*(Y) \xrightarrow{F} \text{Hom}(C_*^{CW}(X) \otimes C_*^{CW}(Y), R) \xrightarrow{\theta^*} \text{Hom}(C_*^{CW}(X \times Y), R) = C_{CW}^*(X \times Y),$$

and is thus a chain map in itself, which is why it descends to a well-defined operation $[\varphi] \times [\psi]$ on cohomology classes.

In singular homology, the chain-level cross product is not an isomorphism, but it is a chain homotopy equivalence, and it therefore makes sense to use its chain homotopy inverse θ in the same way that its cellular counterpart was used above. This leads to a **cochain-level singular cross product**

$$(47.7) \quad C^*(X) \otimes C^*(Y) \xrightarrow{F} \text{Hom}(C_*(X) \otimes C_*(Y), R) \xrightarrow{\theta^*} \text{Hom}(C_*(X \times Y), R) = C^*(X \times Y),$$

where F denotes the canonical chain map

$$C^*(X) \otimes C^*(Y) \xrightarrow{F} \text{Hom}(C_*(X) \otimes C_*(Y), R)$$

defined via the evaluation pairing in (47.5). Since \times is a chain map, it gives rise to a well-defined cross product on singular cohomology with ring coefficients,

$$H^k(X) \otimes H^\ell(Y) \xrightarrow{\times} H^{k+\ell}(X \times Y) : [\varphi] \otimes [\psi] \mapsto [\varphi] \times [\psi] := [\varphi \times \psi].$$

THEOREM 47.8. *The cross product on singular cohomology satisfies the obvious analogues of the associativity, graded commutativity and unit properties stated for the homological cross product in Corollaries 47.2, 47.4 and 47.6, with the role of the unit element played by $[1] \in H^0(\{*\})$ for $1 \in R \cong C^0(\{*\})$.*

PROOF. Since the chain-level cross product is a chain homotopy equivalence, one can invert and then dualize the diagrams in Lemmas 47.1, 47.3 and 47.5 to produce similar diagrams for the cochain-level cross product, namely

$$\begin{array}{ccc}
 C^*(X) \otimes C^*(Y) \otimes C^*(Z) & \xrightarrow{\times \otimes \mathbb{1}} & C^*(X \times Y) \otimes C^*(Z) & & C^*(X) \otimes C^*(Y) & \xrightarrow{\times} & C^*(X \times Y) \\
 \downarrow \mathbb{1} \otimes \times & & \downarrow \times & , & \cong \downarrow \sigma & & \cong \uparrow \tau^* \\
 C^*(X) \otimes C^*(Y \times Z) & \xrightarrow{\times} & C^*(X \times Y \times Z) & & C^*(Y) \otimes C^*(X) & \xrightarrow{\times} & C^*(Y \times X)
 \end{array}$$

$$\begin{array}{ccc}
 C^*(X) & \xleftarrow{u_L^* \circ F} & C^*(\{*\}) \otimes C^*(X) \\
 \searrow \cong & & \downarrow \times \\
 & & C^*(\{*\} \times X)
 \end{array}
 \qquad
 \begin{array}{ccc}
 C^*(X) & \xleftarrow{u_R^* \circ F} & C^*(X) \otimes C^*(\{*\}) \\
 \searrow \cong & & \downarrow \times \\
 & & C^*(X \times \{*\})
 \end{array}
 ,$$

where $\sigma : C^*(X) \otimes C^*(Y) \rightarrow C^*(Y) \otimes C^*(X)$ is the chain map $\varphi \otimes \psi \mapsto (-1)^{|\varphi| \cdot |\psi|} \psi \otimes \varphi$. After descending to cohomology, these diagrams imply associativity, graded commutativity, and the formulas $[1] \times [\varphi] = [\varphi]$ and $[\varphi] \times [1] = [\varphi]$ respectively. \square

In contrast to cellular (co)homology, the relation (47.4) between the homological and cohomological cross products is not strictly true at the chain level—indeed, a precise relation of this form makes no sense when the (co)chain-level cross products are defined only up to chain homotopy. The next lemma says, however, that the relation holds up to chain homotopy. This is where the particular sign choice made in the definition of the cochain complex $C^*(X) = \text{Hom}(C_*(X), R)$ in Lecture 41 becomes important, as it allows us to interpret the evaluation pairing

$$\langle , \rangle : C^*(X) \otimes C_*(X) \rightarrow R_*$$

as a chain map, after viewing $C^*(X)$ as a chain complex with its grading inverted, and defining R_* to be the chain complex that has R in degree 0 and trivial modules in all other degrees (cf. Remark 41.3). Note that since the R -modules R and $R \otimes R$ are canonically isomorphic, the same is true of the chain complexes R_* and $R_* \otimes R_*$.

LEMMA 47.9. *For any two spaces X, Y , the diagram*

$$\begin{array}{ccc}
 C^*(X) \otimes C^*(Y) \otimes C_*(X) \otimes C_*(Y) & \xrightarrow{\times \otimes \times} & C^*(X \times Y) \otimes C_*(X \times Y) \\
 \downarrow \mathbf{1} \otimes \sigma \otimes \mathbf{1} & & \downarrow \langle , \rangle \\
 C^*(X) \otimes C_*(X) \otimes C^*(Y) \otimes C_*(Y) & \xrightarrow{\langle , \rangle \otimes \langle , \rangle} & R_* \otimes R_* \\
 & & \parallel \\
 & & R_*
 \end{array}$$

commutes up to chain homotopy, where $\sigma : C^*(Y) \otimes C_*(X) \rightarrow C_*(X) \otimes C^*(Y)$ denotes the chain map $\psi \otimes a \mapsto (-1)^{|a| \cdot |\psi|} a \otimes \psi$.

PROOF. The proof is a straightforward computation using the definition of the cochain-level cross product in (47.7) and the fact that θ composed with the chain-level cross product is chain homotopic to the identity. \square

COROLLARY 47.10. *The cross products on singular homology and cohomology are related by the formula*

$$\langle [\varphi] \times [\psi], [a] \times [b] \rangle = (-1)^{|a| \cdot |\psi|} \langle [\varphi], [a] \rangle \cdot \langle [\psi], [b] \rangle$$

for homogeneous elements $[\varphi] \in H^*(X)$, $[\psi] \in H^*(Y)$, $[a] \in H_*(X)$ and $[b] \in H_*(Y)$. \square

47.3. The cup product. The cup product on $H^*(X)$ with coefficients in a ring R can be defined in terms of the cross product, but it is more general and more useful to regard it as yet another structure arising from a natural chain map that is canonical up to chain homotopy. The following theorem is proved by a succession of straightforward acyclic models arguments, in which one can (as usual) specialize to the case of integer coefficients and then deduce from it the generalization to coefficients in R . As a preliminary remark, note that for any space X , there is a natural inclusion of R -modules

$$R \hookrightarrow C^0(X) = C^0(X; R),$$

identifying each $r \in R$ with the cochain whose value on every singular 0-simplex is r . In slightly fancier terms, the unique map $\epsilon : X \rightarrow \{*\}$ admits a right-inverse, implying that the induced chain map $\epsilon^* : C^*(\{*\}) \rightarrow C^*(X)$ is injective, and the inclusion described above identifies R with the image of the injection $\epsilon^* : R = C^0(\{*\}) \hookrightarrow C^0(X)$. Regarding each $r \in R$ in this way as a 0-cochain, it is also a 0-cocycle, giving rise to a natural inclusion

$$R \hookrightarrow H^0(X) = H^0(X; R) : r \mapsto [r].$$

In particular, the unit element $1 \in R$ is now identified with the homomorphism

$$C_0(X) \xrightarrow{1} R : \sum_i r_i \sigma_i \mapsto \sum_i r_i \quad \text{for } r_i \in R \text{ and } \sigma_i : \Delta^0 \rightarrow X,$$

which you may recognize as the *augmentation* that is used in constructing the augmented chain complex $\tilde{C}_*(X; R)$. Extending this map trivially to chains of nonzero degrees makes it a chain map

$$C_*(X) \xrightarrow{1} R_*.$$

THEOREM 47.11. *One can associate to every space X a chain map $C_*(X) \xrightarrow{\Psi} C_*(X) \otimes C_*(X)$ that takes the form $\Psi(x) = x \otimes x$ on generators of degree 0 and is natural in the sense that any map $f : X \rightarrow Y$ gives rise to a commutative diagram*

$$\begin{array}{ccc} C_*(X) & \xrightarrow{\Psi} & C_*(X) \otimes C_*(X) \\ \downarrow f_* & & \downarrow f_* \otimes f_* \\ C_*(Y) & \xrightarrow{\Psi} & C_*(Y) \otimes C_*(Y) \end{array} .$$

Moreover, these two properties determine Ψ uniquely up to chain homotopy, and the following diagrams commute up to chain homotopy:

$$\begin{array}{ccc} C_*(X) & \xrightarrow{\Psi} & C_*(X) \otimes C_*(X) \\ \downarrow \Psi & & \downarrow \Psi \otimes 1 \\ C_*(X) \otimes C_*(X) & \xrightarrow{1 \otimes \Psi} & C_*(X) \otimes C_*(X) \otimes C_*(X) \end{array} , \quad \begin{array}{ccc} C_*(X) & \xrightarrow{\Psi} & C_*(X) \otimes C_*(X) \\ & \searrow \Psi & \downarrow \sigma \\ & & C_*(X) \otimes C_*(X) \end{array} ,$$

$$\begin{array}{ccc} C_*(X) & \xrightarrow{\Psi} & C_*(X) \otimes C_*(X) \\ \downarrow \cong & \swarrow 1 \otimes 1 & \\ R_* \otimes C_*(X) & & \end{array} . \quad \begin{array}{ccc} C_*(X) & \xrightarrow{\Psi} & C_*(X) \otimes C_*(X) \\ \downarrow \cong & \swarrow 1 \otimes 1 & \\ C_*(X) \otimes R_* & & \end{array} .$$

Here, the chain map σ is defined as in Lemma 47.3, and the downward arrows in the last two diagrams represent the canonical isomorphisms between $C_*(X)$ and its tensor products with R_* . \square

The chain map $\Psi : C_*(X) \rightarrow C_*(X) \otimes C_*(X)$ in Theorem 47.11 is called a **diagonal approximation**.

DEFINITION 47.12. Given a diagonal approximation Ψ , the **cochain-level cup product**

$$C^*(X) \otimes C^*(X) \xrightarrow{\cup} C^*(X) : \varphi \otimes \psi \mapsto \varphi \cup \psi$$

on the singular cochain complex of a space X with coefficients in a commutative ring R is the chain map defined as the composition

$$C^*(X) \otimes C^*(X) \xrightarrow{F} \text{Hom}(C_*(X) \otimes C_*(X), R) \xrightarrow{\Psi^*} \text{Hom}(C_*(X), R) = C^*(X) .$$

This determines the **cup product** on singular cohomology,

$$[\varphi] \cup [\psi] := [\varphi \cup \psi] \in H^{k+\ell}(X) \quad \text{for} \quad [\varphi] \in H^k(X), [\psi] \in H^\ell(X).$$

As you've hopefully come to expect by now, the cochain-level cup product is dependent on choices, but as a chain map $C^*(X) \otimes C^*(X) \rightarrow C^*(X)$ it is independent of those choices up to chain homotopy, implying that the induced product on cohomology is canonically defined.

As a slightly more direct formula for the cochain-level cup product, we can write

$$(\varphi \cup \psi)(c) = (\varphi \otimes \psi) \circ \Psi(c) \quad \text{for} \quad \varphi \in C^k(X), \psi \in C^\ell(X), c \in C_{k+\ell}(X).$$

This still seems a bit abstract if one only has a general acyclic models argument to provide the diagonal approximation Ψ , but it is also possible to write down explicit examples of diagonal approximations. A favorite choice for this is known as the **Alexander-Whitney** diagonal approximation, and can be 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 $0 \leq j_0 < j_1 < \dots < j_k \leq n$, let

$$[j_0, \dots, j_k] \subset \Delta^n$$

denote the k -dimensional face of Δ^n spanned by the vertices j_0, \dots, j_k , which is identified naturally with the standard k -simplex. For instance, in this notation, the j th 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 $\Psi : C_*(X) \rightarrow C_*(X) \otimes C_*(X)$ via the formula

$$\Psi(\sigma) := \sum_{k+\ell=n} (\sigma|_{[0, \dots, k]}) \otimes (\sigma|_{[k, \dots, n]})$$

for each singular n -simplex $\sigma : \Delta^n \rightarrow X$. It is a straightforward exercise to verify that this satisfies the conditions of a diagonal approximation.

Plugging the Alexander-Whitney approximation into $\varphi \cup \psi = (\varphi \otimes \psi) \circ \Psi$ gives the following formula for the cup product of cochains: for any singular n -simplex $\sigma : \Delta^n \rightarrow X$ with $n = k + \ell$,⁷³

$$(\varphi \cup \psi)(\sigma) = (-1)^{k\ell} \varphi(\sigma|_{[0, \dots, k]}) \psi(\sigma|_{[k, \dots, n]}).$$

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 that 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.

THEOREM 47.13. *The cup product $\cup : H^*(X) \otimes H^*(X) \rightarrow H^*(X)$ on cohomology with coefficients in the ring R has the following properties.*

(1) *It is natural: for all continuous maps $f : Y \rightarrow X$ and $[\varphi], [\psi] \in H^*(X)$,*

$$f^*([\varphi] \cup [\psi]) = f^*[\varphi] \cup f^*[\psi].$$

(2) *It is associative: for all $[\varphi], [\psi], [\eta] \in H^*(X)$,*

$$([\varphi] \cup [\psi]) \cup [\eta] = [\varphi] \cup ([\psi] \cup [\eta]).$$

⁷³The 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 (41.1) for the definition of coboundary maps.

(3) It is graded commutative: for all $[\varphi] \in H^k(X)$ and $[\psi] \in H^\ell(X)$,

$$[\varphi] \cup [\psi] = (-1)^{k\ell} [\psi] \cup [\varphi].$$

(4) It has a unit: under the canonical inclusion $R \hookrightarrow H^0(X) : r \mapsto [r]$, the unit $1 \in R$ and any $[\varphi] \in H^*(X)$ satisfy

$$[1] \cup [\varphi] = [\varphi] \cup [1] = [\varphi].$$

(5) It determines the cohomological cross product for two spaces X, Y via the formula

$$[\varphi] \times [\psi] = \pi_X^*[\varphi] \cup \pi_Y^*[\psi]$$

for $[\varphi] \in H^*(X)$ and $[\psi] \in H^*(Y)$, where $\pi_X : X \times Y \rightarrow X$ and $\pi_Y : X \times Y \rightarrow Y$ are the natural projections.

(6) It is determined by the cohomological cross product via the formula

$$[\varphi] \cup [\psi] = d^*([\varphi] \times [\psi])$$

for $[\varphi], [\psi] \in H^*(X)$, where $d : X \rightarrow X \times X$ denotes the diagonal map $x \mapsto (x, x)$.

PROOF. Naturality is a consequence of the naturality property of diagonal approximations and of the canonical chain map $F : C^*(X) \otimes C^*(X) \rightarrow \text{Hom}(C_*(X) \otimes C_*(X), R)$. Associativity, graded commutativity and the unit property can all be deduced from diagrams of chain maps that commute up to chain homotopy, and are constructed by dualizing the diagrams in Theorem 47.11, namely,

$$\begin{array}{ccc} C^*(X) & \xleftarrow{\cup} & C^*(X) \otimes C^*(X) \\ \uparrow \cup & & \uparrow \cup \mathbb{1} \\ C^*(X) \otimes C^*(X) & \xleftarrow{\mathbb{1} \otimes \cup} & C^*(X) \otimes C^*(X) \otimes C^*(X) \end{array} \qquad \begin{array}{ccc} C^*(X) & \xleftarrow{\cup} & C^*(X) \otimes C^*(X) \\ \swarrow \cup & & \downarrow \sigma \\ & & C^*(X) \otimes C^*(X) \end{array}$$

$$\begin{array}{ccc} C^*(X) & \xrightarrow{\cup \mathbb{1}} & C^*(X) \\ \mathbb{1} & \xrightarrow{\quad} & \mathbb{1} \end{array} \qquad \begin{array}{ccc} C^*(X) & \xrightarrow{\mathbb{1} \cup \cdot} & C^*(X) \\ \mathbb{1} & \xrightarrow{\quad} & \mathbb{1} \end{array}$$

where $\sigma : C^*(X) \otimes C^*(X) \rightarrow C^*(X) \otimes C^*(X)$ denotes the chain map $\varphi \otimes \psi \mapsto (-1)^{|\varphi||\psi|} \psi \otimes \varphi$.

For the two properties relating the cup and cross products, we need two more diagrams involving diagonal approximations, which can be shown to commute up to chain homotopy by straightforward acyclic model arguments, namely

$$\begin{array}{ccc} C_*(X \times Y) & \xrightarrow{\Psi} & C_*(X \times Y) \otimes C_*(X \times Y) \\ & \searrow \theta & \downarrow (\pi_X)_* \otimes (\pi_Y)_* \\ & & C_*(X) \otimes C_*(Y) \end{array} \qquad \begin{array}{ccc} C_*(X) & \xrightarrow{\Psi} & C_*(X) \otimes C_*(X) \\ & \searrow d_* & \uparrow \theta \\ & & C_*(X \times X) \end{array}$$

This is an unimportant detail, but it may be amusing to note that each of these diagrams implies the other (exercise!), so it is only really necessary to carry out one acyclic model argument here. What's more important is that both diagrams can be dualized, producing the diagrams

$$\begin{array}{ccc} C^*(X \times Y) & \xleftarrow{\cup} & C^*(X \times Y) \otimes C^*(X \times Y) \\ \swarrow \times & & \uparrow \pi_X^* \otimes \pi_Y^* \\ & & C^*(X) \otimes C^*(Y) \end{array} \qquad \begin{array}{ccc} C^*(X) & \xleftarrow{\cup} & C^*(X) \otimes C^*(X) \\ \swarrow d^* & & \downarrow \times \\ & & C^*(X \times X) \end{array}$$

which imply the last two properties. □

47.4. Exercises.

EXERCISE 47.1 (*). Prove the formula

$$([\alpha] \cup [\varphi]) \times ([\beta] \cup [\psi]) = (-1)^{|\varphi| \cdot |\beta|} ([\alpha] \times [\beta]) \cup ([\varphi] \times [\psi])$$

for homogeneous elements $[\alpha], [\varphi] \in H^*(X)$ and $[\beta], [\psi] \in H^*(Y)$.

EXERCISE 47.2. In this exercise we compute the cohomology ring $H^*(\mathbb{T}^n; R)$ for every $n \geq 1$, with coefficients in any commutative ring R with unit. The idea is to derive the cup product from information about the homological cross product

$$H_k(\mathbb{T}^m; R) \otimes H_\ell(\mathbb{T}^n; R) \xrightarrow{\times} H_{k+\ell}(\mathbb{T}^{m+n}; R)$$

obtained via the Künneth formula.

The homology of \mathbb{T}^n is fairly easy to compute because $\mathbb{T}^n = S^1 \times \dots \times S^1$ has a natural structure as a product cell complex (cf. Exercise 46.1). 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

$$[*] \in H_0(S^1; \mathbb{Z}), \quad [S^1] \in H_1(S^1; \mathbb{Z}),$$

i.e. $[*]$ is the homology class represented by any singular 0-simplex $\Delta^0 \rightarrow S^1$, and $[S^1]$ is the class represented by the identity map $S^1 \rightarrow S^1$ under the isomorphism $H_1(S^1; \mathbb{Z}) \cong \pi_1(S^1)$.

(a) Derive from the Künneth formula an isomorphism

$$H_m(\mathbb{T}^n; \mathbb{Z}) \cong H_{m-1}(\mathbb{T}^{n-1}; \mathbb{Z}) \oplus H_m(\mathbb{T}^{n-1}; \mathbb{Z})$$

for every $n \geq 2$ and $m \in \mathbb{Z}$. Deduce that $H_m(\mathbb{T}^n; \mathbb{Z})$ is always a finitely-generated abelian group, whose rank is an entry in Pascal's triangle,

$$\text{rank } H_m(\mathbb{T}^n; \mathbb{Z}) = \binom{n}{m}.$$

Hint: Since you already know that $H_(S^1; \mathbb{Z})$ is finitely generated and free, you can prove by induction on $n \in \mathbb{N}$ that the same is true for $H_*(\mathbb{T}^n; \mathbb{Z})$; this should remove any need to worry about Tor terms.*

(b) For each $m \in \mathbb{N}$ and each choice of integers $1 \leq j_1 < \dots < j_m \leq n$, define the homology class

$$e_{j_1, \dots, j_m} := A_1 \times \dots \times A_n \in H_m(\mathbb{T}^n; \mathbb{Z})$$

by setting $A_{j_i} := [S^1]$ for each $i = 1, \dots, m$ and $A_j := [*]$ for all other $j \in \{1, \dots, n\}$. Deduce from the Künneth formula that the set of all such elements forms a basis of $H_*(\mathbb{T}^n; \mathbb{Z})$.

It will be useful to have an alternative description of the degree 1 generators $e_j \in H_1(\mathbb{T}^n; \mathbb{Z})$ that appear in part (b). Pick a base point $t_0 \in S^1$ and consider the embedding

$$(47.8) \quad i_j : S^1 \hookrightarrow \mathbb{T}^n : x \mapsto (\underbrace{t_0 \times \dots \times t_0}_{j-1}, x, \underbrace{t_0 \times \dots \times t_0}_{n-j}).$$

Note that different choices of the base point $t_0 \in S^1$ give homotopic maps $i_j : S^1 \rightarrow \mathbb{T}^n$, thus the induced map $(i_j)_* : H_*(S^1; \mathbb{Z}) \rightarrow H_*(\mathbb{T}^n; \mathbb{Z})$ is independent of this choice.

(c) Show that for each $j = 1, \dots, n$, $(i_j)_*[S^1] = e_j$.

Hint: For the case $j = n$, you can identify S^1 in the obvious way with $\{\} \times S^1$ and then write $i_n : S^1 \hookrightarrow \mathbb{T}^n$ as*

$$i_n = \iota \times \text{Id} : \{*\} \times S^1 \hookrightarrow \mathbb{T}^{n-1} \times S^1,$$

with $\iota : \{*\} \rightarrow \mathbb{T}^{n-1}$ denoting the inclusion of the point (t_0, \dots, t_0) . The naturality of the cross product then gives a commutative diagram

$$\begin{array}{ccc} H_0(\{*\}; \mathbb{Z}) \otimes H_1(S^1; \mathbb{Z}) & \xrightarrow{\times} & H_1(S^1; \mathbb{Z}) \\ \downarrow \iota_* \otimes \mathbf{1} & & \downarrow (i_n)_* \\ H_0(\mathbb{T}^{n-1}; \mathbb{Z}) \otimes H_1(S^1; \mathbb{Z}) & \xrightarrow{\times} & H_1(\mathbb{T}^n; \mathbb{Z}). \end{array}$$

Use the fact that $[*] \times [S^1] = [S^1]$ for the canonical generator $[*] \in H_0(\{*\}; \mathbb{Z})$, after identifying $\{*\} \times S^1 = S^1$, while $\iota_* : H_0(\{*\}; \mathbb{Z}) \rightarrow H_0(\mathbb{T}^{n-1}; \mathbb{Z})$ is an isomorphism relating the canonical generators.

- (d) Use the universal coefficient theorem to upgrade the computation of $H_*(\mathbb{T}^n; \mathbb{Z})$ above to a computation of $H_*(\mathbb{T}^n; G)$ for arbitrary coefficient groups G . Deduce in particular that for any commutative ring R with unit, $H_*(\mathbb{T}^n; R)$ has the structure of a finitely-generated free R -module, with a basis in bijective correspondence with the basis of $H_*(\mathbb{T}^n; \mathbb{Z})$ described above.

Remark: There is no need to assume that R is a principal ideal domain here, as we are only using the universal coefficient theorem and Künneth formula for \mathbb{Z} -modules.

- (e) Show that for each $m \in \mathbb{Z}$, the natural map

$$H^m(\mathbb{T}^n; R) \rightarrow \text{Hom}(H_m(\mathbb{T}^n; R), R) : [\varphi] \mapsto \langle [\varphi], \cdot \rangle$$

is an R -module isomorphism, implying $H^m(\mathbb{T}^n; R) \cong R^{\binom{n}{m}}$.

Hint: Again, you do not need to assume here that R is a principal ideal domain. Just regard R at first as an abelian group, prove that the natural map from $H^m(\mathbb{T}^n; R)$ to the abelian group of group homomorphisms $H_m(\mathbb{T}^n; \mathbb{Z}) \rightarrow R$ is a group isomorphism, and then use what you know about the algebraic structure of $H_m(\mathbb{T}^n; R)$.

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, regarding these groups as R -modules. We can write down a canonical basis for $H^*(\mathbb{T}^n)$ 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 $1 \leq j_1 < \dots < j_m \leq n$, define

$$\lambda_{j_1, \dots, j_m} := \alpha_1 \times \dots \times \alpha_n \in H^m(\mathbb{T}^n),$$

where we choose $\alpha_{j_i} := \lambda$ for each $i = 1, \dots, m$ and $\alpha_j = [1] \in H^0(S^1)$ for all other $j \in \{1, \dots, n\}$. By Corollary 47.10, we have

$$\begin{aligned} \langle \lambda_{j_1, \dots, j_m}, e_{k_1, \dots, k_m} \rangle &= \langle \alpha_1 \times \dots \times \alpha_n, A_1 \times \dots \times A_n \rangle = \pm \langle \alpha_1, A_1 \rangle \dots \langle \alpha_n, A_n \rangle \\ &= \begin{cases} \pm 1 & \text{if } j_i = k_i \text{ for all } i = 1, \dots, m, \\ 0 & \text{otherwise,} \end{cases} \end{aligned}$$

proving that the collection of classes $\lambda_{j_1, \dots, j_m}$ for all choices $1 \leq j_1 < \dots < j_m \leq n$ 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

$$\pi_j : \mathbb{T}^n \rightarrow S^1 : (x_1, \dots, x_n) \mapsto x_j,$$

which is related to the inclusions $i_j : S^1 \hookrightarrow \mathbb{T}^n$ defined in (47.8) above by

$$\pi_j \circ i_k = \begin{cases} \text{Id} : S^1 \rightarrow S^1 & \text{if } j = k, \\ \text{constant} & \text{if } j \neq k. \end{cases}$$

(f) Show that $\pi_j^* \lambda = \lambda_j$ for each $j = 1, \dots, n$.

Hint: Evaluate both $\pi_j^ \lambda$ and λ_j on the generators $e_i \in H_1(\mathbb{T}^n)$.*

(g) Use Theorem 47.13(5) to prove that for any $m \in \mathbb{N}$ and integers $1 \leq j_1 < \dots < j_m \leq n$,

$$\lambda_{j_1} \cup \dots \cup \lambda_{j_m} = \lambda_{j_1, \dots, j_m}.$$

Conclude that the ring $H^*(\mathbb{T}^n; R)$ is isomorphic to the exterior algebra $\Lambda_R[\lambda_1, \dots, \lambda_n]$ over R on n generators of degree 1.

48. Relative cross, cup and cap products

We have two issues to address before we can wrap up the discussion of products: first, there is a product that intertwines homology and cohomology, called the *cap* product, which will be needed in order to define the Poincaré duality isomorphism in a few lectures. Second, the entire discussion of products so far has been restricted to *absolute* singular homology and cohomology, but we will occasionally also need products in relative (co)homology. The latter requires us to dust off the notion of *excisive couples*, which appeared in Lecture 34 in the context of the Mayer-Vietoris sequence.

CONVENTION. As in the previous lecture, the default assumption here will be that we use a fixed commutative ring R with unit for the coefficients in singular homology and cohomology.

48.1. The cap product. The previous lecture stated several results about the existence and uniqueness up to chain homotopy of natural chain maps, but there are still one or two interesting results of this type that we have not seen. As preparation, recall that the evaluation pairing

$$\langle \cdot, \cdot \rangle : C^*(X) \otimes C_*(X) \rightarrow R_* : \varphi \otimes c \mapsto \langle \varphi, c \rangle,$$

defined by $\langle \varphi, c \rangle := \varphi(c)$ when $|\varphi| = |c|$ and $\langle \varphi, c \rangle := 0$ otherwise, is a chain map if we regard $C^*(X)$ as a chain complex by inserting minus signs in front of the degrees of homogeneous elements. This convention will help you remember the correct subscripts to write in the following definition, which generalizes the evaluation pairing.

DEFINITION 48.1. Given a diagonal approximation $\Psi : C_*(X) \rightarrow C_*(X) \otimes C_*(X)$, the **chain-level cap product**

$$C^*(X) \otimes C_*(X) \xrightarrow{\cap} C_*(X) : \varphi \otimes c \mapsto \varphi \cap c$$

on a space X is the chain map defined as the composition

$$\begin{array}{ccc} C^*(X) \otimes C_*(X) & \xrightarrow{\mathbb{1} \otimes \Psi} & C^*(X) \otimes C_*(X) \otimes C_*(X) \xrightarrow{\sigma \otimes \mathbb{1}} & C_*(X) \otimes C^*(X) \otimes C_*(X) \\ & & & \downarrow \mathbb{1} \otimes \langle \cdot, \cdot \rangle \\ & \searrow \cap & & C_*(X) \otimes R_* = C_*(X) \end{array}$$

where $\sigma : C^*(X) \otimes C_*(X) \rightarrow C_*(X) \otimes C^*(X)$ denotes the chain map $\varphi \otimes c \mapsto (-1)^{|c| \cdot |\varphi|} c \otimes \varphi$. This determines the **cap product** on singular cohomology and homology,

$$H^k(X) \otimes H_\ell(X) \xrightarrow{\cap} H_{\ell-k}(X) : [\varphi] \otimes [c] \mapsto [\varphi] \cap [c] := [\varphi \cap c].$$

As usual, the homological cap product is independent of choices because the chain maps appearing in its chain-level variant are canonical up to chain homotopy.

Before continuing, let's unpack the definition of the chain-level cap product. Given a cochain $\varphi \in C^k(X)$ and a chain $c \in C_\ell(X)$, the result of plugging c into a diagonal approximation can be written as a finite sum

$$\Psi(c) = \sum_i a_i \otimes b_i \in C_*(X) \otimes C_*(X),$$

in which the degrees of the individual terms a_i and b_i may vary but always satisfy $|a_i| + |b_i| = \ell$. To obtain $\varphi \cap c$, we carry out the composition

$$\varphi \otimes c \mapsto \varphi \otimes \Psi(c) = \sum_i (\varphi \otimes a_i \otimes b_i) \mapsto \sum_i (-1)^{|a_i| \cdot |\varphi|} a_i \otimes \varphi \otimes b_i \mapsto \sum_i (-1)^{|a_i| \cdot |\varphi|} \langle \varphi, b_i \rangle a_i,$$

where in the last term, $\langle \varphi, b_i \rangle$ can only be nonzero when $|b_i| = k$, implying $|a_i| = \ell - k$ so that the resulting chain is homogeneous with degree $\ell - k$. If we apply the Koszul sign convention to define $\mathbb{1} \otimes \varphi : C_*(X) \otimes C_*(X) \rightarrow C_*(X) \otimes R_* = C_*(X)$ by

$$(\mathbb{1} \otimes \varphi)(a \otimes b) := (-1)^{|a| \cdot |\varphi|} \mathbb{1}(a) \langle \varphi, b \rangle,$$

we obtain the slightly more succinct formula

$$(48.1) \quad \varphi \cap c = (\mathbb{1} \otimes \varphi) \circ \Psi(c).$$

The real motivation to define the cap product in this way is that it satisfies the following chain-level relation with the cup product:

$$(48.2) \quad \langle \psi \cup \varphi, c \rangle = \langle \psi, \varphi \cap c \rangle$$

for any $\psi, \varphi \in C^*(X)$ and $c \in C_*(X)$. Indeed, in the simple case where $\Psi(c)$ takes the form $a \otimes b$ for some $a, b \in C_*(X)$, one computes $\varphi \cap c = (-1)^{|a| \cdot |\varphi|} \langle \varphi, b \rangle a$ and thus

$$\langle \psi, \varphi \cap c \rangle = (-1)^{|a| \cdot |\varphi|} \langle \psi, a \rangle \langle \varphi, b \rangle = (\psi \otimes \varphi)(a \otimes b) = (\psi \otimes \varphi) \circ \Psi(c) = \langle \psi \otimes \varphi, c \rangle.$$

It follows immediately via linearity that the relation also holds in the general case with $\Psi(c)$ of the form $\sum_i a_i \otimes b_i$. One can show in fact that the chain-level cup product uniquely determines the chain-level cap product via the relation (48.2), and of course, the relation also holds at the (co)homology level:

$$(48.3) \quad \langle [\psi] \cup [\varphi], [c] \rangle = \langle [\psi], [\varphi] \cap [c] \rangle \quad \text{for} \quad [\varphi], [\psi] \in H^*(X), [c] \in H_*(X).$$

We can now clarify in what sense the cap product generalizes the evaluation pairing: if the classes $[\varphi] \in H^*(X)$ and $[c] \in H_*(X)$ have the same degree, then taking $[\psi] := [1] \in H^0(X)$ in the latter relation gives

$$(48.4) \quad \langle [1], [\varphi] \cap [c] \rangle = \langle [\varphi], [c] \rangle.$$

Note that in this situation, $[\varphi] \cap [c]$ belongs to $H_0(X)$, which is canonically isomorphic to the coefficients R if X is path-connected, the canonical isomorphism being

$$\langle [1], \cdot \rangle : H_0(X) \rightarrow R,$$

which sends the class represented by any singular 0-simplex $\sigma : \Delta^0 \rightarrow X$ to the unit $1 \in R$.

REMARK 48.2. If desired, one can use the Alexander-Whitney diagonal approximation to turn (48.1) into a more precise chain-level formula for $\varphi \cap c$. A formula of this type appears as the definition of the cap product in [Hat02, §3.3], though with slightly different conventions than we are using here.

LEMMA 48.3. *The following diagrams commute up to chain homotopy:*

$$\begin{array}{ccc}
 C^*(X) \otimes C^*(X) \otimes C_*(X) & \xrightarrow{1 \otimes \cap} & C^*(X) \otimes C_*(X) \\
 \downarrow \cup \otimes \mathbb{1} & & \downarrow \cap \\
 C^*(X) \otimes C_*(X) & \xrightarrow{\cap} & C_*(X)
 \end{array}
 \qquad
 C_*(X) \begin{array}{c} \xrightarrow{1 \cap} \\ \xleftarrow{\mathbb{1}} \end{array} C_*(X) .$$

PROOF. These can be deduced from the diagrams for diagonal approximations in Theorem 47.11. □

COROLLARY 48.4. *The cap and cup products on singular (co)homology are related by the associativity relation*

$$([\psi] \cup [\varphi]) \cap [c] = [\psi] \cap ([\varphi] \cap [c])$$

for $[\psi], [\varphi] \in H^*(X)$ and $[c] \in H_*(X)$, and furthermore,

$$[1] \cap [\varphi] = [\varphi]$$

for all $[\varphi] \in H_*(X)$. □

We observe that in light of (48.4) the associativity relation of Corollary 48.4 generalizes the relation (48.3) for cap and cup products under evaluation.

Two more properties of the cap product will be important to understand. The first is naturality: the mixture of covariance and contravariance makes the statement look slightly more complicated than for the cup product or cross products, but the correct relation is

$$f_*(f^* \varphi \cap c) = \varphi \cap f_*c \quad \text{for} \quad X \xrightarrow{f} Y, \varphi \in C^*(Y) \text{ and } c \in C_*(X).$$

Indeed, one can use the naturality of diagonal approximations to verify this relation directly from the definitions, and since it holds at the chain level, it also holds after descending to (co)homology:

$$(48.5) \quad f_*(f^*[\varphi] \cap [c]) = [\varphi] \cap f_*[c] \quad \text{for} \quad X \xrightarrow{f} Y, [\varphi] \in H^*(Y) \text{ and } [c] \in H_*(X).$$

For the second property, we need one more acyclic model argument, in which the uniqueness up to chain homotopy of natural chain maps $C_*(X) \otimes C_*(Y) \rightarrow C_*(X \times Y) \otimes C_*(X \times Y)$ satisfying certain conditions implies the following:

LEMMA 48.5. *For any spaces X, Y and any diagonal approximation Ψ and chain/cochain-level cross products, the diagram*

$$\begin{array}{ccc}
 C_*(X) \otimes C_*(Y) & \xrightarrow{\times} & C_*(X \times Y) \\
 \downarrow \Psi \otimes \Psi & & \downarrow \Psi \\
 C_*(X) \otimes C_*(X) \otimes C_*(Y) \otimes C_*(Y) & & \\
 \downarrow \mathbb{1} \otimes \sigma \otimes \mathbb{1} & & \\
 C_*(X) \otimes C_*(Y) \otimes C_*(X) \otimes C_*(Y) & \xrightarrow{\times \otimes \times} & C_*(X \times Y) \otimes C_*(X \times Y)
 \end{array}$$

commutes up to chain homotopy, where $\sigma : C_*(X) \otimes C_*(Y) \rightarrow C_*(Y) \otimes C_*(X)$ is the chain map $a \otimes b \mapsto (-1)^{|a| \cdot |b|} b \otimes a$. □

Dualizing the appropriate pieces of this diagram and writing $\sigma : C^*(Y) \otimes C^*(X) \rightarrow C^*(X) \otimes C^*(Y) : \psi \otimes a \mapsto (-1)^{|a| \cdot |\psi|} a \otimes \psi$, one deduces:

COROLLARY 48.6. *The diagram*

$$\begin{array}{ccc}
 C_*(X) \otimes C_*(Y) & \xrightarrow{\quad \times \quad} & C_*(X \times Y) \\
 \uparrow \cap \otimes \cap & & \uparrow \cap \\
 C^*(X) \otimes C_*(X) \otimes C^*(Y) \otimes C_*(Y) & & \\
 \uparrow \mathbf{1} \otimes \sigma \otimes \mathbf{1} & & \\
 C^*(X) \otimes C^*(Y) \otimes C_*(X) \otimes C_*(Y) & \xrightarrow{\quad \times \otimes \times \quad} & C^*(X \times Y) \otimes C_*(X \times Y)
 \end{array}$$

commutes up to chain homotopy. □

COROLLARY 48.7. *The cap product and cross products on homology and cohomology are related by the formula*

$$([\varphi] \times [\psi]) \cap ([a] \times [b]) = (-1)^{|a||\psi|} ([\varphi] \cap [a]) \times ([\psi] \cap [b])$$

for homogeneous elements $[\varphi] \in H^*(X)$, $[\psi] \in H^*(Y)$, $[a] \in H_*(X)$ and $[b] \in H_*(Y)$. □

48.2. Relative 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 relative (co)homology for pairs (X, A) with $A \neq \emptyset$. 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 $\pi_A : A \rightarrow A/A_0$ and $\pi_B : B \rightarrow B/B_0$ 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 $A \oplus B \rightarrow A \otimes B$ with the quotient projection, it descends to a bilinear map $(A/A_0) \oplus (B/B_0) \rightarrow (A \otimes B) / ((A_0 \otimes B) + (A \otimes B_0))$.

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

$$C_*^{\text{CW}}(X, A) \otimes C_*^{\text{CW}}(Y, B) \cong \frac{C_*^{\text{CW}}(X) \otimes C_*^{\text{CW}}(Y)}{(C_*^{\text{CW}}(A) \otimes C_*^{\text{CW}}(Y)) + (C_*^{\text{CW}}(X) \otimes C_*^{\text{CW}}(B))}.$$

As it happens, the denominator on the right hand side is not only a subcomplex of the chain complex $C_*^{\text{CW}}(X) \otimes C_*^{\text{CW}}(Y)$, but its image under the isomorphism $\times : C_*^{\text{CW}}(X) \otimes C_*^{\text{CW}}(Y) \rightarrow C_*^{\text{CW}}(X \times Y)$ 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

$$C_*^{\text{CW}}(X, A) \otimes C_*^{\text{CW}}(Y, B) \xrightarrow{\times} C_*^{\text{CW}}((X, A) \times (Y, B))$$

if we define the product of two CW-pairs to be

$$(48.6) \quad (X, A) \times (Y, B) := (X \times Y, (A \times Y) \cup (X \times B)).$$

The most general version of the cross product on relative cellular homology thus takes the form

$$H_k^{CW}(X, A) \otimes H_\ell^{CW}(Y, B) \xrightarrow{\times} H_{k+\ell}^{CW}((X, A) \times (Y, B)),$$

and the Künneth formula (in the case where R is a principal ideal domain) then becomes

$$0 \longrightarrow \bigoplus_{k+\ell=n} H_k^{CW}(X, A) \otimes H_\ell^{CW}(Y, B) \xrightarrow{\times} H_{k+\ell}^{CW}((X, A) \times (Y, B)) \\ \longrightarrow \bigoplus_{k+\ell=n-1} \text{Tor}(H_k^{CW}(X, A), H_\ell^{CW}(Y, B)) \longrightarrow 0.$$

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 (48.6) as a definition of the product of two pairs of spaces.⁷⁴ Applying the naturality property of the chain-level cross product $C_*(X) \otimes C_*(Y) \xrightarrow{\times} C_*(X \times Y)$ to the inclusions $A \hookrightarrow X$ and $B \hookrightarrow Y$, we see that \times 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

$$C_*(X, A) \otimes C_*(Y, B) \xrightarrow{\times} C_*((X, A) \times (Y, B)),$$

so that the cross product on relative homology is well defined:

$$H_k(X, A) \otimes H_\ell(Y, B) \xrightarrow{\times} H_{k+\ell}((X, A) \times (Y, B)).$$

This part of the story works without any further conditions, and it can also be made to work with arbitrary coefficients.

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 $\theta : C_*(X \times Y) \rightarrow C_*(X) \otimes C_*(Y)$ of the chain-level cross product, and this does *not* always descend to a map

$$C_*((X, A) \times (Y, B)) \rightarrow C_*(X, A) \otimes C_*(Y, B).$$

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 the subcomplex $(C_*(A) \otimes C_*(Y)) + (C_*(X) \otimes C_*(B)) \subset C_*(X) \otimes C_*(Y)$. What we can immediately say instead is that \times and θ descend to chain homotopy inverses between the two quotient complexes

$$(48.7) \quad \frac{C_*(X) \otimes C_*(Y)}{(C_*(A) \otimes C_*(Y)) + (C_*(X) \otimes C_*(B))} \begin{array}{c} \xrightarrow{\times} \\ \xleftarrow{\theta} \end{array} \frac{C_*(X \times Y)}{C_*(A \times Y) + C_*(X \times B)}.$$

⁷⁴This definition of $(X, A) \times (Y, B)$ is the right one for talking about the cross product and Künneth's formula, but for other purposes, it is not a good way of defining the "product" of two objects in the category Top^{rel} . It suffers in particular 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 Top^{rel} , in the spirit of Proposition 43.5.

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

$$\frac{C_*(X \times Y)}{C_*(A \times Y) + C_*(X \times B)} \longrightarrow \frac{C_*(X \times Y)}{C_*((A \times Y) \cup (X \times B))} = C_*((X, A) \times (Y, B)),$$

and it does happen sometimes that this chain map induces an isomorphism on homology. Recall from Proposition 34.16: this is true whenever the two subsets $A \times Y, X \times B \subset X \times Y$ form an *excisive couple* for singular homology.

Now, using the quotient complex at the right hand side of (48.7) 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:

THEOREM 48.8. *For relative singular homology with coefficients in a principal ideal domain R , if (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*

$$0 \rightarrow \bigoplus_{k+\ell=n} H_k(X, A) \otimes H_\ell(Y, B) \xrightarrow{\times} H_{k+\ell}((X, A) \times (Y, B)) \rightarrow \bigoplus_{k+\ell=n-1} \text{Tor}(H_k(X, A), H_\ell(Y, B)) \rightarrow 0,$$

and the sequence splits. □

Adapting this discussion for cohomology requires the following lemma:

LEMMA 48.9. *If $A, B \subset X$ is an excisive couple, then for any abelian group G , the map on cohomology*

$$H^*(X, A \cup B; G) \rightarrow H^*\left(\frac{C_*(X; \mathbb{Z})}{C_*(A + B; \mathbb{Z})}; G\right)$$

induced by the natural chain map $C_*(X; \mathbb{Z})/C_*(A + B; \mathbb{Z}) \rightarrow C_*(X, A \cup B; \mathbb{Z})$ is an isomorphism.

PROOF. Both $C_*(X; \mathbb{Z})/C_*(A + B; \mathbb{Z})$ and $C_*(X, A \cup B; \mathbb{Z})$ are chain complexes of free abelian groups, so Corollary 33.18 and Proposition 34.16 imply that $C_*(X; \mathbb{Z})/C_*(A + B; \mathbb{Z}) \rightarrow C_*(X, A \cup B; \mathbb{Z})$ is a chain homotopy equivalence. It follows via Proposition 41.2 that this remains true after dualizing it. □

If $A \times Y$ and $X \times B$ form an excisive couple in $X \times Y$, the lemma now allows us to identify the cohomology of the quotient complex $C_*(X \times Y)/(C_*(A \times Y) + C_*(X \times B))$ with $H^*((X, A) \times (Y, B))$, so that dualizing the chain map θ in (48.7) gives rise to a cross product on relative singular cohomology

$$H^k(X, A) \otimes H^\ell(Y, B) \xrightarrow{\times} H^{k+\ell}((X, A) \times (Y, B)).$$

Let's identify some concrete situations in which the technical condition required for relative cross products is satisfied:

LEMMA 48.10. *Two subsets $A, B \subset X$ form an excisive couple for singular homology whenever any of the following conditions hold:*

- (1) *A and B are both open in X ;*
- (2) *$A = \emptyset$ or $B = \emptyset$;*
- (3) *X is a CW-complex with A and B as subcomplexes.*

PROOF. The first case is immediate from Proposition 34.14, and the second case is trivial. For the third case, we recall the homological characterization of excisive couples in Proposition 34.16: $A, B \subset X$ form an excisive couple if the inclusion of pairs $(A, A \cap B) \hookrightarrow (A \cup B, B)$ induces

an isomorphism in relative singular homology with integer coefficients. If A and B are both subcomplexes of a CW-complex X , then $(A, A \cap B) \hookrightarrow (A \cup B, B)$ is a map of CW-pairs, and the relative cellular chain complexes of both are freely generated by the set of cells in A that are not also cells in B , implying that the induced chain map is an isomorphism of cellular chain complexes. The result then follows from the isomorphism between cellular and singular homology. \square

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 $f : (X, x_0) \rightarrow (X', x'_0)$ and $g : (Y, y_0) \rightarrow (Y', y'_0)$, the product map $f \times g : X \times Y \rightarrow X' \times Y'$ descends to the quotient as a continuous map

$$f \wedge g : X \wedge Y \rightarrow X' \wedge Y'.$$

EXAMPLE 48.11. For any integers $k, \ell \geq 0$, $S^k \wedge S^\ell \cong S^{k+\ell}$. 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 48.10, the subcomplexes $\{x_0\} \times Y, X \times \{y_0\} \subset X \times Y$ form an excisive couple, and the Künneth formula is thus 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

$$\begin{aligned} 0 \rightarrow \bigoplus_{k+\ell=n} H_k(X, \{x_0\}) \otimes H_\ell(Y, \{y_0\}) &\xrightarrow{\times} H_n(X \times Y, X \vee Y) \\ &\rightarrow \bigoplus_{k+\ell=n-1} \text{Tor}(H_k(X, \{x_0\}), H_\ell(Y, \{y_0\})) \rightarrow 0, \end{aligned}$$

or under the natural isomorphisms $H_*(X, A) = \tilde{H}_*(X/A)$ for good pairs,

$$(48.8) \quad 0 \rightarrow \bigoplus_{k+\ell=n} \tilde{H}_k(X) \otimes \tilde{H}_\ell(Y) \xrightarrow{\times} \tilde{H}_n(X \wedge Y) \rightarrow \bigoplus_{k+\ell=n-1} \text{Tor}(\tilde{H}_k(X), \tilde{H}_\ell(Y)) \rightarrow 0.$$

48.3. Relative cup and cap products. Recall that the cochain-level cup product $C^k(X) \otimes C^\ell(X) \rightarrow C^{k+\ell}(X) : \varphi \otimes \psi \mapsto \varphi \cup \psi$ produces a cochain of the form

$$(48.9) \quad \varphi \cup \psi = (\varphi \otimes \psi) \circ \Psi : C_*(X) \rightarrow R,$$

where Ψ is any choice of diagonal approximation, meaning a natural chain map $C_*(X) \rightarrow C_*(X) \otimes C_*(X)$ that acts on singular 0-simplices as the diagonal map. We claim that whenever $A, B \subset X$ are two subspaces that form an excisive couple for singular homology, this gives rise to a well-defined relative cup product

$$H^k(X, A) \otimes H^\ell(X, B) \xrightarrow{\cup} H^{k+\ell}(X, A \cup B).$$

Indeed, under this assumption, Lemma 48.9 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 $\Psi : C_*(X) \rightarrow C_*(X) \otimes C_*(X)$ and make sense of

$$\varphi \cup \psi = (\varphi \otimes \psi) \circ \Psi : \frac{C_*(X)}{C_*(A) + C_*(B)} \rightarrow R$$

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 47.13, as the homotopy-commutative diagrams underlying those properties still make sense after descending to whichever quotient complexes are permitted by the excisive couple assumption.

As a special case, the product

$$H^k(X, A) \otimes H^\ell(X, A) \xrightarrow{\cup} H^{k+\ell}(X, A)$$

is well defined for every pair (X, A) , as $A, A \subset X$ always trivially forms an excisive couple.

Similarly, the relative cap product takes the form

$$(48.10) \quad H^*(X, A) \otimes H_*(X, A \cup B) \xrightarrow{\cap} H_*(X, B)$$

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

$$C^*(X) \otimes C_*(X) \rightarrow C_*(X) : \varphi \otimes c \mapsto \varphi \cap c = (\mathbf{1} \otimes \varphi) \circ \Psi(c)$$

always descends to a well-defined map on the relative complexes

$$C^*(X, A) \otimes \frac{C_*(X)}{C_*(A) + C_*(B)} \rightarrow C_*(X, B),$$

as $\varphi \in C^*(X, A)$ means $\varphi : C_*(X) \rightarrow R$ 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 (48.10).

48.4. Exercises.

EXERCISE 48.1. Assuming the coefficient ring R to be field \mathbb{K} , find an alternative proof of the formula in Corollary 48.7 relating the cap and cross products, using Exercise 47.1.

EXERCISE 48.2. Show that for the cross product on reduced homology as described in (48.8) and the identification of $S^k \wedge S^\ell$ with $S^{k+\ell}$ as indicated in Example 48.11, if $[S^k] \in \tilde{H}_k(S^k)$ and $[S^\ell] \in \tilde{H}_\ell(S^\ell)$ are generators, then $[S^k] \times [S^\ell] \in \tilde{H}_{k+\ell}(S^{k+\ell})$ is also a generator.

EXERCISE 48.3. Suppose $f : S^k \rightarrow S^k$ and $g : S^\ell \rightarrow S^\ell$ 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 $f : S^k \rightarrow S^k$ 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 $f : X \rightarrow X'$ and $g : Y \rightarrow Y'$ are cellular maps, then the diagram

$$\begin{array}{ccc} C_*^{\text{CW}}(X) \otimes C_*^{\text{CW}}(Y) & \xrightarrow{\times} & C_*^{\text{CW}}(X \times Y) \\ \downarrow f_* \otimes g_* & & \downarrow (f \times g)_* \\ C_*^{\text{CW}}(X') \otimes C_*^{\text{CW}}(Y') & \xrightarrow{\times} & C_*^{\text{CW}}(X' \times Y') \end{array}$$

commutes.

Comment: With significantly more effort, one can proceed from this exercise to a proof that the cellular cross product matches the cross product on singular homology under the natural isomorphisms $H_*^{\text{CW}}(X; R) \cong H_*(X; R)$ 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].

49. The orientation bundle

The next few lectures will focus on a new topic: the global topology of finite-dimensional topological manifolds.⁷⁵

49.1. Finitely-generated homology. There is a basic fact about manifolds that was briefly mentioned in the context of the Lefschetz fixed point theorem (Lecture 40), 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 $r : P \rightarrow M$. Using coefficients in a commutative ring R , the homology of P is a finitely-generated R -module, since its simplicial chain complex is finitely generated, and it follows that the singular homology of M is likewise finitely generated. Since this statement relies on knowledge of the homology of compact CW-complexes, it is true not just for singular homology, but in fact for any axiomatic homology theory with ring coefficients:

THEOREM 49.1. *For every compact manifold M and any axiomatic homology theory $\{h_n : \text{Top}^{\text{rel}} \rightarrow R\text{-Mod}\}_{n \in \mathbb{Z}}$ with coefficients $h_0(\{*\}) \cong R$, the homology $h_*(M) = \bigoplus_{n \in \mathbb{Z}} h_n(M)$ is a finitely-generated R -module. \square*

As usual, the statement that $h_*(M)$ is finitely generated is really two statements in one, namely that $h_k(M)$ is finitely generated for every $k \in \mathbb{Z}$, and also that it vanishes for all k outside of a finite range. One consequence of the main theorem in this lecture will be that, in fact, $h_k(M)$ must always vanish when $k > \dim M$.

49.2. Orientations. I would now like to discuss what it means for a topological manifold to be *orientable*. We discussed this somewhat in Lecture 30 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 overlapping charts are related by a coordinate transformation that

⁷⁵I 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.

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 $\mathbb{R}^n \rightarrow \mathbb{R}^n$ with positive determinant. For maps that are continuous 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 36). Recall that if $\dim M = n$, then for every axiomatic homology theory h_* and every interior point $x \in M \setminus \partial M$, 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

$$(49.1) \quad \begin{aligned} h_k(M, M \setminus \{x\}) &\cong h_k(\mathcal{U}_x, \mathcal{U}_x \setminus \{x\}) \cong h_k(\mathbb{R}^n, \mathbb{R}^n \setminus \{0\}) \cong h_k(\mathbb{D}^n, \partial \mathbb{D}^n) \\ &\cong \tilde{h}_{k-1}(S^{n-1}) \cong \begin{cases} G & \text{if } k = n, \\ 0 & \text{otherwise,} \end{cases} \end{aligned}$$

where $G = h_0(\{*\}) \in R\text{-Mod}$ is the coefficient module. We call

$$h_n(M | x) := h_n(M, M \setminus \{x\}) \cong G$$

the **local homology** group of M at x , and if h_* has coefficients $G = \mathbb{Z}$, a **local orientation** of M at x is defined to be a choice of generator

$$[M]_x \in h_n(M | x) \cong H_n(M | x; \mathbb{Z}) \cong \mathbb{Z}.$$

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 $x \in M \setminus \partial M$, 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 $h_*(M | x) := h_*(M, M \setminus \{x\})$ 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

$$h_*(M | A) := h_*(M, M \setminus A),$$

which we shall call the “homology of M restricted to A ”. In singular homology, for instance, 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 $(M, M \setminus A) \hookrightarrow (M, M \setminus B)$, which therefore induces natural “restriction” homomorphisms

$$j_{B,A} : h_*(M | A) \rightarrow h_*(M | B),$$

or in the case where B is a single point $x \in A$,

$$j_{x,A} : h_*(M | A) \rightarrow h_*(M | x).$$

Note that the absolute homology $h_*(M)$ itself is also an example of a restricted homology group, namely $h_*(M | M)$.

Since the isomorphism $h_n(M | x) \cong G$ only holds when $x \in M$ lies in the interior $M \setminus \partial M$, it will be convenient to assume in most of the following discussion that M has empty boundary,

$$\partial M = \emptyset.$$

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

$$\begin{array}{ccccc} h_n(M|x) & \xleftarrow{\cong} & h_n(\mathcal{U}|x) & \xrightarrow[\cong]{\varphi^*} & h_n(\mathbb{R}^n|\varphi(x)) \\ j_{x,A} \uparrow & & & & \cong \uparrow \\ h_n(M|A) & \xleftarrow{\cong} & h_n(\mathcal{U}|A) & \xrightarrow[\cong]{\varphi^*} & h_n(\mathbb{R}^n|\mathbb{D}^n), \end{array}$$

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

$$j_{x,A}: h_n(M|A) \rightarrow h_n(M|x)$$

is an isomorphism. We shall say in this situation that $A \subset M$ is a **disk-like neighborhood** of $x \in M$.

DEFINITION 49.2. 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 $\partial M \neq \emptyset$, we define an orientation of M to be an orientation of its interior $M \setminus \partial M$.

REMARK 49.3. The choice to use singular homology in Definition 49.2 is arbitrary; any axiomatic homology theory with coefficient group \mathbb{Z} would do just as well.

A manifold equipped with an orientation will be called an **oriented manifold** (*orientierte Mannigfaltigkeit*). In light of (49.1), you can imagine an orientation as a choice for every $x \in M \setminus \partial M$ of a favorite generator $[S_x] \in \tilde{H}_{n-1}(S_x; \mathbb{Z}) \cong \mathbb{Z}$ 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 $\tilde{H}_{n-1}(S_x; \mathbb{Z}) \rightarrow \tilde{H}_{n-1}(S_y; \mathbb{Z})$ sending $[S_x]$ to $[S_y]$. You should take a moment to contemplate why this description matches Definition 49.2 in the case $M = \mathbb{R}^n$.

49.3. Orientation bundles. The notion of orientation described above admits fruitful generalizations, and in order to express them in the most useful language, let us again assume $\partial M = \emptyset$, fix an axiomatic homology theory h_* with coefficient module $G := h_0(\{*\}) \in R\text{-Mod}$, and abbreviate

$$\Theta_x^G := h_n(M|x) \cong H_n(M|x; G) \cong G$$

for each point $x \in M$, or simply

$$\Theta_x := h_n(M|x)$$

for situations where knowledge of the coefficient group is unimportant. This associates to each point $x \in M$ an R -module Θ_x that is isomorphic to the fixed module $G = h_0(\{*\})$, though without a canonical choice of isomorphism. The union of all these modules defines a set that we shall denote by

$$\Theta = \Theta^G := \bigcup_{x \in M} \Theta_x^G.$$

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 $\bigsqcup_{x \in M} \Theta_x$ since this notation normally carries connotations about the topology of the union, and those connotations would be inconsistent with the next definition.

DEFINITION 49.4. 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

$$\mathcal{B} := \{ \mathcal{U}_c \subset \Theta \mid \mathcal{U} \subset M \text{ open and } c \in h_n(M | \bar{\mathcal{U}}) \},$$

where for $\mathcal{U} \subset M$ and $c \in h_n(M | \bar{\mathcal{U}})$ we define

$$\mathcal{U}_c := \{ j_{x, \bar{\mathcal{U}}}(c) \in \Theta_x \mid x \in \mathcal{U} \}.$$


PROPOSITION 49.5. *The collection of subsets $\mathcal{B} = \{ \mathcal{U}_c \}$ appearing in Definition 49.4 is the base of a topology on Θ for which the natural projection map*

$$p : \Theta \rightarrow M$$

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 $j_{x, \bar{\mathcal{U}}} : h_n(M | \bar{\mathcal{U}}) \rightarrow \Theta_x$ 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 $h_n(M | A)$:



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 $p : \Theta \rightarrow M$ 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 $j_{x, A} : h_n(M | A) \rightarrow h_n(M | x)$ factors through $h_n(M | \bar{\mathcal{U}})$ for any smaller open neighborhood $\mathcal{U} \subset \mathring{A}$ of x , implying that $j_{x, \bar{\mathcal{U}}} : h_n(M | \bar{\mathcal{U}}) \rightarrow h_n(M | x)$ 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

$$(p, \psi) : p^{-1}(A) \rightarrow A \times \Theta_x$$

a homeomorphism, where the map $\psi : p^{-1}(A) \rightarrow \Theta_x$ is defined such that its restriction to $\Theta_y \subset p^{-1}(A)$ for each $y \in A$ is $j_{x, A} \circ j_{y, A}^{-1} : \Theta_y \rightarrow \Theta_x$. Using this to identify $p^{-1}(A)$ with $A \times \Theta_x$ turns $p^{-1}(A) \xrightarrow{p} A$ into the trivial covering map $A \times \Theta_x \rightarrow A : (a, c) \mapsto a$. \square

REMARK 49.6. The orientation bundle Θ^G and covering map $p : \Theta^G \rightarrow M$ obviously depend on $G = h_0(\{*\})$, but up to isomorphism, they are otherwise independent of the choice of axiomatic homology theory.

REMARK 49.7. The word “bundle” is borrowed from differential geometry, where fiber bundles $p : E \rightarrow B$ 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 R -modules—this makes $p : \Theta \rightarrow M$ a *sheaf* of R -modules. For readers who may know what this means and find it interesting: $p : \Theta \rightarrow M$ is the *completion* of the *presheaf* that associates to each open subset $U \subset M$ the R -module $h_n(M|U)$.

DEFINITION 49.8. 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 $s : A \rightarrow \Theta$ such that $p \circ s = \text{Id}_A$, 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

$$\Gamma_c(\Theta|_A) \subset \Gamma(\Theta|_A).$$

The sets $\Gamma(\Theta|_A)$ and $\Gamma_c(\Theta|_A)$ defined above are both R -modules in a natural way; see Exercise 49.2 for details.

Our previous definition of orientations can now be recouched in the following terms.

DEFINITION 49.9. 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 $\Theta_x^{\mathbb{Z}} \cong \mathbb{Z}$ 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_x^R as an R -module, i.e. $Rs(x) = \Theta_x^R$. 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 $M \setminus \partial M$.

REMARK 49.10. When $G = h_0(\{*\})$ is a ring R , one should not be misled by the isomorphism $\Theta_x^R \cong \tilde{h}_{n-1}(S^{n-1}) \cong R$ into thinking that Θ_x^R is a ring—it is an R -module, and the isomorphism $\Theta_x^R \cong R$ respects that R -module structure, but Θ_x^R 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 $\Theta_x^{\mathbb{Z}} \cong \mathbb{Z}$, that is, the two local orientations of M at x . Let us write

$$\tilde{M} := \{c \in \Theta^{\mathbb{Z}} \mid c \text{ is a local orientation}\} \subset \Theta^{\mathbb{Z}},$$

in other words, \tilde{M} is the union for all $x \in M$ of the two generators of $\Theta_x^{\mathbb{Z}} \cong \mathbb{Z}$. Assigning to $\tilde{M} \subset \Theta^{\mathbb{Z}}$ the subspace topology, it is easy to see that the restriction of $p : \Theta^{\mathbb{Z}} \rightarrow M$ defines a two-to-one covering map

$$\pi := p|_{\tilde{M}} : \tilde{M} \rightarrow M : [M]_x \mapsto x.$$

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 $s : A \rightarrow \widetilde{M}$ and its opposite, $-s : A \rightarrow \widetilde{M}$, 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, the proof of the following result is an exercise:

LEMMA 49.11. *If $A \subset M$ is connected and $\pi^{-1}(A) \subset \widetilde{M}$ has more than one connected component, then each component intersects $\Theta_x^{\mathbb{Z}}$ for every $x \in A$. It follows in this situation that $\pi^{-1}(A) \subset \widetilde{M}$ has exactly two components, each of which is the image of a section of $\Theta^{\mathbb{Z}}$ along A . \square*

Combining this with the previous remarks proves:

PROPOSITION 49.12. *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 . \square*

EXAMPLE 49.13. For $M = \mathbb{R}P^2$, the orientation double cover is equivalent to the standard covering $S^2 \rightarrow S^2/\mathbb{Z}_2 = \mathbb{R}P^2$ defined via the antipodal map on S^2 . In particular, $\mathbb{R}P^2$ is orientable along a loop $\gamma \subset \mathbb{R}P^2$ 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 49.14. *Every manifold is orientable over \mathbb{Z}_2 .*

PROOF. Each fiber $\Theta_x^{\mathbb{Z}_2}$ 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 . \square

PROPOSITION 49.15. *For every abelian group G , there is a natural group isomorphism*

$$\Phi_x : \Theta_x^{\mathbb{Z}} \otimes G \xrightarrow{\cong} \Theta_x^G$$

for each $x \in M$ 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 $s' \in \Gamma(\Theta^G|_A)$.

PROOF. Since Θ^G depends on G but not on the specific choice of axiomatic homology theory with G as coefficients, we can choose to use singular homology and thus write $\Theta_x^G = H_n(M|x; G)$. The map Φ_x then comes from the universal coefficient theorem, which applies because \mathbb{Z} is a principal ideal domain and $\Theta_x^{\mathbb{Z}} = H_n(M, M \setminus \{x\}; \mathbb{Z})$ is the homology of a chain complex of free \mathbb{Z} -modules $C_*(M, M \setminus \{x\}; \mathbb{Z})$. It is an isomorphism because $\text{Tor}(H_{n-1}(M|x; \mathbb{Z}), G) = 0$, since $H_{n-1}(M|x; \mathbb{Z}) = 0$. \square

COROLLARY 49.16. *If M is orientable along A , then it is also R -orientable along A for every choice of coefficient ring R . \square*

49.4. The compact support axiom. Before the most important result about orientation bundles can be proved axiomatically, we will need to add one additional axiom to the usual list from Eilenberg-Steenrod. Recall that a **compact pair** is a pair of spaces (X, A) such that X is compact and $A \subset X$ is closed. For pairs of spaces, we shall write

$$(X', A') \subset (X, A)$$

when $X' \subset X$ and $A' \subset A$.

DEFINITION 49.17. An axiomatic homology theory h_* is said to be **compactly supported** if for every $(X, A) \in \mathbf{Top}^{\text{rel}}$ and every $c \in h_n(X, A)$, there exists a compact pair $(X', A') \subset (X, A)$ such that c lies in the image of the map $h_n(X', A') \rightarrow h_n(X, A)$ induced by the inclusion.

EXAMPLE 49.18. Singular homology is compactly supported since singular n -chains always have images contained in compact subsets. More precisely, for any singular n -chain $\sum_i a_i \sigma_i \in C_n(X; G)$ that is a relative n -cycle in (X, A) , the same chain is also a relative n -cycle in the compact pair $(X', A') \subset (X, A)$ defined by $X' := \bigcup_i \sigma_i(\Delta^n)$ and $A' := \bigcup_i \sigma_i(\partial\Delta^n)$.

PROPOSITION 49.19. *Suppose h_* is a compactly supported homology theory, $(X', A') \subset (X, A)$ is a compact pair, and $c \in h_*(X', A')$ belongs to the kernel of the map $h_*(X', A') \rightarrow h_*(X, A)$ induced by the inclusion. Then c also belongs to the kernel of the map $h_*(X', A') \rightarrow h_*(X'', A'')$ induced by the inclusion of (X', A') into another compact pair (X'', A'') with*

$$(X', A') \subset (X'', A'') \subset (X, A).$$

Proposition 49.19 is one of those results that are easier to prove directly in the case of singular homology, so let us do that before tackling the general case.

PROOF OF PROP. 49.19 FOR SINGULAR HOMOLOGY. Suppose (X', A') is a compact pair with an inclusion into (X, A) , and $c \in C_n(X')$ is a singular n -chain that is a relative cycle in (X', A') , and therefore also in (X, A) , such that the class it represents in $H_n(X, A)$ vanishes. The latter means that there is an $(n + 1)$ -chain $b = \sum_i b_i \sigma_i \in C_{n+1}(X)$ and an n -chain $a = \sum_j a_j \sigma_j \in C_n(A)$ such that $c = \partial b + a$. Setting $X'' := \bigcup_i \sigma_i(\Delta^{n+1}) \cup X' \subset X$ and $A'' := \bigcup_j \sigma_j(\Delta^n) \cup A' \subset A$ then gives a compact pair (X'', A'') that contains (X', A') and is contained in (X, A) , such that $b \in C_{n+1}(X'')$ and $a \in C_n(A'')$, showing that c represents the trivial class in $H_n(X'', A'')$. \square

PROOF OF PROPOSITION 49.19 (GENERAL CASE). We first establish the absolute case, i.e. assuming $X' \subset X$ is a compact subset and $c \in h_n(X')$ vanishes under the map $i_* : h_n(X') \rightarrow h_n(X)$ induced by the inclusion $i : X' \hookrightarrow X$, we show that c also vanishes under the map induced by the inclusion $j : X' \hookrightarrow X''$ into some compact set $X'' \subset X$ containing X . Indeed, the long exact sequence

$$\dots \longrightarrow h_{n+1}(X, X') \xrightarrow{\partial_*} h_n(X) \xrightarrow{i_*} h_n(X') \longrightarrow \dots$$

produces a class $b \in h_{n+1}(X, X')$ such that $\partial_* b = c$, and the compact support axiom then implies the existence of a compact pair $(A, A') \subset (X, X')$ and an element $a \in h_{n+1}(A, A')$ that maps to b under the homomorphism $h_{n+1}(A, A') \rightarrow h_{n+1}(X, X')$ induced by the inclusion. Let $X'' := A \cup X'$, and note that X'' is compact since A and X' both are. Using the inclusions $\varphi : (A, A') \hookrightarrow (X'', X')$, $k : X'' \hookrightarrow X$ and $j : X' \hookrightarrow X''$, the naturality of long exact sequences produces a commutative diagram

$$\begin{array}{ccccccc} & & h_{n+1}(A, A') & & & & \\ & & \downarrow \varphi_* & & & & \\ \dots & \longrightarrow & h_{n+1}(X'', X') & \xrightarrow{\partial_*} & h_n(X') & \xrightarrow{j_*} & h_n(X'') \longrightarrow \dots \\ & & \downarrow k_* & & \downarrow \mathbb{1} & & \downarrow k_* \\ \dots & \longrightarrow & h_{n+1}(X, X') & \xrightarrow{\partial_*} & h_n(X') & \xrightarrow{i_*} & h_n(X) \longrightarrow \dots \end{array}$$

in which the bottom two rows are exact. Since $k_* \varphi_* a = b$, we deduce $c = \partial_* k_* \varphi_* a = \partial_* \varphi_* a$, and thus $j_* c = j_* \partial_* \varphi_* a = 0$.

Now the relative case: we assume $(X', A') \subset (X, A)$ is a compact pair and $c \in h_n(X', A')$ vanishes under $i_* : h_n(X', A') \rightarrow h_n(X, A)$, with $i : (X', A') \hookrightarrow (X, A)$ denoting the inclusion of

pairs. The obvious inclusions and the naturality of long exact sequences produce the following commutative diagram with exact rows:

$$\begin{array}{ccccccccccc}
 \dots & \longrightarrow & h_n(A') & \xrightarrow{j'_*} & h_n(X') & \xrightarrow{\varphi'_*} & h_n(X', A') & \xrightarrow{\partial'_*} & h_{n-1}(A') & \xrightarrow{j'_*} & h_{n-1}(X') & \longrightarrow & \dots \\
 & & \downarrow i_* & & \downarrow i_* & & \downarrow i_* & & \downarrow i_* & & \downarrow i_* & & \\
 \dots & \longrightarrow & h_n(A) & \xrightarrow{j_*} & h_n(X) & \xrightarrow{\varphi_*} & h_n(X, A) & \xrightarrow{\partial_*} & h_{n-1}(A) & \xrightarrow{j_*} & h_{n-1}(X) & \longrightarrow & \dots
 \end{array}$$

The diagram implies that $\partial'_*c \in h_{n-1}(A')$ is in the kernel of $i_* : h_{n-1}(A') \rightarrow h_{n-1}(A)$, so the absolute case of the result implies the existence of a compact set $A'' \subset A$ containing A' such that ∂'_*c also vanishes under the map $i'_* : h_{n-1}(A') \rightarrow h_{n-1}(A'')$ induced by the inclusion $i' : A' \hookrightarrow A''$. Setting

$$X'' := X' \cup A''$$

then produces a compact pair (X'', A'') with inclusions

$$(X', A') \xleftarrow{i'} (X'', A'') \xleftarrow{i''} (X, A),$$

and thus a more elaborate commutative diagram with three exact rows,

$$\begin{array}{ccccccccccc}
 \dots & \longrightarrow & h_n(A') & \xrightarrow{j'_*} & h_n(X') & \xrightarrow{\varphi'_*} & h_n(X', A') & \xrightarrow{\partial'_*} & h_{n-1}(A') & \xrightarrow{j'_*} & h_{n-1}(X') & \longrightarrow & \dots \\
 & & \downarrow i'_* & & \downarrow i'_* & & \downarrow i'_* & & \downarrow i'_* & & \downarrow i'_* & & \\
 \dots & \longrightarrow & h_n(A'') & \xrightarrow{j''_*} & h_n(X'') & \xrightarrow{\varphi''_*} & h_n(X'', A'') & \xrightarrow{\partial''_*} & h_{n-1}(A'') & \xrightarrow{j''_*} & h_{n-1}(X'') & \longrightarrow & \dots \\
 & & \downarrow i''_* & & \downarrow i''_* & & \downarrow i''_* & & \downarrow i''_* & & \downarrow i''_* & & \\
 \dots & \longrightarrow & h_n(A) & \xrightarrow{j_*} & h_n(X) & \xrightarrow{\varphi_*} & h_n(X, A) & \xrightarrow{\partial_*} & h_{n-1}(A) & \xrightarrow{j_*} & h_{n-1}(X) & \longrightarrow & \dots
 \end{array}$$

Since $0 = i'_*\partial'_*c = \partial''_*i'_*c$, exactness of the middle row implies $i'_*c = \varphi''_*b$ for some $b \in h_n(X'')$, and the hypothesis on c then implies $0 = i''_*i'_*c = i''_*\varphi''_*b = \varphi_*i''_*b$, so that by the exactness of the bottom row, $i''_*b = j_*a$ for some $a \in h_n(A)$. By the compact support axiom, there is a compact subset $K \subset A$ such that the map $h_n(K) \rightarrow h_n(A)$ induced by the inclusion has a in its image, and after replacing both A'' and X'' by their respective unions with K , we can now assume without loss of generality that $a = i''_*a''$ for some $a'' \in h_n(A'')$. It follows that $i''_*b = j_*i''_*a'' = i''_*j''_*a''$, thus $b - j''_*a'' \in \ker i''_* \subset h_n(X'')$. By the absolute case of the result, we can now replace X'' with a larger compact set in order to assume without loss of generality that $b - j''_*a'' = 0 \in h_n(X'')$, or in other words, $b = j''_*a''$. This implies $i'_*c = \varphi''_*j''_*a'' = 0 \in h_n(X'', A'')$. \square

49.5. Sections and fundamental classes. We would now like to formulate a relationship between the R -module of sections $\Gamma(\Theta|_A)$ and the homology $h_n(M|A)$.

LEMMA 49.20. *Assume h_* is a compactly supported axiomatic homology theory and Θ denotes the associated orientation bundle over an n -manifold M with empty boundary. Then for every closed subset $A \subset M$, there exists an R -module homomorphism*

$$J_A : h_n(M|A) \rightarrow \Gamma_c(\Theta|_A) : c \mapsto s_c$$

defined by $s_c(x) := j_{x,A}(c)$ for $x \in A$.

PROOF. We need to show two things about the map $s_c : A \rightarrow \Theta|_A$, 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 $c \in h_n(M|A) = h_n(M, M \setminus A)$, the compact support axiom provides a compact pair (K, L) with an inclusion $i : (K, L) \hookrightarrow (M, M \setminus A)$ such that $c = i_*c'$ for some $c' \in h_n(K, L)$. For $x \in A$, we then have a composition of inclusions

$$(K, L) \xrightarrow{i} (M, M \setminus A) \xrightarrow{j} (M, M \setminus \{x\})$$

such that $s_c(x) = j_*c'$. But if $x \notin K$, then $K \subset M \setminus \{x\}$ and j thus factors through the pair $(M \setminus \{x\}, M \setminus \{x\})$, whose homology is trivial, implying $s_c(x) = 0$, i.e. the support of s_c is contained in K .

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 49.5 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. Given $c \in h_n(M|A)$, write $c = i_*c'$ as in the previous paragraph, with $c' \in h_n(K, L)$ and the inclusion of a compact pair $i : (K, L) \hookrightarrow (M, M \setminus A)$. Since L is compact and disjoint from A , every $x \in A$ admits a disk-like neighborhood $\bar{U} \subset M$ disjoint from L , so that the inclusion $(K, L) \hookrightarrow (M, M \setminus \{x\})$ factors through $(M, M \setminus \bar{U})$, giving rise to a class

$$c'' \in h_n(M|\bar{U}) \quad \text{such that} \quad j_{x,\bar{U}}(c'') = s_c(x) \text{ for all } x \in \mathring{U} \cap A.$$

On a neighborhood of x , our section s_c is therefore the restriction of a section defined on the disk-like neighborhood \mathring{U} , and the latter is continuous for the reasons stated above. \square

Here is the main theorem about the orientation bundle.

THEOREM 49.21. *If M is a topological n -manifold with empty boundary and h_* is a compactly supported homology theory, then for every closed subset $A \subset M$, the map $J_A : h_n(M|A) \rightarrow \Gamma_c(\Theta|_A)$ is an isomorphism, and $h_k(M|A) = 0$ 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 49.22. *Assume M is a topological n -manifold with empty boundary, and h_* is a compactly supported homology theory with coefficients G . Then:*

- (1) $h_k(M) = 0$ for all $k > n$.
- (2) If M is noncompact and connected, then additionally $h_n(M) = 0$.
- (3) If M is compact, connected and orientable, then $h_n(M) \cong G$.
- (4) If M is compact and connected but not orientable, then $h_n(M) \cong \{g \in G \mid 2g = 0\}$.

Moreover, if h_* has coefficients in the ring R and M is compact and R -orientable, then any choice of R -orientation $s \in \Gamma(\Theta^R)$ determines a unique element

$$[M] \in h_n(M) \quad \text{such that} \quad j_{x,M}[M] = s(x) \text{ for all } x \in M,$$

and if M is additionally connected, then the R -module $h_n(M)$ is isomorphic to R and has $[M]$ as a generator.

PROOF. We work through the claims one by one:

- (1) Follows from $h_k(M|A) = 0$ with $A = M$.
- (2) If M is noncompact and connected then $\Gamma_c(\Theta) = 0$, 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 $h_n(M) \cong \Gamma(\Theta)$, where the compact support condition is irrelevant since M is compact. Then given an orientation $s \in \Gamma(\Theta^{\mathbb{Z}})$ and the natural group isomorphisms $\Phi_x : \Theta_x^{\mathbb{Z}} \otimes G \rightarrow \Theta_x^G$ from the universal coefficient theorem (cf. Proposition 49.15), we obtain an injective group homomorphism

$$G \rightarrow \Gamma(\Theta^G) : g \mapsto s_g \quad \text{where} \quad s_g(x) := \Phi_x(s(x) \otimes g).$$

This homomorphism is also surjective if M is connected, because in this case, the value of a section at any one point uniquely determines the section.

- (4) For the submodule $G_0 = \{g \in G \mid 2g = 0\}$, we can again use the isomorphisms $\Phi_x : \Theta_x^{\mathbb{Z}} \otimes G \rightarrow \Theta_x^G$ to define an injective homomorphism

$$G_0 \rightarrow \Gamma(\Theta^G) : g \mapsto s_g \quad \text{where} \quad s_g(x) := \Phi_x(\pm[M]_x \otimes g).$$

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 $[M] \in h_n(M)$ can be written as $J_M^{-1}(s)$. If M is also connected, then the map $\Gamma(\Theta^R) \rightarrow \Theta_x^R : s \mapsto s(x)$ 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 $h_n(M)$. \square

DEFINITION 49.23. The generator $[M] \in h_n(M) \cong R$ associated by Corollary 49.22 to any R -orientation of a closed n -manifold M and a compactly supported homology theory h_* with coefficients R is called the **fundamental class** of M .

49.6. Exercises.

EXERCISE 49.1. Assume M is an n -manifold and $\{h_k : \mathbf{Top}^{\text{rel}} \rightarrow R\text{-Mod}\}_{k \in \mathbb{Z}}$ is an axiomatic homology theory. Given a point $x \in M$, let J denote the set of all open neighborhoods of x , and write $\mathcal{U} < \mathcal{V}$ whenever $\mathcal{V} \subset \mathcal{U}$. This makes $(J, <)$ into a directed set, and whenever $\mathcal{U} < \mathcal{V}$ there is an associated homomorphism $j_{\mathcal{V}, \mathcal{U}} : h_n(M | \mathcal{U}) \rightarrow h_n(M | \mathcal{V})$, so that the collection of R -modules $\{h_n(M | \mathcal{U})\}_{\mathcal{U} \in J}$ forms a direct system. Find a canonical isomorphism

$$\varinjlim \{h_n(M | \mathcal{U})\} \xrightarrow{\cong} h_n(M | x).$$

EXERCISE 49.2. Suppose M is a manifold, G is an R -module, $\Theta = \Theta^G$ is the associated orientation bundle over M , and $A \subset M$ is a subset. Show that $\Gamma(\Theta|_A)$ and $\Gamma_c(\Theta|_A)$ are both naturally R -modules, with addition of sections and scalar multiplication defined pointwise, i.e.

$$(s_1 + s_2)(x) := s_1(x) + s_2(x) \in \Theta_x, \quad (rs)(x) := rs(x) \in \Theta_x$$

for $s_1, s_2, s \in \Gamma(\Theta|_A)$, $r \in R$ and $x \in A$.

EXERCISE 49.3 (*). Prove Lemma 49.11 on the connected components of the orientation double cover.

Hint: Show that the set of $x \in A$ for which $\Theta_x^{\mathbb{Z}}$ intersects a given component of $\pi^{-1}(A)$ is both open and closed.

EXERCISE 49.4. 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 49.5 (*). 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 $\text{Tor}(H_{n-1}(M), \mathbb{Z}_p) = 0$ for various values of $p \geq 2$, and see what you can deduce from it (cf. Exercise 45.13). 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 49.6. Here is an interesting application of Čech cohomology to the question of orientability of manifolds. Fix a space X and R -module G , and recall that the set $\mathcal{O}(X)$ of all open coverings of X admits an ordering relation $<$ that makes it into a directed set: we write $\mathfrak{U} < \mathfrak{U}'$ whenever \mathfrak{U}' is a refinement of \mathfrak{U} . There is a direct system of \mathbb{Z} -graded R -modules over $\mathcal{O}(X)$ whose direct limit is Čech cohomology, namely

$$\check{H}^*(X; G) := \varinjlim \{H_o^*(\mathcal{N}(\mathfrak{U}); G)\}_{\mathfrak{U} \in \mathcal{O}(X)},$$

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 cohomology of a cochain complex $\check{C}^*(\mathfrak{U}; G) := C_o^*(\mathcal{N}(\mathfrak{U}); G)$, where $\check{C}^n(\mathfrak{U}; G) = 0$ for $n < 0$ and, for each $n \geq 0$, $\check{C}^n(\mathfrak{U}; G)$ is the R -module 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:

$$\varphi(\mathcal{U}_0, \dots, \mathcal{U}_n) \in G \quad \text{assuming} \quad \mathcal{U}_0 \cap \dots \cap \mathcal{U}_n \neq \emptyset.$$

The coboundary map $\delta : \check{C}^n(\mathfrak{U}; G) \rightarrow \check{C}^{n+1}(\mathfrak{U}; G)$ is defined by

$$(\delta\varphi)(\mathcal{U}_0, \dots, \mathcal{U}_{n+1}) := (-1)^{n+1} \sum_{k=0}^{n+1} (-1)^k \varphi(\mathcal{U}_0, \dots, \hat{\mathcal{U}}_k, \dots, \mathcal{U}_{n+1}),$$

where the hat over $\hat{\mathcal{U}}_k$ means that that term is skipped. The cohomologies of these cochain complexes form a direct system over $(\mathcal{O}(X), <)$ because, as mentioned in Lecture 43, refinements $\mathfrak{U}' > \mathfrak{U}$ give rise to chain maps $C_o^*(\mathcal{N}(\mathfrak{U}')) \rightarrow C_o^*(\mathcal{N}(\mathfrak{U}))$ that are canonical up to chain homotopy, so dualizing these gives chain maps $\check{C}^*(\mathfrak{U}; G) \rightarrow \check{C}^*(\mathfrak{U}'; G)$ that are also canonical up to chain homotopy and therefore induce canonical maps on the cohomology groups.

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.⁷⁶ We are going to consider covering⁷⁷ maps $f : Y \rightarrow X$ 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 $\varphi : Y_1 \rightarrow Y_2$ such that the diagram

$$\begin{array}{ccc} Y_1 & \xrightarrow{\varphi} & Y_2 \\ & \searrow f_1 & \swarrow f_2 \\ & & X \end{array}$$

⁷⁶Alternatively, one could avoid the need for connected intersections by using Čech cohomology with sheaf coefficients, cf. [Spa95, Chapter 6].

⁷⁷Caution! 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.

commutes. We will say that a covering map (Y, f) is **trivial** if it is isomorphic to the **trivial double cover**

$$X \times \mathbb{Z}_2 \rightarrow X : (x, i) \mapsto x.$$

Given $f : Y \rightarrow X$, any open covering $\mathfrak{U} \in \mathcal{O}(X)$ can be replaced with a refinement such that every $U \in \mathfrak{U}$ is **evenly covered** by $f : Y \rightarrow X$, meaning $f^{-1}(U)$ is the union of two disjoint subsets $\mathcal{V}_0, \mathcal{V}_1 \subset Y$ such that $f|_{\mathcal{V}_i} : \mathcal{V}_i \rightarrow U$ is a homeomorphism for $i = 0, 1$. After a further refinement, assume \mathfrak{U} is also admissible. We can now choose for each $U \in \mathfrak{U}$ a so-called **local trivialization**, meaning a homeomorphism

$$\Phi_U : f^{-1}(U) \rightarrow U \times \mathbb{Z}_2$$

that sends $f^{-1}(x)$ to $\{x\} \times \mathbb{Z}_2$ for each $x \in U$. This determines a set of continuous **transition functions** $g_{\mathcal{U}, \mathcal{V}} : \mathcal{U} \cap \mathcal{V} \rightarrow \mathbb{Z}_2$ for each intersecting pair $\mathcal{U}, \mathcal{V} \in \mathfrak{U}$, defined such that the map

$$(\mathcal{U} \cap \mathcal{V}) \times \mathbb{Z}_2 \xrightarrow{\Phi_{\mathcal{V}} \circ \Phi_{\mathcal{U}}^{-1}} (\mathcal{U} \cap \mathcal{V}) \times \mathbb{Z}_2$$

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 $\mathcal{U} \cap \mathcal{V} \neq \emptyset$ an element $\varphi(\mathcal{U}, \mathcal{V}) := g_{\mathcal{U}, \mathcal{V}} \in \mathbb{Z}_2$. See if you can prove the following:

- $\varphi \in \check{C}^1(\mathfrak{U}; \mathbb{Z}_2)$ is a cocycle, and choosing different local trivializations changes φ by a coboundary.
- Feeding $[\varphi] \in H_o^1(\mathcal{N}(\mathfrak{U}); \mathbb{Z}_2)$ 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.
- If X is an n -manifold and $f : Y \rightarrow X$ 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 .)

50. Existence of the fundamental class

I owe you a proof of Theorem 49.21, on which the general construction of fundamental classes is based. First, we need to finish unpacking its consequences and say what can be said about compact manifolds with boundary.

50.1. Relative fundamental classes. The construction of fundamental classes $[M] \in h_n(M)$ described in the previous lecture works specifically for manifolds that are closed, but only a little bit more effort is required in order to extract from Theorem 49.21 a similar construction for compact manifolds with nonempty boundary. We've already learned from the triangulated manifolds in Lecture 30 what to expect: when M is compact but $\partial M \neq \emptyset$, $[M]$ 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

$$\overset{\circ}{M} := M \setminus \partial M,$$

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 $\overset{\circ}{M}$: 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 $\{0\} \times \partial M$ as the identity map. 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 49.21 can be applied to \mathring{M} , and gives an isomorphism

$$J_A : h_n(\mathring{M}, \mathring{M} \setminus A) \rightarrow \Gamma_c(\Theta|_A)$$

for any closed subset $A \subset \mathring{M}$. Taking $h_* = \{h_k : \mathbf{Top}^{\text{rel}} \rightarrow R\text{-Mod}\}_{k \in \mathbb{Z}}$ to be a homology theory with coefficients $h_0(\{*\}) \cong R$, any R -orientation $s \in \Gamma(\Theta|_{\mathring{M}})$ determines a generator $[M]_x := s(x) \in h_n(M|x) \cong R$ for each interior point $x \in \mathring{M}$. We will refer to a relative homology class

$$[M] \in h_n(M, \partial M)$$

as a **relative fundamental class** for M if the natural map $i^x : h_n(M, \partial M) \rightarrow h_n(M|x)$ defined via the inclusion $(M, \partial M) \hookrightarrow (M, M \setminus \{x\})$ sends $[M]$ to a generator

$$[M]_x := i_*^x[M] \in \Theta_x \cong R \quad \text{such that} \quad R[M]_x = \Theta_x$$

for every $x \in \mathring{M}$. In this way, a relative fundamental class $[M]$ determines an R -orientation $s \in \Gamma(\Theta|_{\mathring{M}})$ with $s(x) := [M]_x$, and we would now like to show that the converse also holds if M is compact, i.e. that every R -orientation determines in this way a relative fundamental class. If $\partial M = \emptyset$, then $[M]$ will be an absolute class in $h_n(M)$, and its existence is a direct consequence of the isomorphism $h_n(M) \cong \Gamma(\Theta)$ in Theorem 49.21. It is not difficult to show (see Exercise 50.1) that if M is equipped with a triangulation, then the constructions of fundamental cycles in Lecture 30 via simplicial homology give rise to relative fundamental classes in this sense, living in $H_n(M, \partial M; \mathbb{Z})$ if the triangulation is oriented, and $H_n(M, \partial M; \mathbb{Z}_2)$ if orientations are ignored. Without triangulations, we have the following more general result:

THEOREM 50.1. *If M is a compact manifold with boundary carrying an R -orientation $s \in \Gamma(\Theta^R|_{\mathring{M}})$, then for any compactly supported homology theory h_* with coefficients R , s determines a unique relative fundamental class $[M] \in h_n(M, \partial M)$ with $i_*^x[M] = s(x)$ for all $x \in \mathring{M}$, and if M is connected, then the R -module $h_n(M, \partial M)$ is isomorphic to R and is generated by $[M]$.*

PROOF. 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

$$\begin{array}{ccccc} h_n(M, \partial M) & \xrightarrow{\cong} & h_n(M, M \setminus M_\epsilon) & \xleftarrow{\cong} & h_n(\mathring{M}, \mathring{M} \setminus M_\epsilon) \\ & \searrow^{i_*^x} & \downarrow j_{x, M_\epsilon} & & \downarrow j_{x, M_\epsilon} \\ & & h_n(M, M \setminus \{x\}) & \xleftarrow{\cong} & h_n(\mathring{M}, \mathring{M} \setminus \{x\}) \end{array}$$

where several maps are labeled as isomorphisms due to homotopy invariance. An R -orientation $s \in \Gamma(\Theta|_{\mathring{M}})$ determines a generator $[M]_x := s(x) \in h_n(M|x) \cong R$ 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 49.21 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 $\epsilon > 0$, which can always be arranged sufficiently small so that any given interior point $x \in \mathring{M}$ lies in M_ϵ . (Hint: The isomorphism of Theorem 49.21 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 $h_n(M, \partial M) \cong h_n(\mathring{M}|M_\epsilon) \cong \Gamma_c(\Theta|_{M_\epsilon}) = \Gamma(\Theta|_{M_\epsilon})$, 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 $\Gamma(\Theta|_{M_\epsilon}) \rightarrow \Theta_x \cong R$ defined by evaluating sections at any point $x \in M_\epsilon$ is then an R -module isomorphism that sends $s|_{M_\epsilon}$ to $[M]_x$. \square

50.2. The proof of Theorem 49.21. The unfinished business of the previous lecture was the proof that the natural map

$$J_A : h_n(M | A) \rightarrow \Gamma_c(\Theta|_A) : c \mapsto s_c$$

is an isomorphism and $h_k(M | A) = 0$ for all $k > n$, assuming M is a topological n -manifold without boundary, $A \subset X$ is a closed subset and h_* is a compactly supported axiomatic homology theory. The proof 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 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 the set $\varphi(A) \subset \mathbb{R}^n$ is convex. The usefulness of this condition lies in the fact 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, since intersections of convex subsets in \mathbb{R}^n are always convex.

The following lemma is a simple exercise in point-set topology: the idea is to form a nested sequence of sets that are finite unions of convex balls.

LEMMA 50.2. *In a topological n -manifold M , every compact subset $A \subset M$ can be written as $A = \bigcap_{i=1}^{\infty} A_i$ for a nested sequence of subsets $M \supset A_1 \supset A_2 \supset A_3 \supset \dots \supset A$ such that each A_i is a finite union of compact convex subsets.* \square

PROOF OF THEOREM 49.21. 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 contained in a disk-like neighborhood $\bar{\mathcal{U}} \subset M$, so that for every $x \in A$, the composition

$$\begin{array}{ccccc}
 & & j_{x, \bar{\mathcal{U}}} & & \\
 & & \cong & & \\
 & \swarrow & & \searrow & \\
 h_k(M | \bar{\mathcal{U}}) & \xrightarrow{j_{A, \bar{\mathcal{U}}}} & h_k(M | A) & \xrightarrow{j_{x, A}} & h_k(M | x)
 \end{array}$$


is an isomorphism. We claim that the map $j_{A, \bar{\mathcal{U}}} : h_k(M | \bar{\mathcal{U}}) \rightarrow h_k(M | A)$ is also an isomorphism. Since A and $\bar{\mathcal{U}}$ are contained in a Euclidean neighborhood, this is equivalent via excision and rescaling to the claim that

$$h_k(\mathbb{R}^n, \mathbb{R}^n \setminus \mathbb{D}^n) \xrightarrow{j_{K, \mathbb{D}^n}} h_k(\mathbb{R}^n, \mathbb{R}^n \setminus K)$$

is an isomorphism for any compact convex subset $K \subset \mathring{\mathbb{D}}^n \subset \mathbb{R}^n$ containing the origin. This is true because radial lines outward from the origin remain outside of K as soon as they leave K , giving rise to a deformation retraction of $\mathbb{R}^n \setminus K$ to $\mathbb{R}^n \setminus \mathbb{D}^n$. The inclusion thus induces an isomorphism $h_*(\mathbb{R}^n \setminus K) \rightarrow h_*(\mathbb{R}^n \setminus \mathbb{D}^n)$, and the usual five-lemma trick using the long exact sequences of the pairs $(\mathbb{R}^n, \mathbb{R}^n \setminus \mathbb{D}^n)$ and $(\mathbb{R}^n, \mathbb{R}^n \setminus K)$ then proves the claim. It follows now that $j_{x, A} : h_k(M | A) \rightarrow h_k(M | x)$ is also an isomorphism for every $x \in A$. For $k > n$, this proves $h_k(M | A) = 0$, while for $k = n$, we observe that Θ admits a unique section on $\bar{\mathcal{U}}$ with any given value at x , and it follows that $J_A : h_n(M | A) \rightarrow \Gamma(\Theta|_A)$ is an isomorphism.

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 §34.6. 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 $M \setminus A$ and $M \setminus B$ thus form an open covering of $M \setminus (A \cap B)$. The obvious inclusions of pairs



then give rise to a long exact sequence of the form

$$\begin{aligned}
 \dots \rightarrow h_{k+1}(M | A) \oplus h_{k+1}(M | B) &\rightarrow h_{k+1}(M | A \cap B) \rightarrow h_k(M | A \cup B) \\
 &\rightarrow h_k(M | A) \oplus h_k(M | B) \rightarrow h_k(M | A \cap B) \rightarrow \dots
 \end{aligned}$$

If $k > n$, then the sequence places $h_k(M | A \cup B)$ in between two vanishing terms and thus proves $h_k(M | A \cup B) = 0$. 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

$$0 \rightarrow \Gamma_c(\Theta|_{A \cup B}) \rightarrow \Gamma_c(\Theta|_A) \oplus \Gamma_c(\Theta|_B) \rightarrow \Gamma_c(\Theta|_{A \cap B}),$$

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

$$h_k^X := h_k(M | X) \quad \text{for any } k \geq 0 \text{ and subset } X \subset M,$$

so that the diagram in question looks like

$$\begin{array}{ccccccc}
 h_{n+1}^A \oplus h_{n+1}^B & \longrightarrow & h_{n+1}^{A \cap B} & \longrightarrow & h_n^{A \cup B} & \longrightarrow & h_n^A \oplus h_n^B & \longrightarrow & h_n^{A \cap B} \\
 \downarrow \cong & & \downarrow \cong & & \downarrow J_{A \cup B} & & \cong \downarrow J_A \oplus J_B & & \cong \downarrow J_{A \cap B} \\
 0 & \longrightarrow & 0 & \longrightarrow & \Gamma_c(\Theta|_{A \cup B}) & \longrightarrow & \Gamma_c(\Theta|_A) \oplus \Gamma_c(\Theta|_B) & \longrightarrow & \Gamma_c(\Theta|_{A \cap B})
 \end{array}$$

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 $A_1 \supset A_2 \supset A_3 \supset \dots$, then it also holds for $A_\infty := \bigcap_{i=1}^\infty A_i \subset M$. This requires a direct limit argument, and for this step we must make explicit use of the assumption that the homology theory h_* is compactly supported—or in the case of singular homology, the fact that singular chains are always confined to compact subsets.

Observe first that the sequence of inclusions

$$(M, M \setminus A_1) \hookrightarrow (M, M \setminus A_2) \hookrightarrow (M, M \setminus A_3) \hookrightarrow \dots (M, M \setminus A_\infty)$$

induces a sequence of homomorphisms

$$h_*(M | A_1) \rightarrow h_*(M | A_2) \rightarrow h_*(M | A_3) \rightarrow \dots \rightarrow h_*(M | A_\infty),$$

so that $\{h_*(M | A_i)\}_{i=1}^\infty$ forms a direct system of \mathbb{Z} -graded R -modules, with $h_*(M | A_\infty)$ as a target. We claim that the sequence of maps $h_*(M | A_i) \rightarrow h_*(M | A_\infty)$ satisfies the universal property so

large. The desired map φ_{∞} is thus defined on any $s \in \Gamma(\Theta|_{A_{\infty}})$ by extending s to A_i and then applying φ_i .

With these preliminaries in place, the hypothesis $h_k(M|A_i) = 0$ for $k > n$ and $i \in \mathbb{N}$ now implies

$$h_k(M|A_{\infty}) \cong \varinjlim \{h_k(M|A_i)\}_{i=1}^{\infty} = 0 \quad \text{for } k > n.$$

For $k = n$, we can combine both of the direct systems above into a commuting diagram

$$\begin{array}{ccccccc} h_n(M|A_1) & \longrightarrow & h_n(M|A_2) & \longrightarrow & h_n(M|A_3) & \longrightarrow & \dots \longrightarrow h_n(M|A_{\infty}) \\ \cong \downarrow J_{A_1} & & \cong \downarrow J_{A_2} & & \cong \downarrow J_{A_3} & & \downarrow J_{A_{\infty}} \\ \Gamma(\Theta|_{A_1}) & \longrightarrow & \Gamma(\Theta|_{A_2}) & \longrightarrow & \Gamma(\Theta|_{A_3}) & \longrightarrow & \dots \longrightarrow \Gamma(\Theta|_{A_{\infty}}), \end{array}$$

so that the sequence of isomorphisms $J_{A_i} : h_n(M|A_i) \rightarrow \Gamma(\Theta|_{A_i})$ 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 $J_{A_{\infty}}$, 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 $\Gamma(\Theta|_{A_{\infty}}) \rightarrow h_n(M|A_{\infty})$ that is the inverse of $J_{A_{\infty}}$.

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 $\{A_{\alpha} \subset M\}_{\alpha \in J}$ **separated** if they admit a collection of open neighborhoods $\{A_{\alpha} \subset \mathcal{U}_{\alpha} \subset M\}_{\alpha \in J}$ such that $\mathcal{U}_{\alpha} \cap \mathcal{U}_{\beta} = \emptyset$ 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 $A := \bigcup_{\alpha \in J} A_{\alpha}$. The point of the separation condition is that if we write $\mathcal{U} := \bigcup_{\alpha \in J} \mathcal{U}_{\alpha}$, then $(\mathcal{U}, \mathcal{U} \setminus A) \cong \coprod_{\alpha} (\mathcal{U}_{\alpha}, \mathcal{U}_{\alpha} \setminus A_{\alpha})$, 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

$$\begin{array}{ccccccc} h_n(M|A) & \xleftarrow{\cong} & h_n(\mathcal{U}|A) & \xleftarrow{\cong} & \bigoplus_{\alpha} h_n(\mathcal{U}_{\alpha}|A_{\alpha}) & \xrightarrow{\cong} & \bigoplus_{\alpha} h_n(M|A_{\alpha}) \\ \downarrow J_A & & \downarrow J_A & & \downarrow \bigoplus_{\alpha} J_{A_{\alpha}} & & \cong \downarrow \bigoplus_{\alpha} J_{A_{\alpha}} \\ \Gamma_c(\Theta|_A) & \xlongequal{\quad} & \Gamma_c(\Theta|_A) & \xleftarrow{\cong} & \bigoplus_{\alpha} \Gamma(\Theta|_{A_{\alpha}}) & \xlongequal{\quad} & \bigoplus_{\alpha} \Gamma(\Theta|_{A_{\alpha}}), \end{array}$$

where the isomorphism $\bigoplus_{\alpha} \Gamma(\Theta|_{A_{\alpha}}) \rightarrow \Gamma_c(\Theta|_A)$ sends each $\sum_{\alpha} s_{\alpha} \in \bigoplus_{\alpha} \Gamma(\Theta|_{A_{\alpha}})$ to the unique section $s \in \Gamma(\Theta|_A)$ such that $s|_{A_{\alpha}} = s_{\alpha}$ for every $\alpha \in J$. Note that s necessarily has compact support since the sets A_{α} are compact and only finitely many of the summands in $\sum_{\alpha} s_{\alpha} \in \bigoplus_{\alpha} \Gamma(\Theta|_{A_{\alpha}})$ can be nonzero. Conversely, a section $s \in \Gamma(\Theta|_A)$ with compact support can be nonzero only on finitely many of the components A_{α} , and is therefore in the image of the map from the direct sum. This proves that $J_A : h_n(M|A) \rightarrow \Gamma_c(\Theta|_A)$ 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 Lemma 50.2, 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 Lemma 50.2, 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 $X^* := X \cup \{\infty\}$ with $\infty \notin X$, where a subset of X^* is considered open if it is either an open set in X or takes the form $(X \setminus K) \cup \{\infty\}$ 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 $\text{dist}(\cdot, \cdot)$ denote a metric on M^* and exhaust A by the countable sequence of subsets

$$\begin{aligned} A_1 &:= \{x \in A \mid 1 \leq \text{dist}(x, \infty) < \infty\}, \\ A_2 &:= \{x \in A \mid 1/2 \leq \text{dist}(x, \infty) \leq 1\}, \\ A_3 &:= \{x \in A \mid 1/3 \leq \text{dist}(x, \infty) \leq 1/2\}, \\ &\dots \end{aligned}$$

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

$$B := \bigcup_{j=1}^{\infty} A_{2j-1}, \quad C := \bigcup_{j=1}^{\infty} A_{2j}, \quad B \cap C = \bigcup_{j=1}^{\infty} \{x \in A \mid \text{dist}(x, \infty) = 1/j\},$$

all of which are unions of separated collections of compact sets. We then conclude from step 2 that the theorem also holds for $A = B \cup C$. \square

50.3. Exercises.

EXERCISE 50.1. Show that under the natural isomorphism $H_*^{\Delta}(K) \cong H_*(|K|)$ between simplicial and singular homologies, the simplicial fundamental classes constructed in Lecture 30 for compact triangulated manifolds also satisfy the conditions required in the present lecture for relative fundamental classes on general compact manifolds.

Hint: For each interior point $x \in \overset{\circ}{M}$, you can assume after a harmless modification of the triangulation that x lies in the interior of one of its n -simplices.

EXERCISE 50.2. In this exercise, assume M is an n -manifold with boundary and h_* is a compactly supported homology theory whose coefficients are the ring R . Let $\Theta^M \rightarrow M$ and $\Theta^{\partial M} \rightarrow \partial M$ denote the associated orientation bundles over $\overset{\circ}{M}$ and ∂M respectively, hence

$$\Theta_y^M = h_n(M|y) \cong R, \quad \text{for } y \in \overset{\circ}{M}, \quad \Theta_x^{\partial M} = h_{n-1}(\partial M|x) \cong R \quad \text{for } x \in \partial M.$$

The goal is to prove the following two statements:

- (1) There is a natural homomorphism $\partial_* : \Gamma(\Theta^M) \rightarrow \Gamma(\Theta^{\partial M})$ that associates to any R -orientation of M an R -orientation ∂M .
- (2) If M is compact with an R -orientation $s \in \Gamma(\Theta^M)$ and $[M] \in h_n(M, \partial M)$ and $[\partial M] \in h_{n-1}(\partial M)$ are the fundamental classes corresponding to the R -orientations s and $\partial_* s \in \Gamma(\Theta^{\partial M})$, then the connecting homomorphism $\partial_* : h_n(M, \partial M) \rightarrow h_{n-1}(\partial M)$ in the long exact sequence of the pair $(M, \partial M)$ satisfies

$$\partial_*[M] = [\partial M].$$

It follows via the long exact sequence that $[\partial M] \in h_{n-1}(\partial M)$ lies in the kernel of the map $h_{n-1}(\partial M) \rightarrow h_{n-1}(M)$ induced by the inclusion $\partial M \hookrightarrow M$. One can use this e.g. to define natural transformations of the form $[(M, \varphi)] \mapsto \varphi_*[M]$ from bordism theory to singular homology,

$$\Omega_n^O(X) \rightarrow H_n(X; \mathbb{Z}_2) \quad \text{and} \quad \Omega_n^{SO}(X) \rightarrow H_n(X; \mathbb{Z}),$$

using fundamental classes $[M] \in H_n(M; \mathbb{Z}_2)$ of unoriented closed manifolds in the first case and integral fundamental classes $[M] \in H_n(M; \mathbb{Z})$ of oriented closed manifolds in the second case.

Consider the following setup. Let $\mathbb{H}^n := [0, \infty) \times \mathbb{R}^{n-1} \subset \mathbb{R}^n$ and $\partial\mathbb{H}^n := \{0\} \times \mathbb{R}^{n-1}$, and for the closed unit disk $\mathbb{D}^n \subset \mathbb{R}^n$, denote

$$\mathbb{D}_+^n := \mathbb{D}^n \cap \mathbb{H}^n, \quad \mathbb{D}^{n-1} := \mathbb{D}^n \cap \partial\mathbb{H}^n \subset \mathbb{D}_+^{n-1}.$$

Now given a point $x \in \partial M$, choose a closed neighborhood $\mathcal{D}_+^n \subset M$ of x in M , making $\mathcal{D}^{n-1} := \mathcal{D}_+^n \cap \partial M$ likewise a closed neighborhood of x in ∂M , such that there exists a homeomorphism of pairs


$$(M, \partial M) \supset (\mathcal{D}_+^n, \mathcal{D}^{n-1}) \cong (\mathbb{D}_+^n, \mathbb{D}^{n-1})$$

identifying $x \in \partial M$ with the origin $0 \in \mathbb{D}^{n-1} \subset \mathbb{D}_+^n$. In this neighborhood, choose also an interior point $y \in \mathcal{D}_+^n \setminus \mathcal{D}^{n-1}$ and let

$$\ell \subset \mathcal{D}_+^n$$

denote the path from x to y that is identified under the homeomorphism $(\mathcal{D}_+^n, \mathcal{D}^{n-1}) \cong (\mathbb{D}_+^n, \mathbb{D}^{n-1})$ with a straight line segment in \mathbb{D}_+^n . Note that topologically, \mathcal{D}_+^n is homeomorphic to \mathbb{D}^n and is thus a compact n -manifold with boundary, whose boundary we will denote as usual by $\partial\mathcal{D}_+^n \cong S^{n-1}$.

(a) Stare for a while at the following diagram



All maps in this diagram are either connecting homomorphisms from long exact sequences of pairs or are induced by obvious inclusions of pairs. Convince yourself that the diagram commutes, and then explain which of the Eilenberg-Steenrod axioms imply that each of the maps labelled with the symbol “ \cong ” is an isomorphism.

- (b) Deduce from the diagram itself that each of the maps labelled with the symbol “ \approx ” is also an isomorphism.
- (c) Use the diagram in part (a) to construct homomorphisms $\Phi : h_n(M, \partial M) \rightarrow h_n(\mathcal{D}_+^n, \partial \mathcal{D}_+^n)$ and $\Psi : \tilde{h}_{n-1}(\partial M) \rightarrow \tilde{h}_{n-1}(\partial \mathcal{D}_+^n)$ such that the following diagram commutes

$$\begin{array}{ccc}
 & & \Theta_y^M \longleftarrow \cong \Theta_y^{\mathcal{D}_+^n} \\
 & \nearrow & \uparrow \cong \\
 h_n(M, \partial M) & \xrightarrow{\Phi} & h_n(\mathcal{D}_+^n, \partial \mathcal{D}_+^n) \\
 \downarrow \hat{c}_* & & \downarrow \hat{c}_* \cong \\
 \tilde{h}_{n-1}(\partial M) & \xrightarrow{\Psi} & \tilde{h}_{n-1}(\partial \mathcal{D}_+^n) \\
 \downarrow & & \downarrow \cong \\
 \Theta_x^{\partial M} \longleftarrow \cong \Theta_x^{\mathcal{D}_+^{n-1}} & \xrightarrow{\cong} & \Theta_x^{\partial \mathcal{D}_+^n}
 \end{array} ,$$

where all maps are once again either connecting homomorphisms or induced by obvious inclusions of pairs. Explain why each of the maps labelled with “ \cong ” is an isomorphism.

- (d) Use the diagram in part (c) to construct an isomorphism $\Theta_y^M \rightarrow \Theta_x^{\partial M}$ that is associated to any boundary point $x \in \partial M$ and any interior point $y \in \overset{\circ}{M}$ in a sufficiently small neighborhood of x in M , and convince yourself that this isomorphism does not depend on the choices involved in the construction. Derive from it a homomorphism

$$\Gamma(\Theta^M) \xrightarrow{\hat{c}_*} \Gamma(\Theta^{\partial M})$$

that sends R -orientations of M to R -orientations of ∂M .

- (e) Assuming M and ∂M are endowed with R -orientations related via the map in part (d), prove that the associated fundamental classes are related by $\partial_*[M] = [\partial M]$.

EXERCISE 50.3. The following is a variant of Exercise 50.2 that applies specifically to singular homology with coefficients in the ring R . Assume M satisfies the hypotheses of Theorem 50.1 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 $\partial_* : H_n(M, \partial M; R) \rightarrow H_{n-1}(\partial M; R)$ 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 49.5(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 $(M, \partial M, \partial M \setminus N)$ and notice that $H_{n-1}(\partial M, \partial M \setminus N) \cong H_{n-1}(N)$ by excision.

Remark: In the setting of triangulated manifolds, a similar result with coefficients \mathbb{Z}_2 or \mathbb{Z} can be deduced from Propositions 30.3 and 30.10 via the natural isomorphism between simplicial and singular homology, in light of Exercise 50.1.

EXERCISE 50.4 (*). The goal of this exercise is to prove that the product $M \times N$ of two R -oriented manifolds inherits a natural R -orientation, and in the compact case, the associated fundamental class $[M \times N]$ is given by the cross product $[M] \times [N]$. Note that if M and N

are topological manifolds of dimensions m and n respectively with boundary, then $M \times N$ is a topological $(m+n)$ -manifold with boundary

$$\partial(M \times N) = (\partial M \times N) \cup (M \times \partial N),$$

thus in terms of the product for pairs of spaces defined in §48.2, we have

$$(M, \partial M) \times (N, \partial N) = (M \times N, \partial(M \times N)).$$

We work in the setting of an axiomatic homology theory h_* with coefficients R , and assume the existence of a cross product

$$h_k(X, A) \otimes h_\ell(Y, B) \rightarrow h_{k+\ell}((X, A) \times (Y, B)) : a \otimes b \mapsto a \times b$$

for arbitrary pairs of spaces (X, A) and (Y, B) .⁷⁸ One can formulate various axioms to guarantee that the cross product in h_* has desirable properties: these should include associativity, graded commutativity and a unit property as in §47.1, plus e.g. a Leibniz rule for feeding products into connecting homomorphisms. We will not go into the details here, but the crucial consequence of these axioms should be that the cross product on h_* matches the *cellular* cross product when restricted to CW-pairs. This implies, for instance, that the map

$$(50.2) \quad h_m(\mathbb{D}^m, \partial\mathbb{D}^m) \otimes h_n(\mathbb{D}^n, \partial\mathbb{D}^n) \xrightarrow{\times} h_{m+n}(\mathbb{D}^m \times \mathbb{D}^n, \partial(\mathbb{D}^m \times \mathbb{D}^n))$$

is an isomorphism; let us take this as an axiom in the following.

- (a) Show that the analogues of the map (50.2) in cellular and singular homology with ring coefficients are isomorphisms.

Remark: The Künneth formula offers one convenient approach to this, but only if R is a principal ideal domain. Try to do without that assumption.

- (b) Given an m -manifold M and an n -manifold N with interior points $x \in \mathring{M}$ and $y \in \mathring{N}$, we have

$$(M, M \setminus \{x\}) \times (N, N \setminus \{y\}) = (M \times N, (M \times N) \setminus \{(x, y)\}),$$

so that the relative cross product defines a map

$$\Theta_x^M \otimes \Theta_y^N \xrightarrow{\times} \Theta_{(x,y)}^{M \times N}.$$

Show that this map is an isomorphism, and that it gives rise to a homomorphism

$$\Gamma(\Theta^M) \otimes \Gamma(\Theta^N) \xrightarrow{\times} \Gamma(\Theta^{M \times N}) : s \otimes t \mapsto s \times t$$

given by $(s \times t)(x, y) = s(x) \times t(y)$. Conclude that if $s \in \Gamma(\Theta^M)$ and $t \in \Gamma(\Theta^N)$ are R -orientations, then so is $s \times t \in \Gamma(\Theta^{M \times N})$.

- (c) Deduce via the naturality of the cross product with respect to maps of the form $(M, \partial M) \rightarrow (M, M \setminus \{x\})$ and $(N, \partial N) \rightarrow (N, N \setminus \{y\})$ that if M and N are compact manifolds with R -orientations and $M \times N$ is equipped with the product R -orientation arising from part (b), then the corresponding fundamental classes $[M] \in h_m(M, \partial M)$, $[N] \in h_n(N, \partial N)$ and $[M \times N] \in h_{m+n}(M \times N, \partial(M \times N))$ are related by

$$[M] \times [N] = [M \times N].$$

⁷⁸If necessary, it would also be possible to work with the less ambitious assumption that the cross product is defined whenever the two subsets $A \times Y, X \times B \subset X \times Y$ form an excisive couple for h_* . One can use excision and collar neighborhoods to show that this is always true when the pairs are $(M, \partial M)$ and $(N, \partial N)$ for compact manifolds M and N .

51. Poincaré duality

51.1. The statement. The classical perspective on Poincaré duality is demonstrated by Figure 27. 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 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 $C_*^{\text{CW}}(M)$ the cellular chain complex for its dual cell complex. Evidently, there is a natural bijection

$$C_k^\Delta(M) \rightarrow C_{n-k}^{\text{CW}}(M),$$

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 $C_{\text{CW}}^*(M)$. Thus it can be more properly interpreted as a bijective chain map

$$C_*^\Delta(M) \rightarrow C_{\text{CW}}^{n-*}(M),$$

therefore giving rise to an isomorphism $H_k(M) \cong H^{n-k}(M)$ for each $k = 0, \dots, n$.

REMARK 51.1. 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 $C_k^\Delta(M) \rightarrow C_{n-k}^{\text{CW}}(M)$, 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 $C_*^\Delta(M) \rightarrow C_{\text{CW}}^{n-*}(M)$. 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 $H^k(M) \rightarrow H_{n-k}(M)$, 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 51.2 (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*

$$H^k(M; R) \xrightarrow{\text{PD}} H_{n-k}(M; R) : \varphi \mapsto \varphi \cap [M]$$




FIGURE 27. A triangulation of a surface and its dual cell decomposition.

is an isomorphism for every $k \in \mathbb{Z}$.

REMARK 51.3. The proof of Theorem 51.2 will be based mostly on the formal (i.e. axiomatic) properties of homology and cohomology, thus it would also be possible—with some effort—to formulate a more general version of Poincaré duality that identifies axiomatic cohomology groups $h^k(M)$ with axiomatic homology groups $h_{n-k}(M)$ under suitable assumptions. The use of the cap product in the definition of the map $\text{PD} : H^k(M; R) \rightarrow H_{n-k}(M; R)$ gives a big hint that something more than just the Eilenberg-Steenrod axioms is needed, as the axioms do not automatically give rise to extra algebraic structures like the cup and cap products—in fact, the axioms do not even provide an evaluation pairing of $h^*(M)$ with $h_*(M)$, as seems natural in the realm of singular (co)homology. It is possible and sometimes desirable, however, to add extra axioms so that the existence of a cup product on h^* satisfying reasonable properties is explicitly assumed, giving rise to the notion of a *multiplicative* cohomology theory. One can similarly introduce axioms for cross products on both h^* and h_* , as well as a so-called *duality pairing* to play the role of a cap product intertwining h^* with h_* . The required definitions can be found e.g. in [tD08, §17.2–3 and §18.2]. In the following exposition of Poincaré duality and intersection theory, we will make our lives easier by explicitly using singular homology and cohomology, but the properties we will make use of in the proof are mostly axiomatic in nature.

CONVENTION. Unless otherwise specified, the coefficients for singular homology and cohomology in this lecture will always be the fixed commutative ring R with unit, and R will usually be omitted from the notation.

51.2. Applications. Before getting into the proof of the duality theorem, 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 51.4. *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 51.5. *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, one can reach the same conclusion using the \mathbb{Z}_2 Betti numbers, thanks to Exercise 45.12(c). □

Here is an application of Corollary 51.5 that plays a fundamental role in bordism theory. 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 [Roh51], 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 51.6. *There is no compact 5-manifold with boundary homeomorphic to $\mathbb{C}\mathbb{P}^2$.*

PROOF. Suppose M is a compact manifold with $\partial M \cong \mathbb{C}\mathbb{P}^2$. We can then construct a closed 5-manifold \widehat{M} by gluing M to a copy of itself along the matching boundary,

$$\widehat{M} := M \cup_{\mathbb{C}\mathbb{P}^2} M,$$

and by Corollary 51.5, $\chi(\widehat{M}) = 0$. But according to Exercise 51.1, we also have

$$\chi(\widehat{M}) = 2\chi(M) - \chi(\mathbb{C}\mathbb{P}^2),$$

a formula that admits an easy interpretation if 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 $-\chi(\mathbb{C}\mathbb{P}^2)$ as a correction term. Thanks to its cell decomposition with one cell in each even dimension, we know the homology of $\mathbb{C}\mathbb{P}^2$ and therefore know its Euler characteristic: it is 3, implying that $2\chi(M) - \chi(\mathbb{C}\mathbb{P}^2)$ can never be 0. □

Poincaré duality also provides considerable information about the ring structure of $H^*(M)$ as a consequence of the relation $\langle \psi \cup \varphi, [M] \rangle = \langle \psi, \varphi \cap [M] \rangle$. For each $k = 0, \dots, n$, consider the bilinear form $Q : H^k(M; R) \times H^{n-k}(M; R) \rightarrow R$ defined via the R -module homomorphism

$$H^k(M) \otimes H^{n-k}(M) \xrightarrow{Q} R$$

$$\varphi \otimes \psi \mapsto Q(\varphi, \psi) := \langle \varphi \cup \psi, [M] \rangle.$$

For reasons that we will discuss in the next two lectures, 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 $H_{\text{free}}^*(M) := H^*(M; \mathbb{Z})/\text{torsion}$,

$$H_{\text{free}}^k(M) \otimes H_{\text{free}}^{n-k}(M) \xrightarrow{Q} \mathbb{Z}.$$

For a general pair of R -modules A and B , a bilinear map $Q : A \times B \rightarrow G$ (or equivalently an R -module homomorphism $Q : A \otimes B \rightarrow G$) is called **nonsingular** if the maps $A \rightarrow \text{Hom}(B, G) : a \mapsto Q(a, \cdot)$ and $B \rightarrow \text{Hom}(A, G) : b \mapsto Q(\cdot, b)$ are both isomorphisms.

COROLLARY 51.7. *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*

$$H^k(M; \mathbb{K}) \otimes_{\mathbb{K}} H^{n-k}(M; \mathbb{K}) \xrightarrow{Q} \mathbb{K}$$

is nonsingular for every $k = 0, \dots, n$, and if M is oriented, Q descends to the free part of $H^(M; \mathbb{Z})$ as a nonsingular bilinear form $H_{\text{free}}^k(M) \otimes H_{\text{free}}^{n-k}(M) \rightarrow \mathbb{Z}$.*

PROOF. With integer coefficients, we saw in Exercise 45.12(a) that the canonical map $h : H_{\text{free}}^{n-k}(M) \rightarrow \text{Hom}(H_{n-k}^{\text{free}}(M), \mathbb{Z}) : \varphi \mapsto \langle \varphi, \cdot \rangle$ is an isomorphism. Since the duality map $\text{PD} : H^k(M) \rightarrow H_{n-k}(M)$ is also an isomorphism, it and its inverse each map torsion to torsion and thus descend to the free parts as isomorphisms $H_{\text{free}}^k(M) \cong H_{n-k}^{\text{free}}(M)$. We can then compose h with the dualization of PD to form an isomorphism

$$H_{\text{free}}^{n-k}(M) \xrightarrow[\cong]{h} \text{Hom}(H_{n-k}^{\text{free}}(M), \mathbb{Z}) \xrightarrow[\cong]{\text{PD}^*} \text{Hom}(H_{\text{free}}^k(M), \mathbb{Z})$$

$$\searrow \Phi \nearrow$$

$$\cong$$

To see what this map actually is, we choose $\psi \in H_{\text{free}}^{n-k}(M)$ and $\varphi \in H_{\text{free}}^k(M)$ 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 $h : H^{n-k}(M; \mathbb{K}) \rightarrow \text{Hom}_{\mathbb{K}}(H_{n-k}(M; \mathbb{K}), \mathbb{K})$ is a vector space isomorphism. \square

COROLLARY 51.8. *If M is a closed oriented n -manifold and $\varphi \in H^k(M; \mathbb{Z})$ is a primitive⁷⁹ element for some $k \in \{0, \dots, n\}$, 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.*

PROOF. The primitivity hypothesis means that the projection of φ to $H_{\text{free}}^k(M)$ is nontrivial and generates a subgroup $H \subset H_{\text{free}}^k(M)$ such that $H_{\text{free}}^k(M)/H$ 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

⁷⁹Recall 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 $m \geq 2$. In particular, this rules out that φ is torsion, since $m\varphi = 0$ would imply $(m+1)\varphi = \varphi$.

element in a basis of $H_{\text{free}}^k(M)$, so that there exists a homomorphism $\Phi : H_{\text{free}}^k(M) \rightarrow \mathbb{Z}$ 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

As important applications of the nonsingularity of the intersection form (see Exercises 51.3 and 51.4), one can fully compute the ring structures of $H^*(\mathbb{C}\mathbb{P}^n; \mathbb{Z})$ and $H^*(\mathbb{R}\mathbb{P}^n; \mathbb{Z}_2)$, giving ring isomorphisms of each to quotients of polynomial rings with one generator,

$$\begin{aligned} H^*(\mathbb{C}\mathbb{P}^n; \mathbb{Z}) &\cong \mathbb{Z}[\alpha]/(\alpha^{n+1}), & |\alpha| &= 2, \\ H^*(\mathbb{R}\mathbb{P}^n; \mathbb{Z}_2) &\cong \mathbb{Z}_2[\alpha]/(\alpha^{n+1}), & |\alpha| &= 1. \end{aligned}$$

These results can also be extended without much difficulty to compute the cohomology rings of the infinite-dimensional cell complexes $\mathbb{C}\mathbb{P}^\infty$ and $\mathbb{R}\mathbb{P}^\infty$, which are not manifolds, but the fact that their finite-dimensional skeleta are manifolds gives enough information to carry out the computation. Note that the computation of $H^*(\mathbb{C}\mathbb{P}^n; \mathbb{Z})$ fills in the last remaining gap in our proof from Lecture 41 (see Theorem 41.1) that all maps $f : \mathbb{C}\mathbb{P}^n \rightarrow \mathbb{C}\mathbb{P}^n$ have fixed points when n is even.

51.3. 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 49.21 to a distinguished homology class

$$[M]_K := J_K^{-1}(s) \in H_n(M|K).$$

Recall from Lecture 48 that the relative cohomology with coefficients in the ring R admits a cap product pairing

$$\cap : H^k(M, M \setminus K) \otimes H_n(M, M \setminus K) \rightarrow H_{n-k}(M),$$

which is well defined in this case because the subsets $M \setminus K$ and \emptyset in M trivially form an excisive couple. We can therefore define a “restricted” duality map by

$$\text{PD}_K : H^k(M|K) \rightarrow H_n(M) : \varphi \mapsto \varphi \cap [M]_K.$$

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

$$i : (M, M \setminus K') \hookrightarrow (M, M \setminus K)$$

satisfies

$$i_*[M]_{K'} = [M]_K.$$

The naturality property (48.5) of the cap product then implies that for all $\varphi \in H^k(M|K)$,

$$i_*(i^*\varphi \cap [M]_{K'}) = \text{PD}_{K'}(i^*\varphi) = \varphi \cap [M]_K = \text{PD}_K(\varphi),$$

where “ i_* ” has disappeared in the second expression since $\text{PD}_{K'}(i^*\varphi)$ is an *absolute* homology class and $i : M \rightarrow M$ is just the identity map. The result is a commutative diagram

$$(51.1) \quad \begin{array}{ccc} H^k(M | K) & \xrightarrow{i^*} & H^k(M | K') \\ & \searrow \text{PD}_K & \downarrow \text{PD}_{K'} \\ & & H_{n-k}(M), \end{array}$$

which means that we can view the maps $\text{PD}_K : H^k(M | K) \rightarrow H_{n-k}(M)$ 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

$$\text{PD} : \varinjlim \{H^k(M | K)\}_K \rightarrow H_{n-k}(M).$$

DEFINITION 51.9. For any space X , we define the **compactly supported cohomology** of X with coefficients in an abelian group G as the direct limit

$$H_c^*(X) = H_c^*(X; G) := \varinjlim \{H^*(X | K; G)\}_K,$$

where K ranges over the set of all compact subsets of X , ordered by inclusion and forming a direct system via the maps $H^*(X | K; G) \rightarrow H^*(X | K'; G)$ induced by inclusions $(X, X \setminus K') \hookrightarrow (X, X \setminus K)$ whenever $K \subset K'$.

With this definition in place, the previous discussion produces natural homomorphisms

$$\text{PD} : H_c^k(M; R) \rightarrow H_{n-k}(M; R)$$

for every $k \in \mathbb{Z}$ whenever M is a (possibly noncompact) manifold of dimension n with a fixed R -orientation. We can now state the noncompact version of Poincaré duality:

THEOREM 51.10. *For every R -oriented topological n -manifold M and every $k \in \mathbb{Z}$, the map*

$$\text{PD} : H_c^k(M; R) \rightarrow H_{n-k}(M; R),$$

defined as the direct limit of the maps $\text{PD}_K : H^k(M | K; R) \rightarrow H_{n-k}(M; R) : \varphi \mapsto \varphi \cap [M]_K$ for all compact subsets $K \subset M$, is an isomorphism.

Before continuing, let us make some observations about the properties of compactly supported cohomology, leaving the proofs as exercises.

First, $H_c^k(M)$ is nothing new if M is compact, as in this case the existence of a maximal element $M \subset M$ in the directed set of compact subsets gives rise to a natural isomorphism $H_c^k(M) \cong H^k(M)$. One can show moreover that this isomorphism identifies the map $\text{PD} : H_c^k(M) \rightarrow H_{n-k}(M)$ defined above with the usual map $\varphi \mapsto \varphi \cap [M]$; see Exercise 51.5. For this reason, Theorem 51.10 implies the compact version of Poincaré duality, Theorem 51.2.

Second, $H_c^*(M)$ is an invariant of M up to homeomorphism, but not homotopy type, and certain simple computations therefore work out a bit differently than one might at first expect. One important example (see Exercise 51.6) that we'll need to make use of is

$$H_c^k(\mathbb{R}^n; G) \cong \begin{cases} G & \text{if } k = n, \\ 0 & \text{if } k \neq n, \end{cases}$$

which bears more resemblance to $H^k(\mathbb{R}^n | \mathbb{D}^n)$ than $H^k(\mathbb{R}^n)$, and this is of course not a coincidence. It turns out that H_c^* does not define a contravariant functor on **Top**, because an arbitrary continuous map $f : X \rightarrow Y$ does not induce a well-defined homomorphism $f^* : H_c^*(Y) \rightarrow H_c^*(X)$ unless it is also *proper*, meaning that preimages of compact sets are compact. Homeomorphisms do have

this property, and that's why $H_c^*(X)$ is a topological invariant, but it fails to be a homotopy type invariant, as demonstrated by the result quoted above, which shows that

$$H_c^*(\mathbb{R}^n) \not\cong H_c^*(\{*\}).$$

Finally, it is interesting to note that $H_c^*(X)$ can be constructed without direct limits as the cohomology of a cochain complex consisting of “compactly supported” cochains; see Exercise 51.6 for the details.

51.4. Proof of the theorem. We proceed toward the proof of Theorem 51.10. The argument that follows will bear a resemblance to the inductive construction of the fundamental class in the previous lecture. We start with a purely local result to begin the induction.

LEMMA 51.11. *For either choice of orientation of \mathbb{R}^n , the map $PD : H_c^k(\mathbb{R}^n) \rightarrow H_{n-k}(\mathbb{R}^n)$ is an isomorphism for every $k \in \mathbb{Z}$.*

PROOF. There is an obvious cofinal set of compact subsets to use in computing $H_c^k(\mathbb{R}^n) = \varinjlim \{H^k(\mathbb{R}^n | K)\}_K$: every compact subset $K \subset \mathbb{R}^n$ is contained in the disk \mathbb{D}_r^n of sufficiently large radius $r > 0$, and the natural maps $H^k(\mathbb{R}^n | \mathbb{D}_r^n) \rightarrow H^k(\mathbb{R}^n | \mathbb{D}_{r'}^n)$ are isomorphisms for all $r' > r$, thus

$$H_c^k(\mathbb{R}^n) \cong H^k(\mathbb{R}^n | \mathbb{D}^n) \cong \begin{cases} R & \text{if } k = n, \\ 0 & \text{if } k \neq n. \end{cases}$$

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 $H^n(\mathbb{R}^n | \mathbb{D}^n) \cong \text{Hom}(H_n(\mathbb{R}^n | \mathbb{D}^n), R)$ by evaluation of cohomology classes on homology classes, so that $\langle \varphi, \cdot \rangle$ generates $\text{Hom}(H_n(\mathbb{R}^n | \mathbb{D}^n), R)$ 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 . □

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 51.6 reveals that continuous maps $f : X \rightarrow Y$ do not always induce homomorphisms $f^* : H_c^*(Y) \rightarrow H_c^*(X)$ unless an additional condition is imposed, i.e. $f : X \rightarrow Y$ 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 $X \setminus A$ is a closed subset contained in the open set $X \setminus K$, hence the inclusion $(A, A \setminus K) \hookrightarrow (X, X \setminus K)$ 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 $H^*(X | K) \rightarrow H^*(X | L)$ induced by the inclusion $(X, X \setminus L) \hookrightarrow (X, X \setminus K)$ produces a map $H^*(A | K) \rightarrow H^*(X | L)$:

$$\begin{array}{ccc} & H^*(X | K) & \\ \cong \swarrow & & \searrow \\ H^*(A | K) & \xrightarrow{\quad} & H^*(X | L). \end{array}$$

If we then compose this with the natural map of $H^*(X | L)$ to the direct limit $H_c^*(X)$, it produces a map $H^*(A | K) \rightarrow H_c^*(X)$ 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

$$\begin{array}{ccc} H^*(A | K) & \longrightarrow & H^*(A | K') \\ & \searrow & \downarrow \\ & & H_c^*(X) \end{array}$$

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

$$H_c^*(A) \rightarrow H_c^*(X).$$

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 51.12. *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*

$$\begin{array}{ccc} H_c^k(A) & \longrightarrow & H_c^k(M) \\ \downarrow \text{PD} & & \downarrow \text{PD} \\ H_{n-k}(A) & \longrightarrow & H_{n-k}(M) \end{array}$$

PROOF. Given a compact set $K \subset A$, pick any compact set $L \subset M$ that contains K , and denote the obvious inclusions

$$A \xrightarrow{i} M, \quad (A, A \setminus K) \xrightarrow{i} (M, M \setminus K), \quad (M, M \setminus L) \xrightarrow{j} (M, M \setminus K).$$

We then claim that the diagram

$$\begin{array}{ccccc} & & H^k(M | K) & & \\ & \swarrow \cong & & \searrow j^* & \\ H^k(A | K) & & & & H^k(M | L) \\ \downarrow \text{PD}_K & & & & \downarrow \text{PD}_L \\ H_{n-k}(A) & \xrightarrow{i_*} & & & H_{n-k}(M) \end{array}$$

commutes. To see this, observe that there is another map we could add to this diagram and sensibly denote by PD_K , namely $H^k(M | K) \rightarrow H_{n-k}(M) : \varphi \mapsto \varphi \cap [M]_K$; let's call this one PD'_K to avoid confusion, and note that by (51.1), it satisfies

$$\text{PD}_L \circ j^* = \text{PD}'_K.$$

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 $\varphi \in H^k(M | K)$,

$$i_* \circ \text{PD}_K \circ i^* \varphi = i_* (i^* \varphi \cap [A]_K) = \varphi \cap i_*[A]_K = \varphi \cap [M]_K = \text{PD}'_K(\varphi),$$

thus proving the claim. This implies in particular that for the natural maps $H^*(A|K) \rightarrow H^*(M|L)$ that determine $H_c^*(A) \rightarrow H_c^*(M)$ via the direct limit, the diagram

$$\begin{array}{ccc} H^k(A|K) & \longrightarrow & H^k(M|L) \\ \downarrow \text{PD}_K & & \downarrow \text{PD}_L \\ H_{n-k}(A) & \longrightarrow & H_{n-k}(M) \end{array}$$

always commutes. The rest is essentially abstract nonsense: if we let $\Psi : H_c^k(A) \rightarrow H_{n-k}(M)$ 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 $H^k(A|K) \rightarrow H_{n-k}(M)$ 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

$$\begin{array}{ccccc} & & (M, M \setminus K) & & \\ & \swarrow & & \searrow & \\ (M, M \setminus (K \cup L)) & & & & (M, M \setminus (K \cap L)) \\ & \swarrow & & \searrow & \\ & & (M, M \setminus L) & & \end{array}$$

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:

$$\begin{array}{cccccccc} \dots & \longrightarrow & H^k(M|K \cap L) & \longrightarrow & H^k(M|K) \oplus H^k(M|L) & \longrightarrow & H^k(M|K \cup L) & \longrightarrow & H^{k+1}(M|K \cap L) & \longrightarrow & \dots \\ & & \downarrow \cong & & \downarrow \cong & & \parallel & & \downarrow \cong & & \\ \dots & \longrightarrow & H^k(A \cap B|K \cap L) & \longrightarrow & H^k(A|K) \oplus H^k(B|L) & \longrightarrow & H^k(M|K \cup L) & \longrightarrow & H^{k+1}(A \cap B|K \cap L) & \longrightarrow & \dots \\ & & \downarrow \text{PD}_{K \cap L} & & \downarrow \text{PD}_K \oplus \text{PD}_L & & \downarrow \text{PD}_{K \cup L} & & \downarrow \text{PD}_{K \cap L} & & \\ \dots & \longrightarrow & H_{n-k}(A \cap B) & \longrightarrow & H_{n-k}(A) \oplus H_{n-k}(B) & \longrightarrow & H_{n-k}(M) & \longrightarrow & H_{n-k-1}(A \cap B) & \longrightarrow & \dots \end{array}$$

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,⁸⁰ then the same arguments as in the proof of Lemma 51.12 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 calculation, so we will skip it and simply refer to [Hat02, pp. 246–247]. The result is:

⁸⁰Recall 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 $H_k(A \cap B) \rightarrow H_k(A) \oplus H_k(B)$ or $H_k(A) \oplus H_k(B) \rightarrow H_k(A \cup B)$. 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 51.13. *The diagram in (51.2) 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*

$$(51.3) \quad \begin{array}{ccccccc} \dots & \rightarrow & H_c^k(A \cap B) & \longrightarrow & H_c^k(A) \oplus H_c^k(B) & \longrightarrow & H_c^k(M) \longrightarrow H_c^{k+1}(A \cap B) \rightarrow \dots \\ & & \downarrow \text{PD} & & \downarrow \text{PD} \oplus \text{PD} & & \downarrow \text{PD} & & \downarrow \text{PD} \\ \dots & \rightarrow & H_{n-k}(A \cap B) & \rightarrow & H_{n-k}(A) \oplus H_{n-k}(B) & \rightarrow & H_{n-k}(M) \rightarrow H_{n-k-1}(A \cap B) \rightarrow \dots \end{array}$$

in which both rows are exact.

SKETCH OF THE PROOF. Aside from the tedious verification that (51.2) commutes, the claim that the top row of (51.3) 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. A special case of this phenomenon was discussed in §43.2 (along with the fact that it does *not* hold for inverse limits), and more generally, it follows as a special case of Proposition 39.21, which states that the functor taking chain complexes to their homologies is continuous under direct limits—an exact sequence is nothing other than a chain complex with trivial homology. \square

Applying the five-lemma now gives:

COROLLARY 51.14. *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 51.11 and 51.13 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 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 $\mathcal{V}_1, \mathcal{V}_2, \mathcal{V}_3, \dots$, one can also arrange them into nested sequences of open subsets

$$\mathcal{U}_1 := \mathcal{V}_1 \subset \mathcal{U}_2 := \mathcal{V}_1 \cup \mathcal{V}_2 \subset \mathcal{U}_3 := \mathcal{V}_1 \cup \mathcal{V}_2 \cup \mathcal{V}_3 \subset \dots$$

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 51.15. *Suppose $\mathcal{U}_1 \subset \mathcal{U}_2 \subset \mathcal{U}_3 \subset \dots \subset M$ 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 39.24 provides a natural isomorphism

$$\varinjlim \{H_*(\mathcal{U}_i)\}_{i=1}^{\infty} \xrightarrow{\cong} H_*(\varinjlim \{\mathcal{U}_i\}_{i=1}^{\infty}) = H_*(M).$$

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

$$H_c^*(\mathcal{U}_1) \rightarrow H_c^*(\mathcal{U}_2) \rightarrow H_c^*(\mathcal{U}_3) \rightarrow \dots \rightarrow H_c^*(M),$$

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 be proved by establishing the universal property: if we have a sequence of morphisms $f_i : H_c^*(\mathcal{U}_i) \rightarrow A$ to some other \mathbb{Z} -graded R -module A such that the diagram

$$\begin{array}{ccccccc}
 H_c^*(\mathcal{U}_1) & \longrightarrow & H_c^*(\mathcal{U}_2) & \longrightarrow & H_c^*(\mathcal{U}_3) & \longrightarrow & \dots \longrightarrow H_c^*(M) \\
 & & & & & \searrow f_3 & \downarrow f_\infty \\
 & & & & & & A \\
 & \searrow f_1 & & & & & \\
 & & & & & &
 \end{array}$$

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 $H^k(M|K) \rightarrow H_c^k(M)$ 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 51.12, we now obtain a commutative diagram

$$(51.4) \quad \begin{array}{ccccccc}
 H_c^k(\mathcal{U}_1) & \longrightarrow & H_c^k(\mathcal{U}_2) & \longrightarrow & H_c^k(\mathcal{U}_3) & \longrightarrow & \dots \\
 \downarrow \text{PD} & & \downarrow \text{PD} & & \downarrow \text{PD} & & \\
 H_{n-k}(\mathcal{U}_1) & \longrightarrow & H_{n-k}(\mathcal{U}_2) & \longrightarrow & H_{n-k}(\mathcal{U}_3) & \longrightarrow & \dots
 \end{array}$$

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:

$$\varinjlim \{H_c^k(\mathcal{U}_i)\}_{i=1}^\infty \cong H_c^k(M) \xrightarrow[\cong]{\text{PD}} H_{n-k}(M) \cong \varinjlim \{H_{n-k}(\mathcal{U}_i)\}_{i=1}^\infty.$$

□

PROOF OF THEOREM 51.10. Lemmas 51.11 and 51.13 prove the theorem for all finite unions of convex open sets \mathbb{R}^n , and feeding this into Lemma 51.15 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 51.13 now proves the theorem for all finite unions of open subsets in Euclidean neighborhoods, and we can then present M as a nested union of such subsets and establish the theorem for M via a second application of Lemma 51.15. □

In subsequent applications, it will be important also to have relative versions of the Poincaré duality isomorphism that apply to an R -oriented compact n -manifold M with boundary. In this situation, the fundamental class is a relative class $[M] \in H_n(M, \partial M; R)$, and the relative cap product pairing

$$H^k(X, A) \otimes H_\ell(X, A \cup B) \xrightarrow{\cap} H_{\ell-k}(X, B),$$

which was defined in (48.10) whenever $A, B \subset X$ are an excisive couple, gives us two obvious options for interpreting the formula $\text{PD}(\varphi) := \varphi \cap [M]$, namely as a map

$$H^k(M, \partial M; R) \xrightarrow{\text{PD}} H_{n-k}(M; R) \quad \text{or} \quad H^k(M; R) \xrightarrow{\text{PD}} H_{n-k}(M, \partial M; R).$$

In fact, both of these are isomorphisms, a result that is sometimes called **Lefschetz duality**. A proof of this fact is sketched in Exercise 51.7.

REMARK 51.16. Here is the promised addendum to Remark 51.1. When M is compact and has an oriented triangulation, $C_{\Delta}^k(M; \mathbb{Z})$ 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 $C_{\Delta}^k(M; \mathbb{Z}) = \bigoplus_{\sigma \in \mathcal{K}_k(M)} \mathbb{Z}$, then the dual elements $\varphi_{\sigma} : C_{\Delta}^k(M; \mathbb{Z}) \rightarrow \mathbb{Z}$ 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 $C_{\Delta}^k(M; \mathbb{Z})$. In this case, we obtain a chain isomorphism

$$C_{\Delta}^k(M; \mathbb{Z}) \rightarrow C_{n-k}^{\text{CW}}(M; \mathbb{Z})$$

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 $C_{\Delta}^k(M; \mathbb{Z})$ is now an infinitely-generated free abelian group, so its dual $C_{\Delta}^k(M; \mathbb{Z})$ is not isomorphic to it, but is actually much larger: the cochains φ_{σ} do not form a basis for $C_{\Delta}^k(M; \mathbb{Z})$ since they only span the subgroup of homomorphisms $C_{\Delta}^k(M; \mathbb{Z}) \rightarrow \mathbb{Z}$ that are nonzero on finitely many simplices. As a consequence, $C_{\Delta}^k(M; \mathbb{Z})$ and $C_{n-k}^{\text{CW}}(M; \mathbb{Z})$ 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 $H_c^*(M; \mathbb{Z})$.

51.5. Exercises.

EXERCISE 51.1. 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 Corollary 40.6 then tell you?*

EXERCISE 51.2. Show that the Klein bottle is homeomorphic to the boundary of a compact (and necessarily non-orientable) 3-manifold, but $\mathbb{R}\mathbb{P}^2$ is not.

EXERCISE 51.3 (*). We can now compute the ring structure of $H^*(\mathbb{C}\mathbb{P}^n; \mathbb{Z})$. Take the usual cell decomposition $\mathbb{C}\mathbb{P}^n = e^0 \cup e^2 \cup \dots \cup e^{2n}$, and for $k = 1, \dots, n$, let $\alpha_k \in H^{2k}(\mathbb{C}\mathbb{P}^n; \mathbb{Z}) \cong \mathbb{Z}$ denote the generator that evaluates to 1 on the generator of $H_{2k}(\mathbb{C}\mathbb{P}^n; \mathbb{Z})$ represented by the $2k$ -cell.

- Use Corollary 51.8 to prove $\alpha_k \cup \alpha_{n-k} = \pm \alpha_n$ for every k .
- Generalize part (a) to show that $\alpha_k \cup \alpha_{\ell} = \pm \alpha_{k+\ell}$ for every $k, \ell \in \mathbb{N}$ with $k + \ell \leq n$.

Hint: There is a natural inclusion $\mathbb{C}\mathbb{P}^{k+\ell} \hookrightarrow \mathbb{C}\mathbb{P}^n$ that is a cellular map. How does it act on cohomology?

This proves that the ring $H^*(\mathbb{C}\mathbb{P}^n; \mathbb{Z})$ is generated by the single element $\alpha := \alpha_1 \in H^2(\mathbb{C}\mathbb{P}^n; \mathbb{Z})$, subject only to the relation $\alpha^{n+1} = 0$ since $H^k(\mathbb{C}\mathbb{P}^n; \mathbb{Z}) = 0$ for all $k > 2n$. We conclude that there

is an isomorphism of \mathbb{Z} -graded rings⁸¹

$$H^*(\mathbb{C}P^n; \mathbb{Z}) \cong \mathbb{Z}[\alpha]/(\alpha^{n+1}), \quad |\alpha| = 2,$$

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 $\mathbb{C}P^n \hookrightarrow \mathbb{C}P^\infty$ to find a graded ring isomorphism $H^*(\mathbb{C}P^\infty; \mathbb{Z}) \cong \mathbb{Z}[\alpha]$, where again $|\alpha| = 2$.

EXERCISE 51.4 (*). Compute each of the following cohomology rings:

- (a) $H^*(\mathbb{R}P^n; \mathbb{Z}_2) \cong \mathbb{Z}_2[\alpha]/(\alpha^{n+1})$ with $|\alpha| = 1$.
 (b) $H^*(\mathbb{R}P^\infty; \mathbb{Z}_2) \cong \mathbb{Z}_2[\alpha]$ with $|\alpha| = 1$.

EXERCISE 51.5 (*). Show that if M is compact, there is a natural isomorphism $H_c^*(M) \cong H^*(M)$ which identifies the map PD : $H_c^k(M) \rightarrow H_{n-k}(M)$ defined via the direct limit with the usual duality map $\varphi \mapsto \varphi \cap [M]$.

EXERCISE 51.6. 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 cohomology of the subcomplex $C_c^*(X; G) \subset C^*(X; G)$ consisting of every cochain $\varphi : C_k(X) \rightarrow G$ 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 $f : X \rightarrow Y$ is called **proper**⁸² if for every compact set $K \subset Y$, $f^{-1}(K) \subset X$ is also compact. Show that proper maps $f : X \rightarrow Y$ induce homomorphisms $f^* : H_c^*(Y; G) \rightarrow H_c^*(X; G)$, 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.
 (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.

EXERCISE 51.7. 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

$$(51.5) \quad \text{PD} : H^k(M, \partial M) \rightarrow H_{n-k}(M),$$

$$(51.6) \quad \text{PD} : H^k(M) \rightarrow H_{n-k}(M, \partial M),$$

both defined by $\text{PD}(\varphi) = \varphi \cap [M]$. We would now like to prove that both are isomorphisms, a result known as *Lefschetz duality*.

⁸¹A \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 $ab \in R_{k+\ell}$.

⁸²This 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 51.6.

- (a) Find a cofinal family of compact subsets $A \subset \overset{\circ}{M} := M \setminus \partial M$ such that the natural maps in the diagram

$$H^*(\overset{\circ}{M} | A) \longleftarrow H^*(M | A) \longrightarrow H^*(M, \partial M)$$

are isomorphisms. Use this to find a natural isomorphism (cf. Exercise 39.2)

$$H_c^*(M) \cong H^*(M, \partial M),$$

and deduce via Theorem 51.10 that (51.5) is an isomorphism.

- (b) Show that the long exact sequences of the pair $(M, \partial M)$ in homology and cohomology fit together into a commutative diagram of the form

$$\begin{array}{ccccccccccc} \dots & \longrightarrow & H^k(M, \partial M) & \xrightarrow{j^*} & H^k(M) & \xrightarrow{i^*} & H^k(\partial M) & \xrightarrow{\delta^*} & H^{k+1}(M, \partial M) & \longrightarrow & \dots \\ & & \downarrow \cdot \cap [M] & & \downarrow \cdot \cap [M] & & \downarrow \cdot \cap [\partial M] & & \downarrow \cdot \cap [M] & & \\ \dots & \longrightarrow & H_{n-k}(M) & \xrightarrow{j_*} & H_{n-k}(M, \partial M) & \xrightarrow{\partial_*} & H_{n-k-1}(\partial M) & \xrightarrow{i_*} & H_{n-k-1}(M) & \longrightarrow & \dots \end{array}$$

where $i : \partial M \hookrightarrow M$ and $j : (M, \emptyset) \hookrightarrow (M, \partial M)$ denote the usual inclusions.

- (c) Deduce from the diagram in part (b) that the map in (51.6) is also an isomorphism.
 (d) If M has a triangulation, interpret the isomorphisms (51.5) and (51.6) in terms of the dual cell decomposition.

EXERCISE 51.8. For two closed, connected and oriented manifolds M, N of dimension $n \in \mathbb{N}$, the **degree** $\deg(f) \in \mathbb{Z}$ of a map $f : M \rightarrow N$ is defined to be the unique integer d such that

$$H_n(M; \mathbb{Z}) \xrightarrow{f_*} H_n(N; \mathbb{Z}) : [M] \mapsto d[N].$$

Use the nonsingularity of the intersection form $Q : H^k(M; \mathbb{K}) \otimes H^{n-k}(M; \mathbb{K}) \rightarrow \mathbb{K}$ with coefficients in a field \mathbb{K} to show that if $f : M \rightarrow N$ has nonzero degree $d := \deg(f)$ and \mathbb{K} is a field whose characteristic does not divide d , the induced maps $f^* : H^k(N; \mathbb{K}) \rightarrow H^k(M; \mathbb{K})$ and $f_* : H_k(M; \mathbb{K}) \rightarrow H_k(N; \mathbb{K})$ are injective and surjective respectively for every k .

52. The Thom class

52.1. A preview of intersection theory. Our next goal is to explain the intersection product on the homology of a closed manifold, which sheds considerable light on the intuitive geometric meaning of both the cup product and the Poincaré duality isomorphism. The following quick sketch is for motivational purposes only; it will be supplemented by precise definitions and proofs in the next lecture.

CONVENTION. In order to simplify notation, for this and the next lecture we will generally work with integer coefficients

$$H_*(M) := H_*(M; \mathbb{Z}), \quad H^*(M) := H^*(M; \mathbb{Z})$$

for homology and cohomology, which requires imposing assumptions about orientations in several places. Note that thanks to the universal coefficient theorem, most things that can be done with integer coefficients can also be done with coefficients in an arbitrary commutative ring R , because orientable manifolds are also R -orientable. If one prefers to drop the orientation assumptions, one still always has the option to work with \mathbb{Z}_2 coefficients.

In the previous lecture, we defined the intersection form on the cohomology of a closed oriented n -manifold M as

$$Q(\alpha, \beta) := \langle \alpha \cup \beta, [M] \rangle \in \mathbb{Z} \quad \text{for} \quad \alpha \in H^k(M), \beta \in H^{n-k}(M).$$

In light of the Poincaré duality isomorphism

$$H^k(M) \xrightarrow{\text{PD}} H_{n-k}(M) : \varphi \mapsto \varphi \cap [M]$$

and the relation $\langle \alpha \cup \beta, [M] \rangle = \langle \alpha, \beta \cap [M] \rangle$, the intersection form can also be written as

$$Q(\alpha, \beta) = \langle \alpha, \text{PD}(\beta) \rangle = (-1)^{k(n-k)} \langle \beta, \text{PD}(\alpha) \rangle,$$

thus giving an interpretation of Poincaré duality: it identifies the cohomology class $\beta \in H^{n-k}(M)$ with the homology class $B := \text{PD}(\beta) \in H_k(M)$ such that feeding B into the evaluation by other cohomology classes $\alpha \in H^k(M)$ is equivalent to feeding β into the intersection form $Q(\alpha, \beta)$. That's nice, but of course it would be nicer if we also understood what the intersection form means, and why it is called what it is.

The answer comes from the following relation, to be proved in the next lecture. For suitable pairs of closed submanifolds $A, B \subset M$ and cohomology classes $\alpha \in H^k(M)$ and $\beta \in H^\ell(M)$ that are Poincaré dual to the homology classes represented by the submanifolds,

$$\text{PD}(\alpha) = [A] \in H_{n-k}(M) \quad \text{and} \quad \text{PD}(\beta) = [B] \in H_{n-\ell}(M)$$

one has

$$(52.1) \quad \text{PD}(\alpha \cup \beta) = [B \cap A] \in H_{n-(k+\ell)}(M).$$

In other words, Poincaré duality identifies the cup product with a so-called *intersection product*, which measures intersections of closed submanifolds homologically. Some conditions are needed in order to fully make sense of this, e.g. you can infer from (52.1) that $B \cap A \subset M$ is expected to be a closed submanifold with codimension equal to the sum of the codimensions of A and B , but this is not automatic—it requires a technical condition that we will specify in the next lecture. This condition necessitates working in the smooth category, so we will eventually add the assumption that all manifolds in this picture are smooth.

Let us elaborate on what happens when the dimensions are complementary, i.e. the case $k + \ell = n$. In this situation, $B \cap A$ is a closed 0-dimensional manifold, meaning a finite set of points. If we are using integer coefficients and assuming everything to be oriented, then $B \cap A$ will also inherit an orientation, which means each point comes with a sign attached to it, and counting these points with signs is equivalent to evaluating the unit cohomology class $1 \in H^0(M)$ on the homology class $[B \cap A] \in H_0(M)$, giving

$$\begin{aligned} \#(B \cap A) &= \langle 1, [B \cap A] \rangle = \langle 1, (\alpha \cup \beta) \cap [M] \rangle = \langle 1 \cup (\alpha \cup \beta), [M] \rangle = \langle \alpha \cup \beta, [M] \rangle \\ &= Q(\alpha, \beta). \end{aligned}$$

That's why it's called the intersection form!

52.2. Tubular neighborhoods and vector bundles. Understanding the intersection product will require understanding what neighborhoods of submanifolds $\Sigma \subset M$ look like, which leads inevitably to the subject of vector bundles. Assume that Σ and M are *smooth* manifolds, so they have tangent spaces: we will denote by $T_x M$ the tangent space of M at a point $x \in M$. If $\Sigma \subset M$ is a smooth submanifold, then its tangent space at each point $x \in \Sigma$ is naturally a linear subspace

$$T_x \Sigma \subset T_x M.$$

The quotient by this subspace is called the **normal space** of Σ in M at x , and will be denoted by

$$N_x \Sigma = N_x^M \Sigma := T_x M / T_x \Sigma.$$

Intuitively, you can imagine the normal space as another subspace of $T_x M$ complementary to $T_x \Sigma$, and indeed, one can realize $N_x \Sigma$ in this way by choosing an inner product on $T_x M$ and identifying the quotient with the orthogonal complement of $T_x \Sigma$ in $T_x M$. The following vague statement

should now seem intuitively plausible: every point in some neighborhood of $\Sigma \subset M$ is obtained by “pushing” some specific point $x \in \Sigma$ in the direction of some specific normal vector in $N_x\Sigma$, after identifying $N_x\Sigma$ with a subspace complementary to $T_x\Sigma \subset T_xM$. In other words, a neighborhood of Σ can be identified with the union of all the normal spaces,

$$N\Sigma = N^M\Sigma := \bigcup_{x \in \Sigma} N_x\Sigma,$$

which is called the **normal bundle** of Σ . Here, the vector spaces $N_x\Sigma$ and $N_y\Sigma$ are regarded as disjoint sets for $x \neq y$, thus $N\Sigma$ is set-theoretically their disjoint union. But one can give $N\Sigma$ a natural topology that is quite different from the disjoint union topology, along with a smooth manifold structure, such that the obvious inclusion

$$\Sigma \hookrightarrow N\Sigma : x \mapsto 0 \in N_x\Sigma$$

is a smooth embedding onto a smooth submanifold, known as the **zero section** of $N\Sigma$, and contracting each of the vector spaces $N_x\Sigma$ to $0 \in N_x\Sigma$ defines a deformation retraction of $N\Sigma$ to its zero section. This is standard material in differential geometry, and we will not go into the details here except to say that none of them are especially deep or surprising. Both the normal bundle and the **tangent bundle**

$$TM := \bigcup_{x \in M} T_xM$$

are examples of vector bundles, a notion that can be defined equally well in the topological or smooth categories. Here are the main definitions.

DEFINITION 52.1. A (real) **vector bundle** $\pi : E \rightarrow X$ of rank $k \geq 0$ over a space X (called the **base** of the bundle) consists of a space E (the **total space** of the bundle) together with a surjective continuous map $\pi : E \rightarrow X$ (the **bundle projection**) such that for each point $x \in X$, the so-called **fiber**

$$E_x := \pi^{-1}(x) \subset E$$

over x is endowed with the structure of a real k -dimensional vector space, and additionally, the following *local triviality* condition is satisfied. Every point $x \in X$ is contained in a neighborhood $\mathcal{U} \subset X$ on which there exists a **local trivialization**, meaning a homeomorphism

$$E|_{\mathcal{U}} := \pi^{-1}(\mathcal{U}) \xrightarrow{\Phi} \mathcal{U} \times \mathbb{R}^k$$

such that for each $y \in \mathcal{U}$, Φ defines a vector space isomorphism from E_y to $\{y\} \times \mathbb{R}^k$. The **zero section** of $\pi : E \rightarrow X$ is the image of the topological embedding

$$X \hookrightarrow E : x \mapsto 0 \in E_x,$$

and we shall therefore often regard the base X as a subspace $X \subset E$ by identifying it with the zero section.

We call $\pi : E \rightarrow X$ a **smooth vector bundle** if, in addition to the conditions above, X and E are smooth manifolds, the bundle projection $\pi : X \rightarrow E$ is a smooth map, and the local trivialization $\Phi : E|_{\mathcal{U}} \rightarrow \mathcal{U} \times \mathbb{R}^k$ can be arranged to be smooth. In this situation, the manifold E has dimension

$$\dim E = \dim X + k,$$

and it contains the zero section $X \subset E$ and the fibers $E_x \subset E$ as smooth submanifolds. If X has nonempty boundary, then E has boundary $\partial E = E|_{\partial X} := \pi^{-1}(\partial X)$.

It is sometimes useful to choose some auxiliary structure on a vector bundle $\pi : E \rightarrow X$, such as a **bundle metric** $\langle \cdot, \cdot \rangle$, meaning a continuous family of inner products defined on the fibers E_x . Bundle metrics always exist if the space X is somewhat reasonable: one can construct them via

partitions of unity, which exist if X is paracompact, and this is always true e.g. if X is a manifold. Once a bundle metric has been chosen, one can define various subspaces of E such as the **unit disk bundle**

$$\mathbb{D}E := \bigcup_{x \in X} \mathbb{D}E_x, \quad \text{where} \quad \mathbb{D}E_x := \{v \in E_x \mid \langle v, v \rangle \leq 1\} \cong \mathbb{D}^k,$$

and if $k \geq 1$, the **unit sphere bundle**

$$SE := \bigcup_{x \in X} SE_x, \quad \text{where} \quad SE_x := \{v \in E_x \mid \langle v, v \rangle = 1\} \cong S^{k-1}.$$

The spaces $\mathbb{D}E$ and SE with their projections π to X are not vector bundles, but are examples of slightly more general objects called **fiber bundles**, and just as we use the notation

$$E|_{\mathcal{U}} := \pi^{-1}(\mathcal{U}) = \bigcup_{x \in \mathcal{U}} E_x \subset E$$

for subsets $\mathcal{U} \subset X$, we will similarly use the notation

$$\mathbb{D}E|_{\mathcal{U}} := E|_{\mathcal{U}} \cap \mathbb{D}E \quad \text{and} \quad SE|_{\mathcal{U}} := E|_{\mathcal{U}} \cap SE.$$

In the situation that will be of greatest interest for us, E is a *smooth* vector bundle over a closed smooth n -manifold $X := M$, and in this case, choosing the bundle metric to be smooth makes $\mathbb{D}E$ into a compact smooth $(n+k)$ -manifold with boundary

$$\partial(\mathbb{D}E) = SE,$$

so that we can speak of relative fundamental classes living in $H_{n+k}(\mathbb{D}E, SE)$.

REMARK 52.2. There is almost never a *canonical* choice of bundle metric, and the definitions of $\mathbb{D}E$ and SE depend on this choice, but the following observation should reassure you that none of the important conclusions we draw from these spaces will depend on it. In the situation described above, relative fundamental classes of $\mathbb{D}E$ live in $H_{n+k}(\mathbb{D}E, SE)$, but if we identify the base M with the zero section $M \subset E$, then thanks to homotopy invariance, the inclusion of pairs $(\mathbb{D}E, SE) \hookrightarrow (E, E \setminus M)$ induces an isomorphism

$$H_*(\mathbb{D}E, SE) \xrightarrow{\cong} H_*(E | M),$$

thus identifying $[\mathbb{D}E] \in H_{n+k}(\mathbb{D}E, SE)$ with a class in $H_{n+k}(E | M)$ that will be independent of choices. It would in fact be possible to formulate everything we have to say about vector bundles and their (co)homology in these terms, without ever choosing a bundle metric. Working with disk bundles and sphere bundles has some intuitive advantages, however, and we will therefore do so.

If we are to work with integer coefficients, we will of course require our vector bundles to be oriented, which requires first clarifying what it means for a real k -dimensional vector space V to be oriented.

Since vector spaces are also manifolds, we could define this in terms of the local homology groups $H_k(V | v; \mathbb{Z})$ at points $v \in V$, but there is a simpler way that is also standard in differential geometry. There, one defines an **orientation** of the vector space V to be an equivalence class of ordered bases (v_1, \dots, v_k) of V , where two such bases are considered equivalent if and only if there exists a continuous deformation of one to the other through a family of ordered bases. The bases in the chosen equivalence class are then called **positively oriented** bases, and the others are called **negatively oriented**. An ordered basis is equivalent to a vector space isomorphism $V \cong \mathbb{R}^k$, so another way of saying this is that an orientation is a deformation class of isomorphisms to \mathbb{R}^k . If we take V to be \mathbb{R}^k itself, then such isomorphisms are literally elements of the general linear group $\text{GL}(k, \mathbb{R})$, and an orientation is thus equivalent to a choice of connected component of $\text{GL}(k, \mathbb{R})$. This group has two components, distinguished by the signs of determinants of matrices. We

conclude that every finite-dimensional real vector space admits exactly two choices of orientation in this sense. To see that this is equivalent to the usual definition we use on topological manifolds, observe that \mathbb{R}^k itself has canonical orientations in both senses: there is the equivalence class of the standard basis, and there are also standard conventions for fixing a canonical generator of $H_k(\mathbb{R}^k | v; \mathbb{Z}) \cong H_k(\mathbb{D}^k, S^{k-1}; \mathbb{Z}) \cong \tilde{H}_{k-1}(S^{k-1}; \mathbb{Z}) \cong \mathbb{Z}$. A deformation class of isomorphisms $\phi : \mathbb{R}^k \rightarrow V$ then determines isomorphisms $\phi_* : H_k(\mathbb{R}^k | v; \mathbb{Z}) \rightarrow H_k(V | \phi(v); \mathbb{Z})$ that determine an orientation of V as a topological k -manifold in terms of the canonical orientation of \mathbb{R}^k . For this reason, the two notions we've defined for orientations of real finite-dimensional vector spaces can be used interchangeably.

DEFINITION 52.3. An **orientation** of the vector bundle $\pi : E \rightarrow X$ is a choice of orientation of the fibers E_x for each $x \in X$ such that for any local trivialization $\Phi : E|_{\mathcal{U}} \rightarrow \mathcal{U} \times \mathbb{R}^k$ over a path-connected region $\mathcal{U} \subset X$, the resulting vector space isomorphisms $E_y \rightarrow \{y\} \times \mathbb{R}^k$ for $y \in \mathcal{U}$ are all either orientation preserving or orientation reversing.

It is not difficult to show that if $E \rightarrow X := M$ is a smooth and oriented vector bundle over a smooth oriented manifold M , then the manifold E is also orientable. We will need to fix a convention for defining this orientation of E , and to be consistent about it. Recall from Exercise 50.4 that for any pair of oriented manifolds M and N , the product $M \times N$ inherits a natural orientation that can be defined via the homological cross product. In the smooth category, there are also easier ways to define product orientations, because orientations of M and N are then equivalent to orientations of their tangent bundles $TM \rightarrow M$ and $TN \rightarrow N$, and one can orient the tangent space

$$T_{(x,y)}(M \times N) \cong T_x M \oplus T_y N$$

at each point $(x, y) \in M \times N$ by insisting that for any pair of positively-oriented bases (v_1, \dots, v_m) of $T_x M$ and (w_1, \dots, w_n) of $T_y N$, the basis $(v_1, \dots, v_m, w_1, \dots, w_n)$ of $T_{(x,y)}(M \times N)$ should be positively oriented. Whether you prefer this definition or the equivalent one in Exercise 50.4, we can use it as follows to define orientations of total spaces of vector bundles in terms of product orientations:

DEFINITION 52.4. Assume $\pi : E \rightarrow M$ is a smooth oriented vector bundle over an oriented manifold M . We then endow the total space E with the unique orientation such that for any smooth local trivialization $\Phi|_{\mathcal{U}} : E|_{\mathcal{U}} \rightarrow \mathcal{U} \times \mathbb{R}^k$ over an open subset $\mathcal{U} \subset M$ for which the isomorphisms $E_y \xrightarrow{\Phi} \{y\} \times \mathbb{R}^k$ are all orientation preserving, the diffeomorphism

$$E|_{\mathcal{U}} \xrightarrow{\cong} \mathbb{R}^k \times \mathcal{U}$$

defined by composing Φ with the obvious bijection $\mathcal{U} \times \mathbb{R}^k \rightarrow \mathbb{R}^k \times \mathcal{U}$ is orientation preserving. Here, $\mathbb{R}^k \times \mathcal{U}$ is endowed with the product orientation as in Exercise 50.4, determined by the canonical orientation of \mathbb{R}^k and the given orientation of $\mathcal{U} \subset M$.

REMARK 52.5. The reversal of the order of \mathcal{U} and \mathbb{R}^k in Definition 52.4 is slightly unfortunate, but it is really just an artefact of the arbitrary convention to define local trivializations as maps to $\mathcal{U} \times \mathbb{R}^k$ rather than $\mathbb{R}^k \times \mathcal{U}$. We will find that this is the most convenient convention to use if we want to avoid unwanted signs in statements of the main results in intersection theory.

We can now turn the vague intuition at the beginning of this section into a precise statement of a theorem:

THEOREM 52.6 (tubular neighborhood theorem). *Assume M is a smooth n -manifold without boundary and $\Sigma \subset M$ is a smooth k -dimensional submanifold without boundary, with normal bundle $\pi : N\Sigma \rightarrow \Sigma$. Then there exists a diffeomorphism of the total space of $N\Sigma$ onto an open*

neighborhood of $\Sigma \subset M$ that restricts to the canonical identification of the zero section with Σ . Moreover, if M and Σ both carry orientations, then the normal bundle $\pi : N\Sigma \rightarrow \Sigma$ inherits a unique orientation such that the diffeomorphism of $N\Sigma$ with a neighborhood of $\Sigma \subset M$ is orientation preserving.

PROOF SKETCH. This is a standard construction in differential geometry, so we provide a brief sketch using a few ideas from that subject. As auxiliary data, one can choose a Riemannian metric on M and use it to identify each fiber $N_x\Sigma$ of the normal bundle with the orthogonal complement $T_x\Sigma^\perp \subset T_xM$ of $T_x\Sigma \subset T_xM$, thus giving an isomorphism of vector bundles $N\Sigma \cong T\Sigma^\perp \subset TM|_\Sigma$. The metric also defines the notions of *geodesics* and the *exponential map*

$$TM \supset \mathcal{O} \xrightarrow{\exp} M,$$

which is defined on an open neighborhood $\mathcal{O} \subset TM$ of the zero section such that for each $p \in M$ and $X \in T_pM$, $\gamma(t) := \exp(tX)$ is the unique geodesic in M satisfying the initial conditions $\gamma(0) = p$ and $\dot{\gamma}(0) = X$. Restricting \exp to the intersection of $T\Sigma^\perp \cong N\Sigma$ with \mathcal{O} then gives a smooth map from some neighborhood of the zero section in $N\Sigma$ to M that restricts to the zero section as the identity map $\Sigma \rightarrow \Sigma \subset M$, and its derivative at each point along the zero section is an isomorphism. The inverse function theorem then implies that this map sends a neighborhood of the zero section diffeomorphically onto a neighborhood of $\Sigma \subset M$. The result now follows by identifying neighborhoods of 0 in each fiber N_xM diffeomorphically with the whole fiber. \square

REMARK 52.7. With a little care, the tubular neighborhood theorem can easily be extended to manifolds with boundary: this works in particular if Σ intersects ∂M transversely, with $\partial\Sigma = \Sigma \cap \partial M$, in which case the normal bundle $N\Sigma$ restricts to $\partial\Sigma$ as the normal bundle of the submanifold $\partial\Sigma \subset \partial M$. The notion of *transversality* needed for this will be explained further in the next lecture.

52.3. The Thom isomorphism theorem. The main results about the intersection product are based on two fundamental ingredients: one is Poincaré duality, and the other is some knowledge of the cohomology of vector bundles, since by the tubular neighborhood theorem, these serve as local models for neighborhoods of smooth submanifolds. For any vector bundle, the projection

$$E \xrightarrow{\pi} X$$

is a homotopy equivalence, making $H^*(E)$ isomorphic to $H^*(X)$. One obtains a more interesting relation between E and X by identifying X with the zero section of E and looking at the restricted cohomology $H^*(E|X)$. Choosing a bundle metric in order to define the unit disk bundle $\mathbb{D}E$ and sphere bundle SE with fibers $SE_x = \partial(\mathbb{D}E_x)$, we recall from Remark 52.2 that the inclusion of pairs $(\mathbb{D}E, SE) \hookrightarrow (E, E \setminus X)$ induces an isomorphism $H^*(E|X) \cong H^*(\mathbb{D}E, SE)$. The following theorem could therefore be stated as a result about $H^*(E|X)$ without needing to choose a bundle metric, but it looks a bit prettier as a result about $H^*(\mathbb{D}E, SE)$, mainly because for each $x \in X$, the fiber $\mathbb{D}E_x \cong \mathbb{D}^k$ has a well-defined relative fundamental class

$$[\mathbb{D}E_x] \in H_k(\mathbb{D}E_x, SE_x),$$

whose image under the homomorphism induced by the inclusion of pairs $(\mathbb{D}E_x, SE_x) \hookrightarrow (\mathbb{D}E, SE)$ we shall denote by

$$[\mathbb{D}E_x]_{\mathbb{D}E} \in H_k(\mathbb{D}E, SE).$$

THEOREM 52.8 (Thom isomorphism theorem). *Assume $\pi : E \rightarrow X$ is an oriented vector bundle of rank $k \in \mathbb{N}$ over a CW-complex X and is endowed with a bundle metric, which defines the associated unit disk bundle $\mathbb{D}E$ and sphere bundle SE . Then:*

- (1) There exists a unique relative cohomology class $\tau(E) \in H^k(\mathbb{D}E, SE)$, called the **Thom class** of the bundle $\pi : E \rightarrow X$, characterized by the condition that

$$\langle \tau(E), [\mathbb{D}E_x]_{\mathbb{D}E} \rangle = 1 \quad \text{for every } x \in X.$$

- (2) For every $m \geq 0$, the relative cup product pairing

$$H^m(\mathbb{D}E) \otimes H^k(\mathbb{D}E, SE) \xrightarrow{\cup} H^{m+k}(\mathbb{D}E, SE)$$

determines an isomorphism

$$H^m(X) \xrightarrow{\cong} H^{m+k}(\mathbb{D}E, SE) : \varphi \mapsto \pi^* \varphi \cup \tau(E),$$

where $\pi^* : H^m(X) \rightarrow H^m(\mathbb{D}E)$ denotes the homomorphism induced by $\mathbb{D}E \xrightarrow{\pi} X$.

REMARK 52.9. The Thom isomorphism theorem can be proved under somewhat more general assumptions, as found e.g. in [tD08, §17.9]. In particular, one can allow X to be something more general than a CW-complex, but keeping this assumption makes the proof especially simple. In keeping with the convention stated at the beginning of this lecture, the assumption that E is oriented can be dropped at the cost of using cohomology with coefficients in \mathbb{Z}_2 instead of \mathbb{Z} . More generally, one can define a notion of *R-orientability* for vector bundles with respect to an arbitrary commutative ring R , and prove that any R -orientation of $\pi : E \rightarrow X$ gives rise to a unique Thom class $\tau(E) \in H^k(\mathbb{D}E, SE; R)$ satisfying $\langle \tau(E), [\mathbb{D}E_x]_{\mathbb{D}E} \rangle = 1$ for every $x \in X$, where the relative fundamental classes $[\mathbb{D}E_x] \in H_k(\mathbb{D}E_x, SE_x; R)$ are determined by a continuous family of R -orientations of the fibers.

The application to intersection theory in the next lecture will not require the full strength of Theorem 52.8, but really only the first statement, which serves simultaneously as a definition of the Thom class $\tau(E)$. The proof given below will also establish that there is an isomorphism $H^m(X) \cong H^{k+m}(\mathbb{D}E, SE)$, but not that it takes the specific form stated in the theorem—in the case where X is a closed oriented manifold, the latter can be deduced via Poincaré duality from Exercise 52.1, using an important fact about Thom classes that will be proved in the next lecture. In the main application we need, X will be a smooth manifold and therefore triangulable, and we can therefore get away with assuming in our proof that X is a polyhedron; this assumption could be relaxed by using a bit more knowledge of the general theory of bundles.

PROOF OF THEOREM 52.8(1), ASSUMING X IS TRIANGULABLE. We are interested in the cohomology of the pair $(\mathbb{D}E, SE)$ in positive degrees, and since this is a “good” pair in the sense of Definition 34.3, we are free to replace its relative cohomology with the absolute cohomology of the quotient space,

$$\text{Th}(E) := \mathbb{D}E/SE,$$

known as the **Thom space**. Topologically, $\text{Th}(E)$ is homeomorphic to the one-point compactification of E , and you can picture it as a family of k -spheres $\mathbb{D}E_x/SE_x \cong \mathbb{D}^k/S^{k-1} \cong S^k$ that all intersect each other at exactly one point

$$\infty \in \text{Th}(E),$$

namely the point obtained by collapsing $SE \subset \mathbb{D}E$. The main idea behind the isomorphism $H^m(X) \cong H^{m+k}(\text{Th}(E)) \cong H^{m+k}(\mathbb{D}E, SE)$ is then to construct out of a cell decomposition of X a closely related cell decomposition of $\text{Th}(E)$, known as the **Thom complex**. Doing this with an arbitrary cell decomposition of X requires the knowledge that vector bundles over disks \mathbb{D}^n are always globally trivializable, which is true but not completely trivial to prove, so we shall now make our lives slightly easier by assuming that X comes with a triangulation. After subdivision, we can then assume without loss of generality that every simplex $\sigma \subset X$ in the triangulation is small enough to be contained in a region $\mathcal{U} \subset X$ on which the disk bundle admits a local trivialization

$\mathbb{D}E|_{\mathcal{U}} \cong \mathcal{U} \times \mathbb{D}^k$, and since $\pi : E \rightarrow X$ is assumed oriented, we are also free to assume that the resulting isomorphisms $\mathbb{D}E_x \cong \mathbb{D}^k$ for points $x \in \sigma$ are all orientation preserving. With this assumption in place, we now construct a cell decomposition of $\text{Th}(E)$ inductively as follows:

- The 0-skeleton of $\text{Th}(E)$ consists of the single point $\infty \in \text{Th}(E)$.
- Each 0-simplex $x \in X$ gives rise to a k -cell in $\text{Th}(E)$, namely the fiber $\mathbb{D}E_x$, which is attached by collapsing its boundary to the 0-skeleton $\infty \in \text{Th}(E)$. The union of ∞ with all the k -cells corresponding in this way to 0-simplices of X forms the k -skeleton of $\text{Th}(E)$.
- For each 1-simplex $\sigma \subset X$, we can use a local trivialization as indicated above to identify $\mathbb{D}E|_{\sigma}$ homeomorphically with $\sigma \times \mathbb{D}^k$ and thus regard $\mathbb{D}E|_{\sigma}$ as the closure of a $(k+1)$ -cell, attached via a map $\partial(\sigma \times \mathbb{D}^k) \rightarrow \text{Th}(E)$ that sends $\sigma \times S^{k-1}$ to ∞ and identifies $\partial\sigma \times \mathbb{D}^k$ with a union of k -cells that were constructed in the previous step. This completes the construction of the $(k+1)$ -skeleton of $\text{Th}(E)$.
- Continuing inductively in this manner, every m -simplex $\sigma \subset X$ gives rise to an $(m+k)$ -cell in $\text{Th}(E)$ whose closure is the region $\mathbb{D}E|_{\sigma}$.

In light of the obvious bijective correspondence between the simplices in X and cells in $\text{Th}(E)$, one obtains isomorphisms

$$C_m^{\Delta}(X) \xrightarrow{\cong} C_{m+k}^{\text{CW}}(\text{Th}(E))$$

for every $m \geq 0$, and one can check that these maps satisfy the chain map relation. A key detail here is that the characteristic maps of the cells in $\text{Th}(E)$ are defined using *orientation-preserving* homeomorphisms $\mathbb{D}E_x \cong \mathbb{D}^k$; if orientations were disregarded, then one would find unwanted signs polluting the chain map relation, though this problem goes away of course if one uses coefficients in \mathbb{Z}_2 instead of \mathbb{Z} . Dualizing the chain map $C_*^{\Delta}(X) \rightarrow C_{*+k}^{\text{CW}}(\text{Th}(E))$ then produces a similar isomorphism of cochain complexes, giving rise to isomorphisms

$$\begin{array}{ccccccc}
 H^m(X) & \xrightarrow{\cong} & H_{\Delta}^m(X) & \xrightarrow{\cong} & H_{\text{CW}}^{m+k}(\text{Th}(E)) & \xrightarrow{\cong} & H^{m+k}(\mathbb{D}E, SE) \\
 & & & & \Phi & & \\
 & & & & \cong & &
 \end{array}$$

for every $m \geq 0$.

For any simplicial 0-cocycle $\alpha \in C_{\Delta}^0(X)$ and any 0-simplex $x \in X$ in the given triangulation, the construction of the isomorphism gives rise to the relation

$$\langle \Phi([\alpha]), [\mathbb{D}E_x]_{\mathbb{D}E} \rangle = \langle [\alpha], [x] \rangle,$$

where the evaluation on the right hand side takes place in the simplicial (co)homology of X , while on the left hand side, it can be understood as the evaluation of a cellular k -cochain on the cellular k -chain in $\text{Th}(E)$ corresponding to x . There is exactly one choice of α that makes this evaluation equal to 1 for every 0-simplex $x \in X$, and we therefore define the Thom class in terms of the unit $1 \in H^0(X)$, hence

$$\tau(E) := \Phi(1) \in H^k(\mathbb{D}E, SE).$$

Since every point in X is connected by a continuous path to a 0-simplex of the triangulation, homotopy invariance then implies that $\langle \tau(E), [\mathbb{D}E_x]_{\mathbb{D}E} \rangle = 1$ also holds for every point $x \in X$. \square

52.4. Exercises.

EXERCISE 52.1. For the unit disk bundle $\mathbb{D}E \xrightarrow{\pi} X$ associated to an oriented vector bundle $\pi : E \rightarrow X$ of rank k over a CW-complex, Theorem 52.8 states an explicit formula for the Thom isomorphism $\Phi : H^m(X) \rightarrow H^{m+k}(\mathbb{D}E, SE)$ in terms of the Thom class $\tau(E) \in H^k(\mathbb{D}E, SE)$, and the relative cup product pairing $H^m(\mathbb{D}E) \otimes H^k(\mathbb{D}E, SE) \rightarrow H^{m+k}(\mathbb{D}E, SE)$, namely

$$\Phi(\alpha) := \pi^* \alpha \cup \tau(E).$$

We will show in the next lecture that in the special case where $X := M$ is a closed oriented n -manifold with fundamental class $[M] \in H_n(M)$ and duality isomorphisms $\text{PD}_M : H^\ell(M) \rightarrow H_{n-\ell}(M)$, the Thom class satisfies

$$\text{PD}_{\mathbb{D}E}(\tau(E)) = i_*[M],$$

where $\text{PD}_{\mathbb{D}E}$ denotes the relative duality isomorphism $H^k(\mathbb{D}E, SE) \rightarrow H_n(\mathbb{D}E)$ and $i : M \hookrightarrow \mathbb{D}E$ is the inclusion of the zero section. Taking this as a black box, prove the formula

$$\text{PD}_{\mathbb{D}E} \circ \Phi = i_* \circ \text{PD}_M,$$

and conclude that Φ is an isomorphism.

53. The intersection product

As in the previous lecture, we continue now under the assumption that \mathbb{Z} is the coefficient group for homology and cohomology, with the understanding that \mathbb{Z}_2 coefficients can also be used if we wish to drop any assumptions about orientations. When orientation assumptions are in place, the universal coefficient theorem also makes it possible to replace \mathbb{Z} with any commutative ring R with unit.

The **intersection product** is defined on any closed oriented n -manifold M as a pairing

$$H_{n-k}(M) \otimes H_{n-\ell}(M) \xrightarrow{\quad} H_{n-(k+\ell)}(M) : A \otimes B \mapsto A \cdot B$$

that is Poincaré dual to the cup product, i.e. it is uniquely determined by the formula

$$\text{PD}(\alpha) \cdot \text{PD}(\beta) = \text{PD}(\beta \cup \alpha) \quad \text{for all } \alpha \in H^k(M), \beta \in H^\ell(M).$$

The reversal of the order of α and β in this definition looks unnatural at first glance, but should be unsurprising if you keep in mind that the passage from homology to cohomology involves a process of *dualization*, which automatically reverses the orders of compositions.

Our main goal in this lecture is to state and prove a precise version of the formula

$$(53.1) \quad [A] \cdot [B] = [A \cap B]$$

for suitable closed oriented submanifolds $A, B \subset M$. The precise statement will appear in §53.3 after clarifying in §53.2 the required transversality assumptions, which will ensure among other things that $A \cap B \subset M$ is a submanifold with $\text{codim}(A \cap B) = \text{codim}(A) + \text{codim}(B)$. In the case $\dim A + \dim B = n$, this makes $A \cap B$ a closed oriented 0-manifold, meaning a finite set of points with signs attached, and counting these points with signs then produces the intersection form of the Poincaré dual cohomology classes: in precise terms, if we assume $\text{PD}(\alpha) = [A] \in H_{n-k}(M)$ and $\text{PD}(\beta) = [B] \in H_k(M)$, then (53.1) and the relation $(\varphi \cup \psi) \cap c = \varphi \cap (\psi \cap c)$ imply

$$(53.2) \quad \begin{aligned} \langle 1, [A \cap B] \rangle &= \langle 1, (\beta \cup \alpha) \cap [M] \rangle = \langle \beta \cup \alpha, [M] \rangle \\ &= Q(\beta, \alpha) = (-1)^{k(n-k)} Q(\alpha, \beta) \\ &= \langle \beta, [A] \rangle = (-1)^{k(n-k)} \langle \alpha, [B] \rangle. \end{aligned}$$

The formula (53.1) will be deduced from a more fundamental result that relates pullbacks of cohomology classes to preimages of submanifolds: roughly speaking, we will show that for smooth maps $f : A \rightarrow M$ and smooth submanifolds $B \subset M$ satisfying suitable conditions such that $f^{-1}(B) \subset A$ is a smooth submanifold with the same codimension as $B \subset M$, one has

$$\text{PD}(\beta) = [B] \in H_*(M) \quad \implies \quad \text{PD}(f^*\beta) = [f^{-1}(B)] \in H_*(A).$$

This gives a transparent geometric interpretation of the Poincaré duality isomorphism: it identifies cohomology classes with homology classes such that pullbacks become preimages.

NOTATION. Here are some useful notational conventions to be adopted throughout this lecture. There will often be multiple compact manifolds in the picture, each having its own Poincaré duality isomorphism, so it will be useful to specify which is which by writing

$$\text{PD}_M : H^k(M) \rightarrow H_{n-k}(M),$$

rather than just “PD”. In some cases, M will also have nonempty boundary, and we recall from Exercise 51.7 that there are two versions of the duality isomorphism for this situation: both will be denoted by

$$\text{PD}_M : H^k(M, \partial M) \rightarrow H_{n-k}(M) \quad \text{and} \quad \text{PD}_M : H^k(M) \rightarrow H_{n-k}(M, \partial M).$$

You will always be able to infer which of these two is meant by observing whether the cohomology class fed into it is an absolute class or a relative class.

We will reserve the notation $[M] \in H_n(M, \partial M)$ for fundamental classes of compact manifolds, thus for a closed k -dimensional submanifold $A \subset M$, $[A] \in H_k(A)$ will be the fundamental class of A , while the class that A represents in M , i.e. the one obtained by feeding $[A]$ into the homomorphism induced by the inclusion $i : A \hookrightarrow M$, will be denoted by

$$[A]_M := i_*[A] \in H_k(M).$$

Under these conventions, the main result about the intersection product can be written in the form

$$\text{PD}_M(\beta \cup \alpha) = [A \cap B]_M \quad \text{where} \quad \text{PD}_M(\alpha) = [A]_M \quad \text{and} \quad \text{PD}_M(\beta) = [B]_M.$$

The notation for classes represented by submanifolds should not be confused with the notation

$$[M]_K := J_K^{-1}(s) \in H_n(\overset{\circ}{M} | K) \cong H_n(M | K)$$

for fundamental classes restricted to compact subsets $K \subset \overset{\circ}{M}$ of the interior, where $s \in \Gamma(\Theta)$ is the orientation corresponding to the fundamental class of M . This notation was used in Lecture 51, and will occasionally still be useful in the following. The context in Lecture 51 was slightly different, because M was assumed to be a (possibly noncompact) manifold without boundary, but if M is compact with boundary and K lies in its interior, then we are free to regard $[M]_K$ as an element of $H_n(M | K)$ because the inclusion of pairs $(\overset{\circ}{M}, \overset{\circ}{M} \setminus K) \hookrightarrow (M, M \setminus K)$ is a homotopy equivalence. In terms of the inclusion of pairs $j^K : (M, \partial M) \hookrightarrow (M, M \setminus K)$, one could equivalently write

$$[M]_K = j^K_*[M] \in H_n(M | K).$$

53.1. The Poincaré dual of a Thom class. The following can be interpreted as a localized special case of the formula $[A] \cdot [B] = [A \cap B]$, and as such, it also serves as one of the main steps in the proof of the general formula. We consider a smooth oriented vector bundle $\pi : E \rightarrow M$ of rank k over a closed, smooth, oriented n -manifold M . The motivation for this setup comes from the tubular neighborhood theorem: a neighborhood of a smooth submanifold can always be identified with a neighborhood of the zero section in the total space of a vector bundle, namely the normal bundle of the submanifold. In the present situation, we are thus identifying M with the zero section $M \subset E$ and regarding it as an n -dimensional submanifold of the $(n+k)$ -dimensional manifold E .

Choose a bundle metric on $\pi : E \rightarrow M$ in order to define the unit disk bundle $\mathbb{D}E$, which is a compact oriented $(n+k)$ -manifold whose boundary is the unit sphere bundle $\partial(\mathbb{D}E) = SE$. Let us pick any point $x \in M$ and count the intersections of the zero section $M \subset \mathbb{D}E$ with the fiber $\mathbb{D}E_x \subset \mathbb{D}E$, which has complementary dimension. The answer is obvious: they intersect precisely at the point $x \in M \subset \mathbb{D}E$, and nowhere else. Since $\mathbb{D}E$ has nonempty boundary, this situation does not fit precisely into the setup for the formula (53.1), though we will discuss in §53.6 relative versions of the intersection product that can accommodate this situation. In any case, if there is a relative version of the correspondence in (53.2) between the intersection product on homology and

the intersection form on cohomology, then for the relative cohomology class $\tau \in H^k(\mathbb{D}E, SE)$ that is Poincaré dual to the zero section $[M]_{\mathbb{D}E} \in H_n(\mathbb{D}E)$, one should expect to find

$$1 = \langle 1, [\mathbb{D}E_x \cap M] \rangle = \langle \tau, [\mathbb{D}E_x]_{\mathbb{D}E} \rangle,$$

where $[\mathbb{D}E_x]_{\mathbb{D}E} \in H_k(\mathbb{D}E, SE)$ is the image of the relative fundamental class of the fiber $\mathbb{D}E_x$ under the map induced by the inclusion of pairs $(\mathbb{D}E_x, SE_x) \hookrightarrow (\mathbb{D}E, SE)$. If this is true for every point $x \in M$, then according to Theorem 52.8, τ is the Thom class of E ; in other words, the cohomology class Poincaré dual to the zero section is the Thom class. That is what we shall now prove.

We can allow a slightly more general setup than what was considered above, in which M is a compact oriented n -manifold that may also have nonempty boundary, and thus has a relative fundamental class $(M, \partial M)$. In this situation, $\mathbb{D}E$ is still a compact $(n+k)$ -manifold with boundary, but its boundary is a union of two pieces

$$\partial(\mathbb{D}E) = SE \cup \mathbb{D}E|_{\partial M}.$$

There is now a relative cap product pairing

$$H^k(\mathbb{D}E, SE) \otimes H_{n+k}(\mathbb{D}E, \partial(\mathbb{D}E)) \xrightarrow{\cap} H_n(\mathbb{D}E, \mathbb{D}E|_{\partial M}),$$

and we let

$$[M]_{\mathbb{D}E} \in H_n(\mathbb{D}E, \mathbb{D}E|_{\partial M})$$

denote the image of $[M] \in H_n(M, \partial M)$ under the homomorphism induced by the inclusion of the zero section, regarded as a map of pairs $(M, \partial M) \hookrightarrow (\mathbb{D}E, \mathbb{D}E|_{\partial M})$.

THEOREM 53.1. *In the setting described above,*

$$\tau(E) \cap [M]_{\mathbb{D}E} = [M]_{\mathbb{D}E}.$$

COROLLARY 53.2. *If M is closed, then the Thom class of $\pi : E \rightarrow M$ corresponds under the Poincaré duality isomorphism $PD : H^k(\mathbb{D}E, SE) \rightarrow H_n(\mathbb{D}E)$ to the homology class represented by the zero section.* □

REMARK 53.3. Combining Corollary 53.2 with Exercise 52.1 proves the portion of the Thom isomorphism theorem (Theorem 52.8) that remained unproved in the previous lecture, in the special case where the base of the bundle is a closed manifold.

The main step in the proof of Theorem 53.1 is to prove by explicit computation that it holds in the special case of a trivial bundle over a disk.

LEMMA 53.4. *Theorem 53.1 holds in the special case of a trivial disk bundle $\mathbb{D}E := \mathbb{D}^k \times \mathbb{D}^n \xrightarrow{\pi} \mathbb{D}^n : (v, x) \mapsto x$.*

PROOF. In the situation at hand, we have

$$SE = S^{k-1} \times \mathbb{D}^n, \quad \mathbb{D}E|_{\partial \mathbb{D}^n} = \mathbb{D}^k \times S^{n-1}, \quad \text{and} \quad \mathbb{D}E_x = \mathbb{D}^k \times \{x\}$$

for each $x \in \mathbb{D}^n$. By Exercise 50.4, $[\mathbb{D}E] = [\mathbb{D}^k] \times [\mathbb{D}^n] \in H_{n+k}(\mathbb{D}^k \times \mathbb{D}^n, \partial(\mathbb{D}^k \times \mathbb{D}^n))$, and one can check that $\tau(E) \in H^k(\mathbb{D}^k \times \mathbb{D}^n, S^{k-1} \times \mathbb{D}^n)$ and $[\mathbb{D}^n]_{\mathbb{D}E} \in H_n(\mathbb{D}^k \times \mathbb{D}^n, \mathbb{D}^k \times S^{n-1})$ are also relative cross products, namely

$$\tau(E) = \alpha \times 1 \quad \text{and} \quad [\mathbb{D}^n]_{\mathbb{D}E} = [*]_{\mathbb{D}^k} \times [\mathbb{D}^n],$$

where $1 \in H^0(\mathbb{D}^n)$ is the unit, $[*]_{\mathbb{D}^k} \in H_0(\mathbb{D}^k)$ is the homology class represented by a single point, and $\alpha \in H^k(\mathbb{D}^k, S^{k-1}) \cong \text{Hom}(H_k(\mathbb{D}^k, S^{k-1}), \mathbb{Z})$ is the unique class such that $\langle \alpha, [\mathbb{D}^k] \rangle = 1$. Indeed, for each $x \in \mathbb{D}^n$, the class $[\mathbb{D}E_x]_{\mathbb{D}E} \in H_k(\mathbb{D}^k \times \mathbb{D}^n, S^{k-1} \times \mathbb{D}^n)$ represented by the fiber

$\mathbb{D}E_x = \mathbb{D}^k \times \{x\}$ is $[\mathbb{D}^k] \times [*]_{\mathbb{D}^n}$, where in this case $[*]_{\mathbb{D}^n} \in H_0(\mathbb{D}^n)$ is again the class represented by a single point, thus

$$\langle \alpha \times 1, [\mathbb{D}^k] \times [*] \rangle = \langle \alpha, [\mathbb{D}^k] \rangle \langle 1, [*]_{\mathbb{D}^n} \rangle = 1,$$

verifying the claim that $\alpha \times 1$ is the Thom class. The identity in Corollary 48.7 now gives

$$\tau(E) \cap [\mathbb{D}E] = (\alpha \times 1) \cap ([\mathbb{D}^k] \times [\mathbb{D}^n]) = (\alpha \cap [\mathbb{D}^k]) \times (1 \cap [\mathbb{D}^n]) = [*]_{\mathbb{D}^k} \times [\mathbb{D}^n] = [\mathbb{D}^n]_{\mathbb{D}E},$$

where the class $\alpha \cap [\mathbb{D}^k] = [*]_{\mathbb{D}^k} \in H_0(\mathbb{D}^k)$ is characterized by the property

$$\langle 1, \alpha \cap [\mathbb{D}^k] \rangle = \langle \alpha, [\mathbb{D}^k] \rangle = 1.$$

□

PROOF OF THEOREM 53.1. The inclusion $(M, \partial M) \hookrightarrow (\mathbb{D}E, \mathbb{D}E|_{\partial M})$ is a homotopy equivalence of pairs and has the projection $\pi : (\mathbb{D}E, \mathbb{D}E|_{\partial M}) \rightarrow (M, \partial M)$ as a homotopy inverse. The goal is thus to show that $\pi_*(\tau(E) \cap [\mathbb{D}E])$ is the fundamental class $[M] \in H_n(M, \partial M)$, or equivalently, that for each $x \in \overset{\circ}{M}$, the image in $H_n(M | x)$ of this class under the map induced by the inclusion of pairs $(M, \partial M) \hookrightarrow (M, M \setminus \{x\})$ is the local orientation $[M]_x$ given by the orientation of M . Choose a closed disk-like neighborhood $\mathcal{U} \subset \overset{\circ}{M}$ of x on which there exists a local trivialization, thus identifying $E|_{\mathcal{U}}$ with a product bundle as in Lemma 53.4. Fix $r \in (0, 1)$ and define the subsets

$$\mathbb{D}_r E := \bigcup_{x \in M} \mathbb{D}_r E_x \quad \text{where} \quad \mathbb{D}_r E_x := \{v \in \mathbb{D}E_x \mid \langle v, v \rangle \leq r^2\}, \quad \text{and} \quad S_r E := \mathbb{D}E \setminus \mathbb{D}_r E.$$

There is then a relative cap product of the form

$$H^k(\mathbb{D}E | \mathbb{D}_r E) \otimes H_{n+k}(\mathbb{D}E | \mathbb{D}_r E_x) \xrightarrow{\cap} H_n(\mathbb{D}E | \mathbb{D}E_x),$$

giving us the middle row of the following diagram

$$\begin{array}{ccccc} H^k(\mathbb{D}E, SE) \otimes H_{n+k}(\mathbb{D}E, \partial(\mathbb{D}E)) & \xrightarrow{\cap} & H_n(\mathbb{D}E, \mathbb{D}E|_{\partial M}) & \xrightarrow[\cong]{\pi_*} & H_n(M, \partial M) \\ \cong \uparrow i^* & & \downarrow i_* & & \downarrow \\ H^k(\mathbb{D}E | \mathbb{D}_r E) \otimes H_{n+k}(\mathbb{D}E | \mathbb{D}_r E_x) & \xrightarrow{\cap} & H_n(\mathbb{D}E | \mathbb{D}E_x) & \xrightarrow[\cong]{\pi_*} & H_n(M | x) \\ \downarrow j^* & & \cong \uparrow j_* & & \cong \uparrow \\ H^k(\mathbb{D}E|_{\mathcal{U}}, SE|_{\mathcal{U}}) \otimes H_{n+k}(\mathbb{D}E|_{\mathcal{U}}, \partial(\mathbb{D}E|_{\mathcal{U}})) & \xrightarrow{\cap} & H_n(\mathbb{D}E|_{\mathcal{U}}, \mathbb{D}E|_{\partial \mathcal{U}}) & \xrightarrow[\cong]{\pi_*} & H_n(\mathcal{U}, \partial \mathcal{U}) \end{array}$$

Here, $i : \mathbb{D}E \rightarrow \mathbb{D}E$ is the identity map, which can be interpreted in various ways as an inclusion of pairs that induces vertical arrows between the top two rows of the diagram, and the inclusion $j : \mathbb{D}E|_{\mathcal{U}} \hookrightarrow \mathbb{D}E$ similar induces arrows between the bottom two rows. The maps on the right hand side are likewise induced by obvious inclusions of pairs, and we observe that some of these maps are isomorphisms due to homotopy invariance and/or excision. The meaning of the arrows in opposing directions between the tensor products is the usual naturality relation for cap products, e.g. the upper left square encodes the fact that for any $\tau \in H^k(\mathbb{D}E | \mathbb{D}_r E)$ and $A \in H_{n+k}(\mathbb{D}E, \partial(\mathbb{D}E))$, pushing forward a cap product from the top row gives a cap product from the middle row,

$$i_*(i^* \tau \cap A) = \tau \cap i_* A.$$

By Lemma 53.4, plugging $\tau(E|_{\mathcal{U}}) \otimes [\mathbb{D}E|_{\mathcal{U}}]$ into the bottom left corner of this diagram and then following the bottom row to the right leads to $[\mathcal{U}] \in H_n(\mathcal{U}, \partial \mathcal{U})$, so that following it one step upward then produces $[M]_x \in H_n(M | x)$. The map i^* in the upper left is an isomorphism, so there is a unique class $\tau \in H^k(\mathbb{D}E | \mathbb{D}_r E)$ such that $i^* \tau$ is the Thom class $\tau(E) \in H^k(\mathbb{D}E | SE)$, and one then deduces from the defining property of Thom classes that $j^* \tau$ is similarly the Thom class of the restricted bundle $E|_{\mathcal{U}} \rightarrow \mathcal{U}$. In the same manner, the defining property of fundamental classes

implies that $[\mathbb{D}E|_{\mathcal{U}}] \in H_{n+k}(\mathbb{D}E|_{\mathcal{U}}, \partial(\mathbb{D}E|_{\mathcal{U}}))$ and $[\mathbb{D}E] \in H_{n+k}(\mathbb{D}E, \partial(\mathbb{D}E))$ have the same image in $H_{n+k}(\mathbb{D}E|_{\mathbb{D}_r E_x})$, namely the restricted fundamental class $[\mathbb{D}E]_{\mathbb{D}_r E_x}$. One deduces from these facts and the commutativity of the diagram that if $\tau(E) \otimes [\mathbb{D}E]$ is plugged in at the upper left corner, following it to the right and then down to $H_n(M|x)$ leads to $[M]_x$. \square

53.2. Transversality. We now discuss the technical conditions needed to ensure that intersections of submanifolds are submanifolds of the expected dimension. It would be possible in theory to formulate the entirety of this discussion in terms of topological manifolds, without mentioning smoothness, but working in the smooth category provides substantial advantages due to the implicit function theorem, and in any case, the most important applications of the intersection product are in the theory of smooth manifolds.

For a smooth manifold M , we denote its tangent space at a point $p \in M$ by $T_p M$, and for a smooth map $f : M \rightarrow N$, its derivative at a point $p \in M$ is a linear map $T_p M \rightarrow T_{f(p)} N$ that we shall denote by

$$T_p M \xrightarrow{T_p f} T_{f(p)} N.$$

The resulting map $Tf : TM \rightarrow TN$ between tangent bundles is also called the **tangent map** of f .

DEFINITION 53.5. Suppose $f : A \rightarrow M$ is a smooth map between smooth manifolds and $B \subset M$ is a smooth submanifold. We say that f is **transverse** to B , written

$$f \pitchfork B,$$

if for every $x \in f^{-1}(B)$ and $y := f(x) \in B$,

$$T_y M = \text{im}(T_x f) + T_y B.$$

If $A \subset M$ is another smooth submanifold, we say that it is transverse to B and write

$$A \pitchfork B$$

if the inclusion map $A \hookrightarrow M$ is transverse to B , which means that for every point $y \in A \cap B$,

$$T_y M = T_y A + T_y B.$$

The transversality condition $f \pitchfork B$ has the following important consequence, due to the implicit function theorem. Assume none of the manifolds in the picture have boundary. Then locally, smooth submanifolds $B \subset M$ can always be presented as regular level sets $g^{-1}(q)$ of smooth functions $g : M \rightarrow \mathbb{R}^k$, where $k := \text{codim}(B)$, and $f \pitchfork B$ then implies that q is also a regular value of the function $g \circ f$, implying:

THEOREM 53.6 (via the implicit function theorem). *Assume $f : A \rightarrow M$ is a smooth map transverse to the submanifold $B \subset M$, where the manifolds A , M and B all have empty boundary. Then $f^{-1}(B)$ is a smooth submanifold of A whose codimension matches the codimension of B in M . In particular, if $A, B \subset M$ are two submanifolds with $A \pitchfork B$, then $A \cap B \subset M$ is a submanifold with*

$$\text{codim}(A \cap B) = \text{codim}(A) + \text{codim}(B).$$

\square

We should also talk about orientations. The key observation here is that if $f \pitchfork B$, then for every $x \in f^{-1}(B) \subset A$, the tangent map $T_x f : T_x A \rightarrow T_{f(x)} M$ descends to an isomorphism of normal spaces

$$N_x^A(f^{-1}(B)) \xrightarrow{\cong} N_{f(x)}^M B,$$

thus producing a commutative diagram

$$\begin{array}{ccc} N^A(f^{-1}(B)) & \xrightarrow{Tf} & N^M B \\ \downarrow \pi & & \downarrow \pi \\ A & \xrightarrow{f} & M \end{array}$$

in which the top arrow is a family of vector space isomorphisms parametrized by $x \in A$. If M and B both carry orientations, then the total space of the normal bundle $N^M B$ inherits an orientation via its identification with a neighborhood of $B \subset M$ from the tubular neighborhood of theorem. This in turn determines an orientation of the bundle $N^M B \rightarrow B$ itself via Definition 52.4. Using the isomorphisms $N_x^A(f^{-1}(B)) \cong N_{f(x)}^M B$, this determines an orientation of the normal bundle $N^A(f^{-1}(B))$ over $f^{-1}(B)$, and using Definition 52.4 again, there is a unique orientation of $f^{-1}(B)$ that is compatible with this orientation of its normal bundle and the given orientation of A . This will be our convention for orienting preimages $f^{-1}(B)$ when $f \pitchfork B$, and in the special case where f is the inclusion of a submanifold $A \hookrightarrow M$, this also determines an orientation of $A \cap B$. Note that if we reverse the order and write $B \cap A$, then the implication is that the orientation may be different, because we are now orienting $A \cap B$ via its normal bundle as a submanifold of B instead of A . One can check that the orientations of $A \cap B$ and $B \cap A$ differ if and only if the codimensions of A and B are both odd.

53.3. Main results. Here is the precise version of the formula for $[A] \cdot [B]$ that was advertised.

THEOREM 53.7. *Assume M is a closed oriented n -manifold, $A, B \subset M$ are closed oriented submanifolds such that $A \pitchfork B$, and the closed submanifold $A \cap B$ is endowed with the induced orientation as explained in the previous section. Then*

$$[A]_M \cdot [B]_M = [A \cap B]_M.$$

We shall deduce this from the following result, which is more fundamental.

THEOREM 53.8. *Assume M is a closed oriented n -manifold, $B \subset M$ is a closed oriented submanifold, A is a closed manifold, $f : A \rightarrow M$ is a smooth map transverse to B , and the submanifold $f^{-1}(B) \subset A$ is endowed with the induced orientation as explained in the previous section. Then if $\beta \in H^*(M)$ is the class Poincaré dual to $[B]_M \in H_*(M)$, its pullback $f^*\beta \in H^*(A)$ is Poincaré dual to $[f^{-1}(B)]_A \in H_*(A)$.*

Let us derive a generalization of Theorem 53.7 from this result. In the situation of Theorem 53.8, suppose $\alpha, \beta \in H^*(M)$ satisfy

$$\text{PD}_M(\alpha) = f_*[A] \quad \text{and} \quad \text{PD}_M(\beta) = [B]_M.$$

Using the associativity and naturality of the cap product, the theorem then implies

$$\begin{aligned} f_*[A] \cdot [B]_M &= \text{PD}_M(\beta \cup \alpha) = (\beta \cup \alpha) \cap [M] = \beta \cap (\alpha \cap [M]) = \beta \cap f_*[A] \\ &= f_*(f^*\beta \cap [A]) = f_*\text{PD}_A(f^*\beta) = f_*[f^{-1}(B)]_A. \end{aligned}$$

COROLLARY 53.9. *In the setting of Theorem 53.8,*

$$f_*[A] \cdot [B] = f_*[f^{-1}(B)]_A.$$

□

Theorem 53.7 is the special case of this corollary where $f : A \rightarrow M$ is the inclusion of a smooth submanifold.




FIGURE 28. The surface and 1-dimensional submanifolds discussed in Example 53.10.

53.4. Applications. In many situations of geometric interest, the intersection product provides an easy criterion for recognizing when a homology class is nontrivial.

EXAMPLE 53.10. Figure 28 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. $[\alpha] = 0 \in H_1(\Sigma)$. 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 oriented surfaces Σ_{\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 $H_1(\beta) \rightarrow H_1(\Sigma)$ is zero because the map $H_1(\beta) \rightarrow H_1(\Sigma_+)$ is zero (see Exercise 50.2),

$$\begin{array}{ccccc}
 H_1(\beta) & \xrightarrow{0} & H_1(\Sigma_+) & \longrightarrow & H_1(\Sigma), \\
 & & \searrow i_* & \nearrow & \\
 & & & &
 \end{array}$$

hence $i_*[\beta] = 0$. The case of $\gamma \subset \Sigma$ is less obvious: it does not split Σ in two pieces, as $\Sigma \setminus \gamma$ 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 $[\gamma] \in H_1(\Sigma)$ 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

$$[\gamma] \cdot [\delta] = \pm 1 \in \mathbb{Z} \cong H_0(\Sigma).$$

This proves that both of the classes $[\gamma], [\delta] \in H_1(\Sigma)$ are not only nontrivial, but also primitive in homology with integer coefficients.

The nonseparating loops in Example 53.10 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 $M \setminus \Sigma$ is disconnected.

THEOREM 53.11. *Suppose M is a closed, connected and oriented smooth n -manifold containing a closed, connected and oriented smooth hypersurface $\Sigma \subset M$. Then the homology class $[\Sigma] \in H_{n-1}(M; \mathbb{Z})$ 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 53.10 implies that $[\Sigma] = 0$. On the other hand, if Σ does not separate M , then $M \setminus \Sigma$ is connected, so we can fix a point $z \in \Sigma$ and two nearby points $z_{\pm} \in M \setminus \Sigma$ that lie in a common Euclidean neighborhood with z identifying Σ with $\mathbb{R}^{n-1} \times \{0\} \subset \mathbb{R}^n$,⁸³ but on opposite sides of Σ , i.e. the n th coordinates of z_{\pm}

⁸³One 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 $\varphi : \mathcal{U} \xrightarrow{\cong} \varphi(\mathcal{U}) \stackrel{\text{open}}{\subset} \mathbb{R}^n$ 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 $k = \dim \Sigma$.

and z_- have opposite signs. We can then find a smooth path γ joining z_+ to z_- in $M \setminus \Sigma$, and then complete it with a path in the Euclidean neighborhood that passes through Σ once, producing (as in Figure 28) a smooth loop $\gamma : S^1 \rightarrow M$ that intersects Σ exactly once and transversely. It follows that

$$\langle 1, \gamma_*[S^1] \cdot [\Sigma] \rangle = \pm 1,$$

hence $[\Sigma] \in H_{n-1}(M)$ and $\gamma_*[S^1] \in H_1(M)$ are both nontrivial. \square

REMARK 53.12. If one drops all orientation assumptions from Theorem 53.11, it remains valid as a statement about homology with \mathbb{Z}_2 coefficients. This observation is relevant in the corollary below.

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 53.13. *For every $n \geq 2$, 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 53.11) a path in $S^n \setminus \Sigma$ that stays within a small neighborhood of Σ but starts and ends on opposite sides of it, thus giving rise to a loop $\gamma : S^1 \rightarrow S^n$ that intersects Σ once transversely. Using \mathbb{Z}_2 coefficients (since Σ is \mathbb{Z}_2 -orientable), the intersection number of Σ with γ is then

$$\langle 1, \gamma_*[S^1] \cdot [\Sigma] \rangle = 1 \in \mathbb{Z}_2,$$

implying $[\Sigma] \neq 0 \in H_{n-1}(S^n; \mathbb{Z}_2)$ and $\gamma_*[S^1] \neq 0 \in H_1(S^n; \mathbb{Z}_2)$. This contradicts the computation of $H_*(S^n; \mathbb{Z}_2)$. \square

REMARK 53.14. The fact that S^n is orientable is not the decisive factor in Corollary 53.13, as there is no obstruction in general to embedding closed non-orientable hypersurfaces into closed orientable manifolds. An easy example is $\mathbb{RP}^2 \hookrightarrow \mathbb{RP}^3$.

Another interesting application of the intersection product is explained in Exercise 53.2 at the end of this lecture: it is a quantitative version of the Lefschetz fixed point theorem that applies to smooth maps on a closed manifold, and interprets the Lefschetz number as an algebraic count of fixed points.

53.5. Pullbacks and preimages. We now prove Theorem 53.8. In the setting of the theorem, denote

$$Q := f^{-1}(B) \subset A,$$

which is necessarily an oriented submanifold with

$$k := \text{codim}(Q) = \text{codim}(B),$$

and let

$$\beta \in H^k(M), \quad \gamma \in H^k(A)$$

denote the cohomology classes satisfying $\text{PD}_M(\beta) = [B]_M$ and $\text{PD}_A(\gamma) = [Q]_A$. Our goal is to prove $f^*\beta = \gamma$.

Consider the inclusions of pairs

$$(M, \emptyset) \xleftarrow{i} (M, M \setminus B) \xleftarrow{j} (\mathbb{D}N^M B, SN^M B),$$

where the second inclusion arises from the tubular neighborhood theorem, and the maps it induces on relative homology and cohomology are isomorphisms due to homotopy invariance and excision.

As in the proof of Theorem 53.1, we can now represent three relative cap products and the naturality relations between them via a diagram

$$\begin{array}{ccc}
 H^k(M) \otimes H_n(M) & \xrightarrow{\quad \cap \quad} & H_{n-k}(M) \\
 i^* \uparrow & & \downarrow i_* = \mathbf{1} \\
 H^k(M|B) \otimes H_n(M|B) & \xrightarrow{\quad \cap \quad} & H_{n-k}(M) \\
 j^* \downarrow \cong & & \uparrow j_* \\
 H^k(\mathbb{D}N^M B, SN^M B) \otimes H_n(\mathbb{D}N^M B, SN^M B) & \xrightarrow{\quad \cap \quad} & H_{n-k}(\mathbb{D}N^M B)
 \end{array}$$

By the defining property of fundamental classes, we have

$$i_*[M] = [M]_B = j_*[\mathbb{D}N^M B].$$

Moreover, since j^* is an isomorphism on relative cohomology, there is a unique class $\hat{\beta} \in H^k(M|B)$ satisfying

$$j^* \hat{\beta} = \tau(N^M B),$$

and in light of Theorem 53.1, the naturality relation represented by the bottom two rows of the diagram implies

$$\begin{aligned}
 [B]_M &= j_*[B]_{\mathbb{D}N^M B} = j_* (\tau(N^M B) \cap [\mathbb{D}N^M B]) = j_* (j^* \hat{\beta} \cap [\mathbb{D}N^M B]) = \hat{\beta} \cap j_*[\mathbb{D}N^M B] \\
 &= \hat{\beta} \cap [M]_B.
 \end{aligned}$$

Applying the relation represented by the top two rows then gives

$$PD_M(i^* \hat{\beta}) = i^* \hat{\beta} \cap [M] = i_* (i^* \hat{\beta} \cap [M]) = \hat{\beta} \cap i_*[M] = \hat{\beta} \cap [M]_B = [B]_M,$$

and since PD_M is an isomorphism, it follows that

$$i^* \hat{\beta} = \beta.$$

This calculation already has a deep and non-obvious consequence: the cohomology class β Poincaré dual to $[B]_M$ is uniquely determined by a class living in the cohomology of M restricted to B , i.e. it depends only on a neighborhood of $B \subset M$, and not on the rest of M . Pulling back via the inclusion of the normal bundle $N^M B$ via the tubular neighborhood theorem, that class becomes precisely the Thom class of $N^M B$.

One can carry out a similar argument with the submanifold $Q = f^{-1}(B) \subset A$ to show that γ is similarly determined by the Thom class of the normal bundle $N^A Q$, with a class $\hat{\gamma} \in H^k(A|Q)$ serving as intermediary between them two. To finish, we note that the embeddings of tubular neighborhoods of $Q \subset A$ and $B \subset M$ can be arranged so that there is a commuting diagram,

$$\begin{array}{ccccc}
 (A, \emptyset) & \xleftarrow{i} & (A, A \setminus Q) & \xleftarrow{j} & (\mathbb{D}N^A Q, SN^A Q) \\
 \downarrow f & & \downarrow f & & \downarrow Tf \\
 (M, \emptyset) & \xleftarrow{i} & (M, M \setminus B) & \xleftarrow{j} & (\mathbb{D}N^M B, SN^M B)
 \end{array}$$

where $Tf : (\mathbb{D}N^A Q, SN^A Q) \rightarrow (\mathbb{D}N^M B, SN^M B)$ restricts for each $x \in A$ to an orientation-preserving homeomorphism $(\mathbb{D}N_x^A Q, SN_x^A Q) \cong (\mathbb{D}N_{f(x)}^M B, SN_{f(x)}^M B)$, defined as a restriction of

the vector space isomorphism $T_x f : N_x^A Q \rightarrow N_{f(x)}^M B$. This induces a diagram

$$\begin{array}{ccccc} H^k(A) & \xleftarrow{i^*} & H^k(A|Q) & \xrightarrow[\cong]{j^*} & H^k(\mathbb{D}N^A Q, SN^A Q) \\ \uparrow f^* & & \uparrow f^* & & \uparrow (Tf)^* \\ H^k(M) & \xleftarrow{i^*} & H^k(M|B) & \xrightarrow[\cong]{j^*} & H^k(\mathbb{D}N^M B, SN^M B) \end{array},$$

and in light of the homeomorphisms $(\mathbb{D}N_x^A Q, SN_x^A Q) \cong (\mathbb{D}N_{f(x)}^M B, SN_{f(x)}^M B)$ defined via Tf , the defining property of Thom classes implies

$$(Tf)^* \tau(N^M B) = \tau(N^A Q).$$

Since the maps j^* are isomorphisms, this implies $f^* \hat{\beta} = \hat{\gamma}$, and pulling back to the left hand side of the diagram then gives $f^* \beta = \gamma$, concluding the proof of Theorem 53.8.

53.6. Relative intersection products. For M a compact, smooth, oriented n -manifold with nonempty boundary, there are multiple ways to generalize the intersection product to relative homology. We will give a quick sketch here and leave the details as exercises.

Just as there are two choices of relative Poincaré duality isomorphism to use, there are essentially two interesting versions of the cup product: in the first variant, one can use the usual cup product on absolute cohomology but identify this via duality with the *relative* homology of the pair $(M, \partial M)$, giving a relative intersection product of the form

$$(53.3) \quad H_{n-k}(M, \partial M) \otimes H_{n-\ell}(M, \partial M) \xrightarrow{\quad} H_{n-(k+\ell)}(\partial M).$$

Two slightly more subtle variants arise by using the relative cup products

$$H^k(M) \otimes H^\ell(M, \partial M) \xrightarrow{\cup} H^{k+\ell}(M, \partial M), \quad H^k(M, \partial M) \otimes H^\ell(M) \xrightarrow{\cup} H^{k+\ell}(M, \partial M)$$

together with the duality isomorphism $PD_M : H^m(M, \partial M) \rightarrow H_{n-m}(M)$. This produces intersection products of the form

$$(53.4) \quad \begin{aligned} H_{n-k}(M, \partial M) \otimes H_{n-\ell}(M) &\xrightarrow{\quad} H_{n-(k+\ell)}(M), \\ H_{n-k}(M) \otimes H_{n-\ell}(M, \partial M) &\xrightarrow{\quad} H_{n-(k+\ell)}(M). \end{aligned}$$

Finally, there is also a relative cup product

$$H^k(M, \partial M) \otimes H^\ell(M, \partial M) \xrightarrow{\cup} H^{k+\ell}(M, \partial M),$$

which is dual to an absolute intersection product

$$(53.5) \quad H_{n-k}(M) \otimes H_{n-\ell}(M) \xrightarrow{\quad} H_{n-(k+\ell)}(M),$$

the meaning of which turns out to be exactly the same as when $\partial M = \emptyset$. This last variant can actually be seen as a special case of (53.4): the product $A \cdot B \in H_*(M)$ of two absolute classes $A, B \in H_*(M)$ becomes the same thing as in (53.4) if one first feeds either A or B into the homomorphism $H_*(M) \rightarrow H_*(M, \partial M)$ induced by the inclusion $(M, \emptyset) \hookrightarrow (M, \partial M)$, which is equivalent to feeding its dual cohomology class into the pullback homomorphism $H^*(M, \partial M) \rightarrow H^*(M)$ and then computing a relative cup product. For these reasons, the absolute intersection product in (53.5) can be ignored in favor of the relative products (53.4).

All of these variants are identical to the usual absolute intersection product if $\partial M = \emptyset$. If M is a connected manifold with $\partial M \neq \emptyset$ and we consider the case of complementary dimensions $k + \ell = n$, then the variants in (53.4) have an advantage over (53.3): these intersection products give an *absolute* homology class in degree 0, which can therefore be paired with $1 \in H^0(M)$ to recover the usual relation between the intersection product and the intersection form:

$$\langle 1, PD_M(\alpha) \cdot PD_M(\beta) \rangle = \langle 1, (\beta \cup \alpha) \cap [M] \rangle = \langle \beta \cup \alpha, [M] \rangle = Q(\beta, \alpha).$$

Note that in general, this version of the intersection form Q pairs an *absolute* cohomology class with a *relative* class in $(M, \partial M)$; the crucial detail is that the product is then in $H_n(M, \partial M)$ instead of $H_n(M)$, so that it can then be evaluated on the relative fundamental class $[M] \in H_n(M, \partial M)$. The usual definition of the intersection form on absolute cohomology classes in this situation gives something trivial, because the interior of M is a connected and noncompact n -manifold, implying $H_n(M) \cong H_n(\overset{\circ}{M}) = 0$.

To understand what the various relative intersection products mean, one can generalize Theorem 53.8 as follows. Assume A, B, M are all compact, oriented smooth manifolds with boundary, $f : (A, \partial A) \rightarrow (M, \partial M)$ is a smooth map of pairs with $f^{-1}(\partial M) = \partial A$, $B \subset M$ is a submanifold with $\partial B = B \cap \partial M$, and we additionally have the following transversality conditions:

- $f \pitchfork \partial M$ and $B \pitchfork \partial M$;
- The map $f|_{\overset{\circ}{A}} : \overset{\circ}{A} \rightarrow \overset{\circ}{M}$ is transverse in $\overset{\circ}{M}$ to the submanifold $\overset{\circ}{B} \subset \overset{\circ}{M}$;
- The map $f|_{\partial A} : \partial A \rightarrow \partial M$ is transverse in ∂M to the submanifold $\partial B \subset \partial M$.

In this situation, if $B \subset M$ has codimension k , then ∂B is similarly a codimension k submanifold of ∂M , and

$$Q := f^{-1}(B) \subset A$$

becomes a smooth codimension k submanifold of A that intersects ∂A transversely such that $\partial Q = Q \cap \partial A$ is a codimension k submanifold of ∂A . The generalization of Theorem 53.8 to this setup is straightforward to state, and its proof only requires extending the arguments in §53.5 to a slightly wider context: the result concerns the absolute cohomology classes in $H^k(M)$ and $H^k(A)$ that are Poincaré dual to the relative homology classes $[B]_M \in H_*(M, \partial M)$ and $[Q]_A \in H_*(A, \partial A)$ respectively, and it simply says

$$\text{PD}_M(\beta) = [B]_M \quad \implies \quad \text{PD}_A(f^*\beta) = [f^{-1}(B)]_A.$$

Taking f to be the inclusion of a submanifold $(A, \partial A) \hookrightarrow (M, \partial M)$ with boundary $\partial A = A \cap \partial M$ such that $A \pitchfork \partial M$ and $\partial A \pitchfork \partial B$ in ∂M , one obtains $A \cap B \subset M$ as a compact oriented submanifold intersecting ∂M transversely in $\partial(A \cap B) = A \cap B \cap \partial M$, and the interpretation of (53.3) then becomes

$$[A]_M \cdot [B]_M = [A \cap B]_M \in H_{n-(k+\ell)}(M, \partial M).$$

If one assumes additionally $\partial A = \emptyset$, so that $A \subset M$ is contained in the interior of M , then one has the freedom to interpret $[A]_M$ as living in $H_*(M)$ instead of $H_*(M, \partial M)$, so that the second of the relative intersection products in (53.4) makes sense, and its interpretation is the same: $A \cap B$ is in this case a *closed* oriented submanifold in the interior of M and thus represents a class in $H_*(M)$. Exploiting the graded commutativity of the cup product, the first product in (53.4) inherits from this a similar interpretation.

53.7. Exercises.

EXERCISE 53.1. Suppose M, N are closed, connected, oriented smooth manifolds of dimension $n \in \mathbb{N}$ and $f : M \rightarrow N$ is a smooth map. Without appealing to the results of Lecture 36, use the intersection product $f_*[M] \cdot [*]_N$ with the homology class of a point $[*]_N \in H_0(N)$ to prove that under a suitable technical assumption on a point $y \in N$, the degree $\deg(f) \in \mathbb{Z}$ of f is a signed count of points in $f^{-1}(y) \subset M$. What exactly is the technical assumption you need?

EXERCISE 53.2. Throughout this exercise, M is a closed, connected and oriented smooth manifold of dimension $n \in \mathbb{N}$. We consider intersection products of homology classes with complementary degree and regard such products as integers

$$A \cdot B \in \mathbb{Z} \cong H_0(M) \quad \text{for } A \in H_k(M), B \in H_{n-k}(M),$$

using the canonical isomorphism $H_0(M) \rightarrow \mathbb{Z} : C \mapsto \langle 1, C \rangle$. For smooth maps $f : M \rightarrow M$, there is a sharp version of the Lefschetz fixed point theorem that can be stated as an intersection product calculation, namely

$$(53.6) \quad [\Gamma_f]_{M \times M} \cdot [\Delta]_{M \times M} = L(f),$$

where Γ_f denotes the **graph** of f ,

$$\Gamma_f := \{(x, f(x)) \mid x \in M\} \subset M \times M,$$

and Δ is the **diagonal** submanifold

$$\Delta := \Gamma_{\text{Id}} = \{(x, x) \mid x \in M\} \subset M \times M.$$

We equip $\Gamma_f \subset M \times M$ with the orientation such that the embedding

$$M \xrightarrow{i} M \times M : x \mapsto (x, f(x))$$

defines an orientation-preserving diffeomorphism of M to Γ_f ; since Δ is a special case of a graph, it inherits an orientation in the same way. In light of the obvious correspondence between $\Gamma_f \cap \Delta$ and the fixed point set of f , (53.6) implies the Lefschetz fixed point theorem for smooth maps on M ; moreover, it gives a quantitative interpretation of the Lefschetz number $L(f)$ as a signed count of fixed points, under the technical condition that Γ_f and Δ intersect transversely. Our goal in this problem is to prove this formula.⁸⁴

- (a) Prove that if $A, B, C, D \in H_*(M)$ are homology classes (with integer coefficients) whose degrees satisfy

$$|A| + |B| = n = |C| + |D|,$$

then we have the following formula for intersections of cross products of degree n in $M \times M$:

$$(A \times B) \cdot (C \times D) = \begin{cases} (-1)^{|B|} (A \cdot C)(B \cdot D) & \text{if } |B| = |C| \text{ and } |A| = |D|, \\ 0 & \text{otherwise.} \end{cases}$$

Remark: You should forgive yourself if you manage to figure out every detail of this problem except the signs, but my advice on signs is this. For most of the calculation, you'll need to keep track of degrees of cohomology classes rather than homology classes. The assumption $|A| + |B| = n = |C| + |D|$ will allow you to view those degrees also as degrees of homology classes if you prefer.

- (b) For an arbitrary map $f : M \rightarrow M$ and homology classes $A, B \in H_*(M)$ of complementary degree $|A| + |B| = n$, prove the formula

$$[\Gamma_f]_{M \times M} \cdot (A \times B) = (-1)^{|A|} f_* A \cdot B \in \mathbb{Z}.$$

Hint: Our definition of the orientation on Γ_f means $[\Gamma_f]_{M \times M} = i_*[M]$. You may at some point find yourself needing to compute $i^*(\alpha \times \beta)$ for some cohomology classes $\alpha, \beta \in H^*(M)$. Try transforming the cross product into a cup product.

Recall that the Lefschetz number $L(f)$ is an alternating sum of the traces of the homomorphisms $f_* : H_k(M; \mathbb{Q}) \rightarrow H_k(M; \mathbb{Q})$ for $k = 0, \dots, n$. Choose a basis $\{e_i\}$ of the rational vector space $H_*(M; \mathbb{Q}) = \bigoplus_{k=0}^n H_k(M; \mathbb{Q})$ consisting only of homogeneous elements. The intersection product on $H_*(M; \mathbb{Q})$ is equivalent to the intersection form on $H^*(M; \mathbb{Q})$ and is therefore nonsingular, so that the basis $\{e_i\}$ uniquely determines a *dual* basis $\{e'_i\}$ satisfying the condition

$$e_i \cdot e'_j = \delta_{ij} := \begin{cases} 1 & \text{if } i = j, \\ 0 & \text{if } i \neq j, \end{cases}$$

⁸⁴As usual, there is also a mod 2 version of (53.6) that holds without any orientability assumptions, in which $L(f)$ is replaced by the \mathbb{Z}_2 Lefschetz number $L_{\mathbb{Z}_2}(f) \in \mathbb{Z}_2$.

with the convention that numerical intersection products $A \cdot B \in \mathbb{Z}$ are understood to be 0 if $|A| + |B| \neq n$. Note that by the Künneth formula, the set of all cross products of the form $e_i \times e'_j$ forms a basis of $H_*(M \times M; \mathbb{Q})$, and so does the set of products of the form $e'_i \times e_j$.

- (c) Prove the formula $[\Delta]_{M \times M} = \sum_k e_k \times e'_k$, where the sum ranges over all the chosen basis elements e_k .

Hint: By the nonsingularity of the intersection product, it suffices to check that both sides have the same intersection pairing with $e'_i \times e_j$ for every i, j .

- (d) Prove (53.6).

54. Higher homotopy groups

The last two lectures in this semester 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.

54.1. Definitions and basic properties. 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

$$[X, Y] := \{\text{continuous maps } X \rightarrow Y\} / \sim,$$

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 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

$$\pi_1(X, x_0) = [(S^1, *), (X, x_0)] = [(I, \partial I), (X, x_0)],$$

where $*$ 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 54.1. For each integer $n \geq 0$, we define the set

$$\pi_n(X, x_0) := [(S^n, *), (X, x_0)].$$

When $n \geq 1$, 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,

$$(I_1 \amalg I_2) / (I_1 \ni 1 \sim 0 \in I_2)$$




FIGURE 29. The homotopy in the proof of Proposition 54.3.

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

$$(I_1^n \amalg I_2^n) / (I_1^n \ni (1, t_2, \dots, t_n) \sim (0, t_2, \dots, t_n) \in I_2^n) \cong I^n,$$

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 $n \geq 2$, since the fundamental 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 54.2. For $n \geq 2$ and two elements $[f], [g] \in \pi_n(X, x_0)$ represented by maps $f, g : (I^n, \partial I^n) \rightarrow (X, x_0)$, we define $[f] + [g] \in \pi_n(X, x_0)$ to be the homotopy class of the map

$$(I^n, \partial I^n) \rightarrow (X, x_0) : (t_1, \dots, t_n) \mapsto \begin{cases} f(2t_1, t_2, \dots, t_n) & \text{if } 0 \leq t_1 \leq 1/2, \\ g(2t_1 - 1, t_2, \dots, t_n) & \text{if } 1/2 \leq t_1 \leq 1. \end{cases}$$

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 54.3. For all $n \geq 2$, the operation in Definition 54.2 makes $\pi_n(X, x_0)$ an abelian group.

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 $n \geq 2$; as we've seen, the fundamental group does not generally have this property. The proof is a homotopy depicted in Figure 29. 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 54.2, 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 $*$ in $\partial\mathbb{D}^n$ and setting

$$\pi_n(X, A, x_0) := [(\mathbb{D}^n, \partial\mathbb{D}^n, *), (X, A, x_0)].$$

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 $n \geq 1$; 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 $n \geq 2$. 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 treating the rest of ∂I^n as if it were a base point: let

$$J_n := I^{n-1} \times \{0\} \subset \partial I^n$$

and redefine $\pi_n(X, A, x_0)$ as

$$\pi_n(X, A, x_0) := [I^n, \partial I^n, \overline{\partial I^n \setminus J_n}, (X, A, x_0)].$$

By this definition, the formula in Definition 54.2 still makes sense for $n \geq 2$ and defines a group structure on $\pi_n(X, A, x_0)$, though Proposition 54.3 no longer works in the $n = 2$ case. You can see why not if you look again at Figure 29 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 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 $n \geq 3$, 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 54.4. *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 $n \geq 2$, and it is abelian for $n \geq 3$. \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 30 shows a recipe for transforming any map $f : (I^n, \partial I^n) \rightarrow (X, y)$ into a map $(I^n, \partial I^n) \rightarrow (X, x)$ 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 $x \xrightarrow{\gamma} y$. 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 $\Phi_\gamma : \pi_1(X, y) \rightarrow \pi_1(X, x)$ 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 \mathbf{Top}_* of pointed spaces to the category \mathbf{Grp} of groups (or \mathbf{Ab} for $n \geq 2$). Namely, every base-point preserving map $f : (X, x_0) \rightarrow (Y, y_0)$ induces a homomorphism

$$f_* : \pi_n(X, x_0) \rightarrow \pi_n(Y, y_0) : [\varphi] \mapsto [f \circ \varphi].$$




FIGURE 30. The isomorphism $\pi_n(X, y) \rightarrow \pi_n(X, x)$ determined by a path $x \rightsquigarrow y$.

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 54.5. *If $f : X \rightarrow Y$ is a homotopy equivalence, then $f_* : \pi_n(X, x_0) \rightarrow \pi_n(Y, f(x_0))$ is an isomorphism for all $n \geq 0$ and $x_0 \in X$.*

Since we've been talking about homology for the rest of this semester, you may have forgotten why Theorem 54.5 is already a nontrivial statement in the $n = 1$ case, which took some effort to prove in *Topology I*. The annoying detail is the base point: if $g : Y \rightarrow X$ 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 $\Phi_\gamma : \pi_n(X, g(f(x_0))) \rightarrow \pi_n(X, x_0)$ 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 $f_* : H_*(X) \rightarrow H_*(Y)$ immediately implied that $H_*(X)$ depends only on the homotopy type of X .

54.2. Sample computations. 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 54.6. For every $k \geq 2$ 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 $p : \mathbb{R}^n \rightarrow \mathbb{T}^n$. Since S^k is simply connected for $k \geq 2$, every map $f : S^k \rightarrow \mathbb{T}^n$ has a lift $\tilde{f} : S^k \rightarrow \mathbb{R}^n$, which is homotopic to a constant map since \mathbb{R}^n is contractible. Composing this homotopy with $p : \mathbb{R}^n \rightarrow \mathbb{T}^n$ then gives a homotopy of f to a constant map $S^k \rightarrow \mathbb{T}^n$. (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 54.7. For $n \in \mathbb{N}$ and $k < n$, $\pi_k(S^n) = 0$. One can see this by proving that every map $f : S^k \rightarrow S^n$ with $n > k$ is homotopic to a map $g : S^k \rightarrow S^n$ that is not surjective: then if $p \in S^n \setminus g(S^k)$, it follows that the image of g is in $S^n \setminus \{p\} \cong \mathbb{R}^n$, 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 §31.5) implies that for suitable choices of triangulations of S^k and S^n , f is homotopic to a simplicial map $g : S^k \rightarrow S^n$, 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]: $f : S^k \rightarrow S^n$ is homotopic to a smooth map $g : S^k \rightarrow S^n$ 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 $T_x g : T_x S^k \rightarrow T_y S^n$ is surjective for every $x \in g^{-1}(y)$, but since that condition can never be satisfied for $n > k$, it follows that $g^{-1}(y) = \emptyset$.

EXAMPLE 54.8. Viewing elements of $\pi_n(S^n)$ as represented by maps $f : S^n \rightarrow S^n$, 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 is a nontrivial result that does not follow directly from anything we've covered in this course. The case $n = 2$ is implied by a long exact sequence appearing in Example 54.17 below, which gives an isomorphism

$$\pi_2(S^2) \cong \pi_1(S^1).$$

There are various approaches for proving the cases $n > 2$. One is to deduce it from the *Freudenthal suspension theorem*, which states that the natural map

$$\pi_n(X) \xrightarrow{\Sigma} \pi_{n+1}(\Sigma X)$$

sending the homotopy class of a map $f : S^n \rightarrow X$ to the homotopy class of its suspension $\Sigma f : \Sigma S^n \cong S^{n+1} \rightarrow \Sigma X$ is an isomorphism under certain conditions; in particular, the theorem applies to $X := S^n$ for $n \geq 2$, giving isomorphisms

$$\pi_n(S^n) \xrightarrow[\cong]{\Sigma} \pi_{n+1}(S^{n+1}) \quad \text{for each } n \geq 2.$$

A minor technical point: since homotopy groups depend on base points, the Freudenthal theorem requires a slight modification to our usual suspension functor $\Sigma : \mathbf{Top} \rightarrow \mathbf{Top}$, called the *reduced suspension*, which defines ΣX slightly differently and gives rise to a functor $\Sigma : \mathbf{Top}_* \rightarrow \mathbf{Top}_*$ on the category of pointed spaces. We will prove the Freudenthal suspension theorem in next semester's *Topologie III* course, as an application of the homotopy excision theorem. There is an alternative approach to proving $\pi_n(S^n) \cong \mathbb{Z}$ via differential topology as in [Mil97], using the so-called *Pontryagin-Thom* construction, which 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 . The latter have a natural correspondence with the integers when $k = n$, so in particular, when $\dim M = n$ this proves that the map $\deg : [M, S^n] \rightarrow \mathbb{Z}$ 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. The Pontryagin-Thom construction can also be used to give an alternative proof of the suspension theorem, see e.g. [Fre12, Lecture 4].

REMARK 54.9. Another approach sometimes cited for the computation $\pi_n(S^n) \cong \mathbb{Z}$ is to deduce it from the Hurewicz theorem, which we will outline in the next lecture: for S^n , it gives a natural isomorphism $\pi_n(S^n) \xrightarrow{\cong} H_n(S^n; \mathbb{Z}) \cong \mathbb{Z}$ for each $n \geq 2$, generalizing the natural isomorphism between $H_1(X; \mathbb{Z})$ and the abelianization of the fundamental group for a path-connected space X . There is a logical drawback to explaining $\pi_n(S^n)$ in this way: the simplest proofs of the Hurewicz

theorem are based on explicit computations in cell complexes that require the result $\pi_n(S^n) \cong \mathbb{Z}$ as a prerequisite.

In Example 54.17 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 54.10. Just to give you a taste of what is studied in modern homotopy theory: the Freudenthal suspension theorem gives natural isomorphisms

$$\pi_{n+k}(S^n) \xrightarrow[\cong]{\Sigma} \pi_{n+k+1}(S^{n+1})$$

for all $k \geq 0$ 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, and can be defined formally as direct limits

$$\pi_k^s := \varinjlim \{\pi_{n+k}(S^n)\}_{n=0}^\infty \cong \pi_{n+k}(S^n) \text{ for } n \gg 0.$$

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.

54.3. A taste of obstruction theory. 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 54.11. For integers $n \geq 0$, a space X is called *n -connected* if $\pi_k(X) = 0$ for all $k \leq n$.

We can now give an example of the kind of problem for which computing higher homotopy groups is useful.

THEOREM 54.12. *If X is a CW-complex of dimension at most n and Y is an n -connected space, then all maps $X \rightarrow Y$ are homotopic.*

PROOF. We need to show that any two given maps $f, g : X \rightarrow Y$ are homotopic. The method of the proof is known as “induction over the skeleta”.⁸⁵ As preparation, one needs to think through the following exercise: if $f|_{X^k} : X^k \rightarrow Y$ is homotopic to $g|_{X^k} : X^k \rightarrow Y$ for some $k \geq 0$, then f is also homotopic on X to a map $f' : X \rightarrow Y$ 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 $e_\alpha^k \subset X$ are homotopic via a homotopy that is fixed at the boundary of the cell, i.e. on the $(k-1)$ -skeleton. Let $\Phi_\alpha : (\mathbb{D}^k, S^{k-1}) \rightarrow (X^k, X^{k-1})$ denote the characteristic map of e_α^k . Then $f \circ \Phi_\alpha$ and $g \circ \Phi_\alpha$ are two maps $\mathbb{D}^k \rightarrow Y$ 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

$$F : S^k \rightarrow Y : x \mapsto \begin{cases} f \circ \Phi_\alpha(x) & \text{if } x \in \mathbb{D}_+^k, \\ g \circ \Phi_\alpha(x) & \text{if } x \in \mathbb{D}_-^k. \end{cases}$$

⁸⁵It 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”.

Since $\pi_k(Y) = 0$, the map $F: S^k \rightarrow Y$ is homotopic to a constant, which is equivalent to saying that it extends to a map $\mathbb{D}^{k+1} \rightarrow Y$, 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. \square

You may notice that Theorem 54.12 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 particular a sphere) such that not all maps $X \rightarrow Y$ 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 49.6). 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. A detailed exposition of this subject can be found e.g. [Ste51], and we will have more to say about it in next semester's *Topologie III* course.

54.4. Long exact sequences. 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

$$(A, x_0) \xrightarrow{i} (X, x_0) \quad \text{and} \quad (X, x_0, x_0) \xrightarrow{j} (X, A, x_0)$$

the obvious inclusions. For each $n \geq 1$ there is also a natural homomorphism

$$\partial_* : \pi_n(X, A, x_0) \rightarrow \pi_{n-1}(A, x_0) : [f] \mapsto [f|_{S^{n-1}}],$$

where we regard elements of $\pi_n(X, A, x_0)$ as represented by maps $f : (\mathbb{D}^n, S^{n-1}, *) \rightarrow (X, A, x_0)$. You can easily check by translating this into the corresponding formula with $f : (I^n, \partial I^n, \partial I^n \setminus \overline{J_n}) \rightarrow (X, A, x_0)$ that it really is a homomorphism.

THEOREM 54.13. *For $x_0 \in A \subset X$, the sequence*

$$\begin{aligned} \dots \rightarrow \pi_{n+1}(X, A, x_0) \xrightarrow{\partial_*} \pi_n(A, x_0) \xrightarrow{i_*} \pi_n(X, x_0) \xrightarrow{j_*} \pi_n(X, A, x_0) \xrightarrow{\partial_*} \pi_{n-1}(A, x_0) \rightarrow \dots \\ \dots \rightarrow \pi_1(X, x_0) \xrightarrow{j_*} \pi_1(X, A, x_0) \xrightarrow{\partial_*} \pi_0(A, x_0) \xrightarrow{i_*} \pi_0(X, x_0). \end{aligned}$$

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 two lectures we saw some examples of vector bundles, which can be imagined 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 54.14. 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 $p : E \rightarrow B$, such that B can be covered by open sets \mathcal{U} that admit **local trivializations**, meaning homeomorphisms

$$\Phi : p^{-1}(\mathcal{U}) \rightarrow \mathcal{U} \times F$$

that send $p^{-1}(b)$ homeomorphically to $\{b\} \times F$ for each $b \in \mathcal{U}$. The **fibers** of the bundle are the subspaces $E_b := p^{-1}(b) \cong F$ 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 $p : E \rightarrow B$ becomes the obvious projection map $B \times F \rightarrow B$.

There's at least one general class of fiber bundles that you've definitely seen plenty of before: a covering map $p : E \rightarrow B$ 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 $p : E \rightarrow B$ can be identified with the trivial covering map, i.e. the one for which E is a disjoint union of copies of B and $p : E \rightarrow B$ 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 54.15. The **Hopf fibration** $p : S^3 \rightarrow S^2$ 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

$$p : S^3 \rightarrow S^2 : (z_1, z_2) \mapsto \frac{z_1}{z_2}.$$

Equivalently, one can identify S^2 with \mathbb{CP}^1 so that this becomes the map

$$p : S^3 \rightarrow \mathbb{CP}^1 : (z_1, z_2) \mapsto [z_1 : z_2].$$

The fiber containing any given point $(z_1, z_2) \in S^3$ is the set

$$\{(e^{i\theta} z_1, e^{i\theta} z_2) \in S^3 \mid \theta \in \mathbb{R}\} \cong S^1.$$

We leave it as an exercise to check that local trivializations exist near every point.

THEOREM 54.16. *Given a fiber bundle $(F, x_0) \hookrightarrow (E, x_0) \xrightarrow{p} (B, b_0)$ with base points, the map $p : (E, F, x_0) \rightarrow (B, b_0, b_0)$ 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

$$\begin{aligned} \dots \longrightarrow \pi_{n+1}(B, b_0) \xrightarrow{\hat{c}_*} \pi_n(F, x_0) \xrightarrow{i_*} \pi_n(E, x_0) \xrightarrow{p_*} \pi_n(B, b_0) \xrightarrow{\hat{c}_*} \pi_{n-1}(F, x_0) \longrightarrow \dots \\ \dots \longrightarrow \pi_1(E, x_0) \xrightarrow{p_*} \pi_1(B, b_0) \xrightarrow{\hat{c}_*} \pi_0(F, x_0) \xrightarrow{i_*} \pi_0(E, x_0), \end{aligned}$$

where the maps $\partial_* : \pi_n(B, b_0) \rightarrow \pi_{n-1}(F, x_0)$ send each $[f]$ to $[\tilde{f}]_{S^{n-1}}$ for $f : (\mathbb{D}^n, S^{n-1}) \rightarrow (B, b_0)$ and $\tilde{f} : (\mathbb{D}^n, S^{n-1}, *) \rightarrow (E, F, x_0)$ solving the lifting problem

$$(54.1) \quad \begin{array}{ccc} & & E \\ & \nearrow \tilde{f} & \downarrow p \\ \mathbb{D}^n & \xrightarrow{f} & B \end{array}$$

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 (54.1). The proper formulation of this condition is something called the *homotopy lifting property*, and Theorem 54.16 is true in fact for any map $p : E \rightarrow B$ 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 (54.1) 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 $\partial_* : \pi_n(B, b_0) \rightarrow \pi_{n-1}(E, x_0)$ described in the theorem is well defined.

EXAMPLE 54.17. Returning to the Hopf fibration of Example 54.15, the homotopy exact sequence has segments of the form

$$0 = \pi_2(S^3) \longrightarrow \pi_2(S^2) \xrightarrow{\hat{c}_*} \pi_1(S^1) \longrightarrow \pi_1(S^3) = 0,$$

and

$$0 = \pi_3(S^1) \longrightarrow \pi_3(S^3) \xrightarrow{p_*} \pi_3(S^2) \longrightarrow \pi_2(S^1) = 0.$$

The first yields the isomorphism $\pi_2(S^2) \cong \pi_1(S^1) \cong \mathbb{Z}$ mentioned in Example 54.8, while the second proves that the map $p_* : \pi_3(S^3) \rightarrow \pi_3(S^2)$ is an isomorphism. Since $\pi_3(S^3) \cong \mathbb{Z}$ is generated by the identity map $S^3 \rightarrow S^3$, this implies that $\pi_3(S^2) \cong \mathbb{Z}$, with the Hopf fibration itself representing a generator.

EXAMPLE 54.18. 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

$$GL_+(n, \mathbb{R}) := \{ \mathbf{A} \in GL(n, \mathbb{R}) \mid \det \mathbf{A} > 0 \}.$$

Here is a trick for computing $\pi_1(GL_+(n, \mathbb{R}))$. Polar decomposition provides a deformation retraction of $GL_+(n, \mathbb{R})$ to $SO(n)$, the special orthogonal group, thus it suffices to compute $\pi_1(SO(n))$. For $n = 1$ and $n = 2$, this is easy because $SO(1) \cong \{*\}$ and $SO(2) \cong S^1$. For $n = 3$, it is not hard to find a homeomorphism of $SO(3)$ to $\mathbb{R}P^3$: this arises from the fact that every element of $SO(3)$ defines a rotation about some axis in \mathbb{R}^3 , so there is a natural map

$$\mathbb{D}^3 \rightarrow SO(3)$$

that sends the origin to $\mathbf{1}$ and sends the point $r\mathbf{x}$ for $0 < r \leq 1$ 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

$$\mathrm{SO}(n) \xrightarrow{i} \mathrm{SO}(n+1) \xrightarrow{p} S^n,$$

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 $e_1 = (1, 0, \dots, 0) \in S^n \subset \mathbb{R}^{n+1}$. The homotopy exact sequence then has segments of the form

$$\dots \longrightarrow \pi_{k+1}(S^n) \longrightarrow \pi_k(\mathrm{SO}(n)) \xrightarrow{i_*} \pi_k(\mathrm{SO}(n+1)) \longrightarrow \pi_k(S^n) \longrightarrow \dots,$$

and taking $k = 1$, both $\pi_2(S^n)$ and $\pi_1(S^n)$ vanish if $n \geq 3$. This produces an infinite sequence of isomorphisms

$$\mathbb{Z}_2 \cong \pi_1(\mathrm{SO}(3)) \cong \pi_1(\mathrm{SO}(4)) \cong \pi_1(\mathrm{SO}(5)) \cong \dots,$$

proving that $\pi_1(\mathrm{GL}_+(n, \mathbb{R})) \cong \mathbb{Z}_2$ for all $n \geq 3$.

55. The theorems of Hurewicz and Whitehead

55.1. Simply connected 3-manifolds. I have more to say about higher homotopy groups, but I want to focus the discussion around a particular application:

THEOREM 55.1. *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 $H_*(M; \mathbb{Z}) \cong H_*(S^3; \mathbb{Z})$ 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 55.1 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 $n \geq 5$ 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. $\mathbb{C}\mathbb{P}^2$ 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 $\mathbb{C}\mathbb{P}^2$ from S^4 via its homology, of course. Part of the interest in Theorem 55.1, 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 semester 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 powerful new theorems involving the higher homotopy groups.

55.2. From simply connected to homology sphere. This part of the argument will be a review of techniques developed in the past semester. Our first main objective is to prove the following lemma:

LEMMA 55.2. *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 55.3. *$H_n(M; \mathbb{Z})$ 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 Euclidean neighborhood retracts; see Theorem 49.1. The groups $H_k(M; \mathbb{Z})$ for $k > n$ vanish by Corollary 49.22. 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 55.4. *$H_1(M; \mathbb{Z}) = 0$.*

PROOF. This is immediate from the isomorphism of $H_1(M; \mathbb{Z})$ with the abelianization of $\pi_1(M)$. \square

LEMMA 55.5. *M is orientable.*

PROOF. If it is not orientable, then its orientation double cover $\pi : \widetilde{M} \rightarrow M$ 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 55.6. *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 49.22. \square

LEMMA 55.7. *$H_2(M; \mathbb{Z})$ is torsion free.*

PROOF. This is true for $H_{n-1}(M; \mathbb{Z})$ whenever M is a closed oriented n -manifold; see Exercise 49.5(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

$$H_n(M; \mathbb{Z}_p) \cong (H_n(M; \mathbb{Z}) \otimes \mathbb{Z}_p) \oplus \text{Tor}(H_{n-1}(M; \mathbb{Z}), \mathbb{Z}_p),$$

where we are working in the setting of \mathbb{Z} -modules and therefore interpret the tensor product \otimes and the Tor functor in that context. Since $H_n(M; \mathbb{Z}_p) \cong \mathbb{Z}_p$ and $H_n(M; \mathbb{Z}) \cong \mathbb{Z}$ by Corollary 49.22, this isomorphism implies the vanishing of $\text{Tor}(H_{n-1}(M; \mathbb{Z}), \mathbb{Z}_p)$. 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 $N \geq 0$, $k_1, \dots, k_N \geq 2$, where $N > 0$ if and only if $H_{n-1}(M; \mathbb{Z})$ has torsion. According to the properties of Tor proved in Lecture 45, we then have

$$0 = \text{Tor}(H_{n-1}(M; \mathbb{Z}), \mathbb{Z}_p) \cong \bigoplus_{i=1}^N \text{Tor}(\mathbb{Z}_{k_i}, \mathbb{Z}_p),$$

implying $\text{Tor}(\mathbb{Z}_{k_i}, \mathbb{Z}_p) = 0$ for every $i = 1, \dots, N$ and every prime p . But if p is chosen to be any prime factor of k_1 , then Theorem 45.1(6) also gives

$$\text{Tor}(\mathbb{Z}_{k_1}, \mathbb{Z}_p) = \ker \left(\mathbb{Z}_p \xrightarrow{\cdot k_1} \mathbb{Z}_p \right) = \ker \left(\mathbb{Z}_p \xrightarrow{0} \mathbb{Z}_p \right) = \mathbb{Z}_p \neq 0,$$

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

$$0 = H^2(M; \mathbb{Z}) \cong \text{Hom}(H_2(M; \mathbb{Z}), \mathbb{Z}) \oplus \text{Ext}(H_1(M; \mathbb{Z}), \mathbb{Z}),$$

hence $\text{Hom}(H_2(M; \mathbb{Z}), \mathbb{Z}) = 0$. 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 $n \geq 3$ by Lemmas 55.3 and 55.6, 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 55.2.

55.3. 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 already discussed in this course. Both will appear as important topics in next semester's *Topologie 3* course.

DEFINITION 55.8. A map $f : X \rightarrow Y$ is called a **weak homotopy equivalence** if for all choices of base points $x_0 \in X$ and $y_0 = f(x_0) \in Y$, $f_* : \pi_n(X, x_0) \rightarrow \pi_n(Y, y_0)$ is an isomorphism for all $n \geq 0$.

Theorem 54.5 in the previous lecture implies that every homotopy equivalence is also a weak homotopy equivalence. We also know of course that if $f : X \rightarrow Y$ 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 55.9 (Whitehead's theorem). *If X and Y are both homotopy equivalent to CW-complexes, then every weak homotopy equivalence $f : X \rightarrow Y$ is a homotopy equivalence.*

While I do not intend to discuss the proof of this theorem here, you will hopefully gain some intuition about it from Theorem 54.12 in the previous lecture. In particular, it should be clear why having cell decompositions of X and Y might be useful in the proof, as e.g. the homotopies needed for showing that $f : X \rightarrow Y$ is a homotopy equivalence could be constructed by induction over the skeleta, one cell at a time.

With Whitehead's theorem added to our toolbox, it would suffice to find a map $f : M \rightarrow S^3$ that induces isomorphisms $\pi_n(M) \rightarrow \pi_n(S^3)$ for all n . This project seems hopeless if we don't yet even know how to compute $\pi_n(M)$ for $n \geq 2$, 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**,

$$h : \pi_n(X, x_0) \rightarrow H_n(X; \mathbb{Z}) : [f] \mapsto f_*[S^n],$$

defined in terms of the fundamental class $[S^n] \in H_n(S^n; \mathbb{Z})$ and maps $f : (S^n, *) \rightarrow (X, x_0)$ 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 $\pi_0(X) = 0$, $h : \pi_1(X) \rightarrow H_1(X; \mathbb{Z})$ descends to an isomorphism on the abelianization of $\pi_1(X)$. For $n \geq 2$, both groups are abelian, so there is some hope of $h : \pi_n(X) \rightarrow H_n(X; \mathbb{Z})$ 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 55.10 (Hurewicz's theorem). *Suppose (X, x_0) is a pointed space that is $(n-1)$ -connected for some $n \geq 2$. Then $\tilde{H}_k(X) = 0$ for all $k \leq n-1$, and the Hurewicz map $h : \pi_n(X, x_0) \rightarrow H_n(X; \mathbb{Z})$ is an isomorphism.*

A rough idea for the proof of this theorem can be summarized as follows. By a construction based on induction over skeleta, one can first replace the space X with a CW-complex that is weakly homotopy equivalent to it, in which case its homotopy groups and singular homology groups are the same.⁸⁶ Moreover, the assumption that X is $(n-1)$ -connected enables one to construct the CW-complex so that its $(n-1)$ -skeleton contains only a single point, i.e. one starts with a 0-cell and then attaches n -cells and cells of higher dimension, skipping over dimensions $1, \dots, n-1$ entirely. For a CW-complex of this form, cellular homology tells us that $\tilde{H}_k(X) = 0$ for $k \leq n-1$, and it is also not difficult to compute both $H_n(X)$ and $\pi_n(X)$ and show that they match. The main ingredients needed for computing $\pi_n(X)$ in this situation are the homotopy exact sequence for pairs of spaces and the computation $\pi_n(S^n) \cong \mathbb{Z}$, mentioned in Example 54.8.

Before we get back to discussing 3-manifolds homotopy equivalent to S^3 , here are a couple of applications of the Hurewicz isomorphism.

COROLLARY 55.11. *If X is path-connected and has universal cover $\tilde{X} \rightarrow X$, then $\pi_2(X) \cong H_2(\tilde{X}; \mathbb{Z})$.*

PROOF. Since S^2 is simply connected, any map $S^2 \rightarrow X$ or homotopy of such maps can be lifted to \tilde{X} , implying $\pi_2(X) \cong \pi_2(\tilde{X})$. Since \tilde{X} is simply connected, the Hurewicz theorem then identifies $\pi_2(\tilde{X})$ with $H_2(\tilde{X}; \mathbb{Z})$. \square

COROLLARY 55.12. *If X is a simply connected CW-complex with $\tilde{H}_*(X; \mathbb{Z}) = 0$, 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 $n \geq 0$. It follows that the unique map $\epsilon : X \rightarrow \{*\}$ induces isomorphisms $\epsilon_* : \pi_n(X) \rightarrow \pi_n(\{*\}) = 0$ for all $n \geq 0$ and is therefore a weak homotopy equivalence. Whitehead's theorem then implies that it is also a homotopy equivalence. \square

You can now imagine at least part of a strategy to complete the proof of Theorem 55.1: instead of the map $\epsilon : X \rightarrow \{*\}$ in the proof of Corollary 55.12, one could take any map $f : M \rightarrow S^3$ of degree 1 and try to prove that $f_* : \pi_n(M) \rightarrow \pi_n(S^3)$ is an isomorphism for all $n \geq 0$. This idea can be carried out for all $n \leq 3$, 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 $n \geq 4$, however, we get stuck, among other reasons

⁸⁶One of these statements is less obvious than the other: it is practically a tautology that homotopy groups are invariants of weak homotopy type, but for singular homology, this invariance is a nontrivial theorem that depends on the choice to use *singular* rather than any other axiomatic homology theory. On the other hand, most important applications of the Hurewicz theorem concern spaces X that are already known to be homotopy equivalent to CW-complexes, in which case Whitehead's theorem eliminates the distinction between strong and weak homotopy equivalences, thus making the invariance of singular homology more obvious.

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 $\tilde{H}_n(S^3; \mathbb{Z}) \neq 0$.

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

$$h : \pi_n(X, A, x_0) \rightarrow H_n(X, A; \mathbb{Z}) : [f] \mapsto f_*[\mathbb{D}^n],$$

where $[f] \in \pi_n(X, A, x_0)$ is represented by a map $f : (\mathbb{D}^n, \partial\mathbb{D}^n, *) \rightarrow (X, A, x_0)$ 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 $n \geq 2$. Let us say that the pair (X, A) is **n -connected** if $\pi_k(X, A) = 0$ for all $k \leq n$. Since $\pi_2(X, A, x_0)$ is not always abelian, we cannot generally expect $h : \pi_2(X, A, x_0) \rightarrow H_2(X, A; \mathbb{Z})$ 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

$$\dots \rightarrow \pi_2(X) \rightarrow \pi_2(X, A) \rightarrow \pi_1(A) = 0,$$

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 55.10:

THEOREM 55.13. *Suppose (X, A) is an $(n - 1)$ -connected pair of spaces for some $n \geq 2$, 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 $k \leq n - 1$, and the relative Hurewicz map $h : \pi_n(X, A, x_0) \rightarrow H_n(X, A; \mathbb{Z})$ is an isomorphism.*

COROLLARY 55.14. *Suppose X and Y are two simply connected spaces that are both homotopy equivalent to CW-complexes, and $f : X \rightarrow Y$ is a map that induces isomorphisms $f_* : H_n(X; \mathbb{Z}) \rightarrow H_n(Y; \mathbb{Z})$ for every $n \geq 0$. 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

$$H_n(Y, X; \mathbb{Z}) = 0 \quad \text{for all } n \geq 0.$$

Similarly, the long exact sequence of relative homotopy groups includes a segment of the form

$$0 = \pi_1(Y) \rightarrow \pi_1(Y, X) \rightarrow \pi_0(X) = 0,$$

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 $n \geq 0$. In light of the exact sequence

$$0 = \pi_{n+1}(Y, X) \rightarrow \pi_n(X) \xrightarrow{f_*} \pi_n(Y) \rightarrow \pi_n(Y, X) = 0,$$

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 $f : X \rightarrow Y$ is an inclusion, we consider the **mapping cylinder** of f , defined as the space

$$M_f := ((X \times I) \amalg Y) / \sim \quad \text{where} \quad (x, 1) \sim f(x) \text{ for all } x \in X.$$

This space contains disjoint homeomorphic copies of X and Y , namely the images of the inclusion maps

$$i_X : X \hookrightarrow M_f : x \mapsto [(x, 0)], \quad \text{and} \quad i_Y : Y \hookrightarrow M_f : y \mapsto [y],$$

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

$$\begin{array}{ccc} Y & \xrightarrow{i_Y} & M_f \\ f \uparrow & \nearrow i_X & \\ X & & \end{array}$$

commutes *up to homotopy*, meaning that while the two maps $i_X : X \rightarrow M_f$ and $i_Y \circ f : X \rightarrow M_f$ have disjoint images and are thus obviously not equal, they are homotopic. As a consequence, $f : X \rightarrow Y$ is a homotopy equivalence if and only if $i_X : X \rightarrow M_f$ is a homotopy equivalence, and the inclusion map $i_X : X \hookrightarrow M_f$ can thus be used as a substitute for $f : X \rightarrow Y$ in arguments that depend only on homotopy type. In particular, if X and Y are simply connected and $f_* : H_*(X; \mathbb{Z}) \rightarrow H_*(Y; \mathbb{Z})$ is an isomorphism, then M_f is also simply connected and $(i_X)_* : H_*(X; \mathbb{Z}) \rightarrow H_*(M_f; \mathbb{Z})$ is also an isomorphism, so the argument of the previous paragraph makes i_X a homotopy equivalence, and so therefore is f . \square

CONCLUSION OF THE PROOF OF THEOREM 55.1. 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 $f : M \rightarrow S^3$ 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 $f : M \rightarrow S^3$ 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 36 implies $\deg(f) = 1$.

It is trivial that $f_* : H_0(M; \mathbb{Z}) \rightarrow H_0(S^3; \mathbb{Z})$ is an isomorphism, and so is $f_* : H_3(M; \mathbb{Z}) \rightarrow H_3(S^3; \mathbb{Z})$ due to the degree assumption. In all other dimensions, both homology groups vanish, so we conclude that $f_* : H_*(M; \mathbb{Z}) \rightarrow H_*(S^3; \mathbb{Z})$ is an isomorphism. Since M and S^3 are both simply connected, Corollary 55.14 now implies that f is a homotopy equivalence. \square

Third semester (Topologie III)

56. Motivation: Detecting exotic 7-spheres

Most of the first lecture is meant to provide motivation for the rest of the semester by sketching a proof of Milnor's famous theorem that there exist smooth manifolds homeomorphic but not diffeomorphic to S^7 . The proof makes use of several fundamental results that will be treated in depth in this course.

Details will appear here in the first week of the semester.

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.
- [Bae40] R. Baer, *Abelian groups that are direct summands of every containing abelian group*, Bull. Amer. Math. Soc. **46** (1940), 800–806.
- [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](https://arxiv.org/abs/1710.05296).
- [Bro94] K. S. Brown, *Cohomology of groups*, Graduate Texts in Mathematics, vol. 87, Springer-Verlag, New York, 1994. Corrected reprint of the 1982 original.
- [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.
- [Fre12] D. S. Freed, *Bordism: Old and New* (2012). lecture notes available at <https://web.ma.utexas.edu/users/dafr/bordism.pdf>.
- [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.
- [Jän05] K. Jänich, *Topologie*, 8th ed., Springer-Verlag, Berlin, 2005 (German).
- [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.]
- [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.
- [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.
- [Рох51] В. А. Рохлин, *А трех-димерсионал манифолд ис тхе боундари оф а фюр-димерсионал оне*, Доклады Акад. Наук СССР (Н.С.) **81** (1951), 355–357 (Руссиан).
- [Rud87] W. Rudin, *Real and complex analysis*, 3rd ed., McGraw-Hill Book Co., New York, 1987.
- [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.
- [tD08] T. tom Dieck, *Algebraic topology*, EMS Textbooks in Mathematics, European Mathematical Society (EMS), Zürich, 2008.
- [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.