

**Topology I and II, 2018–2019, HU Berlin
(revised April 2021)**

Chris Wendl

Contents

First semester (Topologie I)	7
1. Introduction and motivation	7
2. Metric spaces	11
3. Topological spaces	18
4. Products, sequential continuity and nets	23
5. Compactness	29
6. Tychonoff's theorem and the separation axioms	35
7. Connectedness and local compactness	42
8. Paths, homotopy and the fundamental group	50
9. Some properties of the fundamental group	55
10. Retractions and homotopy equivalence	61
11. The easy part of van Kampen's theorem	67
12. Normal subgroups, generators and relations	72
13. The Seifert-van Kampen theorem and π_1 of surfaces	78
14. Torus knots, and an introduction to covering spaces	89
15. The lifting theorem	94
16. Classification of covers	98
17. The universal cover and group actions	103
18. Manifolds	107
19. Surfaces and triangulations	114
20. Orientations and higher homotopy groups	122
21. Bordism groups and simplicial homology	130
22. Singular homology	138
23. Relative homology and long exact sequences	145
24. Homotopy invariance and excision	152
25. The homology of the spheres, and beyond	160
Second semester (Topologie II)	169
26. Categories and functors	169
27. Properties of singular homology	175
28. Excision, suspensions and reduced homology	182
29. Simplicial complexes in singular homology	190
30. The Eilenberg-Steenrod axioms, triples and good pairs	200
31. The Mayer-Vietoris sequence	207
32. Mapping tori and the degree of maps	213
33. Local mapping degree on manifolds	219
34. Degrees, triangulations, and coefficients	225
35. CW-complexes	230
36. Invariance of cellular homology, part 1	237

37. Invariance of cellular homology, part 2	242
38. Direct limits and infinite-dimensional cell complexes	247
39. The Euler characteristic	254
40. The Lefschetz fixed point theorem	260
41. Coefficients and the Tor functor	267
42. Universal coefficient theorem and Künneth formula	274
43. The cross product on homology	281
44. Čech homology and inverse limits	290
45. Singular cohomology	302
46. Axioms for cohomology	310
47. Universal coefficients and the Ext functor	319
48. Products on cohomology, part 1	326
49. Products on cohomology, part 2	332
50. Relative cup and cap products, orientations	339
51. The orientation bundle	345
52. Existence of the fundamental class	350
53. Poincaré duality	355
54. The intersection product	366
55. Higher homotopy groups	377
56. The theorems of Hurewicz and Whitehead	386
Bibliography	391

These notes began in the 2018 Summer Semester as accompaniment for the course *Topologie I* at Humboldt University, Berlin; see

<https://www.mathematik.hu-berlin.de/~wendl/Sommer2018/Topologie1/>

They were then continued as an accompaniment to the followup course in the 2018–19 Winter Semester, *Topologie II*; see

<https://www.mathematik.hu-berlin.de/~wendl/Winter2018/Topologie2/>

They have since been revised a bit further following comments from students in the class, including the incorporation of some assigned homework problems into the notes as exercises within the relevant lectures.

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. Notable examples include:

- The exercises at the end of Lecture 29 concerning the relationship between singular homology and bordisms of triangulated manifolds;
- The proof in Lecture 32 that the relative Mayer-Vietoris sequence is exact;
- The discussion of the *ordered* simplicial complex in Lecture 44;
- The discussion of Alexander-Spanier cohomology in Lecture 46;
- The proof of the Thom isomorphism theorem in Lecture 54.

I would advise anyone teaching a course based on these notes to treat those topics (and possibly several others) as optional unless they have substantially more total teaching hours at their disposal than I did. The reader may also notice that the bijection between class lectures and chapters in these notes occasionally has peculiar consequences, e.g. for reasons of time, Lecture 50 ended up straddling the boundary between two more-or-less unrelated topics that would surely appear in different chapters if this were a textbook.

While the class from which these notes emerged was taught in English, I made an effort to also include the German translations (*geschrieben in dieser Schriftart*) of important terms wherever they were introduced. The reader may notice that this effort is more obvious in the first semester than the second, 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).

Disclaimer: 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 and, especially, Laurenz Upmeyer zu Belzen. (Apologies if I forgot anyone!)

For more detailed treatments of the topics in these notes, I mainly recommend the books by Jänich [Jän05] (or its English translation), Hatcher [Hat02] and Bredon [Bre93].


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) . (Note that as a consequence, all open balls in \mathbb{R}^n are homeomorphic to each other—homeomorphism is an equivalence relation!) 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 a generalization of that theorem later, but it does not say that these two definitions are equivalent for every topological space—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*) 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 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 a finite collection $\mathcal{U}_\alpha^1, \dots, \mathcal{U}_\alpha^{N_\alpha} \in \mathcal{B}$.

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} .

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?)

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 .

If \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 the open set $\mathcal{U} \subset \mathbb{R}^{n+1}$ such that $\mathcal{V} = \mathcal{U} \cap S^n$. One can construct it 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 $I = \{1, \dots, N\}$, $X_1 \sqcup \dots \sqcup X_N$. The topology on this space is called the **disjoint union topology**.

EXERCISE 3.16. Show that the disjoint union topology \mathcal{T} on $X = \bigsqcup_{\alpha \in I} X_\alpha$ 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 $\bigsqcup_{\alpha \in I} X_\alpha$ 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 $X = \bigsqcup_{\alpha \in I} X_\alpha$ 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 . 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 $X_1 \sqcup X_2$ is defined so that its subsets $X_1 \subset X_1 \sqcup X_2$ and $X_2 \subset X_1 \sqcup X_2$ 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

$$\prod_{\alpha \in \mathbb{R}} \{0\} := \{(\alpha, 0) \in \mathbb{R} \times \{0\}\} \rightarrow \mathbb{R},$$

$$(\alpha, 0) \mapsto \alpha.$$

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

4. Products, sequential continuity and nets

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

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

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

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

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

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

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

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

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

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

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

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

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 to $\{x_\alpha\}_{\alpha \in I} \in \prod_{\alpha \in I} X_\alpha$ with 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 exactly 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 $\prod_\alpha X_\alpha$, 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 $\prod_\alpha X_\alpha$ to converge in the box topology? In particular, consider the case where all the X_α are a fixed space X and $\prod_\alpha X$ is identified with the space of all functions $X^I = \{f : I \rightarrow X\}$; 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 basis $\mathcal{U}_1, \mathcal{U}_2, \mathcal{U}_3, \dots$ for x .

It will make our lives slightly easier if the neighborhood basis 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 $n \geq N$,

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

\square

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.* \square

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 ,

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

\square

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} . \square

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 $A \supset A_0 \in I$,

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

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

5. Compactness

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 closure 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.

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) Show 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

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

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

$$\mathcal{A} < \mathcal{B} \iff \mathcal{A} \subset \mathcal{B},$$

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

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

□

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

EXAMPLE 5.24. Though it is not so easy to see this, the space $[0, 1]^{\mathbb{R}}$ of (not necessarily continuous) functions $\mathbb{R} \rightarrow [0, 1]$ 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 $\{\mathcal{U}_\alpha\}_{\alpha \in I}$. □

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 $k_n \geq N$. □

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

6. Tychonoff's theorem and the separation axioms

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 [Kel50], if one assumes that the usual axioms of set theory (not including choice) hold and that Tychonoff's theorem also holds, then the axiom of choice follows, thus the two are actually equivalent.

Speaking only for myself, I had a Ph.D. in mathematics already for several years before I ever started to find the axiom of choice remotely worrying, so if you've never worried about it before, I don't encourage you to start worrying now. If you take functional analysis, you'll encounter a few more fundamental results such as the Hahn-Banach theorem and the existence of bases on Banach spaces that depend on the axiom of choice in similar ways, and you'll need to get used to this unless you prefer to avoid most of modern analysis. As far as this particular course on topology 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]. As with the standard results in functional analysis I just mentioned, it uses the axiom of choice 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 $g_\infty(\alpha)$ defined as $g(\alpha)$ for any $(J, g) \in \mathcal{A}$ such that $\alpha \in J$. The total ordering condition guarantees that (J_∞, g_∞) is independent of choices, but it is not immediately clear whether it is an element of \mathcal{P} , i.e. whether g_∞ is a cluster point of $\{f_\beta|_{J_\infty}\}_{\beta \in K}$. To see this, suppose $\mathcal{U} \subset \prod_{\alpha \in J_\infty} X_\alpha$ is a neighborhood of g_∞ , and recall that by the definition of the product topology, this means

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

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

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

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

We can now apply Zorn's lemma and conclude that \mathcal{P} has a maximal element $(J_M, g_M) \in \mathcal{P}$. We claim $J_M = I$, which means g_M is a cluster point of the original net $\{f_\beta\}_{\beta \in K}$ in $\prod_{\alpha \in I} X_\alpha$. Note that since $g_M \in \prod_{\alpha \in J_M} X_\alpha$ is a cluster point of $\{f_\beta|_{J_M}\}_{\beta \in K}$, Lemma 5.20 provides a subnet $\{f_{\phi(\gamma)}\}_{\gamma \in L}$ of $\{f_\beta\}_{\beta \in K}$ in $\prod_{\alpha \in I} X_\alpha$ whose restriction to J_M converges to g_M . But if $J_M \neq I$, then choosing an element $\alpha_0 \in I \setminus J_M$, 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.

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 $x, y \in X$, there exists an open subset of X that contains either only x or only y .

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 $X \supset \mathcal{V}_1 \supset \mathcal{V}_2 \supset \dots \ni y$. 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 $X \times X$.

- (c) 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.
- (d) Every metrizable space satisfies the axiom T_4 (in particular it is *normal*).
Hint: Given disjoint closed sets $A, A' \subset X$, each $x \in A$ admits a radius $\epsilon_x > 0$ such that the ball $B_{\epsilon_x}(x)$ is disjoint from A' , and similarly for points in A' (why?). The unions of all these balls won't quite produce the disjoint neighborhoods you want, but try cutting their radii in half.

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

- (a) 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.
- (b) Show that if X is also regular and $A \subset X$ is a closed subset, then X/A is Hausdorff.
- (c) Consider $X = \mathbb{R}$ with the non-closed subset $A = (0, 1]$. Which of the separation axioms T_0, \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 7.11 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

PROPOSITION 7.7. *Every connected component $A \subset X$ of a space X is a closed subset.*

PROOF. There is nothing to prove if $A = X$, so assume this is not the case, and thus A is connected but $A \cup \{x\}$ is not connected for every $x \in X \setminus A$. Then there exists a nonconstant continuous function $f : A \cup \{x\} \rightarrow \{0, 1\}$, but it must be constant on A since the latter is connected, so without loss of generality, we may suppose $f|_A = 0$ and $f(x) = 1$. This means $f^{-1}(0) = A$, and since $\{0\}$ is open and closed in $\{0, 1\}$, it follows that A is open and closed as a subset of $A \cup \{x\}$, or equivalently, both A and $\{x\}$ are open subsets of $A \cup \{x\}$.

Since $A \cup \{x\}$ carries the subspace topology as a subset of X , this means that there exist open subsets $\mathcal{U}_x, \mathcal{V}_x \subset X$ whose intersections with $A \cup \{x\}$ are A and $\{x\}$ respectively. Concretely, this means \mathcal{U}_x contains A but not x , while \mathcal{V}_x contains x but is disjoint from A . Since such a set exists for all $x \in X \setminus A$, we can take their union,

$$\mathcal{V} = \bigcup_{x \in X \setminus A} \mathcal{V}_x,$$

which is manifestly an open subset of X and disjoint from A , but also contains $X \setminus A$, so in fact $\mathcal{V} = X \setminus A$. This proves that $X \setminus A$ is open, or equivalently, 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 7.8. *If X is a space with only finitely many connected components, then each of them is both closed and open.* \square

EXERCISE 7.9. Suppose the connected components of X are labeled as a collection of subsets $\{X_\alpha\}_{\alpha \in I}$, endow each with the subspace topology, and consider the continuous inclusion maps $i_\alpha : X_\alpha \hookrightarrow X$. These can be assembled into a continuous map

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

whose restriction to each $X_\beta \subset \coprod_{\alpha \in I} X_\alpha$ is i_β . Show that this map is a homeomorphism if and only if every X_α is an open subset of X . (In particular, Corollary 7.8 implies that this is always true if I is finite, and we will see in Prop. 7.16 below that it is also true if X is locally connected.)

EXAMPLE 7.10. 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 $x \in \mathbb{Q}$. But the disjoint union of these one-point subspaces produces \mathbb{Q} with the discrete topology, not the standard topology of \mathbb{Q} , so (7.1) in this case fails to be a homeomorphism.

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

THEOREM 7.11. *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 7.12. 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 7.11) also connected. But neither X_- nor X_+ is closed, so by Prop. 7.7, neither of these can be a connected component. The maximal connected subset containing X_- , for instance, must be a closed set containing X_- and therefore contains the closure $\overline{X_-}$, which includes points in X_0 . Since X_0 is path-connected, it follows that the connected component containing X_- also contains all of X_0 . But the same argument applies equally well to X_+ , and these two observations together imply that all three path-components are in the same connected component, i.e. X is connected.

The space in Example 7.12 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 7.13. 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 7.14. 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 7.11. 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 7.12 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 7.15. 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 7.16. *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 7.17. *Every space that is connected and locally path-connected is also path-connected.*

PROOF. If X is locally path-connected, then by Prop. 7.16 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 7.18.

- (a) Prove that if X and Y are both connected, then so is $X \times Y$.³
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 $X \times Y \rightarrow \{0, 1\}$.
- (b) Show that for any collection of path-connected spaces $\{X_\alpha\}_{\alpha \in I}$, the space $\prod_{\alpha \in I} X_\alpha$ is path-connected in the usual product topology.
Hint: You might find Exercise 4.5 helpful.
- (c) Consider $\mathbb{R}^{\mathbb{N}}$ with the “box topology” which we discussed in Exercise 4.6. Show that the set of all elements $f \in \mathbb{R}^{\mathbb{N}}$ represented as functions $f : \mathbb{N} \rightarrow \mathbb{R}$ 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).

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

DEFINITION 7.19. 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 7.20. 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 7.21. 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 7.22. 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 7.19 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

³The analogous statement about infinite products is also true, but it takes more work to prove it.

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

- (i) X is locally compact (in the sense of Definition 7.19);
- (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. 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 $y \in \partial\bar{\mathcal{U}}$ there exists a pair of open neighborhoods

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

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

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

Now the set

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

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

EXERCISE 7.24. Extend Theorem 7.23 to the following statement: 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 7.25. 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.

- (a) Verify that \mathcal{T}^* is a topology and that X^* is always compact.

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

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

- (f) Show that for any $x \in \mathbb{Q}$, every neighborhood of x in \mathbb{Q}^* intersects every neighborhood of ∞ , so in particular, \mathbb{Q}^* is not Hausdorff.
Advice: Do not try to argue in terms of sequences with non-unique limits (cf. part (g) below), and do not try to describe precisely what arbitrary compact subsets of \mathbb{Q} can look like (the answer is not nice). One useful thing you can say about arbitrary compact subsets of \mathbb{Q} is that they can never contain the intersection of \mathbb{Q} with any open interval. (Why not?)
- (g) Show that every convergent sequence in \mathbb{Q}^* has a unique limit. (Since \mathbb{Q}^* is not Hausdorff, this implies via Proposition 6.13 that \mathbb{Q}^* is not first countable—in particular, ∞ does not have a countable neighborhood base.)
- (h) Find a point in \mathbb{Q}^* with a neighborhood that does not contain any compact neighborhood.

EXERCISE 7.26. 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 7.27. 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 7.26 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 $X := I = [0, 1]$. 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 $C(Y, Z)$.⁴

- (c) Deduce from part (b) a new proof of the following result from Exercise 7.26(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 7.26(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 7.24 is useful here.

⁴Since $C(X \times Y, Z)$ 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.

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.

- (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.⁵ 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 :

⁵The 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.

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


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 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 $\gamma(1) = y$.⁶ 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

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

defined by $\gamma^{-1}(t) := \gamma(1 - t)$. The reason for this terminology 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 $x \rightsquigarrow^{\alpha} y$ and $y \rightsquigarrow^{\beta} z$, we can define the product path $x \rightsquigarrow^{\alpha \cdot \beta} z$ 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 $x \rightsquigarrow^{e_x} x$, 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 $x \rightsquigarrow^{\alpha} y$, $y \rightsquigarrow^{\beta} z$ and $z \rightsquigarrow^{\gamma} a$, 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).

⁶This seems a good moment to emphasize that all maps in this course are assumed continuous unless otherwise noted.

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 $x \rightsquigarrow y$, the following relations hold:

$$(i) \quad e_x \cdot \gamma \underset{h+}{\sim} \gamma$$

- (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} . \square

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:

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

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

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

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

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 $\pi_1(X)$. We will also prove that for any two homeomorphic spaces X and Y , $\pi_1(X)$ and $\pi_1(Y)$ are isomorphic groups. 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 $\pi_1(S^1)$ is 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

Since the fundamental group $\pi_1(X, p)$ depends on a pair (X, p) consisting of a space X and a point $p \in X$, it will be useful to formalize a few notions concerning such pairs.

DEFINITION 9.1. 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 $s \in I$.

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

THEOREM 9.2. *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).⁷ 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 9.3. *If X is path-connected, then $\pi_1(X, p)$ up to isomorphism is independent of the choice of base point $p \in X$.*

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

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

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

⁷Note 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 9.4. *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*

$$(9.1) \quad \begin{array}{ccc} X & \xrightarrow{\pi} & X/A \\ & \searrow f & \swarrow g \\ & & Y \end{array}$$

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

REMARK 9.5. Commutative diagrams such as (9.1) 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. This was an especially simple example, 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 $\gamma \circ g = g' \circ \beta$.




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 9.6.

Lemma 9.4 gives a canonical bijection between the set of all paths $p \rightsquigarrow p$ 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 $p \rightsquigarrow p$ 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

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

where \sim_{h_+} 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 9.6. *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 9.4, 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 9.4 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 9.7. 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 9.8. 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 9.6, 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 9.6.

EXAMPLES 9.9. 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 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 9.10. *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 9.6. \square

COROLLARY 9.11. *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

We must next discuss in what sense π_1 is a *topological invariant*.

THEOREM 9.12. *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 $(Y, q) \xrightarrow{g} (Z, r)$, $(g \circ f)_* = g_* \circ f_*$.*




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 $I^2 \cong \mathbb{D}^2$.

- (ii) The map associated to the identity map $(X, p) \xrightarrow{\text{Id}} (X, p)$ is the identity homomorphism $\pi_1(X, p) \xrightarrow{\mathbb{1}} \pi_1(X, p)$.
- (iii) Each homomorphism f_* depends only on the pointed homotopy class of f .

PROOF. It is clear that up to homotopy (with fixed end points), the loop $q \xrightarrow{f \circ \gamma} q$ in Y depends only on the pointed homotopy class of $f : (X, p) \rightarrow (Y, q)$ and the homotopy class with fixed end points of the loop $p \xrightarrow{\gamma} p$. This shows that $f_* : \pi_1(X, p) \rightarrow \pi_1(Y, q)$ is a well-defined map. It is similarly easy to check that it is a homomorphism and satisfies the first two stated properties: e.g. for any two loops $p \xrightarrow{\alpha, \beta} p$, 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 9.13. If X and Y are homeomorphic path-connected spaces then $\pi_1(X)$ and $\pi_1(Y)$ are isomorphic.

PROOF. Given a homeomorphism $f : X \rightarrow Y$, choose any base point $p \in X$ and set $q := f(p) \in Y$ so that f is now a pointed map $(X, p) \rightarrow (Y, q)$ and f^{-1} is a pointed map $(Y, q) \rightarrow (X, p)$. Using Theorem 9.12, the commutative diagram of continuous maps

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

then gives rise to a similar commutative diagram of group homomorphisms

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

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. \square

REMARK 9.14. The fancy way to summarize Theorem 9.12 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*.

Let us finally work out a few concrete examples.

EXAMPLE 9.15. 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 $0 \xrightarrow{\sim} 0$ 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 9.16. Since every open ball $B_r(x)$ in \mathbb{R}^n is homeomorphic to \mathbb{R}^n itself, Corollary 9.13 implies that $\pi_1(B_r(x))$ also vanishes, i.e. $B_r(x)$ is simply connected.

EXAMPLE 9.17. 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 $x \xrightarrow{\sim} y$ in S^1 and a point $\tilde{x} \in f^{-1}(x)$, there exists a unique path $\tilde{x} \xrightarrow{\sim} \tilde{y}$ 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 $x \xrightarrow{\sim} y$ (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 $\tilde{x} \xrightarrow{\sim} \tilde{y}$ 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 $1 \xrightarrow{\sim} 1$, there is a unique lift to a path $0 \xrightarrow{\sim} \tilde{\gamma}(1)$ 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 9.18.

- (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).

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 $[p_\alpha] \in \bigvee_{\alpha \in J} X_\alpha$.

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. But while the intersection of X_{β} and X_{γ} in $\coprod_{\alpha} X_{\alpha}$ for $\beta \neq \gamma$ is always empty, in $\bigvee_{\alpha} X_{\alpha}$ they intersect at the base point, and only there. The next example should be understood in this context.

EXAMPLE 10.5. 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 10.6. Convince yourself that the map $f : X \vee Y \rightarrow X$ in Example 10.5 is continuous.

EXAMPLE 10.7. 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 10.5, 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 10.8. 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 10.9. *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 9.12 produces a corresponding commutative diagram of homomorphisms:

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

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 10.8, so Theorem 10.9 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 9.15 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 10.23 below.) But there is no surjective homomorphism from the trivial group to \mathbb{Z} , so this is a contradiction. \square

DEFINITION 10.10. 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 10.11. 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 10.12. This is actually a non-example: while the maps $f : S^1 \vee S^1 \rightarrow S^1$ in Example 10.7 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 10.13. 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 10.9 since $\pi_1(S^1) \cong \mathbb{Z}$ and $\pi_1(\mathbb{R}^2) = 0$.)

EXAMPLE 10.14. 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 10.15. 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 10.16. Verify that homotopy equivalence defines an equivalence relation.

REMARK 10.17. 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 10.18. A homeomorphism $f : X \rightarrow Y$ is obviously also a homotopy equivalence, with homotopy inverse f^{-1} .

EXAMPLE 10.19. 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 10.20. We say that a space X is **contractible** (*zusammenziehbar* or *kontrahierbar*) if it is homotopy equivalent to a one-point space.

REMARK 10.21. The above definitions imply immediately that any space admitting a deformation retraction to a one-point subset (as in Example 10.11) 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 10.10. 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 10.22. *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 10.23. *For every contractible space X , $\pi_1(X) = 0$.* □

PROOF OF THEOREM 10.22. Here is a preliminary remark: if you’re only half paying attention, then you might reasonably think this theorem follows immediately from Theorem 9.12. 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 9.12, 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, we proceed in two steps. 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 $p \xrightarrow{\gamma} r$ by

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

and by Theorem 9.2, 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 $p \xrightarrow{\alpha} p$, 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 $p \xrightarrow{\gamma_s} \gamma(s)$ 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.

If we extend the sequence (10.2) one step further to include $f_* : \pi_1(X, r) \rightarrow \pi_1(Y, f(r))$, a repeat of the same argument shows that

$$f_* \circ g_* : \pi_1(Y, q) \rightarrow \pi_1(Y, f(r))$$

is also an isomorphism, and this concludes step 1 of the proof.

In step 2, we examine the sequence of maps:

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

The conclusion of step 1 was that both of the maps one can form by composing two maps in this sequence are isomorphisms, and since both of these contain the same map g_* (though in different roles), it follows that g_* is injective and surjective, thus an isomorphism. The map f_* in the statement of the theorem can then be written as a composition of two isomorphisms $g_*^{-1} \circ (g_* \circ f_*)$. \square

EXAMPLE 10.24. 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 10.11. 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 10.25. Viewing S^1 as the unit circle in \mathbb{C} , associate to each $z \in \mathbb{C}$ the loop $\gamma_z : S^1 \hookrightarrow \mathbb{C} \setminus \{z\} : e^{i\theta} \mapsto z + e^{i\theta}$. 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 10.13, 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 10.26. 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 9.17, 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 $\gamma(t) = z_0 + re^{2\pi i t}$ 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 inclusions⁸

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

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

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

EXAMPLE 11.7. In Example 1.2 we gave a description of $\mathbb{R}P^2$ 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.

$$S^1 \hookrightarrow \mathbb{M} : e^{i\theta} \mapsto (e^{2i\theta}, \cos(\theta), \sin(\theta)).$$

⁸The 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**.

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 $\pi \circ u : \mathbb{D}^2 \rightarrow \mathbb{RP}^2$ is then an extension over the disk of our loop representing $[\gamma]^2$, which proves via Theorem 9.6 that $[\gamma]^2$ is trivial. This proves that $\pi_1(\mathbb{RP}^2)$ is either the trivial group or is isomorphic to \mathbb{Z}_2 ; we will see that it is the latter when we prove that the generator $[\gamma]$ is nontrivial.

Here is another pair of general constructions that produce many more examples.

DEFINITION 11.8. Given a space X , the **cone** (*Kegel*) of X is the space

$$CX := (X \times I)/(X \times \{1\}).$$

The single point in CX represented by $(x, 1)$ for every $x \in X$ is sometimes called the “summit” or “node” of the cone.

EXERCISE 11.9. Show that CS^{n-1} is homeomorphic to \mathbb{D}^n .

LEMMA 11.10. *For every space X , the cone CX is contractible.*

PROOF. There is an obvious deformation retraction of $X \times I$ to $X \times \{1\}$ defined by pushing every $(x, t) \in X \times I$ upward in the t -coordinate. Writing down this same deformation retraction on the quotient $(X \times I)/(X \times \{1\})$, the result is that everything gets pushed to a single point, the summit of the cone. \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 .⁹ The union of all these open subintervals for all $\alpha \in J$ 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

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

⁹Remember 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 ”.

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.¹⁰ 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

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

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.

¹⁰This 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.

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*

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

sending each reduced word $[\gamma_1] \dots [\gamma_N] \in \ast_{\alpha \in J} \pi_1(A_\alpha, p)$ with $[\gamma_i] \in \pi_1(A_{\alpha_i}, p)$ to the concatenation $[\gamma_1 \cdot \dots \cdot \gamma_N] \in \pi_1(X, p)$, and Φ is surjective. \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. The following result is easy to prove directly from the definitions.

PROPOSITION 12.1. *Assume $\{G_\alpha\}_{\alpha \in J}$ is a collection of groups. Then:*

- (1) *For each $\alpha \in J$, the free product $\ast_{\beta \in J} G_\beta$ contains a distinguished subgroup isomorphic to G_α : it consists of the empty word plus all reduced words of exactly one letter which is in G_α .*
- (2) *If we regard each G_α as a subgroup of $\ast_{\gamma \in J} G_\gamma$ as described above, then for every $\alpha, \beta \in J$ with $\alpha \neq \beta$, $G_\alpha \cap G_\beta = \{e\}$, and any two nontrivial elements $g \in G_\alpha$ and $h \in G_\beta$ satisfy $gh \neq hg$ in $\ast_{\gamma \in J} G_\gamma$.*
- (3) *For any group H with a collection of homomorphisms $\{\Phi_\alpha : G_\alpha \rightarrow H\}_{\alpha \in J}$, there exists a unique homomorphism*

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

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

The third item in this list deserves brief comment: the homomorphism $\Phi : *_{\alpha \in J} G_\alpha \rightarrow H$ 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 descends to a well-defined *injective* homomorphism


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

This 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,



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 $p \xrightarrow{\gamma} p$ 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 $\ast_{\gamma \in J} \pi_1(A_\gamma, p)$, 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 $\ast_{\gamma \in J} \pi_1(A_\gamma, p)$ 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 we can do the usual trick 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 7.25 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 $(\mathbb{R}^3 \setminus \overline{B_R(0)}) \cup \{\infty\}$ where $\overline{B_R(0)} \subset \mathbb{R}^3$ is any closed ball large enough to contain K . Since p might not be contained in B_0 but $\mathbb{R}^3 \setminus K$ 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 homeomorphic 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

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

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. *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 $F_S / \ker \Phi \rightarrow G$. \square

DEFINITION 12.11. 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 12.12. 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 12.10 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 12.13. Given a group G , a choice of generators S and relations R such that $G \cong \{S \mid R\}$ is called a **presentation** of G . We say that G is **finitely presented** if it admits a presentation such that S and R are both finite sets.

EXAMPLE 12.14. 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 12.15. 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 12.16. 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 12.17. 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 12.19 below proves that this group is not abelian. By contrast, 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 12.18 (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 12.19. Let us prove that the finitely-presented group $G = \{x, y \mid x^2 = y^3\}$ mentioned in Example 12.17 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 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 12.20. 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.¹¹
- (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. The Seifert-van Kampen theorem and π_1 of surfaces

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*

$$\ker \Phi = \langle S \rangle_N,$$

i.e. 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

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

¹¹Note 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 $a, b \in S$.

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.

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 $p \xrightarrow{\gamma_i \cdot \gamma_{i+1}} p$. (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 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 $(s, t) \in I^2$,¹²

$$[s - 2\epsilon, s + 2\epsilon] \times [t - 2\epsilon, t + 2\epsilon] \subset H^{-1}(A_\alpha) \quad \text{for some } \alpha \in J.$$

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 $\alpha \in J$, 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

¹²I 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 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.

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 $\gamma_1 \cdot \dots \cdot \gamma_N$ 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 path¹³

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

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

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

¹³This 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].

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

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.

- 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 .
- Call a pointed space (X, x) *nice*¹⁴ 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.
- Here is an example of a space that is not “nice” in the sense of part (b): the so-called *Hawaiian earring* can be defined as the subset of \mathbb{R}^2 consisting of the union for all $n \in \mathbb{N}$ of the circles of radius $1/n$ centered at $(1/n, 0)$. As usual, we assign to this set 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.
- It is tempting to liken the Hawaiian earring 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 $\bigsqcup_{n=1}^{\infty} S^1$. Since both X and



¹⁴Not a standardized term, I made it up.


the Hawaiian earring are unions of infinite collections of circles that all intersect each other at one point, it is not hard to imagine a bijection between them. Show however that such a bijection can never be a homeomorphism; in particular, unlike the Hawaiian earring, X is “nice” for any choice of base point.

Hint: Pay attention to how the topology of X is defined—it is a quotient of a disjoint union.

We will discuss two more applications of the Seifert-van Kampen theorem: one to the study of surfaces, and the other (in the next lecture) to knots. Let’s talk about 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 the compact convex region bounded by some polygon with N edges, each of which has been labeled 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 acts on its boundary by identifying all vertices to a single point and identifying any pair of edges labeled by the same letter via a homeomorphism that matches the directions of the arrows. 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 13.3. 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 13.4. 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 $\mathbb{R}P^2$ that we saw in the first lecture, see Example 1.2.

THEOREM 13.5. 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,

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

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 13.6. Applying the theorem to the torus in Example 13.3 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 9.18(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 13.7. For the picture of \mathbb{RP}^2 in Example 13.4, we obtain

$$\pi_1(\mathbb{RP}^2) \cong \{a \mid a^2 = e\} \cong \mathbb{Z}_2.$$

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 13.8. 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 13.3, i.e. it is the torus \mathbb{T}^2 . The diagram for Σ_2 is shown at the bottom of Figure 5. The projective plane \mathbb{RP}^2 is not an “orientable” surface, so it is not Σ_g for any g , though it is sometimes called a “non-orientable surface of genus 1”. This terminology will make more sense when we later discuss the classification of surfaces.

In order to understand what Σ_g has to do with pictures we’ve seen before, we consider an operation on surfaces called the *connected sum*. It can be defined on any pair of surfaces Σ and Σ' , or more generally, on any pair of n -dimensional topological manifolds, though for now we will consider only the case $n = 2$. Since I haven’t yet actually given you precise definitions of the terms “surface” and “topological manifold,” for now you should just assume Σ and Σ' come from the list of specific examples $\Sigma_0 = S^2$, $\Sigma_1 = \mathbb{T}^2$, $\Sigma_2, \Sigma_3, \dots$ 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 13.9 (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 give the rough idea. The main thing you need to believe is that “up to orientation” (I’ll come




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.

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 \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




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.

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 13.10. *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 13.9 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 13.11 (of Theorem 13.5). *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 12.20, the relation in Corollary 13.11 can be conveniently abbreviated as

$$\prod_{i=1}^g [a_i, b_i] = e.$$

EXERCISE 13.12. Show that the abelianization (cf. Exercise 12.20) 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 the latter 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 13.12 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 13.13. *For two nonnegative integers $g \neq h$, Σ_g and Σ_h are not homeomorphic.* \square

EXERCISE 13.14. 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 13.15. 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 $\pi_1(\Sigma_{g,1})$.¹⁵

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.

¹⁵Terminology: 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 13.15.

- (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.

14. Torus knots, and an introduction to covering spaces

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 $\mathbb{T}^2 = S^1 \times S^1$ and q times around the other. Note that p and q must be relatively prime in order for this loop in \mathbb{T}^2 to be embedded.

EXAMPLE 14.1. $K_{2,3}$ is the trefoil knot (Figure 8, left).

EXAMPLE 14.2. $K_{1,0}$ is the unknot (Figure 8, right).

FIGURE 8. The trefoil knot $K_{2,3}$ and unknot $K_{1,0}$.

Recall that for any knot $K \subset S^3$, its **knot group** 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 $\pi_1(S^3 \setminus K)$. 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 12.18, 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 14.3. $\pi_1(S^3 \setminus K_{p,q}) \cong \{a, b \mid a^p = b^q\}$. □

EXAMPLE 14.4. 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 14.5. 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 12.19 that this group is not abelian, in contrast to Example 14.4, hence $\pi_1(S^3 \setminus K_{2,3})$ and $\pi_1(S^3 \setminus K_{1,0})$ are not isomorphic.

COROLLARY 14.6. *The knot complements $\mathbb{R}^3 \setminus K_{1,0}$ and $\mathbb{R}^3 \setminus K_{2,3}$ are not homeomorphic.* □

Before moving on 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 14.7. *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 easily¹⁶ from the Seifert-van Kampen theorem using the decomposition $X = A \cup B$ where $A = \mathring{\mathbb{D}}^2$ and B is an open neighborhood of X_0 obtained by adding a small annulus near the boundary of $\partial\mathbb{D}^2$. Since the annulus admits

¹⁶I 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.

a deformation retraction to $\partial\mathbb{D}^2$, we have $B \underset{h.e.}{\cong} X_0$, while $A \cap B \underset{h.e.}{\cong} S^1$ and A is contractible. According to Corollary 12.18, $\pi_1(X, p)$ 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

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 14.8. 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 14.9. The map $f : \mathbb{R} \rightarrow S^1 : \theta \mapsto e^{i\theta}$ is a covering map of S^1 .

EXAMPLE 14.10. 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 14.11. 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 14.12. For any space X , the identity map $X \rightarrow X$ is trivially a covering map.

EXAMPLE 14.13. 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 14.14. For each $n \in \mathbb{N}$, the quotient projection $S^n \rightarrow \mathbb{R}\mathbb{P}^n = S^n/\sim$ is a covering map.

THEOREM 14.15. *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 14.16. The cover $S^1 \rightarrow S^1 : z \mapsto z^k$ from Example 14.10 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 14.9 has infinite degree.

REMARK 14.17. 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 14.8, but notice that if X is connected, then it follows immediately from Theorem 14.15. In practice, it is only sensible to consider covers of connected spaces, and we shall always assume connectedness.

Note that in Definition 14.8, one should explicitly require the sets $\mathcal{V}_\alpha \subset f^{-1}(\mathcal{U})$ to be open. This is important, as part of the point of that definition is that X can be covered by open neighborhoods \mathcal{U} whose preimages are homeomorphic to *disjoint unions* of copies of \mathcal{U} , i.e.

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

This is true specifically because each of the sets \mathcal{V}_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 .

15. The lifting theorem

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 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 15.1 (lifting theorem). *Assume X, Y, A are all path-connected spaces, A is also locally path-connected, $p : (Y, y_0) \rightarrow (X, x_0)$ is a base-point preserving covering map and $f : (A, a_0) \rightarrow (X, x_0)$ is a base-point preserving map. Then a (continuous) lift $\tilde{f} : A \rightarrow Y$ of f exists if and only if*

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

and in that case there is exactly one lift satisfying $\tilde{f}(a_0) = y_0$.

Let us discuss some applications before we get to the proof.

COROLLARY 15.2. *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 15.3. *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 15.4. *If X is simply connected, then every path-connected covering space of X is also simply connected.* \square

EXAMPLE 15.5. Corollary 15.4 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-connected¹⁷ 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

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

Theorem 15.1 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 10.26, 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 15.6 (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 15.7 (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 15.1. 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.

¹⁷Since $\mathcal{U} \subset \mathbb{C}^*$ 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 15.6 to a unique path $\widetilde{f \circ \alpha}$ in Y that starts at y_0 . If a lift \widetilde{f} exists, it clearly must satisfy

$$\widetilde{f}(a) = \widetilde{f \circ \alpha}(1).$$

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 $\widetilde{f} : A \rightarrow Y$. 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 $y_0 \xrightarrow{\gamma} y_0$ in Y such that

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

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 15.7 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 15.6 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 $\widetilde{f} : A \rightarrow Y$ as defined by the above procedure is continuous. Given $a \in A$ with $x = f(a) \in X$ and $y = \widetilde{f}(a) \in Y$, 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 $\widetilde{f}(\mathcal{O}) \subset \mathcal{V}$. 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

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

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 $\widetilde{f}(a') \in \mathcal{V}$. \square

EXAMPLE 15.8. If the local path-connectedness assumption on A is dropped, then the proof above gives a procedure for defining a unique lift $\widetilde{f} : A \rightarrow Y$, but it may fail to be continuous. A concrete example is depicted in [Hat02, p. 79], Exercise 7. The idea is to take 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 resulting space A is simply connected but not locally path-connected: to see the latter, note that small neighborhoods of $(0, 1)$ in A are never path-connected because there is no continuous path starting on the infinitely oscillating sine curve and ending at $(0, 1)$. The space is nonetheless path-connected and even simply connected because one can connect two such points to each other by going the other way around the circle, and such paths are unique up to homotopy. 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 $\widetilde{f} : A \rightarrow \mathbb{R}$

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”:¹⁸

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 15.1). 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:

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

We call f an **isomorphism of covers** if it is additionally a homeomorphism, and say in this case 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** (*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 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

¹⁸This is not a universally standard term.

FIGURE 10. A 3-fold cover of $S^1 \vee S^1$ with trivial automorphism group.

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 .

EXERCISE 16.7. Show that if $p : Y \rightarrow X$ is a covering map of path-connected and locally path-connected spaces, then p is also regular if the following slightly weaker condition holds: for some fixed $x \in X$, any two elements $y_1, y_2 \in p^{-1}(x) \subset X$ satisfy $y_2 = f(y_1)$ for some deck transformation $f \in \text{Aut}(p)$.

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, $\tilde{\mathbb{R}P}^n \cong S^n$ for $n \geq 2$, and $\tilde{\mathbb{T}}^n \cong \mathbb{R}^n$.

A substantially less obvious class of examples is given by the surfaces Σ_g of genus $g \geq 2$: these have universal cover $\tilde{\Sigma}_g \cong \mathbb{R}^2$. A standard construction of 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 \mathbb{D}^2 by a suitable group of isometries and then define a covering map $\mathbb{D}^2 \rightarrow \Sigma_g$ 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 $\Phi(y) = H[\gamma]$.

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 $[\alpha \cdot \tilde{\beta}^{-1}] = [\gamma]$, hence $[\alpha] = [\gamma] \cdot [\beta]$, so α is homotopic to $\gamma \cdot \beta$ with fixed end points. Since γ lifts to a loop $\tilde{\gamma}$ and homotopies can also be lifted, we conclude that $\tilde{\alpha}$ is homotopic to $\tilde{\gamma} \cdot \tilde{\beta}$ with fixed end points, implying $y = \tilde{\alpha}(1) = \tilde{\beta}(1) = y'$. □

If the cover is regular so $H \subset G$ is normal, then $\deg(p) = |\text{Aut}(p)|$, 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 $x_0 \xrightarrow{\gamma} x_0$ 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: we will not prove this generalization, but if you're curious, you can read a proof in [Hat02, Prop. 1.39], and the arguments involved will by this point seem very familiar.

THEOREM 17.1. *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 $\text{Aut}(p) \rightarrow N(H)/H$, where $N(H)$ denotes the **normalizer** (Normalisator) of H in G , meaning the largest subgroup of G in which H is a normal subgroup. \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 Theorem 17.1 reduces to Theorem 16.14.

We will not discuss Theorem 17.1 any further, but 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 15.3 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. 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 9.11).
- (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 $p : \tilde{X} \rightarrow X$.

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 15.3 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 17.2. 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 7.27 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 17.3. 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 .

EXAMPLE 17.4. 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 17.5. 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 $1 \in \mathbb{Z}_2$, 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 17.6. Here is a non-discrete example: any subgroup of the orthogonal group $\text{O}(n)$ 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 17.7. The quotient S^n/\mathbb{Z}_2 arising from the action in Example 17.5 is \mathbb{RP}^n .

PROPOSITION 17.8. *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 17.4.

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. \square

EXERCISE 17.9. Show that the two actions of $\pi_1(X, x_0)$ on the universal cover constructed in the above proof are the same.

DEFINITION 17.10. 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 17.11. Show that if a G -action is free and properly discontinuous, then G is discrete.

EXERCISE 17.12. Show that for any covering map $p : Y \rightarrow X$, the action of $\text{Aut}(p)$ on Y as in Example 17.4 is free and properly discontinuous.

The observation that actions of deck transformation groups are free already has some nontrivial consequences, for instance:

PROPOSITION 17.13. *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 17.14. *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.

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 18.2. *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 18.3. *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 7.17 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. 7.16), 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 18.4. 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 18.5. 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 18.6. 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 18.7. 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 $g \geq 0$, $\mathbb{R}P^2$, \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 18.8. Some examples of arbitrary dimension n without boundary are S^n , $\mathbb{R}P^n$, \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.¹⁹ 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

¹⁹Recall 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.

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 18.9. 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*).

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 18.10 (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 18.11. 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 18.12. The special linear group $SL(n, \mathbb{R}) = \det^{-1}(1)$ 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 $SL(n, \mathbb{R})$ is a smooth manifold of dimension $n^2 - 1$.

Now let’s look at a couple of non-examples.

EXAMPLE 18.13. 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 18.14. Here is a space that is locally Euclidean and second countable, but not Hausdorff: the line with two zeroes, i.e. $(\mathbb{R} \times \{0, 1\})/\sim$ with $(x, 0) \sim (x, 1)$ for all $x \neq 0$. We should emphasize that we are considering the quotient topology on this space, not the pseudometric topology (cf. Example 6.12). Now $0_0 := [(0, 0)]$ and $0_1 := [(0, 1)]$ each have neighborhoods that can be identified with open intervals in \mathbb{R} , and so (obviously) does every other point, thus this space would be a 1-manifold if we did not require manifolds 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 18.14 would be a manifold, and we don't like Example 18.14. But there are better reasons. One of them is related to the implicit function theorem, Theorem 18.10 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 18.10, 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 18.15. *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 18.15 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 18.15 truly qualifies as a good theorem. But I want to talk about another one that I think is better.

THEOREM 18.16. *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 18.17. *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 18.16 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 18.16, $\widehat{M} \cong \mathbb{R}$. 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 18.16 is not conceptually very hard, though writing it down correctly with all the details is a little annoying, so I will settle for a sketch.²⁰ My main objective is to point out explicitly where the Hausdorff and second countability conditions are needed. You saw already from Example 18.14 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.

PROOF OF THEOREM 18.16 (SKETCH). Given a nonempty connected 1-manifold M without boundary, let $N \in \mathbb{N}$ denote the smallest possible number (finite or infinite) of charts in any collection $\{\varphi_\alpha : \mathcal{U}_\alpha \xrightarrow{\cong} \mathbb{R}\}_{\alpha \in J}$ such that $M = \bigcup_{\alpha \in J} \mathcal{U}_\alpha$. We consider the following cases.

Case 1: $N = 1$. Then M is homeomorphic to \mathbb{R} , so we are done.

Case 2: $N = 2$. Assume $M = \mathcal{U} \cup \mathcal{V}$, where $\varphi : \mathcal{U} \xrightarrow{\cong} \mathbb{R}$ and $\psi : \mathcal{V} \xrightarrow{\cong} \mathbb{R}$ are charts. Since M is connected, we have $\mathcal{U} \cap \mathcal{V} \neq \emptyset$, and we can also assume that neither of \mathcal{U} or \mathcal{V} fully contains the other since then N would be 1. We claim that for any connected component \mathcal{W} of $\mathcal{U} \cap \mathcal{V}$, $\varphi(\mathcal{W})$ and $\psi(\mathcal{W})$ are each infinite half-intervals in \mathbb{R} , i.e. they are intervals of the form $(-\infty, a)$ or (a, ∞) for $a \in \mathbb{R}$. This requires the assumption that M is Hausdorff. (Exercise: show that the space in Example 18.14 is covered by two charts for which this claim does not hold.) The main point is to show that if $x_n \in \mathcal{W}$ is any sequence such that $x_n \rightarrow x \in \mathcal{V}$ but $x \notin \mathcal{U}$, then $\varphi(x_n)$ diverges to $\pm\infty$. Indeed, suppose to the contrary that x_n can be replaced by a subsequence for which $\varphi(x_n) \in \mathbb{R}$ is bounded. Then after replacing x_n with a further subsequence, we can assume $\varphi(x_n)$ converges to some point $y \in \mathbb{R}$. But then $x_n = \varphi^{-1}(\varphi(x_n)) \rightarrow \varphi^{-1}(y)$ since φ^{-1} is continuous, and the Hausdorff axiom therefore implies $\varphi^{-1}(y) = x$ since limits of convergent sequences are unique. This contradicts the assumption that $x \notin \mathcal{U}$.

Once the claim is established, we observe that at most two disjoint infinite half-intervals can fit into \mathbb{R} , and it follows that $\mathcal{U} \cap \mathcal{V}$ has at most two connected components. If it has only one, then one can easily find a homeomorphism of M to \mathbb{R} , which actually reduces us to the case $N = 1$. If there are instead two components, then there is a homeomorphism of M to S^1 that identifies two points in separate components of $\mathcal{U} \cap \mathcal{V}$ with antipodal points on the circle.

Case 3: $N > 2$ but $N < \infty$. Writing $M = \bigcup_{i=1}^N \mathcal{U}_i$ where the \mathcal{U}_i are domains of charts, the same arguments as above using the Hausdorff property imply that for each $i \neq j$, $\mathcal{U}_i \cap \mathcal{U}_j$ has at most two connected components. One can now imagine several possibilities, but all of them produce homeomorphisms from M to either S^1 or \mathbb{R} . For instance, if some intersection $\mathcal{U}_i \cap \mathcal{U}_j$ has two components, then for the same reason as in case 2, the union of these two sets forms a connected component of M that is homeomorphic to S^1 , meaning (since M is connected) all the other \mathcal{U}_k must be contained in $\mathcal{U}_i \cup \mathcal{U}_j$ and we're actually back in the case $N = 2$. If this does not happen but there exists a set $\{i_1, \dots, i_m\} \in \{1, \dots, N\}$ such that $\mathcal{U}_{i_j} \cap \mathcal{U}_{i_{j+1}}$ for every j and $\mathcal{U}_{i_m} \cap \mathcal{U}_{i_1}$ each have one connected component while all other intersections are empty, then the union of these m sets again forms a connected component homeomorphic to S^1 . In all other cases,

²⁰A more complete version of the proof is explained as a series of straightforward exercises in [Gal87].

M will be homeomorphic to \mathbb{R} , and one can construct a homeomorphism by identifying each \mathcal{U}_i with a (finite or half-infinite) subinterval of \mathbb{R} and piecing these intervals together.

Case 4: $N = \infty$. In this case M cannot be compact, but since it is second countable, we know that any uncountably infinite cover by charts can be replaced by a countable subcover $\mathcal{U}_1, \mathcal{U}_2, \mathcal{U}_3, \dots$, cf. Lemma 5.25. The same argument as usual implies that $\mathcal{U}_i \cap \mathcal{U}_j$ always has at most two components for $i \neq j$, and any case with two components leads as before to a homeomorphism $M \cong S^1$ and thus a contradiction. A similar contradiction arises in the second scenario described in case 3 above. The result is that one can find a homeomorphism of M to some countable union of (finite or half-infinite) intervals with connected overlap regions, and the latter will always be homeomorphic to \mathbb{R} . \square

The second countability axiom was used at the end of this proof by appealing to the fact that \mathbb{R} can be identified with the union of a countable sequence of overlapping open intervals. This would be a lot less obvious if we were required to fit in *uncountably* many open intervals, and I would now like to describe an example of what strange beasts can 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 18.18. *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 18.19. 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 18.20. *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 18.21. 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{RP}^2 , e.g. $\mathbb{RP}^2 \# \mathbb{RP}^2$ is homeomorphic to the Klein bottle. By the same trick that we used in Lecture 13 to understand Σ_g , the g -fold connected sum $\#_{i=1}^g \mathbb{RP}^2$ is homeomorphic to a space obtained from a polygon with $2g$ edges by identifying them in pairs according to the sequence $a_1, a_1, \dots, a_g, a_g$, thus

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

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.* □

You might now be wondering whether new examples can be constructed by taking the connected sum of a surface from Example 19.1 with some surface from Example 19.2. The answer is no:

PROPOSITION 19.5. *$\mathbb{RP}^2 \# \mathbb{T}^2$ is homeomorphic to the connected sum of \mathbb{RP}^2 with the Klein bottle.*²¹

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

$$i_1 \amalg i_2 : \mathbb{D}^2 \amalg \mathbb{D}^2 \hookrightarrow \Sigma,$$

²¹This 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.

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.

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 $\#_{i=1}^g \mathbb{R}P^2$ 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 $\mathbb{R}P^2$ by connected sums. (Note that whenever both \mathbb{T}^2 and $\mathbb{R}P^2$ appear in this collection, Prop. 19.5 allows us to replace \mathbb{T}^2 with two copies of $\mathbb{R}P^2$, as $\mathbb{R}P^2 \# \mathbb{R}P^2$ 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.

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

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.

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

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 $\mathbb{R}P^2 \# \mathbb{R}P^2$, while attaching a cap or cross-cap gives connected sums with S^2 or $\mathbb{R}P^2$ 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 S is a subset of the set of all finite subsets of V , and $\sigma \in S$ is called an **n -simplex** of K if it has $n + 1$ elements. We require the following conditions:

- (1) Every vertex $v \in V$ gives rise to a 0-simplex in K , i.e. $\{v\} \in S$;
- (2) If $\sigma \in S$ then every subset $\sigma' \subset \sigma$ is also an element of S .

For any n -simplex $\sigma \in S$, its subsets are called its **faces** (*Seiten* or *Facetten*), and in particular the subsets that are $(n - 1)$ -simplices are called **boundary faces** (*Seitenflächen*) of σ . The second condition above thus says that for every simplex in the complex, all of its 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. We shall describe it only in the case of a finite complex,²² since that is what we need in our applications. 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

²²The 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, we will not.

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 .

DEFINITION 19.11. A **triangulation** (*Triangulierung*) of a closed 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 , 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.

THEOREM 19.12. *Every closed surface admits a triangulation.*

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 known for a few years—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. informally we can write

$$\Sigma^{(j+1)} = \Sigma^{(j)} / \ell_1 \sim \ell_2.$$

Here ℓ_1 and ℓ_2 are two distinct boundary faces of 2-simplices, so they are each homeomorphic to the compact interval I , and their interiors are disjoint, though they may have boundary points (vertices of the triangulation) in common. One can now imagine various scenarios:

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 boundary 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 $\partial\Sigma^{(j)}$. 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 18.9 and Theorem 18.10). 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 18.10) 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 submanifold²³ 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 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 if the number of points in $f^{-1}(q)$ is even and 1 if it is odd.

- (a) Prove that $\deg_2(f)$ depends only on the homotopy class of $f : M \rightarrow N$.

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 .

- (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 $\deg_2(f) \neq \deg_2(g)$.


²³A subset $Y \subset M$ of a smooth manifold M is called a *smooth submanifold* if it admits the structure of a smooth manifold such that the inclusion $Y \hookrightarrow M$ is smooth.

- (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 and higher homotopy groups

The first portion of 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. In the second portion, we will start expanding our repertoire of algebraic topological invariants with a brief sketch of the higher homotopy groups.

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:



What most of us 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, 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 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 $\mathbb{T}^2 \subset \mathbb{R}^3$. The fancy way of saying this in differential geometric language is that the *normal bundle* of the standard immersion $K^2 \hookrightarrow \mathbb{R}^3$ is nontrivial, whereas the standard embedding $\mathbb{T}^2 \hookrightarrow \mathbb{R}^3$ has trivial normal bundle.

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 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 10.26. 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. This notion can also be defined in all positive dimensions, not only dimension two. 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 from stating the definition for dimensions $n \geq 3$ for now 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 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 circles 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 circle continuously through small circles around other points and end up with a clockwise circle. 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 circle 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 circles 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 toward p , and $\epsilon(p) = -1$ whenever this direction points away from p . (Different authors may define this in slightly different ways, but it hardly matters: 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 group on N objects after choosing a numbering v_1, \dots, v_N for the elements in J . The symmetric group is generated by *flips*, meaning permutations that interchange two elements of J while leaving the rest fixed, and we say that π is an **even** or **odd** permutation if it can be written as a composition of evenly many flips; otherwise it is an **odd** permutation. If we represent π by a matrix, then we can recognize the even/odd permutations as those for which this matrix has positive/negative determinant respectively. 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. Recall that 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




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

come with an ordering, then the vertices of $\partial_{(k)}\sigma$ inherit an ordering from this, though here we 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 circles 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 circles: clearly a choice of “counterclockwise circles” 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 circles can be defined since following a loop 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 we'll discuss the Euler characteristic, and then you'll see why this is true.)

EXERCISE 20.13. In Exercise 13.15 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.

The rest of this 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 "counting holes" in a topological space. Let me describe a few concrete examples in which the fundamental group either does or does not succeed in this task.

EXAMPLE 20.14. 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 20.15. A 3-dimensional generalization of Example 20.14 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 20.16. A different type of generalization of Example 20.14 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 20.17. 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 13.15, 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 *Topology 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 20.18. 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 [[Mi197](#)].

EXAMPLE 20.19. 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.

21. Bordism groups and simplicial homology

The higher homotopy groups do remedy one of the drawbacks of π_1 that I pointed out in the last lecture: e.g. π_2 can be used to detect the hole in $\mathbb{R}^3 \setminus \mathring{\mathbb{D}}^3$ since, by homotopy invariance,

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

with the inclusion $S^2 \hookrightarrow \mathbb{R}^3 \setminus \mathring{\mathbb{D}}^3$ representing a generator. But there's another drawback here: while π_k can detect higher-dimensional holes, they are still holes of a fairly specific type which one might call "sphere-shaped" holes. What kind of hole is not sphere-shaped, you ask? Is there such a thing as a "torus-shaped" hole? How about this one:

EXAMPLE 21.1. Let $X = S^1 \times \mathbb{R}^2$ and $X_0 = S^1 \times \mathring{\mathbb{D}}^2$, so $X \setminus X_0 = S^1 \times (\mathbb{R}^2 \setminus \mathring{\mathbb{D}}^2)$ admits a deformation retraction to $\partial \bar{X}_0 = S^1 \times S^1 = \mathbb{T}^2$. By homotopy invariance, we have $\pi_1(X) \cong \pi_1(S^1) \cong \mathbb{Z}$ and $\pi_1(X \setminus X_0) \cong \pi_1(\mathbb{T}^2) \cong \mathbb{Z}^2$, 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*.²⁴

Remedy 2: Bordism groups. For any space X and each integer $k \geq 0$, let

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

²⁴In 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)²⁵ 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 21.2. $(M, f) \sim (M, g)$ 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 21.3. Recall from Example 20.17 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 21.4. 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 21.1, 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 all elements in $\Omega_k(X)$ are “2-torsion”:

PROPOSITION 21.5. *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) .²⁶ \square

²⁵Note 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.

²⁶One of the slightly confusing things about $\Omega_k(X)$ is that there is always some ambiguity about how to split up the open and closed subsets of ∂W into M_- and M_+ . For the bordism in the proof of Prop. 21.5, 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 21.6. The k -dimensional oriented bordism group of X is²⁷

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

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 21.5 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,

²⁷The ‘‘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 21.7. 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 21.8. $\partial_{n-1} \circ \partial_n = 0$ 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 21.8 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 21.9. 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 eight 2-simplices, twelve 1-simplices and four vertices labeled as follows:

$$\begin{aligned} S_2 &= \{A, B, C, D, E, F, G, H\}, \\ S_1 &= \{a, b, c, d, e, f, g, h, i, j, k, \ell\}, \\ S_0 &= \{\alpha, \beta, \gamma, \delta\}. \end{aligned}$$

In addition to the orientations of the 2-simplices that come from this being an oriented triangulation, the figure shows (via arrows) an arbitrary choice of orientations for all 1-simplices, and we shall assume all the 0-simplices are oriented with a positive sign. One can now begin writing down relations such as

$$\partial A = a - h - c, \quad \partial B = i - k + h, \quad \partial a = \beta - \alpha$$

and so forth, but writing down all such relations would be rather tedious, so let us instead try to reason more geometrically. The computation of $H_0^\Delta(K; \mathbb{Z})$ is not hard in any case: all 0-chains are cycles since $\partial_0 = 0$, including the four generators α, β, γ and δ , but all four of them are also homologous to each other since any pair of them can be connected by an oriented 1-simplex pointing from one to the other, e.g. $\partial a = \beta - \alpha$ implies $[\alpha] = [\beta]$, and $\partial i = \delta - \beta$ implies $[\beta] = [\delta]$. The result is

$$H_0^\Delta(K; \mathbb{Z}) \cong \mathbb{Z},$$

with a canonical generator represented by any of the vertices in the complex. Notice that this matches the oriented bordism group $\Omega_0^{\text{SO}}(\mathbb{T}^2)$ 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

$$\partial(i + \ell - c - b) = 0,$$

so $i + \ell - c - b$ is a 1-cycle, but not a very interesting one since it is also the boundary of the region filled by the 2-simplices C and D : in particular,

$$\partial(-C - D) = i + \ell - c - b,$$

hence $[i + \ell - c - b] = 0 \in H_1^\Delta(K; \mathbb{Z})$. To find more interesting 1-cycles, it helps to remember what we already know about $\pi_1(\mathbb{T}^2) \cong \mathbb{Z}^2$. We can easily find two loops through 1-simplices that




FIGURE 13. A simplicial complex with $|K| = \mathbb{T}^2$.

represent the two distinct generators of this fundamental group: one of them is $i + j$, and we easily see that

$$\partial(i + j) = (\delta - \beta) + (\beta - \delta) = 0.$$

Another is $c + d$, but notice that the loops corresponding to these two 1-cycles are homotopic in \mathbb{T}^2 , and relatedly, they form the boundary of the region filled by the 2-simplices C, D, G and H , so

$$\partial(C + D + G + H) = c + d - (i + j),$$

implying $[c + d] = [i + j] \in H_1^\Delta(K; \mathbb{Z})$. One can show however that this homology class really is nontrivial, and it is not the only one: the other generator of $\pi_1(\mathbb{T}^2)$ corresponds to either of the two homologous 1-cycles $a + b$ or $k + \ell$. The end result is

$$H_1^\Delta(K; \mathbb{Z}) \cong \mathbb{Z}^2,$$

the same as the fundamental group.

As observed at the beginning of this lecture, the fact that \mathbb{T}^2 has a contractible universal cover implies that $\pi_2(\mathbb{T}^2) = 0$, so if there are any interesting 2-cycles in \mathbb{T}^2 , they will not look like spheres. But if you think that $H_2(K; \mathbb{Z})$ should have something to do with the oriented bordism group $\Omega_2^{SO}(\mathbb{T}^2)$, then there is a fairly obvious candidate for a 2-cycle in this picture: \mathbb{T}^2 itself is a closed oriented manifold, and the oriented triangulation we have chosen turns it into a 2-cycle:

$$\partial(A + B + C + D + E + F + G + H) = 0.$$

The point is that since the triangulation is oriented, writing down each individual term in this sum would produce a linear combination of 1-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 eight of the 2-simplices can ever be a cycle, as its boundary will have to include some 1-simplices that have nothing to cancel with. It follows easily that all 2-cycles in this complex are integer multiples of the one found above, and none of them are boundaries since there are no 3-simplices, thus

$$H_2^\Delta(K; \mathbb{Z}) \cong \mathbb{Z}.$$

I can now state a theorem that is really rather amazing, though I'm sorry to say that we will not be able to prove it until next semester:

THEOREM 21.10. *For any simplicial complex K , the simplicial homology groups $H_n^\Delta(K; G)$ depend (up to isomorphism) on the polyhedron $X = |K|$, but not on the complex 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 21.10 was formulated near the beginning of the 20th century and was based on a conjecture called the **Hauptvermutung**: 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 21.10 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 21.10, 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 course.

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. 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

We will also want to prove that the homomorphism induced on homology by a continuous map $f : X \rightarrow Y$ depends only on the homotopy class of f . It would be too much to ask that all homotopic maps induce exactly the same chain map, but a more plausible goal is to establish some algebraic condition on two chain maps that forces them to induce the same homomorphism on homology. This condition is called *chain homotopy*.

DEFINITION 22.6. 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 22.7. If there exists a chain homotopy between two chain maps f and g from (A_*, ∂^A) to (B_*, ∂^B) , then they induce the same sequence of 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

REMARK 22.8. We will not need it, 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 naturally associate 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*, meaning there exists another chain map $T : C_*(K'; G) \rightarrow C_*(K; G)$ such that each of $S \circ T$ and $T \circ S$ are chain homotopic to the identity. It then follows from Proposition 22.7 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 21.10). We will see an analogue of this in singular homology when we need to prove the excision property next week. In the simplicial context, one usually has to consult some of the older textbooks to find adequate discussions of such topics, but e.g. [Spa95] is quite nice.

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 22.9. 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 case. 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 $C_n(X; G)$.²⁸ 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.

²⁸The word "singular" in this context refers to the fact that since there is no condition beyond continuity required for the maps $\sigma : \Delta^n \rightarrow X$, their images might not look "simplex-shaped" at all, but could instead be full of singularities.

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).derstood

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 22.10. $\partial^2 = 0$. □

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 22.11. *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 , so after fixing such a homeomorphism, we can think of each $\sigma \in \mathcal{K}_1(X)$ as a path $\sigma : I \rightarrow X$ and write

$$\partial\sigma = \sigma(1) - \sigma(0) \in C_0(X; \mathbb{Z}).$$

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

The next result is a straightforward exercise based on the definitions, and you should also compare it with our previous discussion of the bordism groups of a point, if only to observe that the result is very different: while bordism groups require some information about the classification of manifolds which has nothing to do with the one-point space, the singular homology of $\{\text{pt}\}$ is much simpler.

EXERCISE 22.12. 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 22.11, 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 22.13. 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 22.14. Recall from Exercise 13.15 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 22.13, 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_2 , e.g. for this same reason, γ also represents the trivial bordism class in $\Omega_1^{\text{SO}}(\Sigma_2)$. 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 $\Sigma_{h,1} \hookrightarrow \Sigma_2$ as a linear combination of singular 2-simplices in Σ_2 .

The proof of Theorem 22.13 is not trivial, but it is simple enough to leave as a guided homework problem (see Exercise 22.15 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 22.15. Let us prove Theorem 22.13. 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 $\Delta^1 \rightarrow I : (t_0, t_1) \mapsto t_0$, 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?
- 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.
- 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 12.20), 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 homomorphism²⁹

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

- (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 13.15.*

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. □

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 $(\text{Id})_* = \text{Id}$. □*

²⁹Since $\pi_1(X) / [\pi_1(X), \pi_1(X)]$ is abelian, we are adopting the convention of writing its group operation as addition, so the multiplication of an integer $m \in \mathbb{Z}$ by an element $\Psi(\sigma) \in \pi_1(X) / [\pi_1(X), \pi_1(X)]$ is defined accordingly.

We will show in the next lecture that the homomorphisms f_* induced by continuous maps f only depend on f up to homotopy, which has the easy consequence that $H_*(X; G)$ only depends on the homotopy type of X .

But first, let us generalize the discussion somewhat. Algebraic gadgets often have the feature that they become easier to compute if you add more structure to them, sometimes at the cost of making the basic definitions slightly more elaborate. We will now do that with singular homology by introducing the *relative homology* groups of pairs. A **pair of spaces** (X, A) , often abbreviated as simply a “pair,” (*topologisches Paar*) consists of a topological space X and a subset $A \subset X$. Given two pairs (X, A) and (Y, B) , a map $f : X \rightarrow Y$ 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 following example will play a major role in our computation of $H_*(S^n; \mathbb{Z})$ next week.

EXAMPLE 23.3. Recall from Lecture 11 that the **suspension** (*Einhangung*) SX of a space X is defined by gluing together two copies of its cone,

$$(23.1) \quad SX = C_+X \cup_X C_-X,$$

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 (23.1) 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 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 $[c] \in C_n(X, A; G)$, and two relative cycles $b, c \in C_n(X; G)$ are homologous (meaning $[b] = [c] \in H_n(X, A; G)$) if and only if

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

In particular, a relative cycle is nullhomologous if and only if it is the sum of a boundary plus a chain contained in A .

This is all abstract nonsense so far, but 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 23.5 (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 23.5 remains true in this case without any need for the multiplicative factor $k \in \mathbb{N}$.

There is an interesting relationship in general between the groups $H_*(X; G)$, $H_*(A; G)$ and $H_*(X, A; G)$ for any pair (X, A) . This begins with another trivial observation: let $i : A \hookrightarrow X$ and $j : X = (X, \emptyset) \hookrightarrow (X, A)$ denote the natural inclusions, and consider the sequence of chain maps

$$(23.2) \quad 0 \longrightarrow C_*(A; G) \xrightarrow{i_*} C_*(X; G) \xrightarrow{j_*} C_*(X, A; G) \longrightarrow 0,$$

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

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

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

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

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 three neighboring groups in the sequence are trivial, then it suffices to look at a sequence of five groups

$$0 \longrightarrow A_1 \xrightarrow{f_1} A_2 \xrightarrow{f_2} A_3 \longrightarrow 0,$$

which is called a **short exact sequence** (*kurze exakte Sequenz*) if it is exact. Being exact 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 (23.2) 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 23.6. *Suppose (A_*, ∂^A) , (B_*, ∂^B) and (C_*, ∂^C) are chain complexes and*

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

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

$$(23.3) \quad \begin{aligned} \dots \xrightarrow{\hat{\partial}_*} H_{n+1}(A_*, \partial^A) &\xrightarrow{f_*} H_{n+1}(B_*, \partial^B) \xrightarrow{g_*} H_{n+1}(C_*, \partial^C) \\ &\xrightarrow{\hat{\partial}_*} H_n(A_*, \partial^A) \xrightarrow{f_*} H_n(B_*, \partial^B) \xrightarrow{g_*} H_n(C_*, \partial^C) \\ &\xrightarrow{\hat{\partial}_*} H_{n-1}(A_*, \partial^A) \xrightarrow{f_*} H_{n-1}(B_*, \partial^B) \xrightarrow{g_*} H_{n-1}(C_*, \partial^C) \xrightarrow{\hat{\partial}_*} \dots \end{aligned}$$

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 (23.2) into the so-called **long exact sequence of the pair** (X, A) ,

$$(23.4) \quad \dots \rightarrow H_{n+1}(X, A; G) \xrightarrow{\partial_*} H_n(A; G) \xrightarrow{i_*} H_n(X; G) \xrightarrow{j_*} H_n(X, A; G) \xrightarrow{\partial_*} H_{n-1}(A; G) \rightarrow \dots$$

To see why this might be useful, consider the pair $(X, A) = (\mathbb{D}^k, S^{k-1})$. We will prove in the next lecture that $H_*(X; G)$ only depends on the homotopy type of X , so since \mathbb{D}^k is contractible, we have

$$H_n(\mathbb{D}^k; \mathbb{Z}) \cong H_n(\{\text{pt}\}; \mathbb{Z}) = 0 \quad \text{for } n > 0$$

by Exercise 22.12. It follows that every third term in the long exact sequence of (\mathbb{D}^k, S^{k-1}) is trivial, producing a whole collection of exact sequences

$$0 \longrightarrow H_{n+1}(\mathbb{D}^k, S^{k-1}; \mathbb{Z}) \xrightarrow{\partial_*} H_n(S^{k-1}; \mathbb{Z}) \longrightarrow 0$$

for $n \geq 1$. Since the first and third maps in this sequence are both trivial, the sequence can only be exact if the map ∂_* is an isomorphism, proving

$$H_{n+1}(\mathbb{D}^k, S^{k-1}; \mathbb{Z}) \cong H_n(S^{k-1}; \mathbb{Z})$$

for every $k, n \in \mathbb{N}$. For example, since we already know $H_1(S^{k-1}; \mathbb{Z}) \cong \pi_1(S^{k-1})$, we have just computed $H_2(\mathbb{D}^k, S^{k-1}; \mathbb{Z})$ for every $k \geq 1$. Combining this with a few more tools in the next lecture will lead to an isomorphism between $H_n(S^{k-1}; \mathbb{Z})$ and $H_{n+1}(S^k; \mathbb{Z})$, allowing us to compute all the homology groups of spheres in arbitrary dimension by induction.

Theorem 23.6 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.³⁰ 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.

³⁰I 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.

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

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

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 $\ker \partial_* = \text{im } g_*$, $\ker g_* = \text{im } f_*$ and $\ker f_* = \text{im } \partial_*$. 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 (23.2), 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:

$$(23.5) \quad \partial_*[c] = [\partial c].$$

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 23.7. Use the formula (23.5) to give a direct proof that the sequence (23.4) is exact.

REMARK 23.8. Exercise 23.7 is straightforward and doable in a much shorter time than the proof of Theorem 23.6, 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*.

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.

LEMMA 24.1. *Homotopies between maps of pairs $f, g : (X, A) \rightarrow (Y, B)$ determine chain homotopies between the induced chain maps $f_*, g_* : C_*(X, A; G) \rightarrow C_*(Y, B; G)$.*

Applying Proposition 22.7, this gives:

COROLLARY 24.2. *The maps $f_* : H_n(X, A; G) \rightarrow H_n(Y, B; G)$ induced by a map of pairs $f : (X, A) \rightarrow (Y, B)$ depend only on the homotopy class of f (as a map of pairs). \square*

The obvious corollary of this corollary 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 10.22). The reason is that fundamental groups and the induced homomorphisms between them require base points for their definition, whereas in homology, base points are completely irrelevant.

COROLLARY 24.3. *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 Corollary 24.2 gives $f_* \circ g_* = \mathbb{1}$ and $g_* \circ f_* = \mathbb{1}$ for the induced maps on homology, implying that both are isomorphisms. \square

You may have thought the definition of a chain homotopy in the algebraic introduction to Lecture 22 seemed a little unnatural—it is not obvious for instance whether a chain homotopy as we defined it is the *only* algebraic condition implying that two chain maps induce the same map on homology. However, the following proof of Lemma 24.1 provides convincing evidence that this definition is the right one. It shows 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.

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 be³¹

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

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:

³¹One 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.

LEMMA 24.4. *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;
- (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

$$h(\sigma) := \sum_i \epsilon_i (\sigma \circ \tau_i) \in C_{n+1}(Y; \mathbb{Z}),$$

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 24.4, 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 24.1 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 24.5 (via Exercise 22.12). 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 24.6 (via Theorem 22.13). 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 24.7 (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 .

Before discussing the proof, here is an example to motivate why we might want such a theorem.

EXAMPLE 24.8. Splitting the sphere S^n into two hemispheres gives a decomposition $S^n \cong \mathbb{D}_+^n \cup_{S^{n-1}} \mathbb{D}_-^n$, where \mathbb{D}_\pm^n are copies of the disk \mathbb{D}^n glued together along their common boundary S^{n-1} , which forms the “equator” of S^n . In Example 23.3, we showed that if the south pole $p_- \in \mathbb{D}_-^n$ is removed, the inclusion $(\mathbb{D}_+^n, S^{n-1}) \hookrightarrow (S^n \setminus \{p_-\}, \mathbb{D}_-^n \setminus \{p_-\})$ becomes a homotopy equivalence of pairs. We claim that, in fact, the inclusion $(\mathbb{D}_+^n, S^{n-1}) \hookrightarrow (S^n, \mathbb{D}_-^n)$ still induces an isomorphism on relative homology, without removing the point p_- . This follows by writing it as a composition of two inclusions

$$(\mathbb{D}_+^n, S^{n-1}) \hookrightarrow (S^n \setminus \{p_-\}, \mathbb{D}_-^n \setminus \{p_-\}) \hookrightarrow (S^n, \mathbb{D}_-^n),$$

inducing a composition of induced homomorphisms

$$H_*(\mathbb{D}_+^n, S^{n-1}; G) \rightarrow H_*(S^n \setminus \{p_-\}, \mathbb{D}_-^n \setminus \{p_-\}; G) \rightarrow H_*(S^n, \mathbb{D}_-^n; G).$$

The first is an isomorphism by Corollary 24.3, and since $\{p_-\}$ is a set with closure contained in the interior of \mathbb{D}_-^n , the second is also an isomorphism due to Theorem 24.7. Recall from the previous lecture that each of the groups $H_*(\mathbb{D}_+^n, S^{n-1}; G)$ and $H_*(S^n, \mathbb{D}_-^n; G)$ fit into long exact sequences that also include the absolute homology groups of S^n , S^{n-1} and \mathbb{D}^n . We know the last one already since \mathbb{D}^n is contractible, so the result of combining all this information will be a precise relation between the homology groups of S^{n-1} and S^n , which we will use in the next lecture to compute all of the groups $H_*(S^n; \mathbb{Z})$.

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 subspace $C_n(A; G) + C_n(X \setminus B; G)$. 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 24.7, 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 from 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 $C_n(\mathcal{U}; G) + C_n(\mathcal{V}; G)$ but represents the same homology class in cases where c is a (relative) cycle. You can imagine why this should be possible if you think of a homology class $f_*[M] \in H_n(X; \mathbb{Z})$ represented by a triangulated oriented n -manifold M with a map $f : M \rightarrow X$, as outlined in the previous lecture (see the discussion surrounding Thom's Theorem 23.5). 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. In particular, we should be free to replace the triangulation by a *finer* one, which has more simplices but all of them are 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 in the barycentric subdivision of Δ^n a parametrization $\tau : \Delta^n \rightarrow \Delta^n$ 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 24.9. $S : C_*(X; G) \rightarrow C_*(X; G)$ 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 24.10. $S : C_*(X; G) \rightarrow C_*(X; G)$ is chain homotopic to the identity map.

PROOF. As in the proof of Lemma 24.1, 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 $S^m c \in C_*(\mathcal{U}; G) + C_*(\mathcal{V}; G)$ for m sufficiently large. This is enough information to prove the excision theorem, but it is worth making a slight extra effort to prove something a bit better which will come in useful for the

further development of homology and cohomology theory next semester. 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 24.7 can be seen as a consequence of the slightly 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 24.11. *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, and that 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 24.10 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 ∂ ,

$$\begin{aligned} S^m - \mathbb{1} &= (S^{m-1} - \mathbb{1})S + (S - \mathbb{1}) = (\partial h_{m-1} + h_{m-1}\partial)S + \partial h_1 + h_1\partial \\ &= \left(\partial h_1 \sum_{k=0}^{m-2} S^k + h_1 \sum_{k=0}^{m-2} S^k \partial \right) S + \partial h_1 + h_1\partial = \partial h_1 \sum_{k=1}^{m-1} S^k + h_1 \sum_{k=1}^{m-1} S^k \partial + \partial h_1 + h_1\partial \\ &= \partial h_1 \sum_{k=0}^{m-1} S^k + h_1 \sum_{k=0}^{m-1} S^k \partial = \partial h_m + h_m \partial. \end{aligned}$$

The iterated subdivision maps S^m can be assumed to satisfy

$$(24.3) \quad S^m(\sigma) \in C_*(\mathcal{U}; G) + C_*(\mathcal{V}; G)$$

for any given $\sigma : \Delta^n \rightarrow X$ if m is large enough, so for each 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 24.10: 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 24.7.

LEMMA 24.12. *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 maps, 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

$$C_*(X \setminus B; G) + C_*(A; G) \xrightarrow{i} C_*(X; G)$$

is a chain homotopy equivalence. Since $X \setminus \bar{B}$ and $\overset{\circ}{A}$ form an open cover of X , Lemma 24.11 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

25. The homology of the spheres, and beyond

Recall that S^{n+1} is homeomorphic to the suspension SS^n of S^n . Let us consider the relationship in general between $H_*(X; G)$ and $H_*(SX; G)$. We write

$$SX = C_+X \cup_X C_-X,$$

where the two cones $C_\pm X$ are contractible spaces that each contain a naturally embedded copy of X . Letting $p_- \in C_-X$ denote the tip of the bottom cone, we saw in Example 23.3 that the natural inclusion

$$(C_+X, X) \xrightarrow{i} (SX \setminus \{p_-\}, C_-X \setminus \{p_-\})$$

is a homotopy equivalence of pairs. Here is an interesting diagram to ponder:

$$(25.1) \quad \begin{array}{ccc} H_k(X; G) & & H_{k+1}(SX; G) \\ \partial_* \uparrow & & \downarrow \varphi_* \\ H_{k+1}(C_+X, X; G) & \xrightarrow{i_*} & H_{k+1}(SX \setminus \{p_-\}, C_-X \setminus \{p_-\}; G) \xrightarrow{j_*} H_{k+1}(SX, C_-X; G) \end{array}$$

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 24.7) 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 (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

$$0 \longrightarrow H_{k+1}(C_+X; G) \xrightarrow{\partial_*} H_k(X; G) \longrightarrow 0$$

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 25.1. *For all spaces X , abelian groups G and integers $k \geq 1$, the diagram (25.1) defines an isomorphism*

$$S_* = \varphi_*^{-1} \circ j_* \circ i_* \circ \partial_*^{-1} : H_k(X; G) \rightarrow H_{k+1}(SX; G).$$

□

EXERCISE 25.2. 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

$$(25.2) \quad \partial c_+ = -\partial c_- = b$$

and

$$(25.3) \quad S_*[b] = [c_+ + c_-].$$

Note that $c_+ + c_- \in C_{k+1}(SX; G)$ is automatically a cycle since $\partial c_+ = -\partial c_-$. Show moreover that (25.3) is satisfied for any pair of chains c_{\pm} satisfying (25.2).

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 22.11), and the latter is the abelianization of $\pi_1(S^n)$ by Theorem 22.13. Since $SS^n \cong S^{n+1}$, we can now compute $H_*(S^n; \mathbb{Z})$ inductively for every $n \geq 1$:

THEOREM 25.3. 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 22.11 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 25.1. 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 25.1 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 22.12 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 25.4. 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 25.4, 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 25.5. 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 $B_\delta^m(0) \subset \mathbb{R}^m$ 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 $H_{n-1}(B_\epsilon^n(0) \setminus \{0\}; \mathbb{Z}) \neq 0$. □

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. For the rest of this lecture, I want to talk about some related topics that will be discussed more at length in *Topologie II*.

First a bit of good news: while the proofs of homotopy invariance and excision in the previous lecture 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. Roughly speaking, an *axiomatic* homology theory h_* with coefficient group G is anything³² that associates to every pair (X, A) a sequence of abelian groups $h_n(X, A)$, together with homomorphisms $f_* : h_n(X, A) \rightarrow h_n(Y, B)$ induced by continuous maps $f : (X, A) \rightarrow (Y, B)$ and connecting homomorphisms $\partial_* : h_n(X, A) \rightarrow h_{n-1}(A)$, such that the following conditions are satisfied:

- (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

$$\dots \rightarrow h_{n+1}(X, A) \xrightarrow{\partial_*} h_n(A) \xrightarrow{i_*} h_n(X) \xrightarrow{j_*} h_n(X, A) \xrightarrow{\partial_*} h_{n-1}(A) \rightarrow \dots$$

is exact.

- (EXCISION) If $B \subset \bar{B} \subset \mathring{A} \subset A \subset X$, 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) $h_0(\{\text{pt}\}) = G$ and $h_n(\{\text{pt}\}) = 0$ for all $n \neq 0$.
- (ADDITIVITY) For any collection of spaces $\{X_\alpha\}_{\alpha \in J}$ with inclusion maps $i^\alpha : X_\alpha \hookrightarrow \coprod_{\beta \in J} X_\beta$, the homomorphisms $i_*^\alpha : h_*(X_\alpha) \rightarrow h_*(\coprod_{\beta \in J} X_\beta)$ determine an isomorphism

$$\bigoplus_{\alpha \in J} h_*(X_\alpha) \rightarrow h_*\left(\prod_{\alpha \in J} X_\alpha\right).$$

We have proved that singular homology $H_*(\cdot; G)$ with coefficient group G satisfies the first four of these axioms, and the proof of the fifth is an easy exercise. If you look again at our computation of $H_*(S^n; \mathbb{Z})$, you'll see that it mostly only used these properties—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 coefficients without knowing anything about the fundamental group. The reason we had to appeal to the fundamental group was that Theorem 25.1 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_*

³²There is some technical terminology from category theory that will be convenient to adopt going forward: we call h_* a **functor** from the category of pairs of spaces to the category of abelian groups. The existence of the induced homomorphisms f_* associated to maps of pairs f is implicit in the use of the word *functor*.

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 ,*

$$h_k(S^n) \cong \begin{cases} G & \text{for } k = 0, n, \\ 0 & \text{for all other } k. \end{cases}$$

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, and we will briefly discuss some of these in *Topologie II*. In the closely related subject of *cohomology*, at least one of the alternative theories (Čech cohomology) is still frequently used, especially in algebraic geometry. Another variant arises frequently in differential geometry: *de Rham* cohomology is defined on smooth manifolds in terms of differential forms, and strictly speaking it is not a cohomology theory in the sense of Eilenberg-Steenrod because it requires a smooth structure, but it satisfies smooth versions of a similar set of axioms. One of the big results we will prove next semester implies that for most spaces that arise in practice, the choice of which axiomatic (co-)homology theory to use will not matter—the results will all be isomorphic. A corollary of the same theorem is that the *simplicial* homology of a triangulated space does not depend on the choice of triangulation, because it is isomorphic to every axiomatic homology theory on that space, in particular singular homology.

I want to quickly sketch the statement of the latter theorem and mention one of its more impressive corollaries: the invariance of the *Euler characteristic*.

One class of “reasonable” spaces we have encountered before are *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—roughly speaking, they are the class of topological spaces for which the standard invariants in algebraic topology are computable. 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 14.7), 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 construction in Theorem 14.7 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 25.6. 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 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 $\partial\mathbb{D}^m \rightarrow X^{m-1}$.

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 are all embeddings. In general, the attaching maps in a cell complex do not need to be embeddings, they only must be continuous, so while the m -cells e_α^m look like open m -disks, their closures 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 25.7. 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 25.8. The cell decomposition of S^n in Example 25.7 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 $C_0^{\text{CW}}(S^n; G) \cong G$ and $C_n^{\text{CW}}(S^n; G)$ 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 21.10). We will work through a complete proof next semester. The first step in that proof will be the computation of $H_*(S^n)$ that began this lecture.

THEOREM. For any cell complex X and any axiomatic homology theory h_* with coefficient group G , $H_*^{\text{CW}}(X; G) \cong h_*(X)$.

Here is a remarkable application. 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.³³ 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

³³One 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.

spaces, the alternating sum of the dimensions of its homology groups can be computed without computing the homology at all: in fact,

$$(25.4) \quad \sum_{n \in \mathbb{Z}} (-1)^n \dim_{\mathbb{K}} H_n(C_*, \partial) = \sum_{n \in \mathbb{Z}} (-1)^n \dim_{\mathbb{K}} C_n.$$

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 (25.4). Now apply this to the cellular chain complex, in which each $C_n^{\text{CW}}(X; \mathbb{K})$ 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 remarkable fact is that the result of this counting game only depends on the space, and not on our choice of how to decompose it into cells:

THEOREM. *For any finite cell complex X ,*

$$\chi(X) = \sum_{n=0}^{\infty} (-1)^n (\text{the number of } n\text{-cells}).$$

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

$$\chi(S^2) = \dim_{\mathbb{R}} H_0(S^2; \mathbb{R}) - \dim_{\mathbb{R}} H_1(S^2; \mathbb{R}) + \dim_{\mathbb{R}} H_2(S^2; \mathbb{R}) = 1 - 0 + 1 = 2.$$

Go ahead. Try it.

Second semester (Topologie II)

26. 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. We saw two important examples of this last semester: the fundamental group (π_1), and the singular homology groups (H_* , with various choices of coefficient group). It is reasonable to think of these in some sense as “functions” with domains consisting of the collection of all topological spaces (possibly with base point), and targets consisting of the collection of all groups (or in the case of homology, all abelian groups). We have not yet developed the right language to make this notion of a “function” precise, so it is time to do so now.

One reason why π_1 and H_* cannot actually be called “functions” is that their domains, strictly speaking, are not sets (*Mengen*). I encourage you to skip the rest of this paragraph if you are not interested in the finer points of axiomatic set theory or the classic set-theoretic paradoxes. . . but for those who are still reading, let us agree that there is no such thing as the “set of all topological spaces”. Indeed, every set can be made into a topological space by assigning it the discrete topology, so if one can talk about the set of all topological spaces, then one must also be able to talk about the set of all sets, and it is a short step from there to the “set of all sets that do not contain themselves”—at which point we find ourselves jumping off the nearest bridge. The architects of abstract set theory solved this dilemma by coming up with a set of axioms that tell you how to construct new sets from old ones, and insisting that *only* collections of objects that arise from these axioms should be called *sets*. Of course, we do sometimes also need to discuss collections of objects that do not arise from the axioms of set theory, and the collection of all topological spaces is an example. Such collections are generally called (proper) **classes** (*Klassen*), but since I do not wish to go any further into the subtleties of set theory in this course, I shall continue to refer to them via the informal word **collections**. You should just keep in mind that while such things can be defined, they are not considered equivalent to sets and cannot be used for all the same purposes that sets can—in particular, an arbitrary “collection” cannot serve as the domain of a function according to the traditional definitions.

Leaving set theory aside, it also must be observed that π_1 and H_* are not just arbitrary “functions” that associate algebraic objects to topological spaces, but they do so in ways that make the algebraic objects into *topological invariants*. In both cases, this results mainly from the fact that continuous maps of spaces induce homomorphisms between the corresponding fundamental groups or homology groups, implying in particular that homeomorphisms induce group isomorphisms. The notion of a *functor* is meant as a form of abstract packaging for this idea.

DEFINITION 26.1. 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 $\text{Mor}(X, Y)$, whose elements are called the **morphisms from X to Y** (*Morphismen von X nach Y*), such that for each $X \in \text{Ob}_{\mathcal{C}}$ there is a distinguished³⁴ element $\text{Id}_X \in \text{Mor}(X, X)$;
- For each $X, Y, Z \in \text{Ob}_{\mathcal{C}}$, a function

$$(26.1) \quad \text{Mor}(X, Y) \times \text{Mor}(Y, Z) \rightarrow \text{Mor}(X, Z) : (f, g) \mapsto g \circ f$$

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, $f \circ \text{Id} = f = \text{Id} \circ f$.

EXAMPLE 26.2. The category Top has $\text{Ob}_{\text{Top}} = \{\text{topological spaces}\}$ and $\text{Mor}(X, Y) = \{f : X \rightarrow Y \mid f \text{ a continuous map}\}$, with Id_X defined for each space X as the identity map and the function (26.1) defined as the usual composition of maps. (Note that this defines a category since the identity map is always continuous and the composition of two continuous maps is also continuous.)

EXAMPLE 26.3. The category Set has $\text{Ob}_{\text{Set}} = \{\text{sets}\}$ and $\text{Mor}(X, Y) = \{f : X \rightarrow Y\}$, with no requirement on continuity of maps since there is no topology.

EXAMPLE 26.4. The objects of Diff are the smooth finite dimensional manifolds, and its morphisms are smooth maps. (As in Example 26.2, the identity is always smooth and the composition of two smooth maps is smooth.)

EXAMPLE 26.5. The category Grp has $\text{Ob}_{\text{Grp}} = \{\text{groups}\}$ and $\text{Mor}(G, H) = \{\text{homomorphisms } G \rightarrow H\}$ for $G, H \in \text{Ob}_{\text{Grp}}$.

EXAMPLE 26.6. There is a **subcategory** (*Unterkategorie*) Ab of Grp whose objects consist of all *abelian* groups, with morphisms defined the same way as in Grp .

The examples above might give you the impression that in every category, a morphism is just a map that may be required to satisfy some specific properties. But nothing in Definition 26.1 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 26.7. Let Top^h denote the category whose objects are the same as in Top , but with $\text{Mor}(X, Y)$ defined as the set of homotopy classes of continuous maps $X \rightarrow Y$ and $\text{Id}_X \in \text{Mor}(X, X)$ as the homotopy class of the identity map. The function (26.1) 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!)

For an interesting example in which objects are not sets and the function (26.1) has nothing to do with composition of maps, see Problem Set 1 #5(b).

DEFINITION 26.8. In any category, a morphism $f \in \text{Mor}(X, Y)$ is called an **isomorphism** (*Isomorphismus*) if there exists a morphism $f^{-1} \in \text{Mor}(Y, X)$ such that $f^{-1} \circ f = \text{Id}_X$ and $f \circ f^{-1} = \text{Id}_Y$. If an isomorphism exists in $\text{Mor}(X, Y)$, we say that the objects X and Y are **isomorphic** (*isomorph*).

³⁴The 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 $\text{Mor}(X, X)$ to contain a specific element that carries that name; there might in theory exist additional morphisms that have the same properties, but only one is called Id_X .

According to this definition, the word “isomorphism” no longer has a strictly algebraic meaning, but will mean whatever is considered to be the notion of “equivalence” in whichever category we are working with. Let’s run through the list: an isomorphism in **Top** is a homeomorphism, in **Set** it is simply a bijection, in **Diff** a diffeomorphism, and in **Grp** or **Ab** it is the usual notion of group isomorphism. The most interesting case so far is **Top^h**: two objects in **Top^h** are isomorphic if and only if they are homotopy equivalent!

DEFINITION 26.9. Given two categories \mathcal{C} and \mathcal{D} , a **functor** (*Funktor*) $\mathcal{F} : \mathcal{C} \rightarrow \mathcal{D}$ assigns to each $X \in \text{Ob}_{\mathcal{C}}$ an object $\mathcal{F}(X) \in \text{Ob}_{\mathcal{D}}$ and to each $f \in \text{Mor}(X, Y)$ for $X, Y \in \text{Ob}_{\mathcal{C}}$ a morphism $\mathcal{F}(f) \in \text{Mor}(\mathcal{F}(X), \mathcal{F}(Y))$ such that:

- (1) $\mathcal{F}(\text{Id}_X) = \text{Id}_{\mathcal{F}(X)}$ for all $X \in \text{Ob}_{\mathcal{C}}$;
- (2) $\mathcal{F}(f \circ g) = \mathcal{F}(f) \circ \mathcal{F}(g)$ for all $g \in \text{Mor}(X, Y)$ and $f \in \text{Mor}(Y, Z)$, $X, Y, Z \in \text{Ob}_{\mathcal{C}}$.

EXAMPLE 26.10. Denote by **Top_{*}** the category whose objects are the pointed spaces (X, p) , i.e. a topological space X together with a point $p \in X$, and

$$\text{Mor}((X, p), (Y, q)) := \{f : X \rightarrow Y \mid f \text{ continuous and } f(p) = q\}.$$

Then the fundamental group 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 26.11. The fundamental group also defines a functor $\pi_1 : \text{Top}_*^h \rightarrow \text{Grp}$ where **Top_{*}^h** is defined to have the same objects as **Top_{*}**, but with $\text{Mor}((X, p), (Y, q))$ defined as the set of pointed homotopy classes of maps $(X, p) \rightarrow (Y, q)$. (See Theorem 9.12 in Lecture 9 from last semester.)

EXAMPLE 26.12. As we will review within the next few lectures, the singular homology group $H_n(X; G)$ for each integer $n \geq 0$ and any fixed (abelian) coefficient group G defines functors

$$H_n(\cdot; G) : \text{Top} \rightarrow \text{Ab} \quad \text{and} \quad \text{Top}^h \rightarrow \text{Ab}.$$

The latter makes sense due to the fact that the homomorphism $f_* : H_n(X; G) \rightarrow H_n(Y; G)$ induced by a continuous map $f : X \rightarrow Y$ depends only on the homotopy class of f .

We will later encounter several algebraic constructions and related topological invariants that satisfy most of the conditions of a functor but differ in one crucial respect: the morphisms they induce go *the other way*. In practice, this phenomenon often arises from the algebraic notion of dualization, and we’ll give an example of this kind immediately after the definition.

DEFINITION 26.13. Given two categories \mathcal{C} and \mathcal{D} , a **contravariant functor** (*kontravarianter Funktor*) $\mathcal{F} : \mathcal{C} \rightarrow \mathcal{D}$ assigns to each $X \in \text{Ob}_{\mathcal{C}}$ some $\mathcal{F}(X) \in \text{Ob}_{\mathcal{D}}$ and to each $f \in \text{Mor}(X, Y)$ for $X, Y \in \text{Ob}_{\mathcal{C}}$ a morphism $\mathcal{F}(f) \in \text{Mor}(\mathcal{F}(Y), \mathcal{F}(X))$ such that

- (1) $\mathcal{F}(\text{Id}_X) = \text{Id}_{\mathcal{F}(X)}$ for all $X \in \text{Ob}_{\mathcal{C}}$;
- (2) $\mathcal{F}(f \circ g) = \mathcal{F}(g) \circ \mathcal{F}(f)$ for all $g \in \text{Mor}(X, Y)$ and $f \in \text{Mor}(Y, Z)$, $X, Y, Z \in \text{Ob}_{\mathcal{C}}$.

A functor that satisfies the original Definition 26.9 instead of Definition 26.13 can be called **covariant** (*kovariant*) when we want to emphasize that it is not contravariant.

EXAMPLE 26.14. Let $\text{Vec}_{\mathbb{K}}$ denote the category of vector spaces over a fixed field \mathbb{K} , so $\text{Mor}(V, W) := \text{Hom}_{\mathbb{K}}(V, W)$ is the space of \mathbb{K} -linear maps $V \rightarrow W$. There is a contravariant functor $\text{Vec}_{\mathbb{K}} \rightarrow \text{Vec}_{\mathbb{K}}$ which sends each vector space V to its dual space $V^* := \text{Hom}_{\mathbb{K}}(V, \mathbb{K})$ and sends each morphism $A : V \rightarrow W$ to its transpose $A^* : W^* \rightarrow V^*$, 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 $V^* \rightarrow V^*$.

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 26.15. 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 $X \in \text{Ob}_{\mathcal{C}}$ a morphism $T_X \in \text{Mor}(\mathcal{F}(X), \mathcal{G}(X))$ such that for all $X, Y \in \text{Ob}_{\mathcal{C}}$ and $f \in \text{Mor}(X, Y)$, the following diagram commutes:

$$\begin{array}{ccc} \mathcal{F}(X) & \xrightarrow{T_X} & \mathcal{G}(X) \\ \downarrow \mathcal{F}(f) & & \downarrow \mathcal{G}(f) \\ \mathcal{F}(Y) & \xrightarrow{T_Y} & \mathcal{G}(Y) \end{array}$$

A natural transformation of contravariant functors can be defined analogously.

EXAMPLE 26.16. Let us abbreviate the singular homology of X with integer coefficients by $H_1(X) := H_1(X; \mathbb{Z})$. Recall that for any pointed space (X, p) , the *Hurewicz homomorphism*

$$h : \pi_1(X, p) \rightarrow H_1(X)$$

sends each equivalence class $[\gamma] \in \pi_1(X, p)$ of paths $\gamma : I \rightarrow X$ starting and ending at p to the homology class represented by the singular 1-cycle $\gamma : \Delta^1 \rightarrow X$ defined by identifying $I = [0, 1]$ with the standard 1-simplex Δ^1 . (See in particular Theorem 22.13 in last semester's Lecture 22.) To view this as a natural transformation, we have to adjust our perspective slightly so that $\pi_1 : \text{Top}_* \rightarrow \text{Grp}$ and $H_1 : \text{Top} \rightarrow \text{Ab}$ become functors between the same pairs of categories: and indeed, we can view H_1 as a functor from Top_* to Grp that simply ignores the base point in each object (X, p) of Top_* and also ignores the fact that the group $H_1(X)$ always happens to be abelian. With this understood, one can show that the Hurewicz homomorphism defines a natural transformation from π_1 to H_1 ; see Problem Set 1 #3.

EXAMPLE 26.17. For any object V in the category $\text{Vec}_{\mathbb{K}}$ of Example 26.14, there is a natural injective homomorphism $\Phi : V \rightarrow (V^*)^*$ from V to the dual of its dual space, defined by

$$\Phi(v)\lambda = \lambda(v) \quad \text{for } v \in V, \lambda \in V^*.$$

The precise meaning of the word "natural" in this case is as follows: there exists a functor $\Delta^2 : \text{Vec}_{\mathbb{K}} \rightarrow \text{Vec}_{\mathbb{K}}$ sending each object V to $(V^*)^*$, and a natural transformation from the "identity functor" to Δ^2 that associates to each object V the morphism $\Phi : V \rightarrow (V^*)^*$. (There are several things to check here: see Problem Set 1 #6.)

REMARK 26.18. Whenever a vector space V is finite dimensional, the map $\Phi : V \rightarrow (V^*)^*$ in Example 26.17 is an isomorphism, and a large part of the reason why it turns out to define a *natural transformation* is that the definition of Φ does not depend on any choices. Contrast this with Example 26.14: restricting again to finite-dimensional spaces, V and V^* are always isomorphic since they have the same dimension, but writing down such an isomorphism requires a choice—it is not canonical. One can show in fact that there does not exist any natural transformation from the identity functor (regarded now as a *contravariant* functor) to the dualization functor of Example 26.14. While this statement sounds purely algebraic, the easiest proof I can think of is topological: it uses vector bundles. (Ignore the rest of this remark if you've never seen vector bundles before). Indeed, if such a natural transformation existed, it would imply that every finite-dimensional vector bundle is isomorphic to its dual bundle. But e.g. for $\mathbb{K} = \mathbb{C}$, there are simple examples of complex vector bundles that are not isomorphic to their dual bundles—one can prove it using Chern classes.

Let us now review the definition of singular homology in this category-theoretic context. For now we will consider only homology with integer coefficients and thus leave the coefficient group out of the notation, writing e.g. $H_*(X)$ instead of $H_*(X; \mathbb{Z})$. It is natural to view singular homology as the composition of *two* covariant functors: one that transforms topological information into algebra (in the form of a chain complex), and another that performs the purely algebraic step of replacing a chain complex with its homology groups. Let us define the algebraic functor first.

DEFINITION 26.19. A \mathbb{Z} -graded abelian group (\mathbb{Z} -graduierte abelsche Gruppe) G_* is an abelian group that is equipped with a direct sum splitting $G_* = \bigoplus_{n \in \mathbb{Z}} G_n$, i.e. G_n for each $n \in \mathbb{Z}$ is a subgroup, $G_n \cap G_m = \{0\}$ whenever $n \neq m$, and every element of G_* can be written as a finite sum of elements that each belong to the various subgroups G_n . An element $g \in G_*$ such that $g \in G_n$ for some $n \in \mathbb{Z}$ is called a **homogeneous** (*homogen*) element of **degree** (*Grad*) n . Let $\text{Ab}_{\mathbb{Z}}$ denote the category whose objects are \mathbb{Z} -graded abelian groups, with morphisms from G_* to H_* defined as group homomorphisms that send G_n into H_n for every $n \in \mathbb{Z}$.

DEFINITION 26.20. A **chain complex** (*Kettenkomplex*) of abelian groups is a \mathbb{Z} -graded abelian group C_* equipped with the additional structure of a homomorphism $\partial : C_* \rightarrow C_*$ that satisfies $\partial(C_n) \subset C_{n-1}$ for every $n \in \mathbb{Z}$ and $\partial^2 = 0$. Given two chain complexes (A_*, ∂^A) and (B_*, ∂^B) , a **chain map** (*Kettenabbildung*) from (A_*, ∂^A) to (B_*, ∂^B) is a morphism $\Phi : A_* \rightarrow B_*$ in the sense of Definition 26.19 such that $\Phi \circ \partial^A = \partial^B \circ \Phi$. Let Chain denote the category that has chain complexes as objects and chain maps as morphisms. (Notice in particular that the identity map is always a chain map, and the composition of two chain maps is also a chain map.)

DEFINITION 26.21. The **homology** (*Homologie*) of a chain complex (C_*, ∂) is the graded abelian group $H_*(C_*, \partial) = \bigoplus_{n \in \mathbb{Z}} H_n(C_*, \partial)$ where

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

with ∂_n denoting the restriction of $\partial : C_* \rightarrow C_*$ to $C_n \rightarrow C_{n-1}$.

PROPOSITION 26.22. *There is a functor $H_* : \text{Chain} \rightarrow \text{Ab}_{\mathbb{Z}}$ that assigns to each chain complex (C_*, ∂) its homology $H_*(C_*, \partial)$ and assigns to each chain map $\Phi : (A_*, \partial^A) \rightarrow (B_*, \partial^B)$ the homomorphism*

$$H_*(A_*, \partial^A) \rightarrow H_*(B_*, \partial^B) : [a] \mapsto [\Phi(a)].$$

PROOF. Straightforward exercise; see Problem Set 1 #1(a). \square

Recall that a **chain homotopy** (*Kettenhomotopie*) between two chain maps $\Phi, \Psi : (A_*, \partial^A) \rightarrow (B_*, \partial^B)$ is a homomorphism $h : A_* \rightarrow B_*$ such that $h(A_n) \subset B_{n+1}$ for each $n \in \mathbb{Z}$ and

$$\partial^B \circ h + h \circ \partial^A = \Phi - \Psi,$$

and we say that Φ and Ψ are **chain homotopic** (*kettenhomotop*) whenever a chain map exists. It is not too hard to show that chain homotopy defines an equivalence relation for chain maps, and the notion of composition for chain maps descends to a well-defined composition of *chain homotopy classes* of chain maps, which is why the following definition makes sense:

DEFINITION 26.23. Let Chain^h denote the category whose objects are chain complexes and whose morphisms are chain homotopy classes of chain maps.

It is an easy algebraic exercise to show that whenever two chain maps $\Phi, \Psi : (A_*, \partial^A) \rightarrow (B_*, \partial^B)$ are chain homotopic, the homomorphisms $H_*(A_*, \partial^A) \rightarrow H_*(B_*, \partial^B)$ that they induce are the same. Proposition 26.22 thus extends as follows:

PROPOSITION 26.24. *The prescription of Proposition 26.22 also defines a functor $\text{Chain}^h \rightarrow \text{Ab}_{\mathbb{Z}}$.* \square

So much for algebra; now back to topology.

DEFINITION 26.25. A **pair of spaces** (*topologisches Paar*) (X, A) consists of a topological space X together with a subset $A \subset X$. A **map of pairs** (*Abbildung von Paaren*) $f : (X, A) \rightarrow (Y, B)$ is a continuous map $f : X \rightarrow Y$ such that $f(A) \subset B$. Two such maps $f, g : (X, A) \rightarrow (Y, B)$ are **homotopic** (*homotop*) if there exists a continuous map $h : I \times X \rightarrow Y$ such that $h(0, \cdot) = f$, $h(1, \cdot) = g$, and $h(I \times A) \subset B$. Let Top_{rel} denote the category whose objects are pairs of spaces and whose morphisms are maps of pairs. Similarly, $\text{Top}_{\text{rel}}^h$ will denote the category with the same objects, but whose morphisms are homotopy classes of maps of pairs.

DEFINITION 26.26. For any integer $n \geq 0$, define the **standard n -simplex** (*Standard n -Simplex*) as the set

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

For $k = 0, \dots, n$, the **k th boundary face** (*k -te Seitenfläche*) of Δ^n is the subset

$$\partial_{(k)} \Delta^n = \{t_k = 0\} \subset \Delta^n,$$

which we will sometimes identify with Δ^{n-1} via the obvious bijection

$$(26.2) \quad \partial_{(k)} \Delta^n \rightarrow \Delta^{n-1} : (t_0, \dots, t_{k-1}, 0, t_{k+1}, \dots, t_n) \mapsto (t_0, \dots, t_{k-1}, t_{k+1}, \dots, t_n).$$

A **singular n -simplex** (*singulärer n -Simplex*) in a space X is defined to be a continuous map $\sigma : \Delta^n \rightarrow X$. Let $\mathcal{K}_n(X)$ denote the set of all singular n -simplices in X .

We can now state the prescription for turning a pair of spaces (X, A) into its **singular chain complex** (*singulärer Kettenkomplex*) $C_*(X, A)$. First define $C_n(X) = 0$ for all $n < 0$, while for $n \geq 0$, define $C_n(X)$ to be the free abelian group generated by the set $\mathcal{K}_n(X)$, i.e.

$$C_n(X) = \bigoplus_{\sigma \in \mathcal{K}_n(X)} \mathbb{Z}.$$

Elements of $C_n(X)$ can be written as finite sums $\sum_j m_j \sigma_j$ where $m_j \in \mathbb{Z}$ and $\sigma_j \in \mathcal{K}_n(X)$. Since the group is freely generated, one can define a homomorphism $\partial : C_n(X) \rightarrow C_{n-1}(X)$ by specifying what it does to each of the generators $\sigma \in \mathcal{K}_n(X)$: the prescription is

$$\partial \sigma = \sum_{k=0}^n (-1)^k (\sigma|_{\partial_{(k)} \Delta^n}),$$

where the identification (26.2) is used in order to view this as an element of $C_{n-1}(X)$. Defining ∂ in this way on $C_n(X)$ for every $n > 0$ and $\partial = 0$ on $C_0(X)$, it is now straightforward to verify that $\partial^2 = 0$, so ∂ equips the graded abelian group $C_*(X) = \bigoplus_{n \in \mathbb{Z}} C_n(X)$ with the structure of a chain complex.

Finally, for any pair (X, A) , one can assign the subspace topology to A and define $C_*(A)$ as above, with the consequence that $C_*(A)$ becomes a subgroup of $C_*(X)$ that is preserved by ∂ , hence the latter descends to the quotients

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

and thus endows $C_*(X, A) := \bigoplus_{n \in \mathbb{Z}} C_n(X, A)$ with the structure of a chain complex. Algebraically, $C_n(X, A)$ is still quite simple: one can identify it with the free abelian group generated by the set of all singular n -simplices in X that are not fully contained in A .

REMARK 26.27. If $A = \emptyset$, then $C_*(X, A)$ is the same thing as $C_*(X)$. It is often convenient to think of Top as the subcategory of Top_{rel} whose objects are all of the form (X, \emptyset) .

In order to view C_* as a functor from $\mathbf{Top}_{\text{rel}}$ to \mathbf{Chain} , we need to explain what it does to morphisms. The answer is again straightforward: for any continuous map $f : X \rightarrow Y$, there is a chain map $f_* : C_*(X) \rightarrow C_*(Y)$ whose action on the generators $\sigma \in \mathcal{K}_n(X)$ is

$$f_*\sigma := f \circ \sigma \in \mathcal{K}_n(Y).$$

If f is also a map of pairs $(X, A) \rightarrow (Y, B)$, then the chain map $C_*(X) \rightarrow C_*(Y)$ sends $C_*(A)$ into $C_*(B)$ and thus descends to a relative chain map $f_* : C_*(X, A) \rightarrow C_*(Y, B)$. The subtler part of the story involves homotopies: as we will review in the next lecture, any homotopy h between maps of pairs $f, g : (X, A) \rightarrow (Y, B)$ induces a chain homotopy h_* between the corresponding chain maps $f_*, g_* : C_*(X, A) \rightarrow C_*(Y, B)$. Putting all of this together gives the following:

PROPOSITION 26.28. *There exist functors $\mathbf{Top}_{\text{rel}} \rightarrow \mathbf{Chain}$ and $\mathbf{Top}_{\text{rel}}^h \rightarrow \mathbf{Chain}^h$ that assign to each pair (X, A) its relative singular chain complex $C_*(X, A)$ and associate to each map of pairs $f : (X, A) \rightarrow (Y, B)$ (or the homotopy class thereof) the induced chain map $f_* : C_*(X, A) \rightarrow C_*(Y, B)$ (or the chain homotopy class thereof).*

The **relative singular homology** (relative singuläre Homologie) of the pair (X, A) is defined as the homology of the chain complex $C_*(X, A)$ and is denoted by

$$H_*(X, A) = \bigoplus_{n \in \mathbb{Z}} H_n(X, A).$$

If $A = \emptyset$ as in Remark 26.27, then we abbreviate it as $H_*(X)$ and call it the **absolute** singular homology (absolute singuläre Homologie) of X . Composing the functors of Propositions 26.22 or 26.24 with those in Proposition 26.28 now gives:

THEOREM 26.29. H_* defines functors $\mathbf{Top}_{\text{rel}} \rightarrow \mathbf{Ab}_{\mathbb{Z}}$ and $\mathbf{Top}_{\text{rel}}^h \rightarrow \mathbf{Ab}_{\mathbb{Z}}$. □

27. Properties of singular homology

In this lecture we shall begin a review of the main properties of singular homology that were established in the last few lectures of *Topologie I*. Let us start by bringing more general coefficient groups back into the picture.

First an algebraic remark: every abelian group G is, in a canonical way, also a module over \mathbb{Z} . Indeed, the obvious multiplication operation $\mathbb{Z} \times G \rightarrow G$ can be defined for positive integers m by

$$mg := \underbrace{g + \dots + g}_m.$$

Combining this with $0g := 0$ and $(-1)g := -g$ gives a definition of mg that makes sense for every $m \in \mathbb{Z}$.

The chain groups $C_n(X, A)$ defined in the previous lecture consist of finite sums $\sum_i m_i \sigma_i$ where each m_i is an integer and each $\sigma_i \in \mathcal{K}_n(X)$ is a singular n -simplex in X whose image is not contained in A . For any abelian group G , one can generalize this to

$$C_n(X, A; G) = C_n(X; G) / C_n(A; G), \quad \text{where} \quad C_n(X; G) = \bigoplus_{\sigma \in \mathcal{K}_n(X)} G.$$

Elements of $C_n(X, A; G)$ can thus be written as finite sums $\sum_i m_i \sigma_i$ where the coefficients m_i now belong to the group G , and addition in $C_n(X, A; G)$ is defined via the obvious relation $k\sigma + m\sigma := (k + m)\sigma$. The natural generalization of $\partial : C_n(X) \rightarrow C_{n-1}(X)$ to a map $C_n(X, A; G) \rightarrow C_{n-1}(X, A; G)$ is determined by the formula

$$\partial(g\sigma) := \sum_{k=0}^n (-1)^k g (\sigma|_{\partial_{(k)} \Delta^n}) \in C_{n-1}(X, A; G), \quad \text{for} \quad g \in G, \sigma \in \mathcal{K}_n(X),$$

which must be understood with the caveat that any term for which $\sigma|_{\partial_{(k)}\Delta^n}$ has image contained in A gets dropped from the sum, as it vanishes in the quotient $C_{n-1}(X;G)/C_{n-1}(A;G)$. For $G = \mathbb{Z}$, we recover the original chain complex $C_*(X, A; \mathbb{Z}) = C_*(X, A)$.

EXAMPLE 27.1. Since every element of \mathbb{Z}_2 is its own inverse, the formula for $\partial\sigma$ on $C_n(X; \mathbb{Z}_2)$ simplifies to $\partial\sigma = \sum_{k=0}^n \sigma|_{\partial_{(k)}\Delta^n}$.

The transformation of $C_*(X, A)$ into $C_*(X, A; G)$ can be reframed in terms of another purely algebraic functor, and this will be a useful perspective going forward. We first need to recall a few standard notions from the theory of abelian groups.

Given a set S , the **free abelian group** (*freie abelsche Gruppe*) on S is defined as a direct sum of copies of \mathbb{Z} , one for each element of S :

$$F^{\text{ab}}(S) := \bigoplus_{s \in S} \mathbb{Z}.$$

We can write elements of $F^{\text{ab}}(S)$ as finite sums $\sum_i m_i s_i$ for $m_i \in \mathbb{Z}$ and the **generators** (*Erzeuger*) $s_i \in S$, with the addition operation determined by $ks + ms := (k+m)s$ for any $k, m \in \mathbb{Z}$ and $s \in S$.

EXERCISE 27.2.

- Show that for any abelian group H , set S , and map $f : S \rightarrow H$, there exists a unique homomorphism $\Phi : F^{\text{ab}}(S) \rightarrow H$ such that $\Phi(s) = f(s)$ for each of the generators $s \in S$.
- Show that there is a natural isomorphism between $F^{\text{ab}}(S)$ and the abelianization of the free (non-abelian) group $F(S)$.

Given abelian groups G, H, K , a map $\Phi : G \oplus H \rightarrow K$ is called **bilinear** if for every fixed $g_0 \in G$ and $h_0 \in H$, the maps $G \rightarrow K : g \mapsto \Phi(g, h_0)$ and $H \rightarrow K : h \mapsto \Phi(g_0, h)$ are both homomorphisms.

The **tensor product** (*Tensorprodukt*) of two abelian groups G and H can be defined as the abelian group

$$G \otimes H := F^{\text{ab}}(G \times H)/N$$

where $N \subset F^{\text{ab}}(G \times H)$ is the smallest subgroup containing all elements of the form $(g + g', h) - (g, h) - (g', h)$ and $(g, h + h') - (g, h) - (g, h')$ for $g, g' \in G$ and $h, h' \in H$. We denote the equivalence class represented by $(g, h) \in F^{\text{ab}}(G \times H)$ in the quotient by

$$g \otimes h \in G \otimes H.$$

EXERCISE 27.3.

- 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 $h \in H$, $0 \otimes h = g \otimes 0 = 0 \in G \otimes H$.
- Show that for any bilinear map $\Phi : G \oplus H \rightarrow K$ of abelian groups, there exists a unique homomorphism $\Psi : G \otimes H \rightarrow K$ such that $\Phi(g, h) = \Psi(g \otimes h)$ for all $(g, h) \in G \oplus H$.
- Show that for any abelian group G , the map $G \rightarrow G \otimes \mathbb{Z} : g \mapsto g \otimes 1$ is a group isomorphism. Write down its inverse.
Hint: Use part (b) to write down homomorphisms in terms of bilinear maps.
- Find a natural isomorphism from $(G \oplus H) \otimes K$ to $(G \otimes K) \oplus (H \otimes K)$.
- Given two sets S and T , find a natural isomorphism from $F^{\text{ab}}(S) \otimes F^{\text{ab}}(T)$ to $F^{\text{ab}}(S \times T)$.
- Let \mathbb{K} be a field, regarded as an abelian group with respect to its addition operation. Show that the abelian group $G \otimes \mathbb{K}$ naturally admits the structure of a vector space over \mathbb{K} such that scalar multiplication takes the form

$$\lambda(g \otimes k) = g \otimes (\lambda k)$$

for every $\lambda, k \in \mathbb{K}$ and $g \in G$, and every group homomorphism $\Phi : G \rightarrow H$ determines a unique \mathbb{K} -linear map $\Psi : G \otimes \mathbb{K} \rightarrow H \otimes \mathbb{K}$ such that $\Psi(g \otimes k) = \Phi(g) \otimes k$ for $g \in G, k \in \mathbb{K}$.

- (g) For any abelian groups A, B, C, D and homomorphisms $f : A \rightarrow B, g : C \rightarrow D$, 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 $c \in C$.

- (h) An element $a \in G$ is said to be **torsion** if $ma = 0$ for some $m \in \mathbb{Z}$. Show that if every element of G is torsion and \mathbb{K} is a field (regarded as an abelian group with respect to addition), then $G \otimes \mathbb{K} = 0$.

The proof of the following result should now be an easy exercise.

PROPOSITION 27.4. *For any fixed abelian group G , there is a covariant functor*

$$\otimes G : \mathbf{Ab} \rightarrow \mathbf{Ab}$$

that sends each abelian group A to $A \otimes G$ and sends each group homomorphism $\Phi : A \rightarrow B$ to $\Phi \otimes \mathbf{1} : A \otimes G \rightarrow B \otimes G$.

Similarly, $\otimes G$ defines functors

$$\mathbf{Ab}_{\mathbb{Z}} \rightarrow \mathbf{Ab}_{\mathbb{Z}}, \quad \mathbf{Chain} \rightarrow \mathbf{Chain}, \quad \text{and} \quad \mathbf{Chain}^h \rightarrow \mathbf{Chain}^h$$

which send \mathbb{Z} -graded abelian groups $C_* = \bigoplus_{n \in \mathbb{Z}} C_n$ to $C_* \otimes G = \bigoplus_{n \in \mathbb{Z}} C_n \otimes G$ and morphisms $\Phi : C_* \rightarrow C_*$ to $\Phi \otimes \mathbf{1} : C_* \otimes G \rightarrow C_* \otimes G$. For a chain complex (C_*, ∂) , the boundary map on $C_* \otimes G$ is defined as $\partial \otimes \mathbf{1}$.

EXERCISE 27.5. Fill in the details of the proof of Proposition 27.4. In particular, check that for any chain map $\Phi : (A_*, \partial^A) \rightarrow (B_*, \partial^B)$, the chain homotopy class of $\Phi \otimes \mathbf{1} : (A_* \otimes G, \partial^A \otimes \mathbf{1}) \rightarrow (B_* \otimes G, \partial^B \otimes \mathbf{1})$ depends only on the chain homotopy class of Φ , so that the functor $\otimes G : \mathbf{Chain}^h \rightarrow \mathbf{Chain}^h$ is well defined.

Hint: If $h_ : A_* \rightarrow B_*$ is a chain homotopy, what can you say about $h_* \otimes \mathbf{1} : A_* \otimes G \rightarrow B_* \otimes G$?*

EXERCISE 27.6. Show that there is a canonical identification between $C_*(X, A; G)$ as defined above and the tensor product chain complex $C_*(X, A) \otimes G$ arising from Proposition 27.4.

In light of this discussion, the **singular homology with coefficients in G**

$$H_*(X, A; G) := H_*(C_*(X, A; G))$$

can now be understood as a composition of *three* covariant functors: first from $\mathbf{Top}_{\text{rel}}$ to \mathbf{Chain} to construct the singular chain complex with \mathbb{Z} -coefficients, then $\otimes G : \mathbf{Chain} \rightarrow \mathbf{Chain}$ to introduce the coefficient group G , and finally the homology functor $H_* : \mathbf{Chain} \rightarrow \mathbf{Ab}_{\mathbb{Z}}$. If we prefer to emphasize homotopy invariance, we can also view this as a composition of three functors $\mathbf{Top}_{\text{rel}}^h \rightarrow \mathbf{Chain}^h \rightarrow \mathbf{Chain}^h \rightarrow \mathbf{Ab}_{\mathbb{Z}}$. The case $A = \emptyset$ will be abbreviated as usual by

$$H_*(X; G) := H_*(X, \emptyset; G),$$

and we will sometimes also drop the coefficient group G from the notation if it is assumed to be \mathbb{Z} or is otherwise clear from context or plays no significant role in the discussion. We will usually continue to denote elements of $C_*(X, A; G)$ by $\sum_i m_i \sigma_i$ instead of $\sum_i \sigma_i \otimes m_i \in C_*(X, A) \otimes G$, but the tensor product perspective will also serve us well, as it can often be used to turn results about homology with integer coefficients into results for general coefficient groups with almost no extra effort.

Let's now run through the list of properties of $H_*(X, A; G)$ that were already proved at the end of *Topologie I*.

Path-components and H_0 (cf. Proposition 22.11). For any space X , 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 comes from the obvious identification between singular 0-simplices (all of which are cycles, i.e. they satisfy $\partial\sigma = 0$) and points in X . Cycles of the form gx and gy for $g \in G$ and $x, y \in X$ represent the same homology class if and only if x and y are in the same path-component, since paths $\gamma : I \rightarrow X$ from x to y can be regarded as singular 1-simplices by identifying $I = [0, 1]$ with Δ^1 , giving $\partial(g\gamma) = gx - gy$.

The fundamental group and H_1 (cf. Theorem 22.13). As mentioned in Example 26.16, the identification $I = \Delta^1$ also gives rise to a natural homomorphism

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

sending the homotopy class of the loop $\gamma : I \rightarrow X$ to the homology class represented by the singular 1-cycle $\gamma : \Delta^1 \rightarrow X$. This map vanishes on the commutator subgroup of $\pi_1(X, p)$ since $H_1(X; \mathbb{Z})$ is abelian, so it descends to a map of the abelianization of $\pi_1(X, p)$ to $H_1(X; \mathbb{Z})$. If X is path-connected, then the latter is an isomorphism, hence

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

The map h is called the **Hurewicz homomorphism**.

I put those two properties first in the list because they are somewhat special: their proofs depend heavily on the precise definition of singular homology in terms of singular simplices. We will later discuss some other functors that are defined very differently but for various reasons deserve to be called “homology theories”: in general they do not have the two properties just mentioned, but they will have the rest of the properties on this list.

The homology of a one-point space (cf. Exercise 22.12). For a space $\{\text{pt}\}$ consisting of only one point, we have

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

This is a straightforward consequence of the definition, using the observation that for each $n \geq 0$, the set $\mathcal{K}_n(X)$ of singular n -simplices in $\{\text{pt}\}$ has only one element, hence $C_n(X; G)$ is naturally isomorphic to G .

Disjoint unions. We didn’t prove this last semester, but it’s a straightforward exercise using the definitions: for any collection of spaces $\{X_\alpha\}_{\alpha \in J}$ with inclusion maps $i^\alpha : X_\alpha \hookrightarrow \coprod_{\beta \in J} X_\beta$, the induced homomorphisms

$$i_*^\alpha : H_*(X_\alpha; G) \rightarrow H_*\left(\coprod_{\alpha \in J} X_\beta; G\right)$$

determine an isomorphism

$$\bigoplus_{\alpha \in J} i_*^\alpha : \bigoplus_{\alpha \in J} H_*(X_\alpha; G) \rightarrow H_*\left(\coprod_{\alpha \in J} X_\beta; G\right).$$

The long exact sequence of a pair (cf. Lecture 23). Recall that a sequence of abelian groups and homomorphisms

$$\dots \longrightarrow G_{n-1} \xrightarrow{f_{n-1}} G_n \xrightarrow{f_n} G_{n+1} \longrightarrow \dots$$

is called **exact** if $\text{im } f_{n-1} = \ker f_n$ for every n . An exact sequence of the form

$$0 \longrightarrow G_1 \longrightarrow G_2 \longrightarrow G_3 \longrightarrow 0$$

is called a **short exact sequence** (*kurze exakte Sequenz*), and one can equally well consider an exact sequence of chain complexes, in which the homomorphisms are all assumed to be chain maps. An obvious example is

$$(27.1) \quad 0 \longrightarrow C_*(A; G) \xrightarrow{i_*} C_*(X; G) \xrightarrow{j_*} C_*(X, A; G) \longrightarrow 0$$

for any pair of spaces (X, A) , where i_* and j_* are the chain maps induced by the natural inclusions $i : A \hookrightarrow X$ and $j : (X, \emptyset) \hookrightarrow (X, A)$. Notice that, in purely algebraic terms, i_* is also the inclusion of the subgroup $C_*(A; G) \hookrightarrow C_*(X; G)$, and j_* is the quotient projection $C_*(X; G) \rightarrow C_*(X, A; G) = C_*(X; G)/C_*(A; G)$. This short exact sequence can then be plugged into the following purely algebraic lemma:

PROPOSITION 27.7 (cf. Theorem 23.6). *Suppose $0 \rightarrow A_* \xrightarrow{f} B_* \xrightarrow{g} C_* \rightarrow 0$ is a short exact sequence of chain complexes. Then for each $n \in \mathbb{Z}$ there exists a so-called **connecting homomorphism** $\partial_* : H_n(C_*) \rightarrow H_{n-1}(A_*)$ such that the sequence*

$$\begin{aligned} \dots \xrightarrow{\partial_*} H_{n+1}(A_*) \xrightarrow{f_*} H_{n+1}(B_*) \xrightarrow{g_*} H_{n+1}(C_*) \\ \xrightarrow{\partial_*} H_n(A_*) \xrightarrow{f_*} H_n(B_*) \xrightarrow{g_*} H_n(C_*) \\ \xrightarrow{\partial_*} H_{n-1}(A_*) \xrightarrow{f_*} H_{n-1}(B_*) \xrightarrow{g_*} H_{n-1}(C_*) \xrightarrow{\partial_*} \dots \end{aligned}$$

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

$$\begin{array}{cccccccccccc} \dots & \longrightarrow & H_{n+1}(C_*) & \xrightarrow{\partial_*} & H_n(A_*) & \xrightarrow{f_*} & H_n(B_*) & \xrightarrow{g_*} & H_n(C_*) & \xrightarrow{\partial_*} & H_{n-1}(A_*) & \longrightarrow & \dots \\ & & \downarrow \gamma_* & & \downarrow \alpha_* & & \downarrow \beta_* & & \downarrow \gamma_* & & \downarrow \alpha_* & & \\ \dots & \longrightarrow & H_{n+1}(C'_*) & \xrightarrow{\partial'_*} & H_n(A'_*) & \xrightarrow{f'_*} & H_n(B'_*) & \xrightarrow{g'_*} & H_n(C'_*) & \xrightarrow{\partial'_*} & H_{n-1}(A'_*) & \longrightarrow & \dots \end{array}$$

also commutes.

The proof of this result uses the standard method known as “diagram chasing”. Let’s do the first step, which is to write down a reasonable candidate for the map $\partial_* : H_n(C_*) \rightarrow H_{n-1}(A_*)$.

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.

We did not mention the “functoriality” aspect of the above statement last semester, but once you've worked out how to define ∂_* , it will not be hard to see why it is true. This amounts to the statement that there exist natural definitions of categories whose objects are short exact sequences of chain maps or long exact sequences of \mathbb{Z} -graded abelian groups, with morphisms defined in each case via commutative diagrams, such that Proposition 27.7 produces a functor from the former category to the latter. See Problem Set 2 for the precise formulation in these terms.

Applying Proposition 27.7 to the short exact sequence (27.1) yields:

THEOREM 27.8. *For every pair of spaces (X, A) , abelian group G and $n \in \mathbb{Z}$, there exists a natural transformation ∂_* from the functor $(X, A) \mapsto H_n(X, A; G)$ to the functor $(X, A) \mapsto H_{n-1}(A; G)$, both regarded as functors $\text{Top}_{\text{rel}} \rightarrow \text{Ab}$, such that the sequence*

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

is exact, where i_* and j_* are induced by the inclusions $i : A \hookrightarrow X$ and $j : (X, \emptyset) \hookrightarrow (X, A)$. □

The fact that ∂_* is a natural transformation concretely means the following: if $f : (X, A) \rightarrow (Y, B)$ is any map of pairs, then the connecting homomorphisms for both pairs fit into the commutative diagram

$$\begin{array}{ccc} H_n(X, A; G) & \xrightarrow{\partial_*} & H_{n-1}(A; G) \\ \downarrow f_* & & \downarrow f_* \\ H_n(Y, B; G) & \xrightarrow{\partial_*} & H_{n-1}(B; G) \end{array}$$

This follows from the functoriality in Proposition 27.7, as f_* also induces a commutative diagram of chain maps

$$\begin{array}{ccccccc} 0 & \longrightarrow & C_*(A; G) & \xrightarrow{i_*} & C_*(X; G) & \xrightarrow{j_*} & C_*(X, A; G) \longrightarrow 0 \\ & & \downarrow f_* & & \downarrow f_* & & \downarrow f_* \\ 0 & \longrightarrow & C_*(B; G) & \xrightarrow{i_*} & C_*(Y; G) & \xrightarrow{j_*} & C_*(Y, B; G) \longrightarrow 0 \end{array}$$

where the rows are simply the short exact sequences of (X, A) and (Y, B) respectively.

One last comment about the connecting homomorphisms $\partial_* : H_n(X, A; G) \rightarrow H_{n-1}(A; G)$. In the above discussion we deduced their existence from an algebraic result, but it is also not hard to write them down with an explicit formula. To express it properly, recall that any relative homology class $[c] \in H_n(X, A; G)$ can be represented by some singular n -chain $c = \sum_i m_i \sigma_i \in C_n(X; G)$, i.e. this is a choice of representative for some element of the quotient $C_n(X, A; G) = C_n(X; G)/C_n(A; G)$, and the fact that that element is a cycle (i.e. is in $\ker \partial$) translates into the condition that ∂c must be an $(n - 1)$ -chain contained in A , i.e. $\partial c \in C_{n-1}(A; G)$. Any n -chain c with this property is called a **relative n -cycle** in (X, A) . But ∂c is manifestly also a cycle in A since $\partial^2 = 0$, so it represents a homology class, and in this way we obtain the simplest possible formula for $\partial_* : H_n(X, A; G) \rightarrow H_{n-1}(A; G)$, namely

$$(27.2) \quad \partial[c] = [\partial c].$$

If you rederive ∂_* from the diagram chase in Proposition 27.7 for the short exact sequence of the pair (X, A) , you'll find that this is what it produces. The simplicity of the formula is deceptive: it looks like the right hand side should be trivial since it is the homology class of a boundary, but you need to keep in mind that while this is understood as a homology class in $H_{n-1}(A; G)$, c is not a chain in A , but in the larger space X .

Homotopy invariance (cf. Lemma 24.1 and Corollary 24.2). It was already mentioned in the last lecture that the homomorphism $f_* : H_*(X, A; G) \rightarrow H_*(Y, B; G)$ induced by any map of pairs $f : (X, A) \rightarrow (Y, B)$ depends only on the homotopy class of f (as a map of pairs), with the important consequence that $H_*(X, A; G)$ up to isomorphism depends only on the homotopy type of (X, A) . In light of the functor $H_* : \text{Chain}^h \rightarrow \text{Ab}_{\mathbb{Z}}$, the main reason for this is the following result.

PROPOSITION 27.9. *Any homotopy $h : (I \times X, I \times A) \rightarrow (Y, B)$ between two maps of pairs $f, g : (X, A) \rightarrow (Y, B)$ induces a chain homotopy $h_* : C_*(X, A; G) \rightarrow C_{*+1}(Y, B; G)$ between the two chain maps $f_*, g_* : C_*(X, A; G) \rightarrow C_*(Y, B; G)$.*

In this statement we are using a popular abuse of notation and writing $h_* : C_*(X, A; G) \rightarrow C_{*+1}(Y, B; G)$ to emphasize the fact that h_* is a map of degree 1 between \mathbb{Z} -graded abelian groups, i.e. it satisfies $h_*(C_n(X, A; G)) \subset C_{n+1}(Y, B; G)$. Let us recall briefly how h_* arises from h , focusing first on the case where $A = \emptyset$ and $G = \mathbb{Z}$. It suffices then to say how h_* is defined on each generator

of $C_n(X)$, i.e. for each singular n -simplex $\sigma : \Delta^n \rightarrow X$, we need to specify $h_*\sigma \in C_{n+1}(Y)$. The key is to look at the map

$$(27.3) \quad I \times \Delta^n \rightarrow Y : (s, t) \mapsto h(s, \sigma(t)).$$

This is not a singular simplex since $I \times \Delta^n$ is not generally a simplex, but if we triangulate $I \times \Delta^n$ in a reasonable way by $(n+1)$ -simplices and identify each one suitably with Δ^{n+1} , then restricting this map to each simplex in the triangulation produces a linear combination (with coefficients ± 1 depending on choices of orientations) of singular $(n+1)$ -simplices in Y , which we define to be $h_*\sigma$. The chain homotopy relation

$$\partial \circ h_* = -f_* + g_* - h_* \circ \partial$$

can then be understood as the algebraic translation of the geometric observation that the boundary of $I \times \Delta^n$ (with suitable orientations) can be written as

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

Indeed, restricting the map (27.3) to $\{0\} \times \Delta^n$ and $\{1\} \times \Delta^n$ produces $f_*\sigma$ and $g_*\sigma$ respectively, while restricting it to $I \times \partial\Delta^n$ or $\partial(I \times \Delta^n)$ produces $h_*(\partial\sigma)$ or $\partial(h_*\sigma)$ respectively.

Extending the above discussion to the case $A \neq \emptyset$ only requires the observation that if $h(I \times A) \subset B$, then the chain homotopy h_* we constructed maps $C_n(A)$ into $C_{n+1}(B)$, thus it descends to a chain homotopy on the quotient complexes $C_*(X, A)$ and $C_*(Y, B)$. Finally, the extension to general coefficient groups G comes for free in light of Exercises 27.5 and 27.6: the existence of the chain homotopy $h_* : C_*(X, A) \rightarrow C_{*+1}(Y, B)$ gives rise to a chain homotopy $h_* \otimes 1 : C_*(X, A) \otimes G \rightarrow C_{*+1}(Y, B) \otimes G$.

28. Excision, suspensions and reduced homology

We were not quite done surveying the properties of singular homology. Here is the last one in the canonical list:

Excision (cf. Theorem 24.7). The **excision** property (*Ausschneidungssatz*) amounts to the statement that the group $H_*(X, A; G)$ does not change if we remove from both A and X a subset with closure in the interior of A . Intuitively this is clear since the definition of $H_*(X, A; G)$ is designed to ignore anything that happens completely inside of A . The proof is nonetheless a bit subtle, and there is a more general phenomenon in the background that we need to review, as it will be useful for a few other things later. Informally, the principle is that singular homology classes can always be represented by cycles whose constituent singular simplices are in some sense “arbitrarily small,” e.g. we lose no generality if we choose an open covering of X and pay attention only to singular simplices that each live entirely in one of the open sets of our covering. The reason this works is that there is a natural algorithm for subdividing simplices, called **barycentric subdivision** (*baryzentrische Zerlegung*), which breaks up each simplex σ into a linear combination of smaller ones whose boundaries add up to a corresponding subdivision of $\partial\sigma$. We will not review the details of barycentric subdivision here, but merely restate the most important lemma that is derived from it. The result concerns the subgroup

$$C_*(\mathcal{U}; G) + C_*(\mathcal{V}; G) \subset C_*(X; G)$$

determined by any covering of X by subsets $X = \mathcal{U} \cup \mathcal{V}$. Observe that since ∂ preserves both $C_*(\mathcal{U}; G)$ and $C_*(\mathcal{V}; G)$, it also preserves their sum, so that $(C_*(\mathcal{U}; G) + C_*(\mathcal{V}; G), \partial)$ defines a *subcomplex* of $(C_*(X; G), \partial)$.

LEMMA 28.1 (cf. Lemma 24.11). *For any subsets $\mathcal{U}, \mathcal{V} \subset X$ with $X = \overset{\circ}{\mathcal{U}} \cup \overset{\circ}{\mathcal{V}}$, the inclusion map*

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

admits a chain homotopy inverse

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

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 $C_*(U; G) + C_*(V; G)$.

With this result in hand, we can consider a pair of spaces (X, A) with a subset

$$B \subset \bar{B} \subset \overset{\circ}{A} \subset A \subset X.$$

The point of this condition is that the sets $X \setminus \bar{B}$ and $\overset{\circ}{A}$ are both open and cover X , so they can be fed into Lemma 28.1. This leads to the following “chain level” version of the excision theorem.

THEOREM 28.2 (cf. Lemma 24.12). *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 $i_* : C_*(X \setminus B, A \setminus B; G) \rightarrow C_*(X, A; G)$.*

SKETCH OF THE PROOF. We start with the observation that the inclusion $C_*(X \setminus B; G) \hookrightarrow C_*(X \setminus B; G) + C_*(A; G)$ descends to a chain map

$$C_*(X \setminus B, A \setminus B; G) \rightarrow (C_*(X \setminus B; G) + C_*(A; G)) / C_*(A; G)$$

which is an isomorphism of chain complexes, as the quotient complexes on the left and right hand side can each be described as free abelian groups with the same set of generators, namely the singular simplices in $X \setminus B$ that are not fully contained in $A \setminus B$. With this identification understood, it suffices to prove that the chain map

$$(28.1) \quad (C_*(X \setminus B; G) + C_*(A; G)) / C_*(A; G) \xrightarrow{j} C_*(X; G) / C_*(A; G) = C_*(X, A; G)$$

induced on the quotients by the inclusion

$$C_*(X \setminus B; G) + C_*(A; G) \xrightarrow{j} C_*(X; G)$$

is a chain homotopy equivalence. Lemma 28.1 provides a chain homotopy inverse $\rho : C_*(X; G) \rightarrow C_*(X \setminus B; G) + C_*(A; G)$ for j , which can be assumed to satisfy $\rho \circ j = \mathbb{1}$, so that ρ restricts to the identity on $C_*(A; G) \subset C_*(X; G)$ and thus descends to the quotient, producing the desired chain homotopy inverse of (28.1). \square

COROLLARY 28.3 (excision). *In the setting of Theorem 28.2, the induced map*

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

is an isomorphism. \square

We have now assembled enough properties of singular homology to compute $H_*(S^n; G)$ for all $n \in \mathbb{N}$ and all coefficient groups G . We carried out this computation at the end of last semester for the case $G = \mathbb{Z}$, and we would like to redo that computation now without making any assumptions about the coefficient group. The main idea is to prove a relation between the homology of an arbitrary space X and the homology one dimension higher of its *suspension* SX , then appealing to the fact that $SS^n \cong S^{n+1}$ to compute $H_*(S^n; G)$ by an inductive argument. The question is then how to begin the induction: in *Topologie I*, we at this point used the fact that $H_1(S^n; \mathbb{Z})$ is isomorphic to the abelianization of $\pi_1(S^n)$, which forced us to set $G = \mathbb{Z}$. Aside from placing a limitation on the coefficient group, there is a more theoretical drawback to this approach: as mentioned in the previous lecture, the relation between $H_1(X; \mathbb{Z})$ and $\pi_1(X)$ is a property distinctive to singular homology that is not shared by other invariants that we will later want to also call “homology theories,” thus the argument in this form would not be valid for those alternative theories. At the

end of this lecture, we will introduce a clever trick that avoids the need to specify the coefficient group or to use any knowledge of $\pi_1(S^n)$.

Recall that for any space X , its **cone** (*Kegel*) is the space

$$C_+X := CX := (X \times [0, 1]) / (X \times \{1\}),$$

and we call the point $p_+ \in C_+X$ represented by the collapsed subset $X \times \{1\}$ the **summit** or **node** of the cone. We will use the notation C_+X to distinguish the so-called “positive” cone from the negative cone

$$C_-X := (X \times [-1, 0]) / (X \times \{-1\}),$$

which has node $p_- \in C_-X$ represented by the collapsed subset $X \times \{-1\}$. These two spaces are obviously homeomorphic, and they each admit deformation retractions to their respective nodes, thus they are contractible. We will often think of X as a subset of $C_\pm X$ by identifying it with $X \times \{0\}$. The **suspension** (*Einhangung*) of X is then the space obtained by gluing the positive and negative cones together along X ,

$$SX := C_+X \cup_X C_-X.$$

For the example of the sphere S^{n-1} , we have $CS^{n-1} \cong \mathbb{D}^n$ and $SS^{n-1} \cong S^n$.

Consider now the following diagram:

$$(28.2) \quad \begin{array}{ccc} H_k(X; G) & & H_{k+1}(SX; G) \\ \partial_* \uparrow & & \downarrow k_* \\ H_{k+1}(C_+X, X; G) & \xrightarrow{i_*} & H_{k+1}(SX \setminus \{p_-\}, C_-X \setminus \{p_-\}; G) \xrightarrow{j_*} H_{k+1}(SX, C_-X; G) \end{array}$$

Three of the maps here are determined by the obvious inclusions of pairs,

$$\begin{aligned} (C_+X, X) &\xrightarrow{i} (SX \setminus \{p_-\}, C_-X \setminus \{p_-\}), \\ (SX \setminus \{p_-\}, C_-X \setminus \{p_-\}) &\xrightarrow{j} (SX, C_-X), \\ (SX, \emptyset) &\xrightarrow{k} (SX, C_-X). \end{aligned}$$

The first of these is a homotopy equivalence, as there exists a deformation retraction of the pair $(SX \setminus \{p_-\}, C_-X \setminus \{p_-\})$ to (C_+X, X) , thus i_* is an isomorphism (cf. Example 23.3). Since $\{p_-\}$ is a closed set contained in the interior of C_-X , excision implies that j_* is also an isomorphism. For the other two maps, we consider the long exact sequences of the pairs (SX, C_-X) and (C_+X, X) , that is

$$\dots \longrightarrow H_{k+1}(C_-X; G) \longrightarrow H_{k+1}(SX; G) \xrightarrow{k_*} H_{k+1}(SX, C_-X; G) \longrightarrow H_k(C_-X; G) \longrightarrow \dots$$

and

$$\dots \longrightarrow H_{k+1}(C_+X; G) \longrightarrow H_{k+1}(C_+X, X; G) \xrightarrow{\partial_*} H_k(X; G) \longrightarrow H_k(C_+X; G) \longrightarrow \dots$$

If $k \geq 1$, then the contractibility of $C_\pm X$ implies via homotopy invariance that

$$H_k(C_\pm X; G) \cong H_k(\{\text{pt}\}; G) = 0, \quad \text{and} \quad H_{k+1}(C_\pm X; G) \cong H_{k+1}(\{\text{pt}\}; G) = 0,$$

thus the exactness of these two sequences implies that k_* and ∂_* are both isomorphisms. We’ve proved:

THEOREM 28.4. *For every space X , abelian group G and integer $k \geq 1$, the diagram (28.2) gives rise to an isomorphism*

$$S_* := k_*^{-1} \circ j_* \circ i_* \circ \partial_*^{-1} : H_k(X; G) \rightarrow H_{k+1}(SX; G).$$

□

To apply this in computing $H_*(S^n; G)$, note first that the problem is easy for $n = 0$: since S^0 is just the disjoint union of two one-point spaces, the computation of $H_*(S^0; G)$ follows immediately from that of $H_*(\{\text{pt}\}; G)$ together with the disjoint union property. In particular, $H_k(S^0; G) = 0$ for all $k > 0$, so if $k > n \geq 1$, we can apply Theorem 28.4 repeatedly to prove

$$H_k(S^n; G) = H_{k-n}(S^0; G) = 0 \quad \text{for } k > n.$$

We have a problem however if $k \leq n$, because before the repeated applications of Theorem 28.4 give us S^0 , we will reach $H_1(S^{n-k+1}; G)$ and not be able to go any further, as the theorem does not give any relation between $H_0(X; G)$ and $H_1(SX; G)$. The reason why not is that even though $C_\pm X$ are contractible, the groups $H_0(C_\pm X; G)$ are not trivial, so that the two exact sequences we considered above fail to prove that k_* and ∂_* are isomorphisms. The problem is thus caused by the fact that $H_*(X; G)$ for a contractible space does not completely vanish—it only *mostly* vanishes.

There is a clever remedy for this problem: one can define a variant of $H_*(X; G)$ that fits into all the same exact sequences but does completely vanish in the case where X is contractible. A brief algebraic digression is in order before we continue.

EXERCISE 28.5. Given a short exact sequence of abelian groups $0 \rightarrow A \xrightarrow{f} B \xrightarrow{g} C \rightarrow 0$, show that the following conditions are equivalent:

- (i) There exists a homomorphism $\pi : B \rightarrow A$ such that $\pi \circ f = 1_A$;
- (ii) There exists a homomorphism $i : C \rightarrow B$ such that $g \circ i = 1_C$;
- (iii) There exists an isomorphism $\Phi : B \rightarrow A \oplus C$ such that $\Phi \circ f(a) = (a, 0)$ and $g \circ \Phi^{-1}(a, c) = c$.

$$\begin{array}{ccccccc}
 & & & & B & & \\
 & & & & \downarrow \Phi & & \\
 0 & \longrightarrow & A & \begin{array}{l} \nearrow f \\ \searrow \end{array} & B & \begin{array}{l} \searrow g \\ \nearrow \end{array} & C \longrightarrow 0 \\
 & & & & A \oplus C & &
 \end{array}$$

DEFINITION 28.6. We say that a short exact sequence **splits** whenever it satisfies any of the three equivalent properties listed in Exercise 28.5.

EXERCISE 28.7. Show that if the groups in Exercise 28.5 are all finite-dimensional vector spaces and the homomorphisms are linear maps, then the sequence always splits. Show also that this is true for any sequence of abelian groups $0 \rightarrow A \rightarrow B \rightarrow C \rightarrow 0$ if C is free.³⁵

Hint: Use a basis of C to write down a right-inverse for $g : B \rightarrow C$.

EXAMPLE 28.8. The sequence $0 \rightarrow 2\mathbb{Z} \rightarrow \mathbb{Z} \rightarrow \mathbb{Z}_2 \rightarrow 0$, with the first and last nontrivial maps defined as the natural inclusion and quotient projection respectively, is exact but does not split. Indeed, a splitting in this case would imply via Exercise 28.5 that \mathbb{Z} is isomorphic to $2\mathbb{Z} \oplus \mathbb{Z}_2 \cong \mathbb{Z} \oplus \mathbb{Z}_2$, which it clearly is not. (In light of Exercise 28.7, we notice of course that \mathbb{Z}_2 is not free.)

End of digression; now we can explain the clever trick. Fix a one-point space $\{\text{pt}\}$, and let

$$\epsilon : X \rightarrow \{\text{pt}\}$$

denote the unique map, which is (trivially) continuous.

³⁵An abelian group G is called **free** whenever it is isomorphic to the free abelian group $F^{\text{ab}}(S)$ on some set S . Equivalently, this means that G admits a **basis**, meaning a subset $S \subset G$ such that every element of G is uniquely representable as a linear combination $\sum_{s \in S} m_s s$ for some coefficients $m_s \in \mathbb{Z}$, only finitely many of which are nonzero.

DEFINITION 28.9. The **reduced singular homology** (*reduzierte singuläre Homologie*) groups of X are defined for each $n \in \mathbb{Z}$ as the subgroup

$$\tilde{H}_n(X; G) = \ker \epsilon_* \subset H_n(X; G),$$

where $\epsilon_* : H_n(X; G) \rightarrow H_n(\{\text{pt}\}; G)$ is the homomorphism induced by the unique map $\epsilon : X \rightarrow \{\text{pt}\}$.

PROPOSITION 28.10. *If X is contractible, then $\tilde{H}_*(X; G) = 0$.*

PROOF. Contractibility means that the map $\epsilon : X \rightarrow \{\text{pt}\}$ is a homotopy equivalence, thus $\epsilon_* : H_*(X; G) \rightarrow H_*(\{\text{pt}\}; G)$ is an isomorphism, and its kernel $\tilde{H}_*(X; G)$ is therefore trivial. \square

PROPOSITION 28.11. *For all spaces X , the map $\epsilon_* : H_*(X; G) \rightarrow H_*(\{\text{pt}\}; G)$ is surjective, and the resulting short exact sequence*

$$0 \rightarrow \tilde{H}_*(X; G) \hookrightarrow H_*(X; G) \xrightarrow{\epsilon_*} H_*(\{\text{pt}\}; G) \rightarrow 0$$

splits. In particular, $H_(X; G)$ is isomorphic to $\tilde{H}_*(X; G) \oplus H_*(\{\text{pt}\}; G)$, thus*

$$H_n(X; G) \cong \begin{cases} \tilde{H}_n(X; G) \oplus G & \text{if } n = 0, \\ \tilde{H}_n(X; G) & \text{if } n \neq 0. \end{cases}$$

PROOF. Choose any map $i : \{\text{pt}\} \hookrightarrow X$ and notice that this is also trivially continuous, though usually not unique. Then $\epsilon \circ i$ is the identity map on $\{\text{pt}\}$, hence $\epsilon_* \circ i_* = \mathbf{1}$ on $H_*(\{\text{pt}\}; G)$, implying that ϵ_* is surjective and (via Exercise 28.5) that the sequence splits. \square

PROPOSITION 28.12. *The homomorphism $f_* : H_*(X; G) \rightarrow H_*(Y; G)$ induced by any continuous map $f : X \rightarrow Y$ sends $\tilde{H}_*(X; G)$ into $\tilde{H}_*(Y; G)$. In particular, $\tilde{H}_*(\cdot; G)$ defines functors $\text{Top} \rightarrow \text{Ab}_{\mathbb{Z}}$ and $\text{Top}^h \rightarrow \text{Ab}_{\mathbb{Z}}$ in the obvious way.*

PROOF. Denote $\epsilon^X : X \rightarrow \{\text{pt}\}$ and $\epsilon^Y : Y \rightarrow \{\text{pt}\}$ for the unique maps, and notice that $\epsilon^Y \circ f = \text{Id} \circ \epsilon^X$, thus the following diagram commutes.

$$\begin{array}{ccc} H_*(X; G) & \xrightarrow{f_*} & H_*(Y; G) \\ \downarrow \epsilon_*^X & & \downarrow \epsilon_*^Y \\ H_*(\{\text{pt}\}; G) & \xrightarrow{\mathbf{1}} & H_*(\{\text{pt}\}; G) \end{array}$$

This implies that $f_*(\ker \epsilon_*^X) \subset \ker \epsilon_*^Y$. \square

The relative version of reduced homology is defined in a trivial way: we set

$$\tilde{H}_*(X, A; G) := H_*(X, A; G) \quad \text{whenever } A \neq \emptyset.$$

This seemingly naive definition is justified by the following considerations. Note first that the functor $\tilde{H}_*(\cdot; G) : \text{Top} \rightarrow \text{Ab}_{\mathbb{Z}}$ now extends to pairs as a functor $\text{Top}_{\text{rel}} \rightarrow \text{Ab}_{\mathbb{Z}}$; 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 $H_*(X, X; G)$ all vanish; one can prove this either directly from the definition of relative singular homology or indirectly via the long exact sequence of the pair. It follows that $\tilde{H}_*(X, A; G)$ for $A \neq \emptyset$ is in fact the kernel of the map

$$H_*(X, A; G) \xrightarrow{\epsilon_*} H_*(\{\text{pt}\}, \{\text{pt}\}; G) = 0$$

induced by the unique map of pairs $\epsilon : (X, A) \rightarrow (\{\text{pt}\}, \{\text{pt}\})$. Moreover, the naturality of connecting homomorphisms gives a commutative diagram

$$\begin{array}{ccc} H_{n+1}(X, A; G) & \xrightarrow{\partial_*} & H_n(A; G) \\ \downarrow \epsilon_* & & \downarrow \epsilon_* \\ H_{n+1}(\{\text{pt}\}, \{\text{pt}\}; G) & \xrightarrow{\partial_*} & H_n(\{\text{pt}\}; G) \end{array}$$

Since the term $H_{n+1}(\{\text{pt}\}, \{\text{pt}\}; G)$ is trivial, this diagram proves that the image of $\partial_* : H_{n+1}(X, A; G) \rightarrow H_n(A; G)$ is always in the subgroup $\tilde{H}_n(A; G)$. We can therefore write down a well-defined sequence of homomorphisms

$$\dots \rightarrow \tilde{H}_{n+1}(X, A; G) \xrightarrow{\partial_*} \tilde{H}_n(A; G) \xrightarrow{i_*} \tilde{H}_n(X; G) \xrightarrow{j_*} \tilde{H}_n(X, A; G) \xrightarrow{\partial_*} \tilde{H}_{n-1}(A; G) \rightarrow \dots$$

using the usual inclusions $i : A \hookrightarrow X$ and $j : (X, \emptyset) \hookrightarrow (X, A)$. It is not immediately obvious, however, whether this sequence is exact. This is where the magic of diagram chasing again comes into play. Consider the commutative diagram

$$\begin{array}{cccccccc} & & 0 & & 0 & & 0 & & 0 & & \\ & & \downarrow & & \downarrow & & \downarrow & & \downarrow & & \\ \dots & \longrightarrow & \tilde{H}_n(A; G) & \xrightarrow{i_*} & \tilde{H}_n(X; G) & \xrightarrow{j_*} & \tilde{H}_n(X, A; G) & \xrightarrow{\partial_*} & \tilde{H}_{n-1}(A; G) & \longrightarrow & \dots \\ & & \downarrow & & \downarrow & & \parallel & & \downarrow & & \\ \dots & \longrightarrow & H_n(A; G) & \xrightarrow{i_*} & H_n(X; G) & \xrightarrow{j_*} & H_n(X, A; G) & \xrightarrow{\partial_*} & H_{n-1}(A; G) & \longrightarrow & \dots \\ & & \downarrow \epsilon_* & & \downarrow \epsilon_* & & \downarrow \epsilon_* & & \downarrow \epsilon_* & & \\ \dots & \longrightarrow & H_n(\{\text{pt}\}; G) & \xrightarrow{=} & H_n(\{\text{pt}\}; G) & \longrightarrow & 0 & \longrightarrow & H_{n-1}(\{\text{pt}\}; G) & \longrightarrow & \dots \\ & & \downarrow & & \downarrow & & \downarrow & & \downarrow & & \\ & & 0 & & 0 & & 0 & & 0 & & \end{array}$$

Here the bottom two nontrivial rows are the long exact sequences of the pairs (X, A) and $(\{\text{pt}\}, \{\text{pt}\})$, and all columns in the diagram are short exact sequences by construction. The rest is algebra:

PROPOSITION 28.13. *Assume the following diagram of abelian groups with homomorphisms commutes, all its columns are exact sequences, and the bottom two nontrivial rows are also exact sequences:*

$$\begin{array}{ccccccccc} & & 0 & & 0 & & 0 & & 0 & & 0 & & \\ & & \downarrow & & \downarrow & & \downarrow & & \downarrow & & \downarrow & & \\ \dots & & A_{n+2} & & A_{n+1} & & A_n & & A_{n-1} & & A_{n-2} & & \dots \\ & & \downarrow \iota_{n+2} & & \downarrow \iota_{n+1} & & \downarrow \iota_n & & \downarrow \iota_{n-1} & & \downarrow \iota_{n-2} & & \\ \dots & \longrightarrow & B_{n+2} & \xrightarrow{g_{n+2}} & B_{n+1} & \xrightarrow{g_{n+1}} & B_n & \xrightarrow{g_n} & B_{n-1} & \xrightarrow{g_{n-1}} & B_{n-2} & \longrightarrow & \dots \\ & & \downarrow \epsilon_{n+2} & & \downarrow \epsilon_{n+1} & & \downarrow \epsilon_n & & \downarrow \epsilon_{n-1} & & \downarrow \epsilon_{n-2} & & \\ \dots & \longrightarrow & C_{n+2} & \xrightarrow{h_{n+2}} & C_{n+1} & \xrightarrow{h_{n+1}} & C_n & \xrightarrow{h_n} & C_{n-1} & \xrightarrow{h_{n-1}} & C_{n-2} & \longrightarrow & \dots \\ & & \downarrow & & \downarrow & & \downarrow & & \downarrow & & \downarrow & & \\ & & 0 & & 0 & & 0 & & 0 & & 0 & & \end{array}$$

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. 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 28.14. *For any pair of spaces (X, A) , there is a long exact sequence of reduced homology groups*

$$\dots \rightarrow \tilde{H}_{n+1}(X, A; G) \xrightarrow{\partial_*} \tilde{H}_n(A; G) \xrightarrow{i_*} \tilde{H}_n(X; G) \xrightarrow{j_*} \tilde{H}_n(X, A; G) \xrightarrow{\partial_*} \tilde{H}_{n-1}(A; G) \rightarrow \dots,$$

where $i : A \hookrightarrow X$ and $j : (X, \emptyset) \hookrightarrow (X, A)$ are the obvious inclusions and $\partial_* : \tilde{H}_n(X, A; G) \rightarrow \tilde{H}_{n-1}(A; G)$ is the same map as the usual connecting homomorphism $H_n(X, A; G) \rightarrow H_{n-1}(A; G)$.

□

Here's the upshot: if we now redo the argument behind Theorem 28.4 but replacing H_* with \tilde{H}_* at every step, then it still works, and it *also* works for $k = 0$ and $k = -1$ since $\tilde{H}_0(C_{\pm} X; G)$ also vanishes. We conclude:

THEOREM 28.15. *For every space X , abelian group G and integer $k \in \mathbb{Z}$, there is an isomorphism*

$$S_* : \tilde{H}_k(X; G) \rightarrow \tilde{H}_{k+1}(SX; G).$$

□

Now apply this to the spheres: when $n > k$, repeating the isomorphism $k + 1$ times gives

$$\tilde{H}_k(S^n; G) \cong \tilde{H}_{-1}(S^{n-k-1}; G) = 0,$$

so in light of Proposition 28.11, we obtain $H_k(S^n; G) = \tilde{H}_k(S^n; G) = 0$ for $0 < k < n$ and $H_0(S^n; G) \cong \tilde{H}_0(S^n; G) \oplus G = G$. In the case $k = n$, we can instead repeat the isomorphism n times and obtain

$$\tilde{H}_n(S^n; G) \cong \tilde{H}_0(S^0; G).$$

The latter is related to $H_0(S^0; G)$ by Proposition 28.11, and we already know $H_0(S^0; G) \cong G \oplus G$ since $S^0 \cong \{\text{pt}\} \sqcup \{\text{pt}\}$, thus

$$(28.3) \quad G \oplus G \cong H_0(S^0; G) \cong \tilde{H}_0(S^0; G) \oplus G,$$

and if G is finitely generated, then the classification of finitely generated abelian groups now implies $\tilde{H}_0(S^0; G) \cong G$. Without any assumption on G , one can instead prove the same result as outlined in the following exercise.

EXERCISE 28.16. For any two spaces X and Y with maps $\epsilon^X : X \rightarrow \{\text{pt}\}$ and $\epsilon^Y : Y \rightarrow \{\text{pt}\}$, show that the natural isomorphism $H_*(X \sqcup Y; G) \cong H_*(X; G) \oplus H_*(Y; G)$ identifies $\tilde{H}_*(X \sqcup Y; G)$ with $\ker(\epsilon_*^X \oplus \epsilon_*^Y) \subset H_*(X; G) \oplus H_*(Y; G)$. Then apply this in the case $X = Y = \{\text{pt}\}$ to identify $\tilde{H}_0(\{\text{pt}\} \sqcup \{\text{pt}\}; G)$ with the kernel of the map

$$\mathbf{1} \oplus \mathbf{1} : G \oplus G \rightarrow G : (g, h) \mapsto g + h,$$

which is isomorphic to G .

Since $H_n(S^n; G) = \tilde{H}_n(S^n; G)$ for all $n > 0$, we've proved:

THEOREM 28.17. *For every $n \in \mathbb{N}$,*

$$H_k(S^n; G) \cong \begin{cases} G & \text{for } k = 0, n, \\ 0 & \text{for all other } k. \end{cases}$$

□

REMARK 28.18. It is sometimes useful to know that $\tilde{H}_*(X; G)$ is also the homology of a chain complex. To see this, note that every element in $C_0(X; G)$ is a cycle, thus $H_0(X; G) = C_0(X; G)/\text{im } \partial_1$ for the restriction $\partial_1 : C_1(X; G) \rightarrow C_0(X; G)$ of ∂ . Thus for the unique map $\epsilon : X \rightarrow \{\text{pt}\}$, the surjective homomorphism $\epsilon_* : H_0(X; G) \rightarrow H_0(\{\text{pt}\}; G) = G$ can be composed with the quotient projection $C_0(X; G) \rightarrow C_0(X; G)/\text{im } \partial_1 = H_0(X; G)$ to define a homomorphism $\epsilon_* : C_0(X; G) \rightarrow G$ such that $\epsilon_* \circ \partial_1 = 0$. This is equivalent to saying that ϵ_* defines a chain map $\epsilon_* : C_*(X; G) \rightarrow G_*$, where G_* is the chain complex with $G_0 := G$ and $G_n := 0$ for all $n \neq 0$, so that its boundary map is necessarily trivial.

In general, if C_* is any chain complex with $C_n = 0$ for all $n < 0$ and G_* is the trivial complex described above, a surjective chain map

$$\epsilon : C_* \rightarrow G_*$$

is called an **augmentation** (*Augmentationsabbildung*) of C_* over G . This is equivalent to a surjective homomorphism $\epsilon : C_0 \rightarrow G$ satisfying $\epsilon \circ \partial_1 = 0$ for the boundary map $\partial_1 : C_1 \rightarrow C_0$. One can therefore define an **augmented chain complex** \tilde{C}_* in the form

$$\dots \longrightarrow C_2 \xrightarrow{\partial_2} C_1 \xrightarrow{\partial_1} C_0 \xrightarrow{\epsilon} G \longrightarrow 0 \longrightarrow \dots$$

in other words, $\tilde{C}_n := C_n$ for all $n \neq -1$ but $\tilde{C}_{-1} := G$, with the new boundary map $\tilde{C}_0 \rightarrow \tilde{C}_{-1}$ defined as $\epsilon : C_0 \rightarrow G$. The homology of this new complex is precisely the kernel of the homomorphism $\epsilon_* : H_*(C_*) \rightarrow G_*$ induced by the chain map $\epsilon : C_* \rightarrow G_*$, thus we can sensibly call it the **reduced homology of the complex** C_* ,

$$\tilde{H}_*(C_*) := H_*(\tilde{C}_*).$$

EXERCISE 28.19. Show that the augmentation $\epsilon_* : C_0(X; G) \rightarrow G$ described in Remark 28.18 is given by the formula

$$\epsilon_* \left(\sum_i g_i \sigma_i \right) = \sum_i g_i$$

for finite sums with $g_i \in G$ and $\sigma_i : \Delta^0 \rightarrow X$.

29. Simplicial complexes in singular homology

Before further developing the theory of singular homology, I would like to pause and address a question that is important for intuition: how can we *visualize* a singular homology class? While this question is not always answerable, there is a standard answer that suffices in most situations: elements of $H_n(X)$ can often be viewed as *triangulated closed n -dimensional submanifolds* of X , where two such submanifolds represent the same homology class whenever their disjoint union (with appropriate orientations) bounds some triangulated compact $(n+1)$ -dimensional submanifold with boundary. To explain in precise terms what this means, I need to digress a little bit into the subject of simplicial complexes and simplicial homology, which will also come in useful for other purposes later. Much of what I will say in this lecture repeats things that I said or hinted at last semester (see especially Lecture 21), but I am now in a position to say those things more precisely.

You will recall that an **n -dimensional (topological) manifold with boundary** is a second countable Hausdorff space M in which every point has a neighborhood homeomorphic to an open subset of the n -dimensional half-space

$$\mathbb{H}^n := [0, \infty) \times \mathbb{R}^{n-1}.$$

The **boundary** $\partial M \subset M$ is defined as the set of points that are sent to $\partial \mathbb{H}^n := \{0\} \times \mathbb{R}^{n-1} \subset \mathbb{H}^n$ under these homeomorphisms, hence ∂M itself is an $(n-1)$ -dimensional topological manifold (with empty boundary: $\partial(\partial M) = \emptyset$). A manifold M is called **closed** (*geschlossen*)³⁶ if it is compact and $\partial M = \emptyset$. All manifolds in this lecture will be compact, so we need not worry about second countability, and we will usually also omit the word “topological”. (The word is often included in order to distinguish topological manifolds from *smooth* manifolds, but smooth structures will not have any role to play in our discussion.) We will usually also omit the words “with boundary” and keep in mind that all manifolds in principle have boundary, but the boundary may be empty.

We usually picture a *simplicial complex* as a space decomposed into a union of simplices. Strictly speaking, a simplicial complex is a purely combinatorial object, and the topological space that we build out of it is called its *polyhedron*. Here are the precise definitions.

³⁶Not to be confused with “*abgeschlossen*,” which is what we call a closed *subset* of a topological space. These two meanings of the English word “closed” are defined in different contexts and are not equivalent.

DEFINITION 29.1. A **simplicial complex** (*Simplizialkomplex*) K consists of two sets V and S , called the sets of **vertices** (*Eckpunkte*) and **simplices** (*Simplizes*) respectively, where S is a subset of the set of all finite subsets of V , and $\sigma \in S$ is called an n -**simplex** of K if it has $n + 1$ elements. We require the following conditions:

- (1) Every vertex $v \in V$ gives rise to a 0-simplex in K , i.e. $\{v\} \in S$;
- (2) If $\sigma \in S$ then every subset $\sigma' \subset \sigma$ is also an element of S .

For any n -simplex $\sigma \in S$, its subsets are called its **faces** (*Seiten* or *Facetten*), and in particular the subsets that are $(n - 1)$ -simplices are called **boundary faces** (*Seitenflächen*) of σ . The second condition above thus says that for every simplex in the complex, all of its faces also belong to the complex. With this condition in place, the first condition is then equivalent to the requirement that every vertex in the set V belongs to at least one simplex.

The complex K is said to be **finite** if V (and therefore also S) is finite, and 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.

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 $V \rightarrow I := [0, 1]$, i.e. each element $t \in I^V$ is determined by a set of real numbers $t_v \in [0, 1]$ associated to the vertices $v \in V$, which we can think of as the *coordinates* of t . For each n -simplex $\sigma = \{v_0, \dots, v_n\}$ 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 is not the most natural choice here. We will instead let the topology of $|K|$ be determined by that of the individual simplices: define a subset $\mathcal{U} \subset |K|$ to be open if and only if for every $\sigma \in S$, $\mathcal{U} \cap |\sigma| \subset |\sigma|$ is open for the topology on $|\sigma|$ defined above. In other words, $|K|$ carries the strongest³⁷ topology for which the inclusions $|\sigma| \hookrightarrow |K|$ are continuous for all σ . You should take a moment to convince yourself that this matches what was already said for the case where K is finite.

EXERCISE 29.2. Show that 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 $\sigma \in S$.

³⁷For 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 29.2.

DEFINITION 29.3. For each integer $n \geq 0$, the n -**skeleton** (n -*Skelett* or n -*Gerüst*) of a polyhedron $X = |K|$ is the subspace $X^n \subset X$ consisting of the union of all $|\sigma| \subset X$ for k -simplices σ in K with $k \leq n$.

By this definition, a polyhedron is n -dimensional (i.e. corresponds to an n -dimensional simplicial complex) if and only if it is equal to its n -skeleton. The 0-skeleton of any polyhedron is just the union of all its vertices—one can show that this is always a discrete set.

While $|K|$ was defined above as a subset of a vector space whose dimension may in general be quite large (or infinite), visualizing $|K|$ in concrete examples is often easier than one might expect.

EXAMPLE 29.4. Suppose $V = \{v_0, v_1, v_2, v_3\}$ and S contains the subsets $A := \{v_0, v_1, v_2\}$ and $B := \{v_1, v_2, v_3\}$, plus all of their respective subsets. Then $|K|$ contains two copies of the triangle Δ^2 , and they intersect each other along a single common edge connecting the vertices labeled v_1 and v_2 . The complex is 2-dimensional, and its 1-skeleton is the union of all the edges of the triangles.

EXAMPLE 29.5. If V has $n + 1$ elements and S consists of all subsets of V except for V itself, then $|K|$ is homeomorphic to $\partial\Delta^n$, i.e. the union of all the boundary faces of Δ^n . In particular, this is homeomorphic to S^{n-1} .

EXAMPLE 29.6. Suppose $V = \{v_0, \dots, v_n\}$ and S is defined to consist of all the one-element subsets $\{v_i\}$ plus the 1-simplices $\{v_i, v_{i+1}\}$ for $i = 0, \dots, n - 1$ and $\{v_n, v_0\}$. Then $|K|$ is a 1-dimensional polyhedron homeomorphic to S^1 .

EXAMPLE 29.7. Taking $V = \mathbb{Z}$ with S as the set of all 0-simplices $\{n\}$ plus 1-simplices of the form $\{n, n + 1\}$ for $n \in \mathbb{Z}$ gives an infinite (but 1-dimensional) simplicial complex whose polyhedron is homeomorphic to \mathbb{R} .

EXAMPLE 29.8. If $V = \mathbb{N}$ and S is the set of all finite subsets of \mathbb{N} , then K is an infinite-dimensional simplicial complex. Every simplex in this complex is a face of $\{1, \dots, n\}$ 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} .

A **subcomplex** (*Unterkomplex* or *Teilkomplex*) of a simplicial complex $K = (V, S)$ is a simplicial complex $K' = (V', S')$ such that $V' \subset V$ and $S' \subset S$. We then call (K, K') a **simplicial pair** (*simpliziales Paar*). The polyhedron $|K'|$ can be regarded naturally as a subspace of $|K|$ via the obvious inclusion $I^{V'} \hookrightarrow I^V$ that sets all coordinates t_v for $v \in V \setminus V'$ to zero. In this way, every simplicial pair (K, K') gives rise to a pair of spaces $(|K|, |K'|)$. Note that the empty set also defines a simplicial complex (whose polyhedron is empty), thus every complex K can be identified with the simplicial pair (K, \emptyset) .

DEFINITION 29.9. Given two simplicial complexes $K_1 = (V_1, S_1)$ and $K_2 = (V_2, S_2)$, a **simplicial map** (*simpliziale Abbildung*) from K_1 to K_2 is a function $f : V_1 \rightarrow V_2$ such that $f(\sigma) \in S_2$ for every $\sigma \in S_1$. A map of simplicial pairs $(K_1, K'_1) \rightarrow (K_2, K'_2)$ is then a simplicial map $K_1 \rightarrow K_2$ that restricts to a simplicial map $K'_1 \rightarrow K'_2$.

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 to 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 formal³⁸ sum

³⁸The 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.

$\sum_{v \in V} t_v e_v$ 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 $v \in \sigma$. If $f : (K_1, K'_1) \rightarrow (K_2, K'_2)$ is a map of simplicial pairs, then it induces in this way a continuous map of pairs $(|K_1|, |K'_1|) \rightarrow (|K_2|, |K'_2|)$. We have thus defined a functor

$$\text{Simp}_{\text{rel}} \rightarrow \text{Top}_{\text{rel}} : (K, K') \mapsto (|K|, |K'|),$$

where Simp_{rel} is the category of simplicial pairs with morphisms defined to be maps of simplicial pairs. 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 mainly be concerned with compact manifolds, the following result enables us to restrict attention to finite simplicial complexes:

PROPOSITION 29.10. *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 29.10 FOR FINITE-DIMENSIONAL COMPLEXES. If K is finite then $|K|$ is a closed and bounded subset of the finite-dimensional vector space \mathbb{R}^V and is therefore compact.

Conversely, if K is infinite but $\dim K < \infty$, there exists an infinite sequence of distinct simplices $\sigma_1, \sigma_2, \dots \in S$ 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

An **orientation** (*Orientierung*) of a simplex $\sigma \in S$ 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. This means that every n -simplex for $n \geq 1$ admits exactly two choices of orientation, whereas only one orientation is possible for a 0-simplex. The term **oriented simplex** (*orientiertes Simplex*) means a simplex together with a choice of orientation. If $\sigma = \{v_0, \dots, v_n\}$, the oriented simplex defined by choosing the ordering v_0, \dots, v_n will be denoted by $[v_0, v_1, \dots, v_n]$, and one can change this to the opposite orientation by writing it as $[v_1, v_0, \dots, v_n]$.

The **simplicial chain complex** (*simplizialer Kettenkomplex*) $C_*^\Delta(K) = \bigoplus_{n \in \mathbb{Z}} C_n^\Delta(K)$ of a simplicial complex $K = (V, S)$ is now defined as follows. For $n \geq 0$, let $\widehat{C}_n^\Delta(K)$ denote the free abelian group generated by the set of all oriented simplices $[v_0, \dots, v_n]$ with $\{v_0, \dots, v_n\} \in S$. Then define

$$C_n^\Delta(K) = \widehat{C}_n^\Delta(K) / Z$$

where $Z \subset \widehat{C}_n^\Delta(K)$ is the subgroup generated by all elements of the form $[v_0, v_1, \dots, v_n] + [v_1, v_0, \dots, v_n]$. In other words, if σ denotes an oriented n -simplex and $\bar{\sigma}$ is the same simplex with the opposite orientation, then both can be regarded as generators of $C_n^\Delta(K)$ subject to the

relation $\bar{\sigma} = -\sigma$. With this understood, we define the boundary map $\partial : C_n^\Delta(K) \rightarrow C_{n-1}^\Delta(K)$ for each $n \geq 1$ by

$$(29.1) \quad \partial[v_0, \dots, v_n] = \sum_{k=0}^n (-1)^k [v_0, \dots, v_{k-1}, v_{k+1}, \dots, v_n].$$

EXERCISE 29.11. Show that $\partial : C_n^\Delta(K) \rightarrow C_{n-1}^\Delta(K)$ given by the above formula is well defined. In particular, prove that if v_0, \dots, v_n on the left hand side of (29.1) is changed by a permutation, then the right hand side is multiplied by ± 1 , with sign determined by the parity of the permutation. *Hint: Notice that in $C_n^\Delta(K)$, $[v_0, \dots, v_n] = (-1)^k [v_k, v_0, \dots, v_{k-1}, v_{k+1}, \dots, v_n]$.*

The chain complex is extended to all $n \in \mathbb{Z}$ by setting $C_n^\Delta(K) = 0$ for $n < 0$, and the formula $\partial^2 = 0$ is now a straightforward consequence of sign cancelations, the same as in singular homology. We can then also define a relative complex for any simplicial pair (K, K') by $C_*^\Delta(K, K') := C_*^\Delta(K)/C_*^\Delta(K')$, and extend the definition to an arbitrary abelian coefficient group G by $C_*^\Delta(K, K'; G) := C_*^\Delta(K, K') \otimes G$. The homologies of these chain complexes define the **simplicial homology** groups (*simpliziale Homologiegruppen*)

$$H_*^\Delta(K), H_*^\Delta(K, K'; G), \text{ etc.}$$

A simplicial map $f : K_1 \rightarrow K_2$ now defines an obvious chain map $f_* : C_*^\Delta(K_1) \rightarrow C_*^\Delta(K_2)$ which acts on the generators of $C_n^\Delta(K_1)$ by

$$f_*[v_0, \dots, v_n] = [f(v_0), \dots, f(v_n)],$$

where some abuse of notation is necessary since f may in general map an n -simplex to a k -simplex for $k < n$: that is what happens if two of the vertices $f(v_i)$ match, and we can accommodate this by extending the notation so that

$$[v_0, \dots, v_n] := 0 \quad \text{whenever} \quad v_i = v_j \text{ for some } i \neq j.$$

With this understood, simplicial maps induce homomorphisms on the simplicial homology groups, so that $H_*^\Delta(\cdot; G)$ defines a functor $\text{Simp}_{\text{rel}} \rightarrow \text{Ab}_{\mathbb{Z}}$.

Notice that while $C_n^\Delta(K, K')$ was defined above as a quotient, it is also a free abelian group, as one can always just choose an orientation for each simplex σ that is in K but not K' , and then write all elements of $C_n(K, K')$ uniquely as finite linear combinations (with integer coefficients) of these chosen oriented simplices. This is essentially the definition of $C_*^\Delta(K)$ that I gave last semester in Lecture 21: it required making an arbitrary choice of orientation for each individual simplex, which is what one has to do in practice to compute $H_*^\Delta(K)$ (see e.g. Figure 14), but the definition in this lecture reveals why the result is not in any way dependent on that arbitrary choice.

We are not yet in a position to prove the most important theorem about simplicial homology, but we can state it. It will follow from a more general theorem about the homology of CW-complexes, which tells us that simplicial and cell complexes provide a practical combinatorial algorithm for computing the singular homology of “reasonable” spaces.

THEOREM 29.12. *For every simplicial pair (K, K') and every coefficient group G ,*

$$H_*^\Delta(K, K'; G) \cong H_*(|K|, |K'|; G).$$

It is not hard to guess how a natural map $H_*^\Delta(K, K'; G) \rightarrow H_*(|K|, |K'|; G)$ might be defined. One obvious procedure is to start by fixing an ordering of the set of all vertices in $K = (V, S)$, so that each n -simplex $\sigma \in S$ now inherits an ordering and thus an orientation. This also determines an identification of $|\sigma| \subset I^\sigma$ with the standard n -simplex $\Delta^n \subset I^{n+1}$ by restricting the linear map




FIGURE 14. The picture shows a simplicial complex K with $|K| \cong \mathbb{T}^2$, and choices of orientations on each simplex indicated via arrows (defining cyclic orderings of three vertices in the case of each 2-simplex). With these orientations fixed, plugging in the definition of $\partial : C_n^\Delta(K) \rightarrow C_{n-1}^\Delta(K)$ gives e.g. $\partial A = a - h - c$, $\partial B = h + i - k$, $\partial a = \beta - \alpha$, $\partial b = \alpha - \beta$ and so forth. The complete computation of $H_*^\Delta(K; \mathbb{Z})$ was carried out near the end of Lecture 21; needless to say, the result matches $H_*(\mathbb{T}^2; \mathbb{Z})$.

that identifies the canonical basis vectors of \mathbb{R}^{n+1} with those of \mathbb{R}^n in the chosen order. Under this identification, the inclusion $|\sigma| \hookrightarrow |K|$ becomes a singular n -simplex in $|K|$, thus defining a map

$$C_*^\Delta(K) \rightarrow C_*(|K|),$$

which is easily seen to be a chain map. It also clearly descends to the quotients $C_*^\Delta(K, K') \rightarrow C_*(|K|, |K'|)$ since each simplex in K' becomes a singular simplex in $|K'|$, and it extends to general coefficient groups $C_*^\Delta(K, K'; G) \rightarrow C_*(|K|, |K'|; G)$ in the standard way to induce a homomorphism of the homology groups $H_*^\Delta(K, K'; G) \rightarrow H_*(|K|, |K'|; G)$.

With this picture in place, we would now like to discuss a special class of singular homology classes that arise in this way from simplicial complexes.

DEFINITION 29.13. A **triangulation** (*Triangulierung*) of a pair of spaces (X, A) is a homeomorphism of (X, A) to $(|K|, |K'|)$ for some simplicial pair (K, K') . If M is a manifold with boundary, then a *triangulation of M* will be understood to mean a triangulation of the pair $(M, \partial M)$.

Given a triangulated space X , we shall often identify X with the associated polyhedron and refer accordingly to the various *skeleta* of X and its constituent simplices. If X is an n -manifold, then its local structure produces the following important observation:

PROPOSITION 29.14. *If M is a triangulated n -dimensional manifold with boundary, then the associated simplicial complex is n -dimensional, and every $(n - 1)$ -simplex σ in the complex is a boundary face of exactly either one or two n -simplices, where the former is the case if and only if σ belongs to the subcomplex triangulating ∂M . \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 14.

To get the most mileage out of this, we need to add a condition involving orientations. Given an oriented n -simplex $[v_0, \dots, v_n]$ with $n \geq 2$ in a complex, let us call the oriented $(n - 1)$ -simplex $[v_1, \dots, v_n]$ an **oriented boundary face** of $[v_0, \dots, v_n]$. Notice that for each $k = 0, \dots, n$, we can use an even permutation of v_0, \dots, v_n to move v_k into the initial position and thus produce an oriented boundary face containing all vertices other than v_k . This definition therefore assigns an orientation to every boundary face of an oriented n -simplex for $n \geq 2$. For the case $n = 1$, it is useful at this point to abuse terminology and redefine the words **oriented 0-simplex** to mean a 0-simplex with an attached sign, i.e. such objects can be written as $\pm\{v\}$ for a vertex v . We then define the oriented boundary faces of an oriented 1-simplex $[v_0, v_1]$ to be the oriented 0-simplices $\{v_1\}$ and $-\{v_0\}$.

DEFINITION 29.15. 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 in ∂M , the two orientations it inherits as an oriented boundary face of two distinct oriented n -simplices (cf. Prop. 29.14) are opposite.

I recommend now taking another look at Figure 14 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 15, 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 29.15. (The trouble is with the 1-simplices labeled c and d .) The problem with the Klein bottle is of course that it is a *non-orientable* manifold, and it turns out that only orientable manifolds can admit oriented triangulations—we sketched a proof of this for surfaces last semester in Lecture 20, and we will be able to prove it for all manifolds later in this course using homology.

EXAMPLE 29.16. The triangulation of S^{n-1} described in Example 29.5 can be oriented by choosing an ordering of the vertex set V , regarding this as an oriented n -simplex σ and then viewing each $(n - 1)$ -simplex of the triangulation as an oriented boundary face of σ . 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.

THEOREM 29.17. *For any compact n -manifold M with an oriented triangulation given by a simplicial pair (K, K') ,³⁹ there is a distinguished nontrivial simplicial homology class*

$$[M] \in H_n^\Delta(K, K'; \mathbb{Z})$$

³⁹It is a pair because M may have boundary, so K' in this case defines a triangulation of ∂M .




FIGURE 15. A triangulation of the Klein bottle that fails to be oriented.

represented by the sum of all the oriented n -simplices in the triangulation. If the triangulation is not assumed oriented, then summing the n -simplices with arbitrary orientations defines a distinguished nontrivial class

$$[M] \in H_n^\Delta(K, K'; \mathbb{Z}_2).$$

PROOF. Let us denote the sum of the oriented n -simplices of K by $\sum_i \sigma_i \in C_n^\Delta(K)$, where the orientations are either determined by the oriented triangulation or are arbitrary. In the latter case, Proposition 29.14 implies that every $(n-1)$ -simplex not in ∂M will appear with an even coefficient in $\partial \sum_i \sigma_i$, thus it vanishes after taking the tensor product with \mathbb{Z}_2 . This leaves only terms that are in the subcomplex K' triangulating ∂M , thus $\sum_i \sigma_i$ defines a relative cycle in $C_n^\Delta(K, K'; \mathbb{Z}_2)$. The resulting homology class $[\sum_i \sigma_i] \in H_n^\Delta(K, K'; \mathbb{Z}_2)$ is obviously nontrivial since the simplicial complex is n -dimensional, implying $C_{n+1}^\Delta(K, K'; \mathbb{Z}_2) = 0$. We can say more if the triangulation is oriented, because the condition on $(n-1)$ -simplices in Definition 29.15 then implies that each one not in ∂M appears in two terms of $\partial \sum_i \sigma_i$ that cancel each other out. Thus there is no longer any need to take a tensor product with \mathbb{Z}_2 , as $\sum_i \sigma_i$ is already a relative cycle in $C_n^\Delta(K, K'; \mathbb{Z})$, and again represents a nontrivial class in $H_n^\Delta(K, K'; \mathbb{Z})$ since $C_{n+1}^\Delta(K, K'; \mathbb{Z}) = 0$. \square

REMARK 29.18. The above theorem also reveals that if the triangulation is oriented, then for any coefficient group G , the sum $g \sum_i \sigma_i \in C_n^\Delta(K, K'; G)$ over the oriented n -simplices σ_i defines a relative cycle and thus a class in $H_n^\Delta(K, K'; G)$. The same is true without orientations if G has the property that all of its nontrivial elements have order 2, e.g. one could take $G = \mathbb{Z}_2 \oplus \mathbb{Z}_2$ instead

of just \mathbb{Z}_2 . In either case, it is not too hard to prove that $H_n^\Delta(K, K'; G) \cong G$ if M is connected, and if G is \mathbb{Z} or \mathbb{Z}_2 , then the distinguished class $[M]$ generates the group.

In light of the map $H_*^\Delta(K, K'; G) \rightarrow H_*(M, \partial M; G)$ discussed previously, Theorem 29.17 gives us distinguished singular homology classes

$$[M] \in H_n(M, \partial M; \mathbb{Z}) \quad \text{or} \quad [M] \in H_n(M, \partial M; \mathbb{Z}_2),$$

where the latter is defined for any compact triangulated n -manifold, and the former if the triangulation is oriented. We call $[M]$ in either case a **fundamental class** (*Fundamentalklasse*) for M . We will typically be most interested in the case where M is a *closed* manifold, so that $[M]$ is an absolute homology class in $H_n(M; G)$ for $G = \mathbb{Z}$ or \mathbb{Z}_2 . Near the end of this semester we will be able to show that $[M]$ does not actually depend on the choice of triangulation except (in the case with \mathbb{Z} coefficients) for its orientation, and in fact, *all* compact topological n -manifolds have distinguished n -dimensional homology classes that can be called fundamental classes—this is true even for manifolds that do not admit triangulations. We are very far at this point from being able to prove that, but there is in any case a large intuitive advantage to thinking about fundamental classes in terms of triangulations: it is not that hard to visualize them.

In fact, if (X, A) is any pair of spaces (not necessarily manifolds), we can now picture fairly general⁴⁰ homology classes in $H_k(X, A)$ in terms of continuous maps $f : (M, \partial M) \rightarrow (X, A)$ where M is a compact k -manifold with an oriented triangulation: any such map defines a class

$$[f] := f_*[M] \in H_k(X, A; \mathbb{Z}).$$

Moreover, we can allow a wider class of (not necessarily orientable) manifolds if we are only interested in \mathbb{Z}_2 coefficients: given any triangulation (oriented or not) of a compact k -manifold M , any continuous map $f : (M, \partial M) \rightarrow (X, A)$ defines a class

$$[f] := f_*[M] \in H_k(X, A; \mathbb{Z}_2).$$

In the most interesting applications, X is typically a smooth n -manifold and we are mainly interested in *absolute* homology classes in $H_k(X)$. An important special case of the above construction is then when $M \subset X$ is a *submanifold* of X and we take $f : M \hookrightarrow X$ to be the inclusion: in this way, triangulated submanifolds $M \subset X$ define homology classes $[M] \in H_k(X; \mathbb{Z}_2)$, or $[M] \in H_k(X; \mathbb{Z})$ whenever the triangulation is oriented.

It is now natural to ask: under what circumstances can we say that two maps $f : M \rightarrow X$ and $g : N \rightarrow X$ defined on triangulated k -manifolds determine the same homology class? The natural answer to this question makes singular homology look a lot like the *bordism* theory that we sketched last semester in Lecture 21. The details are worked out in the following exercise.

EXERCISE 29.19. In each of the following, (X, A) is a pair of spaces, W is a compact triangulated $(k + 1)$ -manifold with boundary, M , M_0 and M_1 are closed triangulated k -manifolds, all maps are continuous and all triangulations are assumed oriented (except in part (e)).

- (a) Show that if $M = \partial W$ and $f = F|_M : M \rightarrow X$ for some map $F : W \rightarrow X$, then $[f] = 0 \in H_k(X; \mathbb{Z})$.

Hint: As in the proof of Theorem 29.17, consider $\partial \sum_{i=1}^N \sigma_i \in C_k(W)$ where $\sigma_1, \dots, \sigma_N$ is the finite set of oriented $(k + 1)$ -simplices in the triangulation of W . Many terms cancel, but some do not.

⁴⁰Even if X is a compact smooth manifold, it is not quite true that *every* class in $H_k(X; \mathbb{Z})$ is of the form $f_*[M]$ for some map $f : M \rightarrow X$ defined on a closed oriented k -manifold. A famous theorem of Thom [Tho54] asserts however that it is almost true: given any $A \in H_k(X; \mathbb{Z})$, there exists a closed oriented k -manifold M , a map $f : M \rightarrow X$ and a number $m \in \mathbb{N}$ such that $mA = f_*[M]$. This implies in particular that every class in $H_k(X; \mathbb{Q})$ can be described via a map from a closed oriented k -manifold into X .




FIGURE 16. The picture shows a scenario as in Exercise 29.19 part (b), where M and W are triangulated submanifolds of X , both with nonempty boundary, and the maps f and F are defined as inclusions. The consequence is that $f : (M, \partial M) \rightarrow (X, A)$ represents the trivial relative homology class in $H_k(X, A; \mathbb{Z})$.

- (b) Generalizing part (a), suppose $f = F|_M : M \rightarrow X$ for some map $F : W \rightarrow X$, but M is a compact subset of ∂W that is both a subcomplex and a k -dimensional submanifold with boundary, such that $F(\overline{\partial W \setminus M}) \subset A$, so in particular $f(\partial M) \subset A$. (See Figure 16.) Show that $[f] = 0 \in H_k(X, A; \mathbb{Z})$.
- (c) Given maps $f_0 : M_0 \rightarrow X$ and $f_1 : M_1 \rightarrow X$, let $f : M_0 \sqcup M_1 \rightarrow X$ denote the map that restricts to M_i as f_i for $i = 0, 1$. Show that $[f] = [f_0] + [f_1] \in H_k(X; \mathbb{Z})$.
- (d) Show that for any map $f : M \rightarrow X$, reversing the orientations of all n -simplices in an oriented triangulation of M reverses the sign of $[f] \in H_k(X; \mathbb{Z})$.
- (e) Show that parts (a), (b) and (c) remain valid with \mathbb{Z}_2 coefficients if all orientation hypotheses are dropped.

EXERCISE 29.20. Let $\Sigma_{1,2}$ denote the 2-torus with two holes cut out, and suppose $\alpha, \beta : S^1 \hookrightarrow \partial \Sigma_{1,2}$ are loops parametrizing its two boundary components, with α following the boundary orientation of $\partial \Sigma_{1,2}$ and β following the opposite orientation (see Figure 17). Show that if we choose an oriented triangulation of $\Sigma_{1,2}$ so that α and β inherit oriented triangulations (with the




FIGURE 17. The surface in Exercise 29.20

orientation on β reversed), then α and β represent the same class in $H_1(\Sigma_{1,2}; \mathbb{Z})$. (One says in this case that α and β are **homologous** (*homolog*). One can show that they are not homotopic.)

There is one major loose end in this discussion: while it is obvious that the fundamental classes $[M]$ we have defined for triangulated n -manifolds are nontrivial elements in *simplicial* homology, we have not yet proved that the latter is isomorphic to $H_*(M)$, thus we have to entertain the possibility that $[M] = 0 \in H_n(M)$, i.e. there could be a singular $(n + 1)$ -chain in M whose boundary gives $[M]$. Actually, it is not too hard to prove using only our present knowledge of singular homology that $[M]$ is always not only nontrivial but is a *primitive* element in $H_n(M)$, meaning it is not a nontrivial integer multiple of any other element. We will return to this subject in a few lectures when we discuss local orientations and the degrees of maps between manifolds.

30. The Eilenberg-Steenrod axioms, triples and good pairs

In the computation of $H_*(S^n; G)$ in Lecture 28, we never had to make any specific reference to singular simplices or any other aspects of the definition of singular homology. We did need to know that singular homology has a particular set of properties, e.g. functoriality, homotopy invariance, long exact sequences and excision, and we needed to understand a variant of the theory known as *reduced* homology, which was defined in terms of the unique map from any space to the one-point space. It will turn out that almost all computations of $H_*(X; G)$ we can carry out for a “reasonable” class of spaces depend on exactly this same list of properties. This realization motivated Eilenberg and Steenrod in the 1950’s to codify a set of axioms for so-called “homology theories”. In the early days of homology, there were in fact several competing homology theories that differed substantially in their definitions but nonetheless seemed *mostly* to be measuring the same information about topological spaces. The Eilenberg-Steenrod axioms provided an explanation for this similarity. While singular homology is now easily the most popular of the original homology theories, the others did not completely die out—I have seen some of the others used from time to time in research papers on subjects I cared about, and there are still authors who argue that other theories are preferable from certain points of view. We will sketch some examples of alternative homology theories later in this course.

DEFINITION 30.1. An **axiomatic homology theory** h_* is a covariant functor

$$\mathrm{Top}_{\mathrm{rel}} \rightarrow \mathrm{Ab}_{\mathbb{Z}} : (X, A) \mapsto h_*(X, A) = \bigoplus_{n \in \mathbb{Z}} h_n(X, A),$$

together with a natural transformation ∂_* from the functor $\mathrm{Top}_{\mathrm{rel}} \rightarrow \mathrm{Ab} : (X, A) \mapsto h_n(X, A)$ to the functor $\mathrm{Top}_{\mathrm{rel}} \rightarrow \mathrm{Ab} : (X, A) \mapsto h_{n-1}(A)$ for each $n \in \mathbb{Z}$ such that the following axioms are satisfied:

- (EXACTNESS) For all pairs (X, A) with inclusion maps $i : A \hookrightarrow X$ and $j : (X, \emptyset) \hookrightarrow (X, A)$, 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.

- (HOMOTOPY) For any two homotopic maps $f, g : (X, A) \rightarrow (Y, B)$, the induced morphisms $f_*, g_* : h_*(X, A) \rightarrow h_*(Y, B)$ are identical.
- (EXCISION) For any pair (X, A) and any subset $B \subset X$ with closure in the interior of A , the inclusion $(X \setminus B, A \setminus B) \hookrightarrow (X, A)$ induces an isomorphism

$$h_*(X \setminus B, A \setminus B) \xrightarrow{\cong} h_*(X, A).$$

- (DIMENSION) For any space $\{\text{pt}\}$ containing only one point, $h_n(\{\text{pt}\}) = 0$ for all $n \neq 0$.
- (ADDITIVITY) For any collection of spaces $\{X_\alpha\}_{\alpha \in J}$ with inclusion maps $i^\alpha : X_\alpha \hookrightarrow \coprod_{\beta \in J} X_\beta$, the induced homomorphisms $i_*^\alpha : h_*(X_\alpha) \rightarrow h_*(\coprod_{\beta \in J} X_\beta)$ determine an isomorphism

$$\bigoplus_{\alpha \in J} i_*^\alpha : \bigoplus_{\alpha \in J} h_*(X_\alpha) \xrightarrow{\cong} h_*\left(\prod_{\beta \in J} X_\beta\right).$$

A few comments are in order.

REMARK 30.2. The original list in [ES52] includes three other axioms before exactness, but the first two of these are equivalent to the statement that $h_* : \mathbf{Top}_{\text{rel}} \rightarrow \mathbf{Ab}_{\mathbb{Z}}$ is a functor, and the third simply requires ∂_* to be a natural transformation. We could equally well have chosen to hide the homotopy axiom by calling h_* a functor $\mathbf{Top}_{\text{rel}}^h \rightarrow \mathbf{Ab}_{\mathbb{Z}}$ instead of $\mathbf{Top}_{\text{rel}} \rightarrow \mathbf{Ab}_{\mathbb{Z}}$.

REMARK 30.3. The additivity axiom did not appear in [ES52] but was added later by Milnor [Mil62]. One can show in fact that for *finite* disjoint unions, additivity follows as a consequence of the other axioms (see Problem Set 3 #3), thus Eilenberg and Steenrod did not need it because they were mainly concerned with computations for compact polyhedra—these come from finite simplicial complexes, so no infinite disjoint unions are allowed.

REMARK 30.4. I am cheating slightly by stating a stronger variant of the excision axiom than appeared in the original list by Eilenberg and Steenrod. The version in [ES52] reads as follows:

- (EXCISION') For any pair (X, A) and any *open* subset $B \subset X$ with closure in the interior of A , the inclusion $(X \setminus B, A \setminus B) \hookrightarrow (X, A)$ induces an isomorphism

$$h_*(X \setminus B, A \setminus B) \xrightarrow{\cong} h_*(X, A).$$

This means that there might in principle exist theories that satisfy the original Eilenberg-Steenrod axioms but not ours, because $h_*(X \setminus B, A \setminus B) \rightarrow h_*(X, A)$ might fail to be an isomorphism in cases where $\bar{B} \subset \overset{\circ}{A}$ but B is not open. Note that we already have applied this axiom in cases where B is not open, e.g. in the argument of Lecture 28 to prove $\tilde{H}_k(X; G) \cong \tilde{H}_{k+1}(SX; G)$, we considered the inclusion $(SX \setminus \{p_-\}, C_-X \setminus \{p_-\}) \hookrightarrow (SX, C_-X)$. This is fine in singular homology because $H_*(\cdot; G)$ satisfies the stronger axiom, but even for some hypothetical theory h_* that satisfies (EXCISION') and not (EXCISION), one could prove that $h_*(SX \setminus \{p_-\}, C_-X \setminus \{p_-\}) \rightarrow h_*(SX, C_-X)$ is an isomorphism by relating it via homotopy invariance to the map

$$h_*(SX \setminus B_\epsilon(p_-), C_-X \setminus B_\epsilon(p_-)) \rightarrow h_*(SX, C_-X)$$

for some small open neighborhood $B_\epsilon(p_-) \subset C_-X$ of p_- , and this map definitely is an isomorphism since $B_\epsilon(p_-)$ is open. In practice, some trick of this sort will be available in every important situation where we need to apply excision, so that it will not really matter which version of the axiom we adopt. We'll opt for the stronger one in this course since several arguments would become slightly longer without it.

REMARK 30.5. It is sometimes useful to expand the definition and allow an axiomatic homology theory to be a functor $\mathcal{C} \rightarrow \mathbf{Ab}_{\mathbb{Z}}$ 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 important example we will see later is the category of **compact pairs**, which are simply pairs of spaces (X, A) such that X is compact Hausdorff and $A \subset X$ is closed. (In reference to Remark 30.4, notice that the category of compact pairs (X, A) requires the weaker version of the excision axiom since $(X \setminus B, A \setminus B)$ will not be an object in the category unless $B \subset X$ is open.) When allowing restrictions of this type, one must take care so that all of the maps needed for expressing the axioms—e.g. the inclusions $A \hookrightarrow X$ and $(X, \emptyset) \hookrightarrow (X, A)$ —are actually morphisms. In [ES52], this concern motivates the definition of the notion of an *admissible* category of pairs.

You may notice that the definition above makes no mention of any coefficient group. It's there, actually, it's just hidden.

DEFINITION 30.6. The **coefficient group**⁴¹ of an axiomatic homology theory is defined to be the group $h_0(\{\text{pt}\})$.

The properties of singular homology listed in Lectures 27 and 28 can now be summarized thus:

THEOREM 30.7. *For any abelian group G , singular homology $H_*(\cdot; G)$ is an axiomatic homology theory with coefficient group G .* \square

REMARK 30.8. The axioms do *not* imply the first two properties of $H_*(\cdot; G)$ that we discussed in Lecture 27: the relation of $H_0(X; G)$ and $H_1(X; \mathbb{Z})$ to $\pi_0(X)$ and $\pi_1(X)$ respectively. We will later see an example of a theory that satisfies all of the Eilenberg-Steenrod axioms, but not these two properties. One has to look at fairly strange spaces in order to see this difference, e.g. spaces that are connected but not path-connected.

EXERCISE 30.9. Show that if the dimension axiom is dropped from Definition 30.1, then for any axiomatic homology theory h_* , the functor $\text{Top} \rightarrow \text{Ab}_{\mathbb{Z}}$ sending each pair (X, A) to the graded abelian group $G_* = \bigoplus_{n \in \mathbb{Z}} G_n$ with $G_n := h_{n+1}(X, A)$ is also an axiomatic homology theory. (This explains why the dimension axiom is called what it is: it guarantees that there is some connection between the subscript n in the notation $h_n(X, A)$ and our intuitive notion of “dimension,” i.e. $\{\text{pt}\}$ is clearly a “zero-dimensional” space.)

REMARK 30.10. One can define singular homology with a trivial coefficient group: it still satisfies the axioms, but it is not very interesting, because $H_*(X, A; G)$ is then trivial for all pairs. Weirdly, it is possible in general for an axiomatic homology theory to be nontrivial on some spaces (though not on “nice” spaces like polyhedra) even if its coefficient group $h_0(\{\text{pt}\})$ is trivial. Look at the axioms again: you'll see that there is no obvious reason why this couldn't be allowed. My first instinct when I learned about these axioms was to try to prove as an exercise that $h_0(\{\text{pt}\}) = 0$ implies $h_0(X, A) = 0$ for all (X, A) , but fortunately I did not spend much time on this exercise—it turns out that someone else thought about it in 1957 and came up with counterexamples [JW58].

DEFINITION 30.11. Any axiomatic homology theory h_* has a corresponding **reduced** theory \tilde{h}_* , defined by $\tilde{h}_*(X, A) := h_*(X, A)$ whenever $A \neq \emptyset$ and

$$\tilde{h}_*(X) := \ker \epsilon_* \subset h_*(X)$$

for the unique map $\epsilon : X \rightarrow \{\text{pt}\}$.

⁴¹There is a slightly awkward semantic issue in this definition: strictly speaking, what we are calling “ $\{\text{pt}\}$ ” is not a unique space, but simply *any* choice of space that happens to contain only one element. It follows that the coefficient group $h_0(\{\text{pt}\})$ is not a uniquely defined group, but is an isomorphism class of groups. Any two choices of one-point spaces P_0 and P_1 are related by a unique homeomorphism $P_0 \rightarrow P_1$, which induces a canonical isomorphism $h_0(P_0) \rightarrow h_0(P_1)$.

The next result follows by exactly the same arguments as in the case of reduced singular homology, cf. Propositions 28.10, 28.11 and 28.12, and Theorem 28.14.

THEOREM 30.12. *For any axiomatic homology theory h_* with coefficient group G :*

- (1) $\tilde{h}_*(X) = 0$ for all contractible spaces X .
- (2) There is a split exact sequence $0 \rightarrow \tilde{h}_*(X) \hookrightarrow h_*(X) \xrightarrow{\epsilon_*} h_*(\{\text{pt}\}) \rightarrow 0$ for all spaces X , giving rise to isomorphisms

$$h_n(X) \cong \begin{cases} \tilde{h}(X) \oplus G & \text{for } n = 0, \\ \tilde{h}(X) & \text{for } n \neq 0. \end{cases}$$

- (3) Morphisms $h_*(X, A) \rightarrow h_*(Y, B)$ induced by maps of pairs $(X, A) \rightarrow (Y, B)$ restrict to $\tilde{h}_*(X, A) \rightarrow \tilde{h}_*(Y, B)$, so that \tilde{h}_* defines a functor $\text{Top}_{\text{rel}}^h \rightarrow \text{Ab}_{\mathbb{Z}}$.
- (4) All connecting homomorphisms $\partial_* : h_n(X, A) \rightarrow h_{n-1}(A)$ have image in $\tilde{h}_{n-1}(A)$ and the resulting sequence

$$\dots \rightarrow \tilde{h}_{n+1}(X, A) \xrightarrow{\partial_*} \tilde{h}_n(A) \xrightarrow{i_*} \tilde{h}_n(X) \xrightarrow{j_*} \tilde{h}_n(X, A) \xrightarrow{\partial_*} \tilde{h}_{n-1}(A) \rightarrow \dots$$

is exact.

□

All together this is enough information to repeat nearly verbatim the computations of Lecture 28 involving suspensions and spheres. We obtain:

THEOREM 30.13. *For any axiomatic homology theory h_* and any space X , there is an isomorphism $\tilde{h}_n(X) \rightarrow \tilde{h}_{n+1}(SX)$ for every $n \in \mathbb{Z}$.* □

THEOREM 30.14. *For every $n \in \mathbb{N}$ and every axiomatic homology theory h_* with coefficient group G ,*

$$h_k(S^n) \cong \begin{cases} G & \text{for } k = 0, n, \\ 0 & \text{for all other } k. \end{cases}$$

□

A large portion of the theorems we prove about singular homology in this course will be based only on the axioms, and will thus be valid for any axiomatic homology theory. Proofs based on the axioms are traditionally considered more elegant than those that require explicit reference to the definition of $H_*(X, A; G)$. On the other hand, there are a few cases in which both types of proof are possible but the one that doesn't use the axioms is much easier. The exact sequence explained below is certainly an example of this.

Suppose $B \subset A \subset X$, so (X, A) , (X, B) and (A, B) are all pairs of spaces, with obvious inclusion maps of pairs

$$i : (A, B) \hookrightarrow (X, B) \quad \text{and} \quad j : (X, B) \hookrightarrow (X, A).$$

These then induce a short exact sequence of relative singular chain complexes

$$0 \longrightarrow C_*(A, B) \xrightarrow{i_*} C_*(X, B) \xrightarrow{j_*} C_*(X, A) \longrightarrow 0.$$

The special case of this with $B = \emptyset$ reproduces the usual short exact sequence for the pair (X, A) . Taking the tensor product with an arbitrary coefficient group G and then applying Proposition 27.7 as before gives the so-called **long exact sequence of the triple** (X, A, B) :

(30.1)

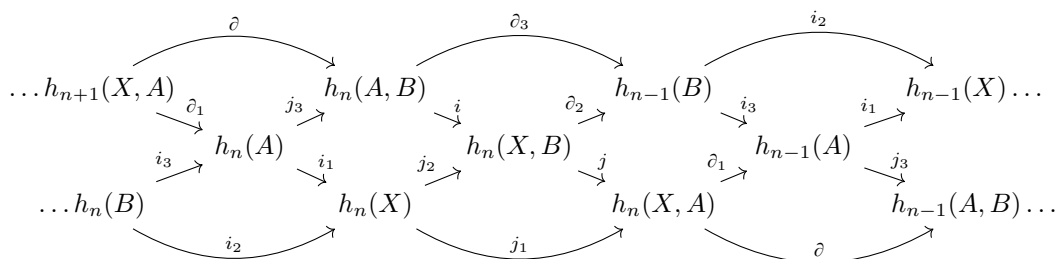
$$\dots \rightarrow H_{n+1}(X, A; G) \xrightarrow{\partial_*} H_n(A, B; G) \xrightarrow{i_*} H_n(X, B; G) \xrightarrow{j_*} H_n(X, A; G) \xrightarrow{\partial_*} H_{n-1}(A, B; G) \rightarrow \dots,$$

which directly generalizes the long exact sequence of (X, A) . It is also not hard to show that the connecting homomorphism $\partial_* : H_n(X, A; G) \rightarrow H_{n-1}(A, B; G)$ satisfies a naturality property, i.e. given any map of pairs $f : (X, A) \rightarrow (X', A')$ and a subset $B' \subset A'$ such that $f(B) \subset B'$, we have commutative diagrams

$$\begin{array}{ccc} H_n(X, A; G) & \xrightarrow{\partial_*} & H_{n-1}(A, B; G) \\ \downarrow f_* & & \downarrow f_* \\ H_n(X', A'; G) & \xrightarrow{\partial_*} & H_{n-1}(A', B'; G) \end{array}$$

for all $n \in \mathbb{Z}$. Moreover, close inspection of the usual diagram chasing argument yields an explicit formula for $\partial_* : H_n(X, A; G) \rightarrow H_{n-1}(A, B; G)$: if $[c] \in H_n(X, A; G)$ is represented by a relative cycle $c \in C_n(X; G)$, then $\partial c \in C_{n-1}(A; G)$ is a cycle in A and therefore also a relative cycle for the pair (A, B) , so that it represents a class in $H_{n-1}(A, B; G)$ and the formula $\partial_*[c] = [\partial c]$ thus makes sense. All of this is proved by nearly the same arguments as in the case $B = \emptyset$, so I will spare you the details.

We will use the exact sequence of the triple several times in our efforts to simplify computations of singular homology, and for this purpose the presentation above is certainly sufficient. Nonetheless, you might now be wondering: is this sequence really a property distinctive to the singular chain complex, or does it work for every axiomatic homology theory, and if so, how can one derive a connecting homomorphism $\partial_* : h_n(X, A) \rightarrow h_{n-1}(A, B)$ from the axioms? The answer to this question requires some cleverness, and I'm at a loss to conjecture how anyone might have come up with it for the first time, but here it is: given an axiomatic homology theory h_* and a triple (X, A, B) with $B \subset A \subset X$, we can consider the following “braid” diagram:



The braid consists of four “strands,” three of which you may recognize as the long exact sequences of the pairs (X, A) , (X, B) and (A, B) . The fourth strand is the sequence

$$(30.2) \quad \dots \rightarrow h_{n+1}(X, A) \xrightarrow{\partial} h_n(A, B) \xrightarrow{i} h_n(X, B) \xrightarrow{j} h_n(X, A) \xrightarrow{\partial} h_{n-1}(A, B) \rightarrow \dots,$$

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.

EXERCISE 30.15. Deduce via the following steps that the sequence (30.2) appearing as the fourth strand in the braid diagram above is exact:

- (a) Use the commutativity of the diagram to show that $i \circ \partial = 0$ and $\partial \circ j = 0$.
Hint: Each can be expressed as a different composition that includes two successive maps in an exact sequence.
- (b) Prove that $j \circ i = 0$ by factoring it through the group $h_*(A, A)$, which is always zero. (Why?)

- (c) Use a purely algebraic diagram-chasing argument to prove that the kernel of each map in the sequence (30.2) is contained in the image of the previous one.

Here is a useful application of the long exact sequence of a triple. Since $H_*(X, A)$ is defined so as to measure the topology of X while ignoring anything that happens entirely in A , it is natural to expect some relationship between this and the *absolute* homology of the space X/A defined by collapsing $A \subset X$ to a point. Here we should restrict attention to the case where $A \subset X$ is closed, since X/A may otherwise be a horrible (e.g. non-Hausdorff) space. It turns out that under a further assumption on the pair (X, A) , the relative homology $H_*(X, A)$ is naturally isomorphic to $\tilde{H}_*(X/A)$. To see this, we start by observing that there is a natural isomorphism between the reduced homology of X/A and the relative homology of the pair $(X/A, A/A)$, in which the subset $A/A \subset X/A$ is actually just a single point. Indeed:

LEMMA 30.16. *For any space X and a point $x \in X$, the inclusion of pairs $(X, \emptyset) \hookrightarrow (X, \{x\})$ induces an isomorphism*

$$\tilde{h}_*(X) \xrightarrow{\cong} \tilde{h}_*(X, \{x\}) = h_*(X, \{x\}),$$

where h_* is any axiomatic homology theory.

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 we're removing needs to have its closure contained in the interior of the smaller set in the pair, which is usually not true if those two sets are the same. Conclusion: we need to impose a condition on (X, A) so that A lies strictly inside of something else that will allow us to apply excision. The following bit of informal terminology is borrowed from [Hat02].

DEFINITION 30.17. A pair of spaces (X, A) will be called **good** if $A \subset X$ is a closed subset and is a deformation retract of some neighborhood $V \subset X$ of itself.

EXAMPLE 30.18. $(\mathbb{D}^n, \partial\mathbb{D}^n)$ 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 30.19. 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 30.21 below, due to the following exercise:

EXERCISE 30.20. Show that for the pair (X, A) in Example 30.19, $H_1(X/A) \not\cong H_1(X, A)$. *Hint: $H_1(X, A)$ 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 $H_1(X/A)$, you might notice that X/A is homeomorphic to the so-called "Hawaiian earring," whose fundamental group is discussed in [Hat02, Example 1.25]. As Hatcher explains, $\pi_1(X/A)$ admits a surjective homomorphism onto the uncountable abelian group $\prod_{i=1}^{\infty} \mathbb{Z}$. Conclude that $H_1(X/A)$ also admits such a surjection and is therefore uncountable.*

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 30.21. *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 30.16 that $h_*(X, A) \cong \tilde{h}_*(X/A)$.

PROOF. Fix a neighborhood $V \subset X$ of A such that A is a deformation retract of V . Then the inclusion $(A, A) \hookrightarrow (V, A)$ is a homotopy equivalence of pairs, thus

$$h_*(V, A) \cong h_*(A, A) = 0,$$

where the latter vanishes due to the long exact sequence of (A, A) . Writing down the long exact sequence of the triple (X, V, A) then gives

$$0 = h_k(V, A) \rightarrow h_k(X, A) \rightarrow h_k(X, V) \rightarrow h_{k-1}(V, A) = 0,$$

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; note that this is where it is crucial for V to be a *neighborhood* of the closed subset A , as the interior of V therefore contains the closure of A . The rightmost map labeled q_* in this diagram is also an isomorphism since it is induced by the map

$$(X \setminus A, V \setminus A) \xrightarrow{q} ((X/A) \setminus (A/A), (V/A) \setminus (A/A)),$$

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

REMARK 30.22. We used the strong version of the excision axiom in the above proof since the interior subsets A or A/A being removed from the pairs (X, V) or $(X/A, V/A)$ respectively are closed and typically not open. To make the argument work with the weaker version of excision mentioned in Remark 30.4, one would need to impose an extra requirement on the pair (X, A) so that instead of removing A we could remove a homotopy equivalent open neighborhood of it that lives inside the neighborhood V . This would make the conditions defining a *good pair* slightly stricter, but they would still be satisfied by almost all pairs we will ever care about.

The following simple example will appear frequently when we compute the homology of CW-complexes.

EXAMPLE 30.23. 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 $G := h_0(\{\text{pt}\})$ is the coefficient group. (Of course it is also not hard to compute this more directly using the reduced long exact sequence of $(\mathbb{D}^n, \partial\mathbb{D}^n)$, in which the connecting homomorphism $h_k(\mathbb{D}^n, \partial\mathbb{D}^n) \rightarrow \tilde{h}_{k-1}(S^{n-1})$ is an isomorphism.)

31. The Mayer-Vietoris sequence

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 $X = A \cup B$ such that the interiors of A and B cover X , and we would like to compute $H_*(X)$ in terms of $H_*(A)$, $H_*(B)$ and $H_*(A \cap B)$. Let us first consider this problem specifically for singular homology, and after that consider the generalization to arbitrary axiomatic homology theories.

In singular homology, we already know from Lemma 28.1 that the inclusion of subgroups

$$C_*(A; G) + C_*(B; G) \hookrightarrow C_*(X; G)$$

is a chain homotopy equivalence, so it induces an isomorphism of homology groups

$$(31.1) \quad H_*(C_*(A; G) + C_*(B; G)) \xrightarrow{\cong} H_*(X; G).$$

The question is then how to relate the homology group on the left hand side of this isomorphism to the homologies of the individual chain complexes $C_*(A; G)$ and $C_*(B; G)$. By now, you will perhaps not be surprised to learn that the answer involves an exact sequence. We can define a short exact sequence of chain complexes

$$0 \longrightarrow C_*(A \cap B; G) \xrightarrow{(i_*^A, -i_*^B)} C_*(A; G) \oplus C_*(B; G) \xrightarrow{k^A \oplus k^B} C_*(A; G) + C_*(B; G) \longrightarrow 0,$$

where $i^A : A \cap B \hookrightarrow A$ and $i^B : A \cap B \hookrightarrow B$ are the obvious inclusions of spaces, and $k^A : C_*(A; G) \hookrightarrow C_*(A; G) + C_*(B; G)$ and $k^B : C_*(B; G) \hookrightarrow C_*(A; G) + C_*(B; G)$ are the inclusions of subgroups of $C_*(X; G)$. It is trivial to verify that this sequence is exact; the crucial detail for this is the sign reversal in the map $C_*(A \cap B; G) \rightarrow C_*(A; G) \oplus C_*(B; G)$, which sends chains $c \in C_n(A \cap B; G)$ to $(c, -c) \in C_n(A; G) \oplus C_n(B; G)$, where we have simplified the notation since singular chains in $A \cap B$ are naturally also singular chains in each of the spaces A and B . The short exact sequence induces a long exact sequence of homology groups by Proposition 27.7, but we need to take a moment to consider what the homology of each of these chain complexes actually is. The first is straightforward: it is simply $H_*(A \cap B; G)$. For the middle term, we are defining the boundary map on the chain complex $C_*(A; G) \oplus C_*(B; G)$ in the obvious way that preserves the direct sum splitting, so the homology of this complex is also a direct sum, namely $H_*(A; G) \oplus H_*(B; G)$. The last term has already been discussed: its homology is isomorphic via the obvious inclusion of chain complexes to $H_*(X; G)$, and the resulting map $H_*(A; G) \oplus H_*(B; G) \rightarrow H_*(X; G)$ is then simply (j_*^A, j_*^B) , induced by the continuous inclusion maps $j^A : A \hookrightarrow X$ and $j^B : B \hookrightarrow X$. We've proved:

THEOREM 31.1. *If $A, B \subset X$ are subsets such that $X = \overset{\circ}{A} \cup \overset{\circ}{B}$ and*

$$i^A : A \cap B \hookrightarrow A, \quad i^B : A \cap B \hookrightarrow B, \quad j^A : A \hookrightarrow X, \quad j^B : B \hookrightarrow X,$$

denote the obvious inclusions, then there exist connecting homomorphisms $\partial_* : H_n(X; G) \rightarrow H_{n-1}(A \cap B; G)$ for every $n \in \mathbb{Z}$ such that the sequence

$$\begin{aligned} \dots \longrightarrow H_{n+1}(X; G) \xrightarrow{\partial_*} H_n(A \cap B; G) \xrightarrow{(i_*^A, -i_*^B)} H_n(A; G) \oplus H_n(B; G) \\ \xrightarrow{j_*^A \oplus j_*^B} H_n(X; G) \xrightarrow{\partial_*} H_{n-1}(A \cap B; G) \longrightarrow \dots \end{aligned}$$

is exact. Moreover, the connecting homomorphisms are natural in the sense that for any map $f : X \rightarrow X'$ such that $X' = \overset{\circ}{A}' \cup \overset{\circ}{B}'$ for subsets $A', B' \subset X'$ with $f(A) \subset A'$ and $f(B) \subset B'$, the diagram

$$\begin{array}{ccc} H_n(X; G) & \xrightarrow{\partial_*} & H_{n-1}(A \cap B; G) \\ \downarrow f_* & & \downarrow f_* \\ H_n(X'; G) & \xrightarrow{\partial_*} & H_{n-1}(A' \cap B'; G) \end{array}$$

commutes. □

The exact sequence in this theorem is known as the **Mayer-Vietoris sequence** for the homology of $X = \overset{\circ}{A} \cup \overset{\circ}{B}$. Note that the naturality part of the statement follows from the functoriality of the short \rightarrow long exact sequence construction, cf. Proposition 27.7 and Problem Set 2 #1.

EXERCISE 31.2. Prove that a formula for $\partial_* : H_n(X; G) \rightarrow H_{n-1}(A \cap B; G)$ can be written as follows: every element of $H_n(X; G)$ can be written in the form $[a + b]$ for singular n -chains $a \in C_n(A; G)$ and $b \in C_n(B; G)$, so the condition that $a + b$ is a cycle implies $\partial a = -\partial b \in C_{n-1}(A \cap B; G)$, and

$$\partial_*[a + b] = [\partial a] = -[\partial b] \in H_{n-1}(A \cap B; G).$$

I promise to look at an example or two, but first we have a theoretical issue to deal with. The construction above seems to depend heavily on the details of the singular chain complex rather than just the axioms of a homology theory—but does it really? As with the exact sequence of triples in the previous lecture, the answer is no, but a rather clever diagram-chase is required in order to prove it. I will show you the diagram and let you work out the details as an exercise.

Assume h_* is an arbitrary axiomatic homology theory and $X = \overset{\circ}{A} \cup \overset{\circ}{B}$. Notice first that the inclusion of pairs $(B, A \cap B) \hookrightarrow (X, A)$ induces an excision isomorphism

$$h_*(B, A \cap B) \xrightarrow{\cong} h_*(X, A).$$

Indeed, $(B, A \cap B)$ is obtained from (X, A) by removing the subset $A \setminus (A \cap B) = A \cap (X \setminus B)$, whose closure is in $\overset{\circ}{A}$ because there must otherwise exist a point x that is not in $\overset{\circ}{A}$ and therefore is in $\overset{\circ}{B}$, while simultaneously $x \in \overline{A \setminus (A \cap B)}$, meaning every neighborhood of x contains points that are in A and not in B , contradicting $x \in \overset{\circ}{B}$. We can now place the long exact sequences of the pairs $(B, A \cap B)$ and (X, A) into the top and bottom rows of the diagram

$$\begin{array}{ccccccccc} \dots & \longrightarrow & h_{n+1}(B) & \longrightarrow & h_{n+1}(B, A \cap B) & \longrightarrow & h_n(A \cap B) & \xrightarrow{i_*^B} & h_n(B) & \longrightarrow & h_n(B, A \cap B) & \longrightarrow & \dots \\ & & \downarrow j_*^B & & \downarrow \cong & & \downarrow i_*^A & & \downarrow j_*^B & & \downarrow \cong & & \\ \dots & \longrightarrow & h_{n+1}(X) & \longrightarrow & h_{n+1}(X, A) & \longrightarrow & h_n(A) & \xrightarrow{j_*^A} & h_n(X) & \longrightarrow & h_n(X, A) & \longrightarrow & \dots \end{array}$$

Here the rows are both exact, all maps in the diagram are either the connecting homomorphisms from long exact sequences of pairs or are induced by obvious inclusion maps, and as usual the diagram commutes due to the commutativity of those inclusion maps plus the naturality of connecting homomorphisms. The following diagram-chasing exercise transforms this diagram into a Mayer-Vietoris sequence for $h_*(X)$:

EXERCISE 31.3. Assume the following diagram commutes and that its top and bottom rows (both including the A_* terms) are exact.

$$\begin{array}{ccccccccccccccc}
 \cdots & \longrightarrow & C_{n+1} & & & & B_n & \xrightarrow{j} & C_n & & & & B_{n-1} & \longrightarrow & \cdots \\
 & & \downarrow \ell & \searrow f & & & \downarrow m & & \downarrow \ell & \searrow f & & & \downarrow m & & \\
 \cdots & \longrightarrow & E_{n+1} & & A_n & \xrightarrow{h} & D_n & \xrightarrow{k} & E_n & & A_{n-1} & \xrightarrow{h} & D_{n-1} & \longrightarrow & \cdots \\
 & & \uparrow g & & \downarrow i & & \uparrow g & & \downarrow i & & & & & &
 \end{array}$$

Prove that

$$\cdots \longrightarrow E_{n+1} \xrightarrow{h \circ g} B_n \xrightarrow{(m, -j)} D_n \oplus C_n \xrightarrow{k \oplus \ell} E_n \xrightarrow{h \circ g} B_{n-1} \longrightarrow \cdots$$

is then an exact sequence.

The exercise produces not only an axiomatic version of Theorem 31.1, but also a formula for $\partial_* : h_n(X) \rightarrow h_{n-1}(A \cap B)$. In the abstract diagram of the exercise it is the composition $h \circ g : E_n \rightarrow B_{n-1}$. In our situation, $g : E_n \rightarrow A_{n-1}$ is just the map $h_n(X) \rightarrow h_n(X, A)$ induced by the inclusion $(X, \emptyset) \hookrightarrow (X, A)$. We then need to replace $h_n(X, A)$ by $h_n(B, A \cap B)$ using the inverse of the excision map $h_n(B, A \cap B) \xrightarrow{\cong} h_n(X, A)$, so that $h : A_{n-1} \rightarrow B_{n-1}$ becomes the connecting homomorphism $h_n(B, A \cap B) \rightarrow h_{n-1}(A \cap B)$ from the exact sequence of $(B, A \cap B)$. To summarize:

THEOREM 31.4. *In the setting of Theorem 31.1, any axiomatic homology theory h_* also admits a Mayer-Vietoris exact sequence*

$$\cdots \longrightarrow h_n(A \cap B) \xrightarrow{(i_*^A, -i_*^B)} h_n(A) \oplus h_n(B) \xrightarrow{j_*^A \oplus j_*^B} h_n(X) \xrightarrow{\partial_*} h_{n-1}(A \cap B) \longrightarrow \cdots,$$

with natural connecting homomorphisms $\partial_* : h_n(X) \rightarrow h_{n-1}(A \cap B)$ determined by the diagram

$$(31.2) \quad \begin{array}{ccc}
 h_n(X) & \xrightarrow{\partial_*} & h_{n-1}(A \cap B) \\
 \downarrow & & \uparrow \\
 h_n(X, A) & \xrightarrow{\cong} & h_n(B, A \cap B)
 \end{array}$$

which is formed out of the obvious inclusions $(B, A \cap B) \hookrightarrow (X, A)$ and $(X, \emptyset) \hookrightarrow (X, A)$ together with the connecting homomorphism from the long exact sequence of the pair $(B, A \cap B)$. \square

REMARK 31.5. The naturality of the connecting homomorphism $\partial_* : h_n(X) \rightarrow h_{n-1}(A \cap B)$ can be deduced directly from the diagram (31.2).

There is also an analogue for reduced homology, which follows from a similar argument to the one we used for the long exact sequence of the pair. Indeed, we can put the Mayer-Vietoris sequence for $X = \hat{A} \cup \hat{B}$ into a commutative diagram together with the Mayer-Vietoris sequence

for $\{\text{pt}\} = \{\text{pt}\} \cup \{\text{pt}\}$ as follows:

$$\begin{array}{cccccccc}
 & & 0 & & 0 & & 0 & & 0 & & \\
 & & \downarrow & & \downarrow & & \downarrow & & \downarrow & & \\
 \dots & & \tilde{h}_n(A \cap B) & & \tilde{h}_n(A) \oplus \tilde{h}_n(B) & & \tilde{h}_n(X) & & \tilde{h}_{n-1}(A \cap B) & & \dots \\
 & & \downarrow & & \downarrow & & \downarrow & & \downarrow & & \\
 \dots & \longrightarrow & h_n(A \cap B) & \longrightarrow & h_n(A) \oplus h_n(B) & \longrightarrow & h_n(X) & \longrightarrow & h_{n-1}(A \cap B) & \longrightarrow & \dots \\
 & & \downarrow \epsilon_* & & \downarrow \epsilon_* \oplus \epsilon_* & & \downarrow \epsilon_* & & \downarrow \epsilon_* & & \\
 \dots & \longrightarrow & h_n(\{\text{pt}\} \cap \{\text{pt}\}) & \longrightarrow & h_n(\{\text{pt}\}) \oplus h_n(\{\text{pt}\}) & \longrightarrow & h_n(\{\text{pt}\}) & \longrightarrow & h_{n-1}(\{\text{pt}\} \cap \{\text{pt}\}) & \longrightarrow & \dots \\
 & & \downarrow & & \downarrow & & \downarrow & & \downarrow & & \\
 & & 0 & & 0 & & 0 & & 0 & &
 \end{array}$$

Since all columns of this diagram are exact and so are the bottom two nontrivial rows, Proposition 28.13 provides uniquely determined maps on the top row that preserve the commutativity of the diagram and make the top row exact:

THEOREM 31.6. *In the setting of Theorem 31.4, all maps can be restricted to the respective reduced homology groups to produce an exact sequence*

$$\dots \longrightarrow \tilde{h}_n(A \cap B) \xrightarrow{(i_*^A, -i_*^B)} \tilde{h}_n(A) \oplus \tilde{h}_n(B) \xrightarrow{j_*^A \oplus j_*^B} \tilde{h}_n(X) \xrightarrow{\partial_*} \tilde{h}_{n-1}(A \cap B) \longrightarrow \dots,$$

□

REMARK 31.7. If one prefers to assume only the weaker version of the excision axiom as mentioned in Remark 30.4, then some additional condition must be imposed on the subsets $A, B \subset X$ in Theorem 31.4 to make sure that $(B, A \cap B) \hookrightarrow (X, A)$ is a valid excision map, or equivalently, that $A \setminus (A \cap B) = A \cap (X \setminus B)$ is open. It suffices for instance to assume that B is closed, so in particular, everything is fine in the category of compact pairs, where only closed subsets $A, B \subset X$ are allowed.

Now let's see the Mayer-Vietoris sequence in action.

EXAMPLE 31.8. Here's the *easy* way to see the isomorphism $\tilde{h}_n(X) \cong \tilde{h}_{n+1}(SX)$. Write $SX = C_+X \cup_X C_-X$ as usual where $C_+X = (X \times [0, 1]) / (X \times \{1\})$ and $C_-X = (X \times [-1, 0]) / (X \times \{-1\})$, pick $\epsilon > 0$ small and define two new subsets of SX by

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

In other words, $A \subset SX$ is a neighborhood of C_+X that also deformation retracts to C_+X , and B is similarly a neighborhood of C_-X , so that the interiors of A and B cover SX . Notice also that the intersection $A \cap B = X \times [-\epsilon, \epsilon]$ deformation retracts to X . We can then write down the Mayer-Vietoris sequence

$$\dots \longrightarrow \tilde{h}_{n+1}(A) \oplus \tilde{h}_{n+1}(B) \longrightarrow \tilde{h}_{n+1}(SX) \longrightarrow \tilde{h}_n(A \cap B) \longrightarrow \tilde{h}_n(A) \oplus \tilde{h}_n(B) \longrightarrow \dots$$

By homotopy invariance, $\tilde{h}_*(A) \cong \tilde{h}_*(C_+X) = 0$ and $\tilde{h}_*(B) \cong \tilde{h}_*(C_-X) = 0$, while $\tilde{h}_*(A \cap B) \cong \tilde{h}_*(X)$, so we obtain from this an exact sequence $0 \rightarrow \tilde{h}_{n+1}(SX) \rightarrow \tilde{h}_n(X) \rightarrow 0$ and thus an isomorphism $\tilde{h}_{n+1}(SX) \cong \tilde{h}_n(X)$.

EXAMPLE 31.9. Let's compute the singular homology $H_*(\mathbb{T}^2)$ of the 2-torus, with integer coefficients. We can decompose $\mathbb{T}^2 = S^1 \times S^1$ as the union of two overlapping annuli: simply write S^1 as the union of two intervals $I_+, I_- \subset S^1$ whose interiors cover the whole circle, then set

$A = I_+ \times S^1$ and $B = I_- \times S^1$. This makes both A and B homotopy equivalent to S^1 , and since the intersection $I_+ \cap I_-$ necessarily contains two disjoint intervals, $A \cap B$ is homotopy equivalent to $S^1 \amalg S^1$. We thus have

$$H_k(A) \cong H_k(B) \cong H_k(S^1) \cong \mathbb{Z}, \quad \text{and} \quad H_k(A \cap B) \cong H_k(S^1) \oplus H_k(S^1) = \mathbb{Z} \oplus \mathbb{Z} \quad \text{for } k = 0, 1,$$

where the two factors in $H_k(A \cap B)$ correspond to the two connected components of the overlap $A \cap B$. For $k \geq 2$, we have $H_k(A) = H_k(B) = H_k(A \cap B) = 0$, so the Mayer-Vietoris sequence contains a segment of the form

$$0 = H_k(A) \oplus H_k(B) \rightarrow H_k(\mathbb{T}^2) \rightarrow H_{k-1}(A \cap B) = 0 \quad \text{for } k \geq 3,$$

implying $H_k(\mathbb{T}^2) = 0$ for all $k \geq 3$.

To understand $H_k(\mathbb{T}^2)$ for $k \leq 2$, we need to know more about the actual maps in the Mayer-Vietoris sequence. A crucial observation here is that the inclusions $i^A : A \cap B \hookrightarrow A$ and $i^B : A \cap B \hookrightarrow B$ restrict to each of the connected components of $A \cap B$ as homotopy equivalences, thus they induce isomorphisms

$$\mathbb{Z} \xrightarrow{\cong} \mathbb{Z} \cong H_k(A) \text{ or } H_k(B)$$

when restricted to each of the factors in $H_k(A \cap B) \cong \mathbb{Z} \oplus \mathbb{Z}$. We are free to choose the isomorphisms of $H_k(A)$ and $H_k(B)$ with \mathbb{Z} so that each of these maps $\mathbb{Z} \rightarrow \mathbb{Z}$ is the identity, and we can therefore write the map $\Phi_k := (i_*^A, -i_*^B) : H_k(A \cap B) \rightarrow H_k(A) \oplus H_k(B)$ for $k = 0, 1$ in the form

$$(31.3) \quad \Phi_k : \mathbb{Z} \oplus \mathbb{Z} \rightarrow \mathbb{Z} \oplus \mathbb{Z} : (m, n) \mapsto (m + n, -m - n).$$

Since $H_2(A) = H_2(B) = H_2(S^1) = 0$, let us use this term to begin the sequence and write it as

$$\begin{aligned} 0 = H_2(A) \oplus H_2(B) &\longrightarrow H_2(\mathbb{T}^2) \xrightarrow{\partial_2} H_1(A \cap B) \xrightarrow{\Phi_1} H_1(A) \oplus H_1(B) \longrightarrow H_1(\mathbb{T}^2) \\ &\xrightarrow{\partial_1} H_0(A \cap B) \xrightarrow{\Phi_0} H_0(A) \oplus H_0(B) \longrightarrow H_0(\mathbb{T}^2) \longrightarrow 0. \end{aligned}$$

The first thing we can deduce from this is that the map labeled ∂_2 is injective, and its image is $\ker \Phi_1$, which by the formula in (31.3), is

$$\ker \Phi_1 \cong \{(m, n) \in \mathbb{Z} \oplus \mathbb{Z} \mid m + n = 0\} \cong \mathbb{Z}.$$

This proves $H_2(\mathbb{T}^2) \cong \mathbb{Z}$.

To compute $H_1(\mathbb{T}^2)$, we can use a convenient trick for turning long exact sequences into short ones: observe first that while the map $H_1(A) \oplus H_1(B) \rightarrow H_1(\mathbb{T}^2)$ cannot be assumed injective, it will become injective if we quotient its domain by its kernel, which is precisely the image of Φ_1 . Similarly, $\partial_1 : H_1(\mathbb{T}^2) \rightarrow H_0(A \cap B)$ may fail to be surjective, but it trivially becomes surjective if we replace its target space with $\text{im } \partial_1$, which equals $\ker \Phi_0$. We therefore have a short exact sequence

$$0 \rightarrow \text{coker } \Phi_1 \rightarrow H_1(\mathbb{T}^2) \rightarrow \ker \Phi_0 \rightarrow 0,$$

where the **cokernel** of Φ_1 is defined as the quotient of its target by its image, i.e.

$$\text{coker } \Phi_1 := (H_1(A) \oplus H_1(B)) / \text{im } \Phi_1.$$

Using (31.3) again, $\text{coker } \Phi_1$ is then the quotient of $\mathbb{Z} \oplus \mathbb{Z}$ by the subgroup $\{(m, -m) \mid m \in \mathbb{Z}\}$, and this quotient is isomorphic to \mathbb{Z} . As luck would have it, the same subgroup is also $\ker \Phi_0$, and it is also isomorphic to \mathbb{Z} , so our short exact sequence now looks like

$$(31.4) \quad 0 \rightarrow \mathbb{Z} \rightarrow H_1(\mathbb{T}^2) \rightarrow \mathbb{Z} \rightarrow 0.$$

At this point there is a cheap trick available to finish the job: \mathbb{Z} is a free group, thus the sequence splits by Exercise 28.7, and Exercise 28.5 then provides an isomorphism $H_1(\mathbb{T}^2) \cong \mathbb{Z} \oplus \mathbb{Z}$.




FIGURE 18. The square at the left (with opposite edges identified) represents \mathbb{T}^2 , which is split by two parallel circles $S_+, S_- \subset \mathbb{T}^2$ into a pair of annuli \mathbb{T}^2_+ and \mathbb{T}^2_- such that $A \subset \mathbb{T}^2$ is a neighborhood of \mathbb{T}^2_+ and $B \subset \mathbb{T}^2$ is a neighborhood of \mathbb{T}^2_- . The chosen triangulation of \mathbb{T}^2 respects this splitting and produces a 2-cycle of the form $a+b \in C_2(\mathbb{T}^2)$ where $a \in C_2(A)$ and $b \in C_2(B)$, such that $[a+b] \in H_2(\mathbb{T}^2) \cong \mathbb{Z}$ is a generator.

Finally, the sequence also gives us a surjective map $H_0(A) \oplus H_0(B) \rightarrow H_0(\mathbb{T}^2)$ whose kernel is $\text{im } \Phi_0$, so it descends to an isomorphism $\text{coker } \Phi_0 \rightarrow H_0(\mathbb{T}^2)$. Here (31.3) implies once again that $\text{coker } \Phi_0 \cong \mathbb{Z}$.

The end result is:

$$H_k(\mathbb{T}^2) \cong \begin{cases} \mathbb{Z} & \text{for } k = 0, 2, \\ \mathbb{Z} \oplus \mathbb{Z} & \text{for } k = 1, \\ 0 & \text{for all other } k \in \mathbb{Z}. \end{cases}$$

It's worth noting that we did not use any properties of singular homology in the above computation beyond the axioms, so the result is equally valid for any axiomatic homology theory with coefficient group \mathbb{Z} . We did use a specific property of the coefficient group at one step: we assumed it was *free* in order to conclude that the sequence (31.4) splits. More generally, if h_* is any homology theory with a free coefficient group G , then since $h_0(S^1) \cong h_1(S^1) \cong G$, the same argument gives the result $h_0(\mathbb{T}^2) \cong h_2(\mathbb{T}^2) \cong G$ and $h_1(\mathbb{T}^2) \cong G \oplus G$. Actually, this is true for *all* coefficient groups, but we cannot see it so easily from this argument—we will see a much easier proof of this once we've learned how to compute cellular homology, and for singular homology it will also follow by combining the above computation with the universal coefficient theorem.

One drawback of the above method for computing $H_*(\mathbb{T}^2)$ is that Mayer-Vietoris does not make it very easy to say e.g. what the two generators of $H_1(\mathbb{T}^2) \cong \mathbb{Z} \oplus \mathbb{Z}$ look like *geometrically*. Of course you already know the answer to this question for other reasons: its generators are the same as those of $\pi_1(\mathbb{T}^2)$, i.e. they can be represented by loops of the form $S^1 \times \{\text{const}\}$ and $\{\text{const}\} \times S^1$ in \mathbb{T}^2 . It is not impossible but by no means straightforward to find both of those loops hiding in

the exact sequence argument we used above. The situation is slightly better for $H_2(\mathbb{T}^2)$. Choose a pair of parallel circles $S_+, S_- \subset A \cap B$ such that each connected component of $A \cap B$ admits a deformation retraction onto one of them. Since $H_1(S_\pm) \cong \pi_1(S_\pm)$, any loop parametrizing S_\pm then represents a generator of one of the factors in $H_1(A \cap B) \cong \mathbb{Z} \oplus \mathbb{Z}$; more generally, any oriented triangulation of S_\pm does the same thing, since a triangulation in this case just means a string of 1-simplices that can be concatenated to form a loop parametrizing the whole circle. Let us denote the generators defined in this way by $[S_+], [S_-] \in H_1(A \cap B)$. The Mayer-Vietoris sequence told us that the connecting homomorphism

$$\partial_2 : H_2(\mathbb{T}^2) \rightarrow H_1(A \cap B)$$

maps $H_2(\mathbb{T}^2) \cong \mathbb{Z}$ isomorphically to the kernel of $\Phi_1 : H_1(A \cap B) \rightarrow H_1(A) \oplus H_1(B)$, which we can deduce from (31.3) is generated by $[S_+] - [S_-]$. According to Exercise 31.2, a generator of $H_2(\mathbb{T}^2)$ can thus be represented by any singular 2-cycle of the form $a + b$ where $a \in C_2(A)$, $b \in C_2(B)$, and $\partial a = -\partial b$ is a 1-cycle in $A \cap B$ representing $[S_+] - [S_-]$. To find such a 2-cycle, observe that \mathbb{T}^2 is the union of a pair of annuli $\mathbb{T}_+^2, \mathbb{T}_-^2 \subset \mathbb{T}^2$ that each have boundary $S_+ \amalg S_-$ and intersect each other only along that boundary. Choose an oriented triangulation of \mathbb{T}^2 in which every 2-simplex is in either \mathbb{T}_+^2 or \mathbb{T}_-^2 ; an example is shown in Figure 18. As explained in Lecture 29, the sum of these oriented simplices defines a singular 2-cycle after choosing an ordering of all the vertices; that cycle represents a so-called *fundamental class* $[\mathbb{T}^2] \in H_2(\mathbb{T}^2)$. That cycle clearly also has the property that (up to a sign) $\partial_2[\mathbb{T}^2] = [S_+] - [S_-]$, so we've proved:

PROPOSITION 31.10. *There exists an oriented triangulation of \mathbb{T}^2 such that the resulting fundamental class $[\mathbb{T}^2]$ as defined in Lecture 29 generates $H_2(\mathbb{T}^2; \mathbb{Z}) \cong \mathbb{Z}$.* \square

32. Mapping tori and the degree of maps

Topic 1: The relative Mayer-Vietoris sequence. We have a few things to discuss in this lecture, but before moving on from the Mayer-Vietoris sequence, I should mention that there is also a relative version. We will not need it until much later in this course, and we will only need it for singular homology in particular, so let us sketch where it comes from in this case. Assume (X, Y) , (A, C) and (B, D) are pairs of spaces such that $X = \mathring{A} \cup \mathring{B}$ and $Y = \mathring{C} \cup \mathring{D}$. To simplify notation, let us omit the coefficient group G and also abbreviate the chain complexes

$$C_*(A + B) := C_*(A) + C_*(B) \subset C_*(X), \quad C_*(C + D) := C_*(C) + C_*(D) \subset C_*(Y),$$

$$C_*(A + B, C + D) := C_*(A + B) / C_*(C + D).$$

We can then write down a commutative diagram of chain complexes and chain maps in the form (32.1)

$$\begin{array}{ccccccc} & & 0 & & 0 & & 0 \\ & & \downarrow & & \downarrow & & \downarrow \\ 0 & \longrightarrow & C_*(C \cap D) & \longrightarrow & C_*(C) \oplus C_*(D) & \longrightarrow & C_*(C + D) \longrightarrow 0 \\ & & \downarrow & & \downarrow & & \downarrow \\ 0 & \longrightarrow & C_*(A \cap B) & \longrightarrow & C_*(A) \oplus C_*(B) & \longrightarrow & C_*(A + B) \longrightarrow 0 \\ & & \downarrow & & \downarrow & & \downarrow \\ 0 & \longrightarrow & C_*(A \cap B, C \cap D) & \longrightarrow & C_*(A, C) \oplus C_*(B, D) & \longrightarrow & C_*(A + B, C + D) \longrightarrow 0 \\ & & \downarrow & & \downarrow & & \downarrow \\ & & 0 & & 0 & & 0 \end{array}$$

The horizontal maps in this diagram are assumed to be the same ones that appear in the Mayer-Vietoris sequence, so the top two rows are just the short exact sequences that underlie the Mayer-Vietoris sequences for $Y = \mathring{C} \cup \mathring{D}$ and $X = \mathring{A} \cup \mathring{B}$ respectively. All three columns are exact for the usual reasons: the top vertical maps are actually just inclusions, and the bottom vertical maps are quotient projections (or direct sums of two such maps, in the case of the middle column). To understand the quotient complex at the lower right, we need to consider the obvious chain map

$$C_*(A + B, C + D) \rightarrow C_*(X, Y)$$

that is induced by the inclusion $C_*(A + B) \hookrightarrow C_*(X)$. The latter is a chain homotopy equivalence, and so is its restriction to the subgroup $C_*(C + D) \hookrightarrow C_*(Y)$. Moreover, there is a short exact sequence of chain complexes

$$0 \rightarrow C_*(C + D) \hookrightarrow C_*(A + B) \rightarrow C_*(A + B, C + D) \rightarrow 0,$$

which produces a long exact sequence of their corresponding homology groups. These form the top row of the following commutative diagram,

$$\begin{array}{ccccccccc} H_n(C_*(C + D)) & \rightarrow & H_n(C_*(A + B)) & \rightarrow & H_n(C_*(A + B, C + D)) & \rightarrow & H_{n-1}(C_*(C + D)) & \rightarrow & H_{n-1}(C_*(A + B)) \\ \downarrow \cong & & \downarrow \cong & & \downarrow & & \downarrow \cong & & \downarrow \cong \\ H_n(Y) & \longrightarrow & H_n(X) & \longrightarrow & H_n(X, Y) & \longrightarrow & H_{n-1}(Y) & \longrightarrow & H_{n-1}(X) \end{array}$$

in which the bottom row is the long exact sequence of (X, Y) . Since both rows are exact and four out of the five vertical maps are isomorphisms, a standard diagram-chasing exercise known as the **five-lemma** implies that the remaining vertical map is also an isomorphism, i.e.

$$H_*(C_*(A + B, C + D)) \cong H_*(X, Y).$$

For completeness, here is the statement of the five-lemma, whose proof we leave as an exercise (see Problem Set 3 #4):

LEMMA 32.1. *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:*

$$\begin{array}{ccccccccc} A & \xrightarrow{f} & B & \xrightarrow{g} & C & \xrightarrow{h} & D & \xrightarrow{i} & E \\ \downarrow \alpha & & \downarrow \beta & & \downarrow \gamma & & \downarrow \delta & & \downarrow \varepsilon \\ A' & \xrightarrow{f'} & B' & \xrightarrow{g'} & C' & \xrightarrow{h'} & D' & \xrightarrow{i'} & E' \end{array}$$

If α , β , δ and ε are all isomorphisms, then so is γ . □

Returning to the large diagram (32.1), it is obvious that the top two nontrivial rows are exact sequences, as are all three of the columns, and it is straightforward to deduce from this that the third row is also a chain complex, though it is not immediately obvious whether it is exact. If you recall Proposition 28.13 however, it should not surprise you to learn that the exactness of two of the rows implies the exactness of the third. Here is the clever way to see this: think of each row in the diagram as a chain complex, so that the vertical chain maps from top to bottom define a short exact sequence of chain complexes. (We normally draw exact sequences of chain complexes with the maps oriented horizontally and not vertically, but there's no law against doing it this way instead.) From this perspective, the short exact sequence induces a long exact sequence of homology groups, but here's the point: since the top two rows are exact, the homology groups they induce are all trivial, which means that two out of every three groups in our long exact sequence will be trivial. Exactness then forces all the other homology groups to be trivial as well, which is equivalent to the statement that the third row in the diagram is exact. Cute, no?

Putting all this together, since we can now recognize the third row in (32.1) as a short exact sequence of chain complexes, and we can identify the homology of the last term in it with $H_*(X, Y)$, the induced long exact sequence is what we will call the **relative Mayer-Vietoris sequence**

$$\begin{aligned} \dots \rightarrow H_{n+1}(X, Y; G) \rightarrow H_n(A \cap B, C \cap D; G) \rightarrow H_n(A, C; G) \oplus H_n(B, D; G) \\ \rightarrow H_n(X, Y; G) \rightarrow H_{n-1}(A \cap B, C \cap D; G) \rightarrow \dots \end{aligned}$$

We'll need this in the proof of Poincaré duality, but until then you are free to forget about it.

Topic 2: Mapping tori. I'd next like to explain another way to compute 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 map⁴²

$$i : X \hookrightarrow X_f : x \mapsto [(x, 1)].$$

THEOREM 32.2. *For any map $f : X \rightarrow X$ and its mapping torus X_f , every axiomatic homology theory h_* admits a long exact sequence*

$$\dots \longrightarrow h_{k+1}(X_f) \longrightarrow h_k(X) \xrightarrow{\mathbb{1}-f_*} h_k(X) \xrightarrow{i_*} h_k(X_f) \longrightarrow h_{k-1}(X) \longrightarrow \dots$$

Let's do an example before talking about the proof.

EXAMPLE 32.3. 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

$$\dots \longrightarrow h_k(\mathbb{T}^{n-1}) \xrightarrow{0} h_k(\mathbb{T}^{n-1}) \xrightarrow{i_*} h_k(\mathbb{T}^n) \xrightarrow{\Phi} h_{k-1}(\mathbb{T}^{n-1}) \xrightarrow{0} h_{k-1}(\mathbb{T}^{n-1}) \longrightarrow \dots$$

The triviality of the two maps $\mathbb{1} - \text{Id}_* = 0$ here means that we actually have a short exact sequence

$$(32.2) \quad 0 \longrightarrow h_k(\mathbb{T}^{n-1}) \xrightarrow{i_*} h_k(\mathbb{T}^n) \xrightarrow{\Phi} h_{k-1}(\mathbb{T}^{n-1}) \longrightarrow 0.$$

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 (32.2) 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:

$$\text{rank } H_k(\mathbb{T}^n; \mathbb{Z}) = \binom{n}{k} \quad \text{for } 0 \leq k \leq n, \quad H_k(\mathbb{T}^n; \mathbb{Z}) = 0 \text{ for } k > n.$$

EXERCISE 32.4. 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 32.2 to compute $H_*(K^2; \mathbb{Z})$ and $H_*(K^2; \mathbb{Z}_2)$.

⁴²Of 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 $f(x) = f(y)$.

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 32.2 follows immediately from the next statement:

THEOREM 32.5. *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 32.6. 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 32.2. 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

(32.3)

$$\begin{array}{ccccccc} \dots & \longrightarrow & h_{k+1}(X \times I) & \xrightarrow{\beta_*} & h_{k+1}(X \times I, X \times \partial I) & \xrightarrow{\partial_*} & h_k(X \times \partial I) & \xrightarrow{\alpha_*} & h_k(X \times I) & \xrightarrow{\beta_*} & \dots \\ & & \downarrow q_* & & \downarrow q_* & & \downarrow q_* & & \downarrow q_* & & \\ \dots & \longrightarrow & h_{k+1}(Z) & \longrightarrow & h_{k+1}(Z, Y) & \xrightarrow{\partial_*^-} & h_k(Y) & \xrightarrow{i_*} & h_k(Z) & \longrightarrow & \dots \end{array}$$

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

$$(32.4) \quad j_*^0 \oplus j_*^1 : h_k(X) \oplus h_k(X) \xrightarrow{\cong} h_k(X \times \partial I),$$

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

$$(32.5) \quad \Psi : h_k(X) \xrightarrow{\cong} \ker \alpha_* : c \mapsto (j_*^0 \oplus j_*^1)(-c, c) = j_*^1 c - j_*^0 c.$$

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

$$(32.6) \quad q_* \circ \Psi : h_k(X) \rightarrow h_k(Y) : c \mapsto g_* c - f_* c = (g_* - f_*)c.$$

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 30.17, Theorem 30.21 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 (32.3); 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 32.2 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.

Topic 3: The mapping degree. For a final topic today, let's talk about the homomorphism $f_* : H_n(S^n; \mathbb{Z}) \rightarrow H_n(S^n; \mathbb{Z})$, a case in which the homology groups are indeed especially simple yet still nontrivial. We'll have a lot more to say about this in the next lecture, but we'll start with a straightforward definition based on the algebraic fact that every group homomorphism $\mathbb{Z} \rightarrow \mathbb{Z}$ is defined via multiplication with some integer.

DEFINITION 32.7. The **mapping degree** (*Abbildungsgrad*) of a map $f : S^n \rightarrow S^n$ for $n \in \mathbb{N}$ is the unique integer $\deg(f) = k \in \mathbb{Z}$ such that⁴³

$$f_* : H_n(S^n; \mathbb{Z}) \rightarrow H_n(S^n; \mathbb{Z}) : c \mapsto kc.$$

The basic properties of singular homology already imply a few quite nontrivial results about this definition.

PROPOSITION 32.8. *The degree for maps $S^n \rightarrow S^n$ has the following properties.*

- (1) *If $f, g : S^n \rightarrow S^n$ are homotopic, then $\deg(f) = \deg(g)$.*

⁴³We will give a more general definition in the next lecture that also applies for $n = 0$; see Definition 33.1.

- (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 $H_*(\cdot; \mathbb{Z})$ 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 $\epsilon : S^n \rightarrow \{\text{pt}\}$ and a suitable inclusion $i : \{\text{pt}\} \hookrightarrow S^n$, thus $f_* : H_n(S^n; \mathbb{Z}) \rightarrow H_n(S^n; \mathbb{Z})$ factors through $H_n(\{\text{pt}\}; \mathbb{Z}) = 0$. Finally, the fifth property follows from standard facts about $\pi_1(S^1)$ and the natural isomorphism $H_1(S^1; \mathbb{Z}) \cong \pi_1(S^1)$. \square

Recall from Problem Set 2 #2 that the suspension $SX = C_+X \cup_X C_-X$ of a space X can be regarded as a functor $\text{Top} \rightarrow \text{Top}$ sending objects X to SX , where maps $f : X \rightarrow Y$ are transformed to maps

$$Sf : SX \rightarrow SY : [(x, t)] \mapsto [(f(x), t)].$$

In particular, any map $f : S^n \rightarrow S^n$ gives rise to a map $Sf : S^{n+1} \rightarrow S^{n+1}$ using the identification $SS^n \cong S^{n+1}$.

PROPOSITION 32.9. For any $f : S^n \rightarrow S^n$, $\deg(f) = \deg(Sf)$.

PROOF. Recall from Example 31.8 that the isomorphism $H_{n+1}(SX; \mathbb{Z}) \rightarrow H_n(X; \mathbb{Z})$ can always be constructed as the connecting homomorphism in a Mayer-Vietoris exact sequence for SX . Given a map $f : S^n \rightarrow S^n$, the naturality of this connecting homomorphism produces a commuting diagram

$$\begin{array}{ccc} H_{n+1}(S^{n+1}; \mathbb{Z}) & \xrightarrow{\partial_*} & H_n(S^n; \mathbb{Z}) \\ \downarrow (Sf)_* & & \downarrow f_* \\ H_{n+1}(S^{n+1}; \mathbb{Z}) & \xrightarrow{\partial_*} & H_n(S^n; \mathbb{Z}) \end{array}$$

where the two maps labeled ∂_* are the same isomorphism. Now if $(Sf)_*c = kc$ for some nontrivial $c \in H_{n+1}(S^{n+1}; \mathbb{Z})$, it follows that $\partial_*(Sf)_*c = k\partial_*c = f_*\partial_*c$, where $\partial_*c \in H_n(S^n; \mathbb{Z})$ is also nontrivial, hence $\deg(Sf) = k = \deg(f)$. \square

PROPOSITION 32.10. 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 $\mathbf{A} \in \text{O}(n+1)$, then $\deg(f) = \det(\mathbf{A}) = \pm 1$.

PROOF. Recall that $\text{O}(n+1)$ has exactly two path-components, which can be labeled according to whether their elements have determinant $+1$ or -1 . A given $\mathbf{A} \in \text{O}(n+1)$ thus admits a continuous path in $\text{O}(n+1)$ 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 . We prove it by induction on the dimension: for $n = 1$, reflection produces a map $S^1 \rightarrow S^1$ with winding number -1 , so the claim follows from Proposition 32.8.

Now if we assume the claim is true for reflections $f : S^n \rightarrow S^n$, it suffices to observe that $Sf : SS^n \rightarrow SS^n$ is also a reflection under a suitable identification $SS^n \cong S^{n+1}$, so the result follows from Proposition 32.9. \square

33. Local mapping degree on manifolds

The degree of a map $S^n \rightarrow S^n$ was defined in the previous lecture for each $n \geq 1$ using the fact that $H_n(S^n; \mathbb{Z}) \cong \mathbb{Z}$. That definition does not make sense for $n = 0$, since $H_0(S^0; \mathbb{Z}) \cong \mathbb{Z} \oplus \mathbb{Z}$, but we can rectify this by using the reduced homology $\tilde{H}_0(S^0; \mathbb{Z}) \cong \mathbb{Z}$. It will take a little while to see why this is a reasonable thing to do and what it means geometrically, but we can in any case now state a more general definition:

DEFINITION 33.1. The **mapping degree** (*Abbildungsgrad*) of a map $f : S^n \rightarrow S^n$ for integers $n \geq 0$ is the unique integer $\deg(f) = k \in \mathbb{Z}$ such that

$$f_* : \tilde{H}_n(S^n; \mathbb{Z}) \rightarrow \tilde{H}_n(S^n; \mathbb{Z}) : c \mapsto kc.$$

One further variation on this definition will be useful:

DEFINITION 33.2. The **mod 2 degree** of a map $f : S^n \rightarrow S^n$ for $n \geq 0$ is the number $\deg_2(f) = k \in \mathbb{Z}_2$ such that

$$f_* : \tilde{H}_n(S^n; \mathbb{Z}_2) \rightarrow \tilde{H}_n(S^n; \mathbb{Z}_2) : c \mapsto kc.$$

Since $\tilde{H}_n(S^n; \mathbb{Z}_2) \cong \mathbb{Z}_2$, the mod 2 degree is algebraically much simpler than the integer-valued degree: its value is 0 if $f_* : \tilde{H}_n(S^n; \mathbb{Z}_2) \rightarrow \tilde{H}_n(S^n; \mathbb{Z}_2)$ is the trivial map, and 1 if f_* is the identity, which is the only other possibility. It is easy to see that $\deg_2(f)$ satisfies the obvious analogues of the first four properties of $\deg(f)$ listed in Proposition 32.8, as these depend only on the fact that $H_*(\cdot; \mathbb{Z}_2)$ is a functor and satisfies the homotopy axiom. In this lecture we will focus mainly on the integer-valued degree, but most of what we say will have simpler analogues for the mod 2 degree, and will sometimes also work in greater generality for $\deg_2(f) \in \mathbb{Z}_2$ than for $\deg(f) \in \mathbb{Z}$.

The facts established in the previous lecture about $\deg(f)$ already have some quite nontrivial consequences about the topology of spheres. Here are two such results.

THEOREM 33.3. *Every map $f : S^n \rightarrow S^n$ with $\deg(f) \neq (-1)^{n+1}$ has a fixed point.*

PROOF. 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 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

$$h(t, x) = \frac{g_t(x)}{|g_t(x)|},$$

and it follows that $\deg(f)$ must match the degree of the antipodal map. The latter is the restriction to S^n of an orthogonal transformation on \mathbb{R}^{n+1} given by minus the identity matrix, so its determinant is $(-1)^{n+1}$, and this is then $\deg(g_1)$ by Proposition 32.10. \square

THEOREM 33.4 (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 we can normalize it so that without loss of generality $|V(x)| = 1$ for all x , making V a map $S^n \rightarrow S^n$ that satisfies $V(x) \perp x$ for all x . Now for each $x \in S^n$, let $P_x \subset \mathbb{R}^{n+1}$ denote the 2-dimensional plane spanned by x and $V(x)$, so $P_x \cap S^n$ is a circle in S^n , and we can imagine following a path along this circle from x through $V(x)$ and ending at $-x$. Concretely, such a path is given by the formula

$$t \mapsto f_t(x) := (\cos \pi t)x + (\sin \pi t)V(x) \in S^n \quad \text{for } t \in [0, 1],$$

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 $1 = \text{deg}(f_0) = \text{deg}(f_1) = (-1)^{n+1}$, implying n must be odd. \square

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 14.15). 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.

Suppose M is a topological manifold of dimension $n \in \mathbb{N}$ with $\partial M = \emptyset$. For each $x \in M$, we can then find an open neighborhood $\mathcal{U} \subset M$ of x with a homeomorphism $\varphi : \mathcal{U} \rightarrow \mathbb{R}^n$, and by adding a shift we can assume without loss of generality that $\varphi(x) = 0$, so φ also restricts to a homeomorphism $\mathcal{U} \setminus \{x\} \rightarrow \mathbb{R}^n \setminus \{0\}$. Choose an axiomatic homology theory h_* and consider the following string of maps:

$$(33.1) \quad h_n(M, M \setminus \{x\}) \xleftarrow{i_*} h_n(\mathcal{U}, \mathcal{U} \setminus \{x\}) \xrightarrow{\varphi_*} h_n(\mathbb{R}^n, \mathbb{R}^n \setminus \{0\}) \xleftarrow{j_*} h_n(\mathbb{D}^n, S^{n-1}) \xrightarrow{\partial_*} \tilde{h}_{n-1}(S^{n-1}),$$

where $i : (\mathcal{U}, \mathcal{U} \setminus \{x\}) \hookrightarrow (M, M \setminus \{x\})$ and $j : (\mathbb{D}^n, \partial S^{n-1}) \hookrightarrow (\mathbb{R}^n, \mathbb{R}^n \setminus \{0\})$ are the obvious inclusions, and ∂_* denotes the connecting homomorphism in the reduced long exact sequence of the pair (\mathbb{D}^n, S^{n-1}) . We claim that all of these maps are isomorphisms. For φ_* this is immediate since φ is a homeomorphism, and for i_* it follows from excision since $M \setminus \mathcal{U}$ is a closed subset contained in the open subset $M \setminus \{x\}$.⁴⁴ For j_* , you hopefully have some intuition telling you that the reason involves homotopy invariance, though we have to be a bit careful: there is no deformation retraction of pairs from $(\mathbb{R}^n, \mathbb{R}^n \setminus \{0\})$ to (\mathbb{D}^n, S^{n-1}) , even though there are deformation retractions individually from \mathbb{R}^n to \mathbb{D}^n and $\mathbb{R}^n \setminus \{0\}$ to S^{n-1} . What is true however is that since the inclusion j defines homotopy equivalences $\mathbb{D}^n \hookrightarrow \mathbb{R}^n$ and $S^{n-1} \hookrightarrow \mathbb{R}^n \setminus \{0\}$, we can put the long exact sequences of (\mathbb{D}^n, S^{n-1}) and $(\mathbb{R}^n, \mathbb{R}^n \setminus \{0\})$ together in a commutative diagram

$$\begin{array}{ccccccccc} \dots & \longrightarrow & h_n(S^{n-1}) & \longrightarrow & h_n(\mathbb{D}^n) & \longrightarrow & h_n(\mathbb{D}^n, S^{n-1}) & \longrightarrow & h_{n-1}(S^{n-1}) & \longrightarrow & h_{n-1}(\mathbb{D}^n) & \longrightarrow & \dots \\ & & \downarrow j_* & & \downarrow j_* & & \downarrow j_* & & \downarrow j_* & & \downarrow j_* & & \\ \dots & \longrightarrow & h_n(\mathbb{R}^n \setminus \{0\}) & \longrightarrow & h_n(\mathbb{R}^n) & \longrightarrow & h_n(\mathbb{R}^n, \mathbb{R}^n \setminus \{0\}) & \longrightarrow & h_{n-1}(\mathbb{R}^n \setminus \{0\}) & \longrightarrow & h_{n-1}(\mathbb{R}^n) & \longrightarrow & \dots \end{array}$$

⁴⁴Notice how we just used the assumption that manifolds are Hausdorff?

where both rows are exact and all of the five vertical maps except the middle one are already known to be isomorphisms, thus the five-lemma (Lemma 32.1) implies that so is $j_* : h_n(\mathbb{D}^n, S^{n-1}) \rightarrow h_n(\mathbb{R}^n, \mathbb{R}^n \setminus \{0\})$. Finally, the reduced exact sequence

$$0 = \tilde{h}_n(\mathbb{D}^n) \longrightarrow h_n(\mathbb{D}^n, S^{n-1}) \xrightarrow{\partial_*} \tilde{h}_{n-1}(S^{n-1}) \longrightarrow \tilde{h}_{n-1}(\mathbb{D}^n) = 0$$

implies that ∂_* is also an isomorphism. We've proved:

PROPOSITION 33.5. *For any manifold M of dimension $n \in \mathbb{N}$ with empty boundary, any point $x \in M$ and any axiomatic homology theory h_* with coefficient group G , $h_n(M, M \setminus \{x\}) \cong G$. \square*

The relative homology group $h_n(M, M \setminus \{x\})$ is sometimes called the **local homology** of M at the point x . Specializing to singular homology with integer coefficients, we can now use this to define an n -dimensional notion of *orientations* for manifolds, without mentioning triangulations.

DEFINITION 33.6. A **local orientation** of an n -manifold M without boundary at a point $x \in M$ 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 open neighborhood $\mathcal{U} \subset M$ of x that is homeomorphic to \mathbb{R}^n , 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 33.7. If M is a surface without boundary and $x \in M$, 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 (33.1): they map $[\sigma]$ to the homology class of a 1-cycle in $S^1 \cong \partial\Delta^2$ consisting of the three edges of the triangle, a loop that clearly generates $\pi_1(\partial\Delta^2) \cong H_1(\partial\Delta; \mathbb{Z})$. In this picture, we can think of a local orientation at x as a choice (up to homotopy) of a small embedded loop in M about x : since there are two directions that such a loop can wind around x , there are two choices of local orientation.

DEFINITION 33.8. Suppose M and N are manifolds of dimension $n \in \mathbb{N}$ without boundary, $f : M \rightarrow N$ is a map and $x \in M$ and $y = f(x) \in N$ are 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 this neighborhood \mathcal{U} is homeomorphic to \mathbb{R}^n . 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 (33.1), any choice of local coordinates on a Euclidean neighborhood $\mathcal{U} \subset M$ of x gives rise to an isomorphism of $H_n(\mathcal{U}, \mathcal{U} \setminus \{x\}; G)$ with $\tilde{H}_{n-1}(S^{n-1}; G) \cong G$, and we can use this isomorphism to transform the above definition into a condition about maps between spheres:

PROPOSITION 33.9. *In the setting of Definition 33.8, fix a generator $[S^{n-1}] \in H_{n-1}(S^{n-1}; \mathbb{Z})$ and local coordinates on neighborhoods $x \in \mathcal{U} \subset M$ and $y \in \mathcal{V} \subset N$ that identify both points with $0 \in \mathbb{R}^n$ 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 $H_{n-1}(S^{n-1}; \mathbb{Z})$ send $[M]_x$ and $[N]_y$ to $[S^{n-1}]$. Then if \hat{f} denotes the map f written in these coordinates as a map between neighborhoods of 0 in \mathbb{R}^n , we have*

$$\deg(f; x) = \deg \left(\frac{\hat{f}}{|\hat{f}|} \Big|_{\partial \mathbb{D}_\epsilon^n} : \partial \mathbb{D}_\epsilon^n \rightarrow S^{n-1} \right)$$

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 $\partial \mathbb{D}_\epsilon^n \rightarrow S^{n-1}$. \square

COROLLARY 33.10. *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 points $x \in M$ and $y \in N$ 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 33.9 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. \square

EXAMPLE 33.11. Continuing the discussion of Example 33.7, suppose $f : M \rightarrow N$ is a map between surfaces such that x is an isolated point in $f^{-1}(y)$ and local orientations $[M]_x \in H_2(M, M \setminus \{x\}; \mathbb{Z})$ and $[N]_y \in H_2(N, N \setminus \{y\}; \mathbb{Z})$ are fixed. 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 (33.1) a string of natural isomorphisms

$$H_n(\mathbb{R}^n, \mathbb{R}^n \setminus \{x\}; \mathbb{Z}) \cong H_n(\mathbb{D}_x^n, \partial \mathbb{D}_x^n; \mathbb{Z}) \cong H_{n-1}(S^{n-1}; \mathbb{Z}),$$

so that any choice of generator $[S^{n-1}] \in H_{n-1}(S^{n-1}; \mathbb{Z})$ 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 33.12. *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 33.10, 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 33.10, 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 32.10, 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 exercise, as it provides a criterion for existence of solutions to equations of the form $f(x) = y$ that are *stable* under small perturbations of f :

EXERCISE 33.13. Prove that 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 .

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 33.14. 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.

In order to define the global degree of maps $f : M \rightarrow N$ in general, we will need to impose a condition on the manifolds that we consider. It will later turn out that this condition is satisfied for all closed and connected manifolds that are also *orientable*,⁴⁵ though it will be a while before we are in a position to prove this.

DEFINITION 33.15. Given an axiomatic homology theory h_* , a topological manifold M of dimension $n \in \mathbb{N}$ will be called h_* -**admissible**⁴⁶ if M is closed and the obvious inclusion $i^x : (M, \emptyset) \hookrightarrow (M, M \setminus \{x\})$ induces an isomorphism

$$i_*^x : h_n(M) \xrightarrow{\cong} h_n(M, M \setminus \{x\})$$

⁴⁵We will define later what it means in general for a topological n -manifold to be orientable; a definition for the case $n = 2$ was given last semester in Lecture 20.

⁴⁶This is not a universally standard term, but it is convenient for our purposes at the moment.

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. We have already done a few such computations, so let us say what can immediately be said:

PROPOSITION 33.16. *For any axiomatic homology theory h_* , the spheres S^n are h_* -admissible for all $n \in \mathbb{N}$.*

PROOF. For any $x \in S^n$, $S^n \setminus \{x\} \cong \mathbb{R}^n$ is contractible, thus the reduced long exact sequence of $(S^n, S^n \setminus \{x\})$ takes the form

$$0 = \tilde{h}_n(S^n \setminus \{x\}) \longrightarrow \tilde{h}_n(S^n) \xrightarrow{i_*^x} h_n(S^n, S^n \setminus \{x\}) \longrightarrow \tilde{h}_{n-1}(S^n \setminus \{x\}) = 0,$$

implying that $i_*^x : h_n(S^n) = \tilde{h}_n(S^n) \rightarrow h_n(S^n, S^n \setminus \{x\})$ is an isomorphism. \square

We also proved in the last lecture that $H_n(\mathbb{T}^n; \mathbb{Z}) \cong \mathbb{Z}$ for every $n \in \mathbb{N}$, so it seems reasonable to expect that \mathbb{T}^n is \mathbb{Z} -admissible. The projective plane $\mathbb{R}\mathbb{P}^2$ and Klein bottle K^2 are definitely not \mathbb{Z} -admissible, as one can use Mayer-Vietoris sequences (see Problem Set 4 #1) to show that $H_2(\mathbb{R}\mathbb{P}^2; \mathbb{Z}) = H_2(K^2; \mathbb{Z}) = 0$, but one also obtains $H_2(\mathbb{R}\mathbb{P}^2; \mathbb{Z}_2) \cong H_2(K^2; \mathbb{Z}_2) \cong \mathbb{Z}_2$, so that both may plausibly be \mathbb{Z}_2 -admissible. We will prove all of these plausible claims in the next lecture using triangulations, and near the end of this course we will also see that *every* closed and connected manifold M is \mathbb{Z}_2 -admissible, and \mathbb{Z} -admissibility additionally holds if and only if M is orientable. For now I will ask you to take these facts on faith so that you can imagine a wide range of applicability for the following definition.

DEFINITION 33.17. Assume M and N are \mathbb{Z} -admissible manifolds of dimension $n \in \mathbb{N}$, and choose generators $[M] \in H_n(M; \mathbb{Z})$ and $[N] \in H_n(N; \mathbb{Z})$. 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 32.8 are satisfied for this new definition.

We can now state the main result relating global and local degrees.

THEOREM 33.18. *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 finite,*

$$(33.2) \quad \deg(f) = \sum_{x \in f^{-1}(y)} \deg(f; x).$$

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 (33.2) 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.

Theorem 33.18 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 always perturb f or the point $y \in N$ so that $df(x)$ is invertible for all $x \in f^{-1}(y)$, in which case Proposition 33.12 tells us that $\deg(f; x)$ is always ± 1 . One then defines $\deg(f)$ essentially as the right hand side of (33.2) 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. The latter can also be proven as a consequence of transversality results, without any knowledge of homology—the main point is that if f_0 and f_1 are homotopic, then the homotopy can be used to construct a compact oriented 1-manifold with boundary such that $\#f_0^{-1}(y) - \#f_1^{-1}(y)$ is interpreted as a count (with suitable signs) of the number of points in the boundary of this 1-manifold. But 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].⁴⁷ 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 (33.2) matches the left hand side. We will prove this in the next lecture.

34. Degrees, triangulations, and coefficients

The previous lecture left several issues unresolved: most notably, I need to convince you that the condition we called “admissibility” in Definition 33.15 is actually satisfied by a large number of manifolds beyond spheres, and I also need to prove Theorem 33.18 on the relation between global and local degrees. Let’s first observe the easiest corollary of that theorem:

COROLLARY 34.1. *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 33.18 to identify $\deg(f)$ or $\deg_2(f)$ with a suitable count of points in $f^{-1}(y)$ where $y \notin f(M)$. \square

REMARK 34.2. For maps $f : S^n \rightarrow S^n$, Corollary 34.1 follows more easily from the observation that $S^n \setminus \{y\}$ is contractible, so if $y \notin f(S^n)$ then f is homotopic to a constant.

Before proving Theorem 33.18, let us discuss the abundance of manifolds that are admissible in the sense of Definition 33.15. We saw that S^n for each $n \in \mathbb{N}$ is h_* -admissible for every axiomatic homology theory. For more general topological n -manifolds it is less easy to see whether this is true, so we will focus specifically on singular homology with coefficients in \mathbb{Z} or \mathbb{Z}_2 .

⁴⁷The differentiable approach to the mapping degree was also sketched in Exercise 19.14.

LEMMA 34.3. *Suppose M is a closed topological manifold of dimension $n \in \mathbb{N}$ with a triangulation, and $[M] \in H_n(M; \mathbb{Z}_2)$ denotes the resulting fundamental class, represented by the sum of the n -simplices in the triangulation (see Lecture 29). Then for every $x \in M$ and the natural inclusion of pairs $i^x : (M, \emptyset) \hookrightarrow (M, M \setminus \{x\})$, $i_*^x[M] \neq 0 \in H_n(M, M \setminus \{x\}; \mathbb{Z}_2)$. Moreover, if the triangulation is oriented and $[M]$ denotes the resulting fundamental class in $H_n(M; \mathbb{Z})$, then $i_*^x[M]$ is a generator of $H_n(M, M \setminus \{x\}; \mathbb{Z})$.*

PROOF. In the following, G may denote either \mathbb{Z}_2 or (only if the triangulation is oriented) \mathbb{Z} . Given $x \in M$, we claim first that the triangulation can always be adjusted without changing the fundamental class $[M]$ so that its $(n-1)$ -skeleton does not contain x . Indeed, the complement of the $(n-1)$ -skeleton in M is dense, so if x lies in the $(n-1)$ -skeleton, then we can choose coordinates in some neighborhood of x and a smooth vector field⁴⁸ supported in this neighborhood so that flowing a small distance along the vector field produces a family of homeomorphisms $\Phi^t : M \rightarrow M$ that push the $(n-1)$ -skeleton away from x . This changes the triangulation and thus changes the fundamental class $[M]$ to $\Phi_*^t[M]$, but the latter matches $[M]$ since Φ^t is homotopic to the identity.

Now assuming x is in the interior of an n -simplex σ_1 of the triangulation, we have

$$i_*^x[M] = [\sigma_1] \in H_n(M, M \setminus \{x\}; G)$$

since all the other n -simplices of the triangulation lie in $M \setminus \{x\}$. Just as in Example 33.7, we can now identify $\sigma_1(\partial\Delta^n)$ with S^{n-1} so that the natural isomorphism $H_n(M, M \setminus \{x\}; G) \rightarrow \tilde{H}_{n-1}(S^{n-1}; G)$ sends $[\sigma_1]$ to the fundamental class $[S^{n-1}]$ represented by the obvious oriented triangulation of $\partial\Delta^n$ (cf. Example 29.5). In the case $n = 1$, the triangulation of S^0 just consists of the obvious two 0-simplices (i.e. points), with opposite signs in the oriented case, and one can check explicitly in this case that $[S^0]$ is a generator of $\tilde{H}_0(S^0; G)$. Now argue by induction on the dimension: if the lemma is already proved for manifolds of dimension at most $n-1$, then in particular it is true for S^{n-1} , implying that $[S^{n-1}]$ maps to generators of the local homology $\tilde{H}_{n-1}(S^{n-1}, S^{n-1} \setminus \{y\}; G)$ for every $y \in S^{n-1}$ and is therefore a primitive element, i.e. it is not a nontrivial integer multiple of another element. Since $\tilde{H}_{n-1}(S^{n-1}; G) \cong G$, this proves that $[S^{n-1}]$ is a generator, and therefore so is $i_*^x[M]$. \square

The lemma implies a result that I promised you at the end of Lecture 29: the fundamental class $[M] \in H_n(M; G)$ defined via a triangulation is always a primitive element—indeed, if it were a nontrivial multiple of some other element, then the same would necessarily be true of $i_*^x[M]$. In cases where we also know $H_n(M; G) \cong G$ for $G = \mathbb{Z}$ or \mathbb{Z}_2 , this implies that $[M]$ generates it. At this point we have a slowly growing list of examples for which this result has been established, e.g. we can prove via Mayer-Vietoris sequences that $H_n(\mathbb{T}^n; \mathbb{Z}) \cong \mathbb{Z}$, $H_2(\mathbb{RP}^2; \mathbb{Z}_2) \cong \mathbb{Z}_2$ and $H_2(K^2; \mathbb{Z}_2) \cong \mathbb{Z}_2$, where K^2 denotes the Klein bottle. We have thus proved:

PROPOSITION 34.4. *The torus \mathbb{T}^n is \mathbb{Z} -admissible for every $n \in \mathbb{N}$, and \mathbb{RP}^2 and the Klein bottle are \mathbb{Z}_2 -admissible.* \square

REMARK 34.5. The projective plane and the Klein bottle are not \mathbb{Z} -admissible since $H_2(\mathbb{RP}^2; \mathbb{Z}) = H_2(K^2; \mathbb{Z}) = 0$. We will eventually see that this has to do with the fact that they are not orientable. We will also eventually see that \mathbb{T}^n is \mathbb{Z}_2 -admissible, but we have not proved it yet.

It's worth stating one more result on this subject that we cannot quite prove yet, but the proof will present itself once we have understood cellular homology over the next few lectures. In the setting of Lemma 34.3, suppose M is connected, and let K denote the simplicial complex

⁴⁸It is natural to object at this point that choosing a smooth vector field requires having a smooth structure, and M might not have one. But we are only doing this in a small ball and using a single coordinate chart—it may happen that our vector field does not look smooth in different coordinate charts, but this does not matter.

corresponding to its triangulation. It is then fairly easy to compute the simplicial homology of this complex in dimension n : the answer for \mathbb{Z}_2 and \mathbb{Z} coefficients is

$$H_n^\Delta(K; \mathbb{Z}_2) \cong \mathbb{Z}_2, \quad H_n^\Delta(K; \mathbb{Z}) \cong \begin{cases} \mathbb{Z} & \text{if the triangulation is orientable,} \\ 0 & \text{otherwise.} \end{cases}$$

Indeed, the $(n + 1)$ -chain group is trivial, so the homology in dimension n is the same as the group of n -cycles. In the oriented case, suppose the oriented n -simplices are denoted by $\sigma_1, \dots, \sigma_N$ and we write down an arbitrary simplicial n -chain $\sum_i m_i \sigma_i$. Then $\partial \sum_i m_i \sigma_i$ contains exactly two terms involving each individual $(n - 1)$ -simplex, and these two will cancel each other out if and only if the two coefficients in front of the corresponding n -simplices match, thus proving that $\sum_i m_i \sigma_i$ can only be a cycle if the coefficients in front of every pair of neighboring n -simplices match. Since M is connected, this can only hold if *all* the coefficients match, implying $\sum_i m_i \sigma_i$ must be an integer multiple of $\sum_i \sigma_i$, so the n -cycle group is \mathbb{Z} . If the triangulation is not orientable, then the cancelation must always fail at some $(n - 1)$ -simplex and one instead deduces that the n -cycle group with integer coefficients is trivial. But with \mathbb{Z}_2 coefficients, every n -chain can be written as $\sum_i m_i \sigma_i$ with $m_i \in \{0, 1\}$, and since $2 = 0 \in \mathbb{Z}_2$, cancelation will happen without any need for orientations for any $(n - 1)$ -simplex whose two neighboring n -simplices have the same coefficient. The conclusion is that $\sum_i \sigma_i$ is a cycle in $C_n^\Delta(K; \mathbb{Z}_2)$, and it is the only nontrivial one. If you are already willing to believe the theorem that simplicial homology is isomorphic to singular homology, this implies that $H_n(M; \mathbb{Z}_2) \cong \mathbb{Z}_2$, and $H_n(M; \mathbb{Z}) \cong \mathbb{Z}$ whenever the triangulation is oriented, while $H_n(M; \mathbb{Z}) = 0$ otherwise. The conclusion is:

PROPOSITION 34.6. *Every closed and connected n -dimensional manifold with a triangulation is \mathbb{Z}_2 -admissible, and it is also \mathbb{Z} -admissible if and only if the triangulation is orientable. \square*

We now proceed to prove the theorem about local and global degrees. Here's the statement again.

THEOREM 34.7. *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 finite,*

$$(34.1) \quad \deg(f) = \sum_{x \in f^{-1}(y)} \deg(f; x).$$

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

$$\deg_2(f) = \sum_{x \in f^{-1}(y)} \deg_2(f; x).$$

PROOF. For later convenience, we shall carry out most of the proof in the framework of an arbitrary axiomatic homology theory h_* , assuming M and N to be h_* -admissible. Write

$$f^{-1}(y) = \{x_1, \dots, x_\ell\},$$

fix a Euclidean neighborhood $\mathcal{V} \subset N$ of y , along with Euclidean neighborhoods $\mathcal{U}_k \subset M$ of the individual points x_k for $k = 1, \dots, \ell$ such that

$$f(\mathcal{U}_k) \subset \mathcal{V} \quad \text{and} \quad \mathcal{U}_k \cap \mathcal{U}_j = \emptyset \text{ for } j \neq k.$$

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

$$(34.2) \quad \begin{array}{ccccc} & & h_n(\mathcal{U}_k, \mathcal{U}_k \setminus \{x_k\}) & \xrightarrow{f_*} & h_n(\mathcal{V}, \mathcal{V} \setminus \{y\}) \\ & \swarrow & \downarrow \gamma_*^k & & \downarrow \beta \\ h_n(M, M \setminus \{x_k\}) & \xleftarrow{\alpha_*^k} & h_n(M, M \setminus f^{-1}(y)) & \xrightarrow{f_*} & h_n(N, N \setminus \{y\}) \\ & \swarrow p_*^k & \uparrow j_* & & \uparrow i_*^y \\ & & h_n(M) & \xrightarrow{f_*} & h_n(N) \end{array}$$

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 Problem Set 3 #4), thus p is also an isomorphism. If we use this to replace $h_n(M, M \setminus f^{-1}(y))$ in (34.2) 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 (34.2) then gives the formula

$$(34.3) \quad f_* j_* = \sum_{k=1}^{\ell} f_* i_*^{x_k} = i_*^y f_* : h_n(M) \rightarrow h_n(N, N \setminus \{y\}).$$

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

$$\sum_{k=1}^{\ell} f_* [M]_{x_k} = \sum_{k=1}^{\ell} \deg(f; x_k) [N]_y = i_*^y f_* [M] = \deg(f) i_*^y [N] = \deg(f) [N]_y,$$

from which the formula for the integer degree follows. The formula for the mod 2 degree follows in the same way using $h_* = H_*(\cdot; \mathbb{Z}_2)$. \square

The following result can also be extracted from this proof. We will mainly care about this result in the case $M = N = S^n$, but it can be proved somewhat more generally:

THEOREM 34.8. *Suppose M and N are smooth manifolds⁴⁹ of dimension $n \in \mathbb{N}$ that are both \mathbb{Z} -admissible and h_* -admissible for some axiomatic homology theory h_* . Then for any map $f : M \rightarrow N$ of degree $k \in \mathbb{Z}$, the homomorphism $f_* : h_n(M) \rightarrow h_n(N)$ takes the form $f_*c = kc$.*

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 (see Problem Set 3 #2). 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 f . Since M and N are smooth manifolds, one can use standard results from differential topology as in [Hir94] to perturb $f : M \rightarrow N$ to a smooth map without changing its homotopy class, and Sard’s theorem then implies that almost every point $y \in N$ is a regular value, meaning the derivative of f at every $x \in f^{-1}(y)$ is an isomorphism. The latter implies that each $x \in f^{-1}(y)$ is isolated in $f^{-1}(y)$, hence $f^{-1}(y) \subset M$ is a discrete set and therefore (since M is compact) also finite. Now write $f^{-1}(y) = \{x_1, \dots, x_\ell\}$ and, given $c \in h_n(M)$, use (34.3) to write

$$i_*^y f_* c = \sum_{k=1}^{\ell} f_* i_*^{x_k} c,$$

where the individual terms on the right hand side involve the homomorphisms

$$f_* : h_n(M, M \setminus \{x_k\}) \rightarrow h_n(N, N \setminus \{y\}).$$

Using excision and connecting homomorphisms as in Proposition 33.9, 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

$$\begin{array}{ccc} h_n(M, M \setminus \{x_k\}) & \xrightarrow{f_*} & h_n(N, N \setminus \{y\}) \\ \downarrow \cong & & \downarrow \cong \\ h_{n-1}(S^{n-1}) & \xrightarrow{\cdot \deg(f; x_k)} & h_{n-1}(S^{n-1}) \end{array}$$

Adding up these contributions for every $x_k \in f^{-1}(y)$ produces multiplication by $\deg(f)$ according to Theorem 34.7. □

One consequence of this result is that the definition of $\deg(f)$ does not actually depend on the choice to use *singular* homology in particular—we could have replaced $H_*(\cdot; \mathbb{Z})$ in Definition 33.17 with any other axiomatic homology theory with coefficient group \mathbb{Z} and would thus obtain an equivalent definition. We also can now state a straightforward relationship between $\deg(f)$ and $\deg_2(f)$:

COROLLARY 34.9. *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 $\mathbb{Z} \rightarrow \mathbb{Z}_2$.* □

⁴⁹I am assuming smoothness here because my proof uses Sard’s theorem, but I am slightly unhappy about this and think there should be another way. The argument basically works whenever one can prove that $f : M \rightarrow N$ is homotopic to a map for which some point $y \in N$ has a finite preimage. Even for continuous maps $S^n \rightarrow S^n$, it is hard to see why this should be true in general, but appealing to a couple of standard results from differential topology makes it obvious.

The more important consequence of Theorem 34.8 will be the role it plays in computing arbitrary axiomatic homology theories $h_*(X)$ on spaces X endowed with cell decompositions. The next major theorem we need to prove is that if h_* has coefficient group G and X is a so-called *CW-complex*, then $h_*(X)$ is always isomorphic to an object that is typically much easier to compute: the homology of the *cellular chain complex* with coefficients in G . The definition of the boundary map in this chain complex requires a good understanding of the homomorphisms $\tilde{h}_*(S^n) \rightarrow \tilde{h}_*(S^n)$ induced by certain continuous maps that form part of the structure of the CW-complex. Theorem 34.8 then allows us to express this data purely in terms of the mapping degree, at which point it no longer matters which axiomatic homology theory we are trying to compute—the only detail that makes any difference is the coefficient group. The theorem thus provides a large portion of the underlying reason why *all* axiomatic homology theories turn out to be isomorphic if we choose to compute them only on relatively “reasonable” spaces. We will start the discussion of CW-complexes in earnest in the next lecture.

35. CW-complexes

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 35.1. A **CW-complex** (*CW-Komplex*) or **cell complex** (*Zellkomplex*) is a topological space X that is the union of a sequence of nested subspaces

$$X^0 \subset X^1 \subset X^2 \subset \dots \subset X$$

constructed by the following inductive procedure:

- X^0 is a space with the discrete topology;
- For each $n \in \mathbb{N}$, there exists 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.

$$(35.1) \quad X^n = X^{n-1} \cup_{\varphi^n} \coprod_{\alpha \in \mathcal{K}^n} \mathbb{D}^n, \quad \text{where} \quad \varphi^n := \coprod_{\alpha \in \mathcal{K}^n} \varphi_\alpha : \coprod_{\alpha \in \mathcal{K}^n} \partial \mathbb{D}^n \rightarrow X^{n-1}.$$

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 (35.1) 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 35.2. 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 $n \geq 0$.

EXERCISE 35.3. Show that 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.

EXERCISE 35.4. Show that 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 $\alpha \in \mathcal{K}^n$.

REMARK 35.5. 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 35.2 is not weak at all, but is the *strongest* with a given property. This discrepancy is apparently the fault of J.H.C. Whitehead, whose influence on the subject was so substantial that many authors still refer to the topology of CW-complexes as “the weak topology” in the literature. Exercise 35.4 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 35.6. A **cell decomposition** (*Zellenzerlegung*) of a space X is a choice of homeomorphism from X to a CW-complex.

EXAMPLE 35.7. 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 the above example, the attaching map of the n -cell is very far from being injective, thus its characteristic map is also not injective at $\partial\mathbb{D}^n$, though the restriction of a characteristic map to the interior is always injective.

EXAMPLE 35.8. 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 injective.

EXAMPLE 35.9. 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 35.8 also gives rise to a cell decomposition of $\mathbb{R}\mathbb{P}^n = S^n/\mathbb{Z}_2$ with exactly one k -cell for each $k = 0, \dots, n$. The k -skeleton of $\mathbb{R}\mathbb{P}^n$ is then a submanifold of the form

$$X^k = \{(x_0, \dots, x_n) \in \mathbb{R}\mathbb{P}^n = S^n/\mathbb{Z}_2 \mid x_{k+1} = \dots = x_n = 0\} \cong \mathbb{R}\mathbb{P}^k.$$

In contrast to Example 35.8, the characteristic maps $\mathbb{D}^k \rightarrow \mathbb{R}\mathbb{P}^n$ for this cell decomposition are not injective: indeed, the k -cells in Example 35.8 are attached to the $(k-1)$ -skeleton S^{k-1} via a homeomorphism $S^{k-1} \rightarrow S^{k-1}$, but in $\mathbb{R}\mathbb{P}^n$ this must be understood as a map to $X^{k-1} = \mathbb{R}\mathbb{P}^{k-1} = S^{k-1}/\mathbb{Z}_2$, thus the homeomorphism $S^{k-1} \rightarrow S^{k-1}$ from Example 35.8 gets composed with the quotient projection $S^{k-1} \rightarrow \mathbb{R}\mathbb{P}^{k-1}$ and thus becomes a covering map of degree 2.

EXAMPLE 35.10. This will be harder to picture, but one can adjust Example 35.8 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

$$S^0 \subset S^1 \subset S^2 \subset S^3 \subset \dots \subset S^\infty.$$

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. As the next exercise shows, there is something a bit subtle about the topology of S^∞ .

EXERCISE 35.11. Show that if $x_k \in S^\infty$ is a convergent sequence, then there exists $n \in \mathbb{N}$ such that $x_k \in S^n$ for every k .

Hint: Given $x \in S^n \subset S^\infty$ and a sequence $x_k \in S^\infty$ such that $x_k \notin S^k$ for all k , construct a neighborhood $\mathcal{U} \subset S^\infty$ of x such that $x_k \notin \mathcal{U}$ for all k .

REMARK 35.12. The exercise reveals that S^∞ is in some sense fundamentally 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 35.13. Combining Examples 35.9 and 35.10 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 35.14. 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 13.8 in last semester's Lecture 13). 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.

EXAMPLE 35.15. Since the standard n -simplex Δ^n is homeomorphic to \mathbb{D}^n , the polyhedron of any simplicial complex is also a CW-complex whose n -cells are precisely the n -simplices, and the characteristic maps are all inclusions. The use of the term “ n -skeleton” in this context coincides with our usage in Lecture 29.

DEFINITION 35.16. A **subcomplex** of a CW-complex X is a subset $A \subset X$ that is also a CW-complex with n -skeleton $A^n = A \cap X^n$ for all $n \geq 0$, such that every cell in A is also a cell in X with the same characteristic map.

Our goal in this lecture is to get as quickly as possible to the definition of cellular homology so that we can compute some examples. For this definition to make sense in full generality, we need an observation about the point-set topology of CW-complexes that is vacuous in the case of finite complexes but nontrivial for infinite complexes. We will postpone its proof until the next lecture, and since most of the interesting examples we consider will be finite complexes anyway, you are safe in ignoring it most of the time.

PROPOSITION 35.17. *For any CW-complex X , any compact subspace $K \subset X$ is contained in a finite subcomplex of X , i.e. in a subcomplex with only finitely many cells.*

The following consequence is the reason for the term “closure-finite”:

COROLLARY 35.18. *The closure of each cell in a CW-complex intersects only finitely many other cells.* \square

REMARK 35.19. Notice that Proposition 35.17 immediately implies the result of Exercise 35.11. It is worth trying to do the exercise independently of this in order to develop some intuition as to why Proposition 35.17 is true.

We can now define the **cellular chain complex** (*zellulärer Kettenkomplex*) associated to a CW-complex X . For $n \in \mathbb{Z}$, define $C_n^{\text{CW}}(X)$ to be the trivial group if $n < 0$ and otherwise

$$C_n^{\text{CW}}(X) := \bigoplus_{\alpha \in \mathcal{K}^n} \mathbb{Z},$$

i.e. $C_n^{\text{CW}}(X)$ is the free abelian group generated by the set of n -cells e_α^n in our given cell decomposition of X . We shall denote the generators of this group as cells e_α^n , thus writing elements of $C_n^{\text{CW}}(X)$ as finite sums

$$\sum_i m_i e_{\alpha_i}^n \in C_n^{\text{CW}}(X)$$

for coefficients $m_i \in \mathbb{Z}$ and indices $\alpha_i \in \mathcal{K}^n$. The direct sum of all these groups produces a \mathbb{Z} -graded abelian group

$$C_*^{\text{CW}}(X) = \bigoplus_{n \in \mathbb{Z}} C_n^{\text{CW}}(X),$$

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:

$$(35.2) \quad \begin{array}{ccc} X^n & \xrightarrow{p_\alpha} & \mathbb{D}^n / \partial \mathbb{D}^n = S^n \\ \downarrow \text{pr} & \swarrow \Phi_\alpha & \\ X^n / (X^n \setminus e_\alpha^n) & & \end{array}$$

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 in most cases a natural bijection of $X^0 / (X^0 \setminus e_\alpha^0)$ with $S^0 = \{1, -1\}$ sending the cell e_α^0 to 1 and the equivalence class represented by every other 0-cell to -1 . The only case that still requires special consideration is when $X^0 \cong \{\text{pt}\}$, so there is only one 0-cell e_α^0 and thus $X^0 \setminus e_\alpha^0 = \emptyset$. As a convention for this case we shall define $p_\alpha : X^0 \rightarrow S^0$ so that it sends e_α^0 to $1 \in S^0$, and with this in place, we now have a definition of $p_\alpha : X^n \rightarrow S^n$ for every n .

DEFINITION 35.20. 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 (35.2). We may sometimes abbreviate the incidence number by $[\beta : \alpha]$.

Observe that whenever $\overline{e_\alpha^n} \cap e_\beta^{n-1} = \emptyset$, 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 35.18, this implies that the sum in the following definition makes sense, because it can only have finitely-many nonzero terms.

DEFINITION 35.21. For each $n \in \mathbb{N}$, the boundary map on the cellular chain complex $C_*^{\text{CW}}(X)$ is defined at degree n as the unique homomorphism $\partial : C_n^{\text{CW}}(X) \rightarrow C_{n-1}^{\text{CW}}(X)$ satisfying

$$\partial e_\alpha^n = \sum_{\beta \in \mathcal{K}^{n-1}} [e_\beta^{n-1} : e_\alpha^n] e_\beta^{n-1}$$

for each $\alpha \in \mathcal{K}^n$.

Implicit in this definition is that $\partial : C_n^{\text{CW}}(X) \rightarrow C_{n-1}^{\text{CW}}(X)$ is the trivial map for every $n \leq 0$, as it must be since its target is then the trivial group. We shall now state two important theorems whose proofs will be postponed: the first states simply that $(C_*^{\text{CW}}(X), \partial)$ is a chain complex.

THEOREM 35.22. *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 can now be defined as

$$H_*^{\text{CW}}(X) := H_*(C_*^{\text{CW}}(X), \partial).$$

Similarly, we can introduce any abelian coefficient group G and define

$$H_*^{\text{CW}}(X; G) := H_*(C_*^{\text{CW}}(X) \otimes G, \partial \otimes \mathbf{1}),$$

with the case $G = \mathbb{Z}$ reproducing the previous definition. The notation $H_*^{\text{CW}}(X; G)$ is in some sense slightly non-ideal, as it hides the fact that the definition of $H_*^{\text{CW}}(X; G)$ depends 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 35.23. *For any CW-complex X and any axiomatic homology theory h_* with coefficient group G , there is an isomorphism $H_*^{\text{CW}}(X; G) \cong h_*(X)$.*

We will improve this statement slightly in the next lecture by talking about CW-pairs and morphisms of CW-pairs, so that the isomorphism between $H_*^{\text{CW}}(\cdot; G)$ and h_* can be regarded as a natural transformation. Theorem 35.23 has several remarkable consequences that can be recognized immediately: one is that $H_*^{\text{CW}}(X; G)$ depends (up to isomorphism) only on the topology of X and not on its cell decomposition, and another is that all axiomatic homology theories are isomorphic if we restrict them to spaces that are nice enough to have cell decompositions. In light of Example 35.15, this also tells us why the *simplicial* homology of a polyhedron depends only on its topology and not on its simplicial decomposition—simplicial homology is just the special case of cellular homology for CW-complexes that take the form of simplicial complexes.

Before trying to explain why all this is true, let's look at a couple of examples that will make Theorem 35.23 look more plausible.

EXAMPLE 35.24. We saw in Example 35.7 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 $X^0 = X^1 = \dots = X^{n-1} \cong \{\text{pt}\}$ and $X^n = S^n$. These two cells are thus the only generators of $C_*^{\text{CW}}(S^n; G)$, 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 $H_*^{\text{CW}}(S^n; G) = C_*^{\text{CW}}(S^n; G)$, 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 $\partial : C_k^{\text{CW}}(S^n; G) \rightarrow C_{k-1}^{\text{CW}}(S^n; G)$ is trivial. When $n = 1$ there is still something to check: $\partial : C_1^{\text{CW}}(S^1; G) \rightarrow C_0^{\text{CW}}(S^1; G)$ 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 $\varphi : S^0 \rightarrow X^0 \cong \{\text{pt}\}$ 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 35.25. We consider S^2 with the alternative cell decomposition described in Example 35.8, 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 $C_k^{\text{CW}}(S^2; G) = 0$ for $k < 0$ or $k > 2$, while $C_k^{\text{CW}}(S^2; G) = G \oplus G$ for each $k = 0, 1, 2$, with the two factors of G corresponding to the two

generators $e_+^k, e_-^k \in C_k^{\text{CW}}(S^n; \mathbb{Z})$. Denote the attaching map for e_\pm^k by $\varphi_\pm^k : S^{k-1} \rightarrow X^{k-1}$, and denote the projection map as defined in (35.2) by $p_\pm^k : X^k \rightarrow S^k$, so $\partial : C_k^{\text{CW}}(S^n; G) \rightarrow C_{k-1}^{\text{CW}}(S^n; G)$ is now determined by

$$(35.3) \quad \begin{aligned} \partial e_+^k &= \deg(p_+^{k-1} \circ \varphi_+^k) e_+^{k-1} + \deg(p_-^{k-1} \circ \varphi_+^k) e_-^{k-1}, \\ \partial e_-^k &= \deg(p_+^{k-1} \circ \varphi_-^k) e_+^{k-1} + \deg(p_-^{k-1} \circ \varphi_-^k) e_-^{k-1}. \end{aligned}$$

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 (35.3) 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 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 $\deg(p_+^1 \circ \varphi_+^2) = -1$. 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 $\deg(p_-^1 \circ \varphi_-^2) = 1$. We thus have

$$\partial e_+^2 = \partial e_-^2 = -e_+^1 + e_-^1.$$

With these formulas in place, we can compute the homology of $C_*^{\text{CW}}(S^2; G)$ 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 $C_2^{\text{CW}}(S^2; G)$ with $G \oplus G$, the group of 2-cycles takes the form

$$\ker \partial_2 = \{(g, -g) \in G \oplus G \mid g \in G\} \subset C_2^{\text{CW}}(S^2; G),$$

which is isomorphic to G . Since $C_3^{\text{CW}}(S^2; G) = 0$, we conclude $H_2^{\text{CW}}(S^2; G) \cong G$. 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 $H_1^{\text{CW}}(S^2; G) = 0$.

Finally, all 0-chains $ge_+^0 + he_-^0$ are cycles since $C_{-1}^{\text{CW}}(S^2; G) = 0$, but under the obvious isomorphism $C_0^{\text{CW}}(S^2; G) = G \oplus G$ we have

$$\text{im } \partial_1 = \{(g, -g) \in G \oplus G \mid g \in G\} \subset C_0^{\text{CW}}(S^2; G),$$

so $H_0^{\text{CW}}(S^2; G)$ 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 35.24.

It is not too hard to extend Example 35.25 to a computation of $H_*^{\text{CW}}(S^n; G)$ 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 $C_0^{\text{CW}}(S^n; G)/\text{im } \partial_1$ 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 35.24, so we will invest considerable effort over the next couple of lectures into proving that this is allowed, because the isomorphism class of $H_*^{\text{CW}}(X; G)$ depends in general only on the topology of X and not on its cell decomposition.

36. Invariance of cellular homology, part 1

Having defined the cellular homology $H_*^{\text{CW}}(X; G)$ for a CW-complex X in the previous lecture, we would now like to begin working toward the proof that it is isomorphic to $h_*(X)$ for any axiomatic homology theory h_* with coefficient group G . The proper statement of that result is Theorem 36.9 below. But first, there was a more basic result about CW-complexes that we left unproved in the previous lecture, without which the cellular chain complex is not generally well defined.

PROPOSITION 36.1. *For any CW-complex X , any compact subspace $K \subset X$ is contained in a finite subcomplex of X , i.e. in a subcomplex with only finitely many cells.*

PROOF. *Step 1:* Suppose $A \subset X$ is a subset with the property that for every pair of distinct elements $x, y \in A$, x and y belong to different cells of the complex. We claim then that $A \cap X^n$ is a closed subset of X^n for every integer $n \geq 0$. 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, Exercise 35.3 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

Let's briefly recall how the cellular chain complex $C_*^{CW}(X)$ is defined. Each chain group $C_n^{CW}(X)$ is freely generated by the set of n -cells e_α^n in X , and $\partial : C_n^{CW}(X) \rightarrow C_{n-1}^{CW}(X)$ is determined by the formula

$$\partial e_\alpha^n = \sum_{e_\beta^{n-1}} [e_\beta^{n-1} : e_\alpha^n] e_\beta^{n-1},$$

where the sum is over all the $(n-1)$ -cells e_β^{n-1} in X , and Proposition 36.1 implies that only finitely many terms are nonzero. The most important detail here is the incidence number $[e_\beta^{n-1} : e_\alpha^n] \in \mathbb{Z}$, which is the degree of the composition of two maps

$$S^{n-1} \xrightarrow{\varphi_\alpha} X^{n-1} \xrightarrow{p_\beta} S^{n-1},$$

where φ_α is the attaching map for e_α^n , and p_β is defined by collapsing everything outside of e_β^{n-1} to a point and using the characteristic map $\Phi_\beta : \mathbb{D}^{n-1} \rightarrow X^{n-1}$ to identify the resulting quotient with $\mathbb{D}^{n-1}/\partial\mathbb{D}^{n-1} = S^{n-1}$.

As mentioned in the previous lecture, this description of p_β doesn't quite work when $n = 1$, so let us work out a more useful formula for $\partial_1 : C_1^{CW}(X) \rightarrow C_0^{CW}(X)$. If $X^0 \cong \{\text{pt}\}$, 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 $\partial_1 = 0$. 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 $C_0^{CW}(X)$ and thus deduce from the remarks above the following:

PROPOSITION 36.2. *The map $\partial : C_1^{CW}(X) \rightarrow C_0^{CW}(X)$ is determined by the formula*

$$\partial e_\alpha^1 = \varphi_\alpha(1) - \varphi_\alpha(-1).$$

\square

One immediately derives from this result a formula for $\partial : C_1^{CW}(X; G) \rightarrow C_0^{CW}(X; G)$ with any coefficient group G .

Let's do another easy example.

EXAMPLE 36.3. We saw in Example 35.14 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 $C_1^{CW}(\Sigma_g)$:

$$\partial e_{a_j}^1 = \partial e_{b_j}^1 = 0 \quad \text{for } j = 1, \dots, g.$$

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

$$H_k^{\text{CW}}(\Sigma_g; G) = C_k^{\text{CW}}(\Sigma_g; G) \cong \begin{cases} G & \text{for } k = 0, 2, \\ G^{2g} & \text{for } k = 1, \\ 0 & \text{for } k < 0 \text{ and } k > 2. \end{cases}$$

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 $C_*^{\text{CW}}(A; G)$ is a subcomplex of $C_*^{\text{CW}}(X; G)$, i.e. it is a subgroup preserved by the boundary map, giving rise to a quotient chain complex

$$C_*^{\text{CW}}(X, A; G) := C_*^{\text{CW}}(X; G) / C_*^{\text{CW}}(A; G).$$

The homology of this complex is the **relative cellular homology**

$$H_*^{\text{CW}}(X, A; G) := H_*(C_*^{\text{CW}}(X, A; G)).$$

By this point you should not be surprised to learn that one can define a category CW_{rel} whose objects are CW-pairs, but I still need to tell you what its morphisms are.

DEFINITION 36.4. 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 36.5. If X and Y are polyhedra (and therefore also CW-complexes as explained in Example 35.15), then any simplicial map $f : X \rightarrow Y$ is also a cellular map.

Unlike 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 ∂ . The key point 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

$$(36.1) \quad S^n = \mathbb{D}^n / \partial \mathbb{D}^n \xrightarrow{\Phi_\alpha} X^n / X^{n-1} \xrightarrow{f} Y^n / Y^{n-1} \xrightarrow{\text{pr}} Y^n / (Y^n \setminus e_\beta^n) \xrightarrow{\Phi_\beta^{-1}} \mathbb{D}^n / \partial \mathbb{D}^n = S^n,$$

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

$$[e_\beta^n : e_\alpha^n] \in \mathbb{Z},$$

a number that can be defined for any pair of n -cells $e_\beta^n \subset Y$ and $e_\alpha^n \subset X$ if $f : X \rightarrow Y$ is a cellular map. It vanishes whenever $e_\beta^n \cap \overline{f(e_\alpha^n)} = \emptyset$ since the map in (36.1) is in this case constant, so Proposition 36.1 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

$$f_* : C_*^{\text{CW}}(X; G) \rightarrow C_*^{\text{CW}}(Y; G)$$

acting on the generators $e_\alpha^n \in C_n^{\text{CW}}(X)$ as

$$(36.2) \quad f_* e_\alpha^n = \sum_{e_\beta^n} [e_\beta^n : e_\alpha^n] e_\beta^n \in C_n^{\text{CW}}(Y),$$

where the sum ranges over all n -cells $e_\beta^n \subset Y$ and has only finitely many nonzero terms.

EXERCISE 36.6. 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, $f_* : C_*^{\text{CW}}(X; G) \rightarrow C_*^{\text{CW}}(Y; G)$ is the identity homomorphism.

EXERCISE 36.7. Show that if $f : X \rightarrow Y$ and $g : Y \rightarrow Z$ are cellular maps then $(g \circ f)_* = g_* \circ f_* : C_*^{\text{CW}}(X; G) \rightarrow C_*^{\text{CW}}(Z; G)$.

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 $C_*^{\text{CW}}(A; G)$ into $C_*^{\text{CW}}(B; G)$ and thus descends to a homomorphism

$$f_* : C_*^{\text{CW}}(X, A; G) \rightarrow C_*^{\text{CW}}(Y, B; G).$$

The proof of the next theorem will arise naturally from the proof of the much bigger theorem that follows it:

THEOREM 36.8. For any map of CW-pairs $f : (X, A) \rightarrow (Y, B)$, $f_* : C_*^{\text{CW}}(X, A; G) \rightarrow C_*^{\text{CW}}(Y, B; G)$ is a chain map and thus induces homomorphisms $f_* : H_n^{\text{CW}}(X, A; G) \rightarrow H_n^{\text{CW}}(Y, B; G)$ for every n . In particular, cellular homology with coefficients in G defines a functor

$$H_*^{\text{CW}}(\cdot; G) : \text{CW}_{\text{rel}} \rightarrow \text{Ab}_{\mathbb{Z}},$$

where CW_{rel} denotes the category of CW-pairs, with morphisms defined as maps of CW-pairs.

We can now state the complete version of the theorem about cellular homology and axiomatic homology theories.

THEOREM 36.9. Suppose h_* is an axiomatic homology theory with coefficient group G . Then one can associate to any CW-pair (X, A) isomorphisms

$$\Psi_{(X,A)} : H_n^{\text{CW}}(X, A; G) \xrightarrow{\cong} h_n(X, A)$$

for every n , which are natural in the sense that for any map of CW-pairs $f : (X, A) \rightarrow (Y, B)$, the following diagram commutes:

$$\begin{array}{ccc} H_n^{\text{CW}}(X, A; G) & \xrightarrow{\Psi_{(X,A)}} & h_n(X, A) \\ \downarrow f_* & & \downarrow f_* \\ H_n^{\text{CW}}(Y, B; G) & \xrightarrow{\Psi_{(Y,B)}} & h_n(Y, B) \end{array}$$

In the language of category theory, this theorem says the following. There is a functor $\text{CW}_{\text{rel}} \rightarrow \text{Top}_{\text{rel}}$ 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_* with this functor produces a functor $\text{CW}_{\text{rel}} \rightarrow \text{Ab}_{\mathbb{Z}}$. The theorem defines a natural transformation from $H_*^{\text{CW}}(\cdot; G)$ to the latter functor, associating to every CW-pair (X, A) the isomorphism $\Psi_{(X,A)}$.

Let us begin setting up the proof of the theorem. We shall focus here on the case of absolute homology, i.e. pairs (X, A) with $A = \emptyset$, leaving the relative case as a (worthwhile!) exercise. The key idea is to establish a relationship between $h_*(X)$ and the homology of a chain complex built out of the long exact sequences of the pairs (X^n, X^{n-1}) and (X^{n+1}, X^n) , as it will turn out that the latter chain complex can be naturally identified with $C_*^{\text{CW}}(X; G)$.

LEMMA 36.10. For all $n \in \mathbb{N}$, (X^n, X^{n-1}) is a good pair in the sense of Definition 30.17, i.e. X^{n-1} is a deformation retract of some neighborhood $V \subset X^n$ of X^{n-1} .

PROOF. Since $X^n = X^{n-1} \cup_{\varphi^n} \coprod_{\alpha \in \mathcal{K}^n} \mathbb{D}^n$, it suffices to set $V := X^{n-1} \cup_{\varphi^n} \coprod_{\alpha \in \mathcal{K}^n} (\mathbb{D}^n \setminus \{0\})$. \square

By Theorem 30.21, we now have a natural isomorphism $h_*(X^n, X^{n-1}) \cong \tilde{h}_*(X^n/X^{n-1})$. 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 $\mathring{\mathbb{D}}^n$, and the quotient collapses the boundaries of all these disks to a point. With this understood, it follows that the map Φ_*^n at the bottom of the following diagram is an isomorphism, and so therefore is the map at the top:

$$\begin{array}{ccc} h_*\left(\coprod_{\alpha \in \mathcal{K}^n} (\mathbb{D}^n, \partial\mathbb{D}^n)\right) & \xrightarrow{\Phi_*^n} & h_*(X^n, X^{n-1}) \\ \downarrow \cong & & \downarrow \cong \\ \tilde{h}_*\left(\coprod_{\alpha \in \mathcal{K}^n} \mathbb{D}^n / \coprod_{\alpha \in \mathcal{K}^n} \partial\mathbb{D}^n\right) & \xrightarrow{\Phi_*^n} & \tilde{h}_*(X^n/X^{n-1}). \end{array}$$

Applying the additivity axiom (in conjunction with the five-lemma as in Problem Set 3 #4) 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 36.11. 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 $n \in \mathbb{N}$. \square

The long exact sequence of $(\mathbb{D}^n, \partial\mathbb{D}^n)$ in reduced homology implies that the connecting homomorphisms

$$h_k(\mathbb{D}^n, \partial\mathbb{D}^n) \xrightarrow{\partial_*} \tilde{h}_{k-1}(S^{n-1}) \cong \begin{cases} G & \text{if } k = n, \\ 0 & \text{if } k \neq n \end{cases}$$

are isomorphisms for all k and n , thus we've proved

$$(36.3) \quad h_k(X^n, X^{n-1}) \cong \begin{cases} C_n^{\text{CW}}(X; G) = \bigoplus_{\alpha \in \mathcal{K}^n} G & \text{if } k = n, \\ 0 & \text{if } k \neq n. \end{cases}$$

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

$$X^0 = \coprod_{\alpha \in \mathcal{K}^0} \{\text{pt}\},$$

so that (36.3) is also correct in this case due to the dimension and additivity axioms. The group $h_n(X^n, X^{n-1})$ can therefore serve as a stand-in for $C_n^{\text{CW}}(X; G)$ in our proof of Theorem 36.9. This proof will be the main topic of the next lecture.

37. Invariance of cellular homology, part 2

Let's quickly rephrase what we've done so far toward the proof of Theorem 36.9. We are assuming h_* is an axiomatic homology theory with coefficient group G . The latter means (via the additivity axiom) that $h_0(S^0)$ has a canonical isomorphism with $G \oplus G$, where the first factor corresponds to the point $1 \in S^0$ and the second to the other point $-1 \in S^0$. Applying Exercise 28.16 to h_* instead of singular homology gives

$$\tilde{h}_0(S^0) = \{(g, -g) \in G \oplus G \mid g \in G\},$$

which we can identify with G via the injection $G \hookrightarrow G \oplus G : g \mapsto (g, -g)$, and $\tilde{h}_k(S^0) = 0$ for all $k \neq 0$ due to the additivity and dimension axioms. Applying the suspension isomorphisms $S_* : \tilde{h}_{k-1}(S^{n-1}) \xrightarrow{\cong} \tilde{h}_k(S^n)$ repeatedly, we can then identify G with $\tilde{h}_{n-1}(S^{n-1})$ for each $n \in \mathbb{N}$, and the connecting homomorphism in the reduced long exact sequence of (\mathbb{D}^n, S^{n-1}) then identifies G in turn with $h_n(\mathbb{D}^n, \partial\mathbb{D}^n)$, while simultaneously proving $h_k(\mathbb{D}^n, \partial\mathbb{D}^n) = 0$ for all $k \neq n$. Now if X is a CW-complex, the argument at the end of the previous lecture showed that

$$(37.1) \quad \bigoplus_{\alpha \in \mathcal{K}^n} (\Phi_\alpha)_* : \bigoplus_{\alpha \in \mathcal{K}^n} h_k(\mathbb{D}^n, \partial\mathbb{D}^n) \rightarrow h_k(X^n, X^{n-1})$$

is an isomorphism for every $n \geq 1$ and $k \in \mathbb{Z}$, which proves $h_k(X^n, X^{n-1}) = 0$ for $k \neq n$ and identifies $h_n(X^n, X^{n-1})$ via our isomorphism $h_n(\mathbb{D}^n, \partial\mathbb{D}^n) = G$ with the cellular n -chain group $C_n^{\text{CW}}(X; G)$. The plan going forward is to use the Eilenberg-Steenrod axioms to construct a boundary map on $\bigoplus_{n \in \mathbb{Z}} h_n(X^n, X^{n-1})$ and prove that the homology of the resulting chain complex is isomorphic to $h_*(X)$. The last step will then be to show that our boundary map on $\bigoplus_{n \in \mathbb{Z}} h_n(X^n, X^{n-1})$ matches the cellular boundary map $\partial : C_*^{\text{CW}}(X; G) \rightarrow C_{*-1}^{\text{CW}}(X; G)$ 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

$$(37.2) \quad 0 = h_{k+1}(X^n, X^{n-1}) \rightarrow h_k(X^{n-1}) \rightarrow h_k(X^n) \rightarrow h_k(X^n, X^{n-1}) = 0,$$

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:

$$h_k(X^n) \cong h_k(X^{n-1}) \cong \dots \cong h_k(X^0) \cong \bigoplus_{\alpha \in \mathcal{K}^0} h_k(\{\text{pt}\}) = 0,$$

where at the last step we have applied the additivity and dimension axioms, using the fact that X^0 is a discrete space. This already proves a quite nontrivial fact that we did not yet know, though you may have expected it: for any homology theory, the homology groups of an n -dimensional CW-complex vanish in dimensions greater than n .

LEMMA 37.1. *For every $k > n$, $h_k(X^n) = 0$.* □

Similarly, starting with $k < n$ and applying (37.2) repeatedly to increase n gives:

LEMMA 37.2. For every $k < n$, the inclusions $X^n \hookrightarrow X^{n+1} \hookrightarrow X^{n+2} \hookrightarrow \dots$ induce isomorphisms $h_k(X^n) \cong h_k(X^{n+1}) \cong h_k(X^{n+2}) \cong \dots$ \square

We can now proceed to the heart of the proof of Theorem 36.9. 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:

(37.3)

$$\begin{array}{ccccccc}
 & & & & h_n(X^{n-1}) & \xlongequal{\quad} & 0 \\
 & & & & \downarrow & & \\
 0 & \parallel & & & & & \\
 \parallel & & & & & & \\
 h_{n+1}(X^n) & \rightarrow & h_{n+1}(X^{n+1}) & \xrightarrow{j_{n+1}} & h_{n+1}(X^{n+1}, X^n) & \xrightarrow{\partial_{n+1}} & h_n(X^n) & \xrightarrow{i_n} & h_n(X^{n+1}) & \rightarrow & h_n(X^{n+1}, X^n) \\
 & & & \searrow & \beta_{n+1} & & \downarrow j_n & & & & \parallel \\
 & & & & & & h_n(X^n, X^{n-1}) & & & & 0 \\
 & & & & & & \downarrow \partial_n & & & & \\
 & & & & & & h_{n-1}(X^{n-1}) & & & &
 \end{array}$$

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

$$(37.4) \quad \dots \rightarrow h_n(X^n, X^{n-1}) \xrightarrow{\beta_n} h_{n-1}(X^{n-1}, X^{n-2}) \rightarrow \dots \rightarrow h_1(X^1, X^0) \xrightarrow{\beta_1} h_0(X^0) \xrightarrow{\beta_0} 0 \rightarrow \dots$$

as a chain complex whose individual chain groups are canonically isomorphic to the chain groups in $C_*^{CW}(X; G)$. 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

$$(37.5) \quad h_n(X^n) / \ker i_n \xrightarrow{\cong} h_n(X^{n+1}).$$

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

$$(37.6) \quad h_n(X^n) / \ker i_n \xrightarrow{\cong} \ker \beta_n / \text{im } \beta_{n+1}.$$

The latter is of course the n th homology group of the chain complex (37.4). 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 37.2 to conclude via (37.5) and (37.6) 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 $h_*(X) \cong H_*^{CW}(X; G)$ will be complete as

soon as we can show that the boundary maps β_n in (37.4) are the same as our usual cellular boundary maps. In other words, we need to prove that the diagram

$$\begin{array}{ccc} C_n^{\text{CW}}(X; G) & \xrightarrow{\cong} & h_n(X^n, X^{n-1}) \\ \downarrow \partial & & \downarrow \beta_n \\ C_{n-1}^{\text{CW}}(X; G) & \xrightarrow{\cong} & h_{n-1}(X^{n-1}, X^{n-2}) \end{array}$$

commutes for every n , where the horizontal maps are the canonical isomorphisms that we discussed at the beginning of this lecture. The theorem that $\partial^2 = 0$ 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 maps on reduced homology

$$(37.7) \quad \bigoplus_{\alpha \in \mathcal{K}^n} (\Phi_\alpha)_* : \bigoplus_{\alpha \in \mathcal{K}^n} \tilde{h}_n(\mathbb{D}^n/\partial\mathbb{D}^n) \rightarrow \tilde{h}_n(X^n/X^{n-1})$$

is an isomorphism. Indeed, under the natural isomorphisms between relative homology for good pairs and reduced homology of quotients, this is equivalent to the fact that (37.1) is an isomorphism. The advantage of rewriting this map in terms of quotients is, however, that we can explicitly write down its inverse. We recall the projections $p_\alpha : X^n \rightarrow X^n/(X^n \setminus e_\alpha^n) = \mathbb{D}^n/\partial\mathbb{D}^n$ 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 37.3. *The inverse of the map (37.7) 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 (37.7) 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:

$$(37.8) \quad \begin{array}{ccccc} & & C_n^{\text{CW}}(X; G) & & C_{n-1}^{\text{CW}}(X; G) \\ & & \parallel & & \parallel \\ \bigoplus_{\alpha \in \mathcal{K}^n} h_n(\mathbb{D}^n, \partial\mathbb{D}^n) & \xrightarrow[\cong]{\bigoplus_{\alpha} \partial_*} & \bigoplus_{\alpha \in \mathcal{K}^n} \tilde{h}_{n-1}(\partial\mathbb{D}^n) & \xrightarrow{\partial^{\text{CW}}} & \bigoplus_{\beta \in \mathcal{K}^{n-1}} \tilde{h}_{n-1}(\mathbb{D}^{n-1}/\partial\mathbb{D}^{n-1}) \\ \cong \downarrow \bigoplus_{\alpha} (\Phi_\alpha)_* & & \downarrow \bigoplus_{\alpha} (\varphi_\alpha)_* & & \uparrow \prod_{\beta} (p_\beta)_* \cong \\ h_n(X^n, X^{n-1}) & \xrightarrow{\partial_n} & \tilde{h}_{n-1}(X^{n-1}) & \xrightarrow{q_*} & \tilde{h}_{n-1}(X^{n-1}/X^{n-2}) \\ & \searrow \beta_n & \downarrow j_{n-1} & & \downarrow \cong \\ & & h_{n-1}(X^{n-1}, X^{n-2}) & \xrightarrow{\cong} & h_{n-1}(X^{n-1}/X^{n-2}, X^{n-2}/X^{n-2}) \end{array}$$

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 30.21 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 $h_{n-1}(X^{n-1}, X^{n-2}) \cong \tilde{h}_{n-1}(X^{n-1}/X^{n-2})$.
- 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 $\partial^{CW} : C_n^{CW}(X; G) \rightarrow C_{n-1}^{CW}(X; G)$, 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 group $h_{n-1}(X^{n-1}, X^{n-2})$ of β_n with $\tilde{h}_{n-1}(X^{n-1}/X^{n-2})$ so that we can then Lemma 37.3 to identify the latter with $C_{n-1}^{CW}(X; G)$ 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 34.8, each of these maps is just multiplication by the degree of $p_\beta \circ \varphi_\alpha$, also known as the incidence number $[e_\beta^{n-1} : e_\alpha^n]$. This proves that ∂^{CW} is indeed simply the cellular boundary map ∂ , and in particular, the latter satisfies $\partial^2 = 0$.

To complete the proof of Theorem 36.9 in the absolute case, we still need to understand how a cellular map $f : X \rightarrow Y$ between two CW-complexes interacts with the two isomorphisms $H_*^{CW}(X; G) \cong h_*(X)$ and $H_*^{CW}(Y; G) \cong h_*(Y)$. Being a cellular map implies that f 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 (37.3) 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 (37.4) to the corresponding chain complex for Y , and therefore induces a homomorphism $H_*^{\text{CW}}(X; G) \rightarrow H_*^{\text{CW}}(Y; G)$. To relate this to the map $f_* : h_*(X) \rightarrow h_*(Y)$, recall that the isomorphism $H_n^{\text{CW}}(X; G) = \ker \beta_n / \text{im } \beta_{n+1} \cong h_n(X)$ 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

$$\begin{array}{ccc} H_n^{\text{CW}}(X; G) & \xrightarrow{\cong} & h_n(X) \\ \downarrow f_* & & \downarrow f_* \\ H_n^{\text{CW}}(Y; G) & \xrightarrow{\cong} & h_n(Y). \end{array}$$

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 $C_n^{\text{CW}}(X; G)$ and $C_n^{\text{CW}}(Y; G)$ 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 (36.2) for maps $C_n^{\text{CW}}(X; G) \rightarrow C_n^{\text{CW}}(Y; G)$ induced by cellular maps. This will prove simultaneously the theorem that the map in (36.2) is a chain map. Here is the analogue of the diagram (37.8) for the situation at hand:

$$(37.9) \quad \begin{array}{ccc} C_n^{\text{CW}}(X; G) & & C_n^{\text{CW}}(Y; G) \\ \parallel & & \parallel \\ \bigoplus_{e_\alpha^n \subset X} \tilde{h}_n(\mathbb{D}^n / \partial \mathbb{D}^n) & \xrightarrow{f_*^{\text{CW}}} & \bigoplus_{e_\beta^n \subset Y} \tilde{h}_n(\mathbb{D}^n / \partial \mathbb{D}^n) \\ \cong \downarrow \bigoplus_\alpha (\Phi_\alpha)_* & & \uparrow \prod_\beta (p_\beta)_* \cong \\ \tilde{h}_n(X^n / X^{n-1}) & \xrightarrow{f_*} & \tilde{h}_n(Y^n / Y^{n-1}) \end{array}$$

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 $f_*^{\text{CW}} : C_n^{\text{CW}}(X; G) \rightarrow C_n^{\text{CW}}(Y; G)$, 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 34.8 again, this map is multiplication by $\text{deg}(p_\beta \circ f \circ \Phi_\alpha) = [e_\beta^n : e_\alpha^n]$, thus f_*^{CW} does indeed match the formula given in (36.2) for $f_* : C_n^{\text{CW}}(X; G) \rightarrow C_n^{\text{CW}}(Y; G)$.

The proof of Theorem 36.9 is now complete except for three details, the first two of which will be left as exercises.

EXERCISE 37.4. Some portions of the discussion above do not make sense for $n = 0$, especially when $\mathbb{D}^n / \partial \mathbb{D}^n$ is mentioned. Adapt the discussion as needed for that particular case.

EXERCISE 37.5. Extend the entire discussion to the case of a CW-pair (X, A) with $A \neq \emptyset$.
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 $(X^n \cup A, X^{n-1} \cup A, A)$.

Comment: This exercise is a bit lengthy, but it is not fundamentally difficult—every step is simply a minor generalization of something that we discussed in this lecture. Working through it is one of the best ways to achieve a deeper understanding of the isomorphism $H_*^{\text{CW}}(X; G) \cong h_*(X)$.

The third unresolved issue is the simplifying assumption $\dim X < \infty$ that we imposed in order to argue that $h_n(X^{n+1}) \cong h_n(X)$. We will discuss in the next lecture how to lift this assumption for the specific homology theory $h_* := H_*(\cdot; G)$.

38. 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 , so that if the functor h_* can be shown to be “continuous” with respect to such limits, we would conclude

$$h_n(X) = h_n\left(\lim_{k \rightarrow \infty} X^k\right) = \lim_{k \rightarrow \infty} h_n(X^k),$$

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 $\{X^n\}_{n=0}^\infty$, and similarly for a sequence of abelian groups such as $\{h_n(X^k)\}_{k=0}^\infty$.

Suppose I is a set with a pre-order $<$, 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 $(I, <)$ is called a **directed set** (*gerichtete Menge*) if for every pair $\alpha, \beta \in I$, there exists $\gamma \in I$ 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 38.17 below.

In the following, we use the notation $X \xrightarrow{f} Y$ 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 38.1. Given a category \mathcal{C} , a **direct system** (*induktives System*) $\{X_\alpha, \varphi_{\beta\alpha}\}$ in \mathcal{C} over $(I, <)$ associates to each $\alpha \in I$ an object X_α of \mathcal{C} , along with morphisms

$$\varphi_{\beta\alpha} \in \text{Mor}(X_\alpha, X_\beta) \quad \text{for each } \alpha < \beta$$

such that

$$\varphi_{\alpha\alpha} = \text{Id}_{X_\alpha}$$

and the diagram

$$\begin{array}{ccccc} X_\alpha & \xrightarrow{\varphi_{\beta\alpha}} & X_\beta & \xrightarrow{\varphi_{\gamma\beta}} & X_\gamma \\ & & \searrow & \nearrow & \\ & & & \varphi_{\gamma\alpha} & \end{array}$$

commutes for every triple $\alpha, \beta, \gamma \in I$ with $\alpha < \beta < \gamma$.

REMARK 38.2. Recall from Problem Set 1 #4 that a pre-order $<$ on a set I can be encoded by calling I the collection of objects in a category \mathcal{I} , such that for each pair $x, y \in I$, the set of morphisms $\text{Mor}(x, y)$ contains exactly one element whenever $x < y$ and is otherwise empty. A direct system in \mathcal{C} over $(I, <)$ is then nothing other than a (covariant) functor $\mathcal{I} \rightarrow \mathcal{C}$.

EXAMPLE 38.3. For any CW-complex X , its collection of skeleta $\{X^n\}_{n=0}^\infty$ forms a direct system in Top over (\mathbb{N}_0, \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 Top_{rel} .

EXAMPLE 38.4. For any axiomatic homology theory h_* , the homology groups of the skeleta of a CW-complex form a direct system in $\text{Ab}_{\mathbb{Z}}$ over (\mathbb{N}_0, \leq) : it consists of the graded abelian groups $\{h_*(X^n)\}_{n=0}^\infty$ and for each $m \geq n$ the map $h_*(X^n) \rightarrow h_*(X^m)$ induced by the inclusion $X^n \hookrightarrow X^m$. For each individual $k \in \mathbb{Z}$, we can also extract from this a direct system in Ab over (\mathbb{N}_0, \leq) , formed by the sequence of abelian groups $\{h_k(X^n)\}_{n=0}^\infty$.

The last example illustrates the following general observation, which is immediate from the definitions:

PROPOSITION 38.5. *If $\{X_\alpha, \varphi_{\beta\alpha}\}$ is a direct system in \mathcal{A} over $(I, <)$, 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 $(I, <)$. \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 $\alpha \in I$ becomes large. What we do have in every category is the notion of morphisms and the composition function $(f, g) \mapsto f \circ g$, so this is the structure that we will use. The idea is to measure the convergence of a direct system $\{X_\alpha, \varphi_{\beta\alpha}\}$ in terms of the morphisms from each X_α to other fixed objects in the category.

DEFINITION 38.6. For a direct system $\{X_\alpha, \varphi_{\beta\alpha}\}$ in \mathcal{C} over $(I, <)$, a **target** $\{Y, f_\alpha\}$ of the system consists of an object Y of \mathcal{C} together with associated morphisms $f_\alpha \in \text{Mor}(X_\alpha, Y)$ for each $\alpha \in I$ such that the diagram

$$\begin{array}{ccc} X_\alpha & \xrightarrow{\varphi_{\beta\alpha}} & X_\beta \\ & \searrow f_\alpha & \nearrow f_\beta \\ & & Y \end{array}$$

commutes for every pair $\alpha, \beta \in I$ with $\alpha < \beta$.

DEFINITION 38.7. A target $\{X_\infty, \varphi_\alpha\}$ of the direct system $\{X_\alpha, \varphi_{\beta\alpha}\}$ is called a **direct limit**⁵⁰ (*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 $f_\infty \in \text{Mor}(X_\infty, Y)$ 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 $\alpha \in I$.

The essential meaning of a direct limit can be encoded in the diagram

$$\begin{array}{ccccccc} X_\alpha & \xrightarrow{\varphi_{\beta\alpha}} & X_\beta & \xrightarrow{\varphi_{\gamma\beta}} & X_\gamma & \longrightarrow & \dots \longrightarrow \varinjlim \{X_\alpha\} \\ & & & & & & \downarrow \text{dashed} \\ & & & & & & Y \end{array}$$

where we assume $\alpha < \beta < \gamma < \dots \in I$. 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:

EXERCISE 38.8. 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 $\psi \in \text{Mor}(X, Y)$, show that $\{Y, \psi \circ f_\alpha\}$ is also a direct limit of $\{X_\alpha, \varphi_{\beta\alpha}\}$. *Remark: The invertibility of ψ is needed only for showing that $\{Y, \psi \circ f_\alpha\}$ satisfies the universal property; it is already a target without this.*

The non-uniqueness exhibited by the exercise above is however the worst thing that can happen: if $\{X, f_\alpha\}$ and $\{Y, g_\alpha\}$ are any two direct limits of the same system $\{X_\alpha, \varphi_{\beta\alpha}\}$, then the universal property provides unique morphisms $g_\infty \in \text{Mor}(X, Y)$ and $f_\infty \in \text{Mor}(Y, X)$ satisfying $g_\infty \circ f_\alpha = g_\alpha$ and $f_\infty \circ g_\alpha = f_\alpha$ for every $\alpha \in I$. It follows that $f_\infty \circ g_\infty$ is the unique morphism from X to X satisfying $(f_\infty \circ g_\infty) \circ f_\alpha = f_\alpha$ for every $\alpha \in I$, which implies $f_\infty \circ g_\infty = \text{Id}_X$. 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 exercise computes direct limits in a situation that is of concrete interest for the homology of a CW-complex X : recall from the previous lecture that for each $k \in \mathbb{Z}$, the sequence of homology groups $h_k(X^0) \rightarrow h_k(X^1) \rightarrow \dots \rightarrow h_k(X^n) \rightarrow h_k(X^{n-1}) \rightarrow \dots$ 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.

EXERCISE 38.9. Suppose $\{X_\alpha, \varphi_{\beta\alpha}\}$ is a direct system in \mathcal{C} over $(I, <)$ with the property that for some $\alpha_0 \in I$, $\varphi_{\gamma\beta} \in \text{Mor}(X_\beta, X_\gamma)$ is an isomorphism for every $\beta, \gamma \in I$ with $\beta > \alpha_0$ and $\gamma > \alpha_0$. For each $\alpha \in I$, choose $\gamma \in I$ such that $\gamma > \alpha$ and $\gamma > \alpha_0$, and define

$$\varphi_\alpha := \varphi_{\gamma\alpha_0}^{-1} \circ \varphi_{\gamma\alpha} \in \text{Mor}(X_\alpha, X_{\alpha_0}).$$

⁵⁰Direct limits are also sometimes called **inductive limits** or **colimits** (*Kolimes*).

- (a) Prove that the morphism φ_α does not depend on the choice of the element $\gamma \in I$.
- (b) Prove that $\{X_{\alpha_0}, \varphi_\alpha\}$ is a target of the system.
- (c) Prove that $\{X_{\alpha_0}, \varphi_\alpha\}$ also satisfies the universal property in Definition 38.7, hence $X_{\alpha_0} = \varinjlim\{X_\alpha\}$.

For the categories that we are most interested in, we will see presently that direct limits always exist and can be described in more concrete terms.

EXERCISE 38.10. Suppose $\{X_\alpha, \varphi_{\beta\alpha}\}$ is a direct system in \mathcal{C} over $(I, <)$, where \mathcal{C} is any category in which objects are sets (possibly with extra structure) and morphisms are maps between them. For any $\alpha, \beta \in I$, $x \in X_\alpha$ and $y \in X_\beta$, define the relation $x \sim y$ to mean

$$x \sim y \iff \varphi_{\gamma\alpha}(x) = \varphi_{\gamma\beta}(y) \text{ for some } \gamma \in I \text{ with } \gamma > \alpha \text{ and } \gamma > \beta.$$

Prove that \sim is an equivalence relation on the set-theoretic disjoint union $\coprod_{\alpha \in I} X_\alpha$.⁵¹

PROPOSITION 38.11. *If $\{X_\alpha, \varphi_{\beta\alpha}\}$ is a direct system in \mathbf{Top} over $(I, <)$, then its direct limit is the space*

$$\varinjlim\{X_\alpha\} = \coprod_{\alpha \in I} X_\alpha / \sim,$$

where the equivalence relation is defined as in Exercise 38.10, and the associated morphisms $\varphi_\alpha : X_\alpha \rightarrow \varinjlim\{X_\alpha\}$ are the compositions of the inclusions $X_\alpha \hookrightarrow \coprod_{\beta \in I} X_\beta$ with the quotient projection. Moreover, the topology on $\varinjlim\{X_\alpha\}$ is the strongest topology for which the maps $\varphi_\alpha : X_\alpha \rightarrow \varinjlim\{X_\alpha\}$ are continuous for all $\alpha \in I$.

PROOF. Abbreviate $X_\infty = \coprod_{\alpha \in I} X_\alpha / \sim$. 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 $q^{-1}(\mathcal{U}) \subset \coprod_{\beta \in I} X_\beta$ via the quotient projection $q : \coprod_{\beta \in I} X_\beta \rightarrow X_\infty$ 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 $\alpha \in I$. Since $q^{-1}(\mathcal{U}) \cap X_\alpha = \varphi_\alpha^{-1}(\mathcal{U})$, this means that $\mathcal{U} \subset X_\infty$ is open if and only if every $\varphi_\alpha^{-1}(\mathcal{U}) \subset X_\alpha$ is open, thus characterizing the topology of X_∞ as the strongest for which every map $\varphi_\alpha : X_\alpha \rightarrow X_\infty$ 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 $\alpha \in I$ (cf. Exercise 35.4).

It is clear that $\{X_\infty, \varphi_\alpha\}$ is a target since for any $\alpha, \beta \in I$ 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 $\alpha \in I$. To write down $f_\infty(x)$ for an arbitrary element $x \in X_\infty$, observe that since the quotient projection $q : \coprod_{\beta \in I} X_\beta \rightarrow X_\infty$ is surjective, we have $x = q(x_\alpha) = \varphi_\alpha(x_\alpha)$ for some $\alpha \in I$ and $x_\alpha \in X_\alpha \subset \coprod_{\beta \in I} X_\beta$, so in order to achieve $f_\infty \circ \varphi_\alpha = f_\alpha$, we are forced to define

$$f_\infty(x) := f_\alpha(x_\alpha).$$

⁵¹The **set-theoretic disjoint union** of a collection of sets $\{X_\alpha\}_{\alpha \in I}$ can be defined in general as the set $\{(\alpha, x) \mid \alpha \in I, x \in X_\alpha\}$, i.e. it is a union of all the sets X_α , but defined such that even if some pair of the sets X_α and X_β for $\alpha \neq \beta$ have elements in common, they are each identified with *disjoint* subsets of $\coprod_{\gamma \in I} X_\gamma$. 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 38.10 is not needed.

We claim that $f_\infty(x)$ is then independent of the choice of element $x_\alpha \in q^{-1}(x)$. Indeed, suppose $\beta \in I$ and $x_\beta \in X_\beta \subset \coprod_\gamma X_\gamma$ such that $\varphi_\beta(x_\beta) = q(x_\beta) = x$. The equivalence $x_\alpha \sim x_\beta$ then means that for some $\gamma \in I$ satisfying $\gamma > \alpha$ 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 $\alpha \in I$. □

REMARK 38.12. Proposition 38.11 extends in an obvious way to give a concrete description of any direct limit in the category Top_{rel} of pairs of spaces.

Consider the specific direct system of topological spaces $\{X^n\}_{n=0}^\infty$ from Example 38.3, 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 38.11 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 38.13. For the direct system of Example 38.3 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 38.11 in the category Ab of abelian groups.

PROPOSITION 38.14. If $\{G_\alpha, \varphi_{\beta\alpha}\}$ is a direct system in Ab over $(I, <)$, then its direct limit is the group

$$\varinjlim \{G_\alpha\} = \bigoplus_{\alpha \in I} G_\alpha / H,$$

where $H \subset \bigoplus_\alpha G_\alpha$ is the subgroup generated by all elements of the form $g - \varphi_{\beta\alpha}(g)$ for $g \in G_\alpha$ and $\beta > \alpha$, 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

$$\bigoplus_{\alpha \in I} \psi_\alpha : \bigoplus_{\alpha \in I} G_\alpha \rightarrow A$$

vanishes on the subgroup 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 . □

EXERCISE 38.15. Prove the obvious analogues of Proposition 38.14 for direct systems in the categories $\text{Ab}_\mathbb{Z}$ of \mathbb{Z} -graded abelian groups and Chain of chain complexes.

The following consequence of Proposition 38.14 makes proving things about direct limits of abelian groups (or the other algebraic categories mentioned in the exercise above) considerably easier.

COROLLARY 38.16. *The following statements hold for any direct system $\{G_\alpha, \varphi_{\beta\alpha}\}$ in \mathbf{Ab} , $\mathbf{Ab}_{\mathbb{Z}}$ or \mathbf{Chain} over a directed set $(I, <)$:*

- (i) *For every $x \in \varinjlim \{G_\alpha\}$, there exists $\beta \in I$ and $x_\beta \in G_\beta$ such that $x = \varphi_\beta(x_\beta)$.*
- (ii) *For every $\beta \in I$ and $x_\beta \in G_\beta$ satisfying $\varphi_\beta(x_\beta) = 0 \in \varinjlim \{G_\alpha\}$, there exists $\gamma > \beta$ such that $\varphi_{\gamma\beta}(x_\beta) = 0 \in G_\gamma$.*

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

$$\sum_{\alpha \in I_0} g_\alpha \in \bigoplus_{\alpha \in I} G_\alpha$$

for some finite subset $I_0 \subset I$. Since $(I, <)$ is a directed set, we can then find an element $\beta \in I$ satisfying $\beta > \alpha$ for every $\alpha \in I_0$, so

$$\sum_{\alpha \in I_0} g_\alpha - \sum_{\alpha \in I_0} \varphi_{\beta\alpha}(g_\alpha) \in H,$$

implying that $x_\beta := \sum_{\alpha \in I_0} \varphi_{\beta\alpha}(g_\alpha) \in G_\beta$ satisfies $\varphi_\beta(x_\beta) = x$.

For the second statement, we observe that $\varphi_\beta(x_\beta) = 0$ holds if and only if $x_\beta \in G_\beta \subset \bigoplus_\alpha G_\alpha$ belongs to the subgroup H , meaning

$$(38.1) \quad x_\beta = \sum_{i=1}^N (g_i - \varphi_{\beta_i\alpha_i}(g_i))$$

for some finite collection of elements $\beta_i > \alpha_i \in I$ and $g_i \in G_{\alpha_i}$, $i = 1, \dots, N$. Choose a finite subset $I_0 \subset I$ that contains all the α_i, β_i for $i = 1, \dots, N$, along with an element $\gamma \in I$ such that $\gamma > \alpha$ for all $\alpha \in I_0$. Applying the homomorphism $\bigoplus_{\alpha \in I_0} \varphi_{\gamma\alpha}$ to both sides of (38.1) then produces $\varphi_{\gamma\beta}(x_\beta) \in G_\gamma$ on the left hand side and kills the right hand side since for each i ,

$$\left(\bigoplus_{\alpha \in I_0} \varphi_{\gamma\alpha} \right) (g_i - \varphi_{\beta_i\alpha_i}(g_i)) = \varphi_{\gamma\alpha_i}(g_i) - \varphi_{\gamma\beta_i} \circ \varphi_{\beta_i\alpha_i}(g_i) = \varphi_{\gamma\alpha_i}(g_i) - \varphi_{\gamma\alpha_i}(g_i) = 0.$$

We have thus proved $\varphi_{\gamma\beta}(x_\beta) = 0$. □

We have seen above that any CW-complex X can be identified with the direct limit of its skeleta. Combining Exercise 38.9 with the computations of the previous lecture proves moreover that for any axiomatic homology theory h_* and any $k \in \mathbb{Z}$, the direct system of abelian groups $\{h_k(X^n)\}_{n=0}^\infty$ stabilizes as $n \rightarrow \infty$ and thus has direct limit $h_k(X)$ for any n sufficiently large, which matches $H_k^{\text{CW}}(X^n) = H_k^{\text{CW}}(X)$. This gives an isomorphism of \mathbb{Z} -graded abelian groups

$$H_*^{\text{CW}}(X) \cong \varinjlim \{h_*(X^n)\}.$$

The isomorphism $H_*^{\text{CW}}(X) \cong h_*(X)$ will therefore follow if we can prove that the functor h_* behaves “continuously” under this direct limit, i.e. the question becomes

$$\varinjlim \{h_*(X^n)\} \cong h_*(\varinjlim \{X^n\})?$$

It is time to insert a word of caution: the next exercise shows that singular homology does *not* always behave as nicely as one would hope with respect to direct limits.

EXERCISE 38.17. Define $\{X_\alpha\}_{\alpha \in I}$ to be the collection of all countable subspaces of S^1 , with a partial order assigned to the index set such that

$$\alpha < \beta \quad \Leftrightarrow \quad X_\alpha \subset X_\beta.$$

In this case we can define $\varphi_{\beta\alpha} : X_\alpha \hookrightarrow X_\beta$ to be the inclusion map and regard $\{X_\alpha, \varphi_{\beta\alpha}\}$ as a direct system of topological spaces. Prove that $\varinjlim X_\alpha$ is homeomorphic to S^1 , but $\varinjlim H_*(X_\alpha; \mathbb{Z})$ is not isomorphic to $H_*(S^1; \mathbb{Z})$.

Hint 1: Describing $\varinjlim X_\alpha$ as in Proposition 38.11, it is not hard to find a natural bijection between $\varinjlim X_\alpha$ and $\bigcup_{\alpha \in I} X_\alpha = S^1$, but you need to check that the topology of this direct limit matches the standard topology of S^1 .

Hint 2: What can you say about $H_1(X_\alpha; \mathbb{Z})$ for each α ?

To see nonetheless why it might *sometimes* be true that $\varinjlim h_*(X_\alpha) \cong h_*(\varinjlim X_\alpha)$, 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 $(I, <)$, 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} . The natural morphisms

$$X_\alpha \xrightarrow{\varphi_\alpha} \varinjlim X_\alpha$$

for every $\alpha \in I$ then induce morphisms

$$\mathcal{F}(X_\alpha) \xrightarrow{\Phi_\alpha := \mathcal{F}(\varphi_\alpha)} \mathcal{F}(\varinjlim X_\alpha)$$

which satisfy

$$\Phi_\beta \circ \mathcal{F}(\varphi_{\beta\alpha}) = \mathcal{F}(\varphi_\beta) \circ \mathcal{F}(\varphi_{\beta\alpha}) = \mathcal{F}(\varphi_\beta \circ \varphi_{\beta\alpha}) = \mathcal{F}(\varphi_\alpha) = \Phi_\alpha$$

for all $\beta > \alpha$ and thus make $\{\mathcal{F}(\varinjlim X_\alpha), \Phi_\alpha\}$ a target of the system $\{\mathcal{F}(X_\alpha), \mathcal{F}(\varphi_{\beta\alpha})\}$. It follows via the universal property of the direct limit that there is a limiting morphism

$$(38.2) \quad \varinjlim \{\mathcal{F}(X_\alpha)\} \xrightarrow{\Phi_\infty} \mathcal{F}(\varinjlim X_\alpha).$$

We would now like to identify some situations in which Φ_∞ is guaranteed to be an isomorphism—in particular, we shall prove that this is true when $\{X_\alpha, \varphi_{\beta\alpha}\}$ is the direct system in \mathbf{Top} formed by the skeleta of a CW-complex and \mathcal{F} is the singular homology functor $H_*(\cdot; G)$.

Recall that $H_*(\cdot; G) : \mathbf{Top} \rightarrow \mathbf{Ab}_\mathbb{Z}$ is in fact the composition of two functors: the first is $C_*(\cdot; G) : \mathbf{Top} \rightarrow \mathbf{Chain}$, which sends each space X to its singular chain complex with coefficients in G , and the second is $H_* : \mathbf{Chain} \rightarrow \mathbf{Ab}_\mathbb{Z}$, sending a chain complex to its homology. The second of these two functors turns out to be extremely well behaved with respect to direct limits.

PROPOSITION 38.18. *Suppose $(I, <)$ is a directed set, with a chain complex C_*^α associated to each $\alpha \in I$ and a chain map $\varphi_{\beta\alpha} : C_*^\alpha \rightarrow C_*^\beta$ associated to each pair $\alpha < \beta \in I$ such that $\{C_*^\alpha, \varphi_{\beta\alpha}\}$ is a direct system in \mathbf{Chain} over $(I, <)$. Then choosing \mathcal{F} to be the functor $H_* : \mathbf{Chain} \rightarrow \mathbf{Ab}_\mathbb{Z}$, the map*

$$\Phi_\infty : \varinjlim \{H_*(C_*^\alpha)\} \rightarrow H_*(\varinjlim \{C_*^\alpha\})$$

defined as in (38.2) is an isomorphism of \mathbb{Z} -graded abelian groups.

We will prove this in the next lecture; the most essential ingredient in the proof is Corollary 38.16, which allows us to replace statements about direct limits with statements about individual objects in the direct system.

If one now tries to prove the analogous statement about the functor $C_*(\cdot; G) : \mathbf{Top} \rightarrow \mathbf{Chain}$, one finds sooner or later that one needs to know whether a singular simplex $\sigma : \Delta^n \rightarrow \varinjlim X_\alpha$ in a space that is a direct limit can always be factored through the map $X_\beta \rightarrow \varinjlim X_\alpha$ for some

“sufficiently large” index β . This would be something of an analogue to Corollary 38.16 in the setting of the singular chain complex, but it is far from obvious whether it is true. Exercise 38.17 shows in fact that the result we’d ideally like to prove is not true in general without some extra condition. We already know however that if X is a CW-complex and $\varphi_n : X^n \hookrightarrow X$ denotes the inclusion for each $n \geq 0$, then every singular simplex $\sigma : \Delta^k \rightarrow X$ is of the form $\varphi_n \circ \sigma'$ for some $\sigma' : \Delta^k \rightarrow X^n$ if n is large enough; this follows from the fact that compact subsets are always contained in finite subcomplexes (Proposition 36.1). This turns out to be enough:

PROPOSITION 38.19. *Suppose $\{X_\alpha, \varphi_{\beta\alpha}\}$ is a direct system of topological spaces over $(I, <)$ satisfying the following conditions:*

- (1) *For every $\alpha \in I$, X_α is a subspace of $X := \varinjlim \{X_\alpha\}$ and the maps $\varphi_{\beta\alpha} : X_\alpha \rightarrow X_\beta$ 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 $\alpha \in I$.*

Then choosing \mathcal{F} to be the singular chain complex functor $C_(\cdot; G) : \mathbf{Top} \rightarrow \mathbf{Chain}$ with an arbitrary coefficient group G , the chain map*

$$\Phi_\infty : \varinjlim \{C_*(X_\alpha; G)\} \rightarrow C_*(\varinjlim \{X_\alpha\}; G)$$

defined as in (38.2) is an isomorphism of chain complexes.

Once we have proved these last two propositions, applying them together proves:

THEOREM 38.20. *Under the same hypotheses as in Proposition 38.19, there is a natural isomorphism of \mathbb{Z} -graded abelian groups*

$$\varinjlim \{H_*(X_\alpha; G)\} \xrightarrow{\cong} H_*(\varinjlim \{X_\alpha\}; G)$$

for every coefficient group G . □

COROLLARY 38.21. *For any CW-complex X , there is a natural isomorphism $H_*^{\text{CW}}(X; G) \cong H_*(X; G)$.* □

The corollary is also true if singular homology $H_*(\cdot; G)$ is replaced by an arbitrary axiomatic homology theory h_* , but proving this would take us into somewhat more abstract territory than we have time for right now, so we will settle for the special case of singular homology. The original treatment of the homology axioms by Eilenberg and Steenrod [ES52] dealt mainly with finite complexes, for which our proof of the isomorphism $H_*^{\text{CW}}(X; G) \cong h_*(X)$ was already completed in the previous lecture. The extension of this result to infinite-dimensional complexes was accomplished originally by Milnor in [Mil62], who introduced the additivity axiom for this purpose.⁵² Milnor’s proof via the “mapping telescope” construction is reproduced in [Hat02, pp. 138–139].

We will begin the next lecture with the proofs of Propositions 38.18 and 38.19.

39. The Euler characteristic

Before we get to the main topic for today, I owe you two proofs from the end of the last lecture. In the interest of time, I will give only the first half of each proof and leave the remainders as exercises. The first result says that the functor $H_* : \mathbf{Chain} \rightarrow \mathbf{Ab}_{\mathbb{Z}}$ that takes each chain complex to its homology is continuous with respect to direct limits.⁵³

⁵²Recall that for finite disjoint unions, the additivity axiom follows from the other axioms, and one does not need any infinite disjoint unions to compute the homology of a finite CW-complex.

⁵³The technical term for this in category theory is that H_* is *cocontinuous*; this is consistent with the standard practice in homotopy theory of calling direct limits “colimits”. In that context, a “limit” (without the “co-”) is what we would call an *inverse limit*—we will encounter these later in the context of cohomology and other contravariant functors.

PROPOSITION 39.1. *Suppose $(I, <)$ is a directed set, with a chain complex C_*^α associated to each $\alpha \in I$ and a chain map $\varphi_{\beta\alpha} : C_*^\alpha \rightarrow C_*^\beta$ associated to each pair $\alpha < \beta \in I$ such that $\{C_*^\alpha, \varphi_{\beta\alpha}\}$ is a direct system in Chain over $(I, <)$. Then choosing \mathcal{F} to be the functor $H_* : \text{Chain} \rightarrow \text{Ab}_\mathbb{Z}$, the map*

$$\Phi_\infty : \varinjlim \{H_*(C_*^\alpha)\} \rightarrow H_*(\varinjlim \{C_*^\alpha\})$$

defined as in (38.2) is an isomorphism of \mathbb{Z} -graded abelian groups.

PROOF. We prove first that Φ_∞ is surjective. Given a homology class $[c] \in H_*(\varinjlim \{C_*^\alpha\})$ represented by a cycle $c \in \varinjlim \{C_*^\alpha\}$, Corollary 38.16 implies $c = \varphi_\beta(c_\beta)$ for some $\beta \in I$ and $c_\beta \in C_*^\beta$, 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 Corollary 38.16, we can find some $\gamma > \beta$ and replace c_β with $c_\gamma := \varphi_{\gamma\beta}(c_\beta) \in C_*^\gamma$ such that $\varphi_\gamma(c_\gamma) = \varphi_\beta(c_\beta) = c$ but also $\partial c_\gamma = 0$, and c_γ thus represents a homology class $[c_\gamma] \in H_*(C_*^\gamma)$. Now let

$$\Psi_\gamma : H_*(C_*^\gamma) \rightarrow \varinjlim \{H_*(C_*^\alpha)\}$$

denote the natural morphism associated to the direct limit of the system $\{H_*(C_*^\alpha), \Phi_{\gamma\alpha}\}$, where $\Phi_{\gamma\alpha} := (\varphi_{\gamma\alpha})_* : H_*(C_*^\alpha) \rightarrow H_*(C_*^\gamma)$ for $\gamma > \alpha$. Writing $\Phi_\gamma := (\varphi_\gamma)_* : H_*(C_*^\gamma) \rightarrow H_*(\varinjlim \{C_*^\alpha\})$, the diagram

$$\begin{array}{ccc} H_*(C_*^\gamma) & \xrightarrow{\Psi_\gamma} & \varinjlim \{H_*(C_*^\alpha)\} \\ & \searrow \Phi_\gamma & \downarrow \Phi_\infty \\ & & H_*(\varinjlim \{C_*^\alpha\}) \end{array}$$

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 Corollary 38.16, which is a tool for replacing statements about direct limits with corresponding statements about individual objects in the direct system. We saw in Exercise 38.17 that the singular homology functor $H_*(\cdot; G) : \text{Top} \rightarrow \text{Ab}_\mathbb{Z}$ is not always continuous with respect to direct limits; since this is the composition of two functors $H_* : \text{Chain} \rightarrow \text{Ab}_\mathbb{Z}$ and $C_*(\cdot; G) : \text{Top} \rightarrow \text{Chain}$, Proposition 39.1 implies that something must go wrong in general with the continuity of $C_*(\cdot; G) : \text{Top} \rightarrow \text{Chain}$. But if $\{X_\alpha, \varphi_{\beta\alpha}\}$ is a direct system of spaces with certain properties, then it is still possible to prove continuity by an argument very much analogous to the proof above. The key point is that we need something playing the role of Corollary 38.16 to tell us when a singular n -simplex $\sigma : \Delta^n \rightarrow \varinjlim \{X_\alpha\}$ can be written as $\varphi_\beta \circ \sigma_\beta$ for some $\beta \in I$ and a singular n -simplex $\sigma_\beta : \Delta^n \rightarrow X_\beta$.

PROPOSITION 39.2. *Suppose $\{X_\alpha, \varphi_{\beta\alpha}\}$ is a direct system of topological spaces over $(I, <)$ satisfying the following conditions:*

- (1) *For every $\alpha \in I$, X_α is a subspace of $X := \varinjlim \{X_\alpha\}$, $X_\alpha \subset X_\beta$ whenever $\alpha < \beta$ 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 $\alpha \in I$.*

Then choosing \mathcal{F} to be the singular chain complex functor $C_(\cdot; G) : \text{Top} \rightarrow \text{Chain}$ with an arbitrary coefficient group G , the chain map*

$$\Phi_\infty : \varinjlim \{C_*(X_\alpha; G)\} \rightarrow C_*(\varinjlim \{X_\alpha\}; G)$$

defined as in (38.2) is an isomorphism of chain complexes.

PROOF. For surjectivity, given $c = \sum_i g_i \sigma_i \in C_n(\varinjlim\{X_\alpha\}; G)$, 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 $\alpha_i \in I$ and $\sigma'_i : \Delta^n \rightarrow X_{\alpha_i}$ since Δ^n is compact. We can then find $\beta \in I$ 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 $c_\beta := \sum_i g_i \sigma''_i \in C_n(X_\beta; G)$ such that $(\varphi_\beta)_* c_\beta = c$. Writing $\Psi_\beta : C_*(X_\beta; G) \rightarrow \varinjlim\{C_*(X_\alpha; G)\}$ for the natural map associated to the direct limit, the diagram

$$\begin{array}{ccc} C_*(X_\beta; G) & \xrightarrow{\Psi_\beta} & \varinjlim\{C_*(X_\alpha; G)\} \\ & \searrow (\varphi_\beta)_* & \downarrow \Phi_\infty \\ & & C_*(\varinjlim\{X_\alpha\}; G) \end{array}$$

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. \square

As explained in the previous lecture, these two results in combination with the computations in Lecture 37 imply that for every CW-complex X and coefficient group G , there is a natural isomorphism

$$H_*^{\text{CW}}(X; G) \cong \varinjlim\{H_*(X^n; G)\} \cong H_*(\varinjlim\{X^n\}; G) = H_*(X; G).$$

It is a straightforward matter to extend that entire discussion to the case of a CW-pair (X, A) and prove $H_*^{\text{CW}}(X, A; G) \cong H_*(X, A; G)$, and also to prove that these isomorphisms commute with the maps induced by any map of CW-pairs $f : (X, A) \rightarrow (Y, B)$. We will leave the further details of these extensions as exercises and thus regard the proof of Theorem 36.9 (at least for singular homology) as complete.

We would now like to discuss a few applications of the isomorphism $H_*^{\text{CW}}(X, A; G) \cong H_*(X, A; G)$. One of the advantages of cellular homology is that for compact spaces, cell decompositions are always finite, in which case the cellular chain complex itself is finitely generated, and so therefore is its homology. This proves:

COROLLARY 39.3. *If (X, A) is a compact CW-pair, then its singular homology $H_*(X, A; \mathbb{Z})$ is finitely generated.* \square

Recall that if G is any abelian group and \mathbb{K} is a field, regarded as an abelian group with respect to its addition operation, then the tensor product $G \otimes \mathbb{K}$ inherits the structure of a vector space over \mathbb{K} : indeed, to define this we just need to say what scalar multiplication $\mathbb{K} \times (G \otimes \mathbb{K}) \rightarrow G \otimes \mathbb{K}$ means, and the obvious definition determined by the formula

$$\lambda(g \otimes k) := g \otimes (\lambda k) \quad \text{for } \lambda, k \in \mathbb{K}, g \in G$$

satisfies the required axioms. Moreover, if $\Phi : G \rightarrow H$ is a homomorphism between two abelian groups, then $\Phi \otimes 1 : G \otimes \mathbb{K} \rightarrow H \otimes \mathbb{K}$ becomes a \mathbb{K} -linear map with respect to the natural vector space structures on its domain and target. This means that if we choose a field \mathbb{K} as a coefficient group, $C_*(X, A; \mathbb{K}) = C_*(X, A) \otimes \mathbb{K}$ becomes a vector space over \mathbb{K} and the boundary map on $C_*(X, A; \mathbb{K})$ becomes \mathbb{K} -linear, so that $H_*(X, A; \mathbb{K})$ is also a vector space over \mathbb{K} . The compactness of (X, A) then makes $C_*(X, A; \mathbb{K})$ a finite-dimensional vector space, and Corollary 39.3 becomes a statement of linear algebra:

COROLLARY 39.4. *If (X, A) is a compact CW-pair, then for any field \mathbb{K} , $H_*(X, A; \mathbb{K})$ is a finite-dimensional vector space over \mathbb{K} .* \square

Note that each of these corollaries is actually two statements in one: they say on the one hand that $H_n(X, A; \mathbb{Z})$ is finitely generated or $H_n(X, A; \mathbb{K})$ is finite dimensional for every $n \in \mathbb{Z}$, but also that both are *trivial* for all but finitely many values of n . It is similarly obvious that $C_k^{\text{CW}}(X; G)$ and therefore also $H_k^{\text{CW}}(X; G)$ must vanish for any CW-pair that has no k -cells:

COROLLARY 39.5. *If (X, A) is an n -dimensional CW-pair, then $H_k(X, A; G) = 0$ for all $k > n$ and every coefficient group G . \square*

REMARK 39.6. 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 39.5 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 39.5 still applies.

For a closed n -manifold, we will see another proof that $H_k(M) = 0$ for all $k > n$ when we talk about Poincaré duality later in the course, and that proof requires no knowledge of cell decompositions. It's worth mentioning that homology is in this sense very different from the higher homotopy groups: there are plenty of n -dimensional manifolds M that have $\pi_k(M) \neq 0$ for some $k > n$, e.g. the simplest example is $\pi_3(S^2) \cong \mathbb{Z}$. This is one of the details that makes homology generally easier than homotopy theory.

REMARK 39.7. By results of Palais [Pal66] proved in 1966, it is also known that every smooth (but not necessarily finite-dimensional) Fréchet manifold is homotopy equivalent to a (not necessarily finite-dimensional) CW-complex. Fréchet manifolds are spaces that can be covered by charts identifying them locally with Fréchet spaces, a class of complete metrizable topological vector space that includes all Banach spaces, plus popular non-Banach examples like the space of C^∞ -functions on a compact smooth manifold. For example, if M and N are two smooth finite-dimensional manifolds and M is compact, then $C^\infty(M, N)$ is naturally a Fréchet manifold. Since many results of algebraic topology hold only for CW-complexes, Palais's theorem makes the techniques of the subject applicable in many of the functional-analytic settings that are used to study nonlinear PDEs.

Associating a sequence of abelian groups to every topological space is a nice thing to do, but sometimes one would prefer something simpler, e.g. a *number*. There are several numerical invariants that we can now associate to spaces in terms of their homology. Recall that according to the classification of finitely generated abelian groups, every such group G is isomorphic to

$$G \cong \mathbb{Z}^n \oplus T,$$

for a unique integer $n \geq 0$ 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 integer n is called the **rank** (*Rang*) of G , and we will denote it by

$$\text{rank } G \geq 0.$$

If like many people you prefer linear algebra to group theory, then you might prefer the following way of repackaging this definition: suppose \mathbb{K} is a field of characteristic 0. Then $G \otimes \mathbb{K} \cong (\mathbb{Z}^n \otimes \mathbb{K}) \oplus (T \otimes \mathbb{K})$, but for every $g \in T$, we have $mg = 0$ for some $m \in \mathbb{N}$ and thus for $q \in \mathbb{K}$,

$$g \otimes q = g \otimes \underbrace{\left(\frac{q}{m} + \dots + \frac{q}{m} \right)}_m = m \left(g \otimes \frac{q}{m} \right) = mg \otimes \frac{q}{m} = 0,$$

implying $T \otimes \mathbb{K} = 0$.⁵⁴ Since $\mathbb{Z} \otimes \mathbb{K} \cong \mathbb{K}$, this gives $G \otimes \mathbb{K} \cong \mathbb{K}^n$, so in terms of the natural \mathbb{K} -vector space structure on $G \otimes \mathbb{K}$, we have

$$\text{rank } G = \dim_{\mathbb{K}}(G \otimes \mathbb{K}).$$

We cannot justify this quite yet, but in a couple of lectures we will prove the *universal coefficient theorem*, a purely algebraic result that implies among other things that for singular homology with coefficients in any field \mathbb{K} of characteristic zero,

$$(39.1) \quad H_*(X; \mathbb{K}) \cong H_*(X) \otimes \mathbb{K}.$$

We will assume this result for the rest of the present lecture, as it makes it possible to erase one or two slightly subtle algebraic issues by converting them into linear algebra.

DEFINITION 39.8. For any space X and integer $k \geq 0$, the **k th Betti number** of X is the nonnegative (or possibly infinite) integer

$$b_k(X) := \text{rank } H_k(X; \mathbb{Z}).$$

In light of (39.1), $b_k(X)$ could equivalently be defined as

$$b_k(X) = \dim_{\mathbb{K}} H_k(X; \mathbb{K})$$

if \mathbb{K} is any field of characteristic zero. The most popular choice for this purpose is \mathbb{Q} , though \mathbb{R} and \mathbb{C} work just as well.

DEFINITION 39.9. For any space X with finitely-generated singular homology, the **Euler characteristic** (*Eulercharakteristik*) of X is the integer⁵⁵

$$\chi(X) = \sum_{k=0}^{\infty} (-1)^k b_k(X) \in \mathbb{Z}.$$

The usefulness of $\chi(X)$ as an invariant derives from a simple phenomenon in homological algebra that has remarkable consequences for topology:

PROPOSITION 39.10. *If C_* is a finitely-generated chain complex of free abelian groups, then*

$$\sum_{n \in \mathbb{Z}} (-1)^n \text{rank } H_n(C_*) = \sum_{n \in \mathbb{Z}} (-1)^n \text{rank } C_n = \sum_{n \in \mathbb{Z}} (-1)^n \dim_{\mathbb{K}} H_n(C_* \otimes \mathbb{K})$$

for any field \mathbb{K} .

In the following proof, we shall make use of the same special case of the universal coefficient theorem that implies (39.1), namely that for any chain complex of free abelian groups C_* and any field \mathbb{K} of characteristic zero,

$$(39.2) \quad H_*(C_* \otimes \mathbb{K}) \cong H_*(C_*) \otimes \mathbb{K}.$$

It is not strictly necessary to assume this right now, but doing so makes the following proof algebraically a bit easier.

PROOF OF PROPOSITION 39.10. Since each C_n is a free abelian group for every n , we have

$$\text{rank } C_n = \dim_{\mathbb{K}}(C_n \otimes \mathbb{K})$$

⁵⁴Did you notice how we used the assumption here that \mathbb{K} has characteristic 0? It implies in particular that \mathbb{K} contains the rational numbers \mathbb{Q} . You cannot do the same trick e.g. with $\mathbb{K} = \mathbb{Z}_2$.

⁵⁵Let'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".

for any field \mathbb{K} . (Here there is no need for \mathbb{K} to have characteristic zero because C_n by assumption has no torsion.) With this in mind, let us suppose instead for the moment that C_* is a finite-dimensional chain complex of \mathbb{K} -vector spaces with \mathbb{K} -linear boundary maps $\partial_n : C_n \rightarrow C_{n-1}$, and abbreviate

$$Z_n := \ker \partial_n, \quad B_n := \operatorname{im} \partial_{n+1}, \quad H_n := Z_n/B_n.$$

The dimensions of these vector spaces over \mathbb{K} are then related by

$$\dim H_n = \dim Z_n - \dim B_n$$

for every $n \in \mathbb{Z}$. Since $\partial_n : C_n \rightarrow C_{n-1}$ descends to an isomorphism $C_n/Z_n \rightarrow B_{n-1}$, we also have

$$\dim C_n = \dim Z_n + \dim B_{n-1}.$$

Combining these two relations gives

$$\begin{aligned} \sum_{n \in \mathbb{Z}} (-1)^n \dim H_n &= \dots - (\dim Z_{-1} - \dim B_{-1}) + (\dim Z_0 - \dim B_0) - (\dim Z_1 - \dim B_1) + \dots \\ &= \dots - \dim Z_{-1} + (\dim B_{-1} + \dim Z_0) - (\dim B_0 + \dim Z_1) + \dim B_1 + \dots \\ &= \dots - \dim C_{-1} + \dim C_0 - \dim C_1 + \dots = \sum_{n \in \mathbb{Z}} (-1)^n \dim C_n. \end{aligned}$$

Moving back to the original hypothesis where each C_n is a free abelian group, this computation implies

$$\sum_{n \in \mathbb{Z}} (-1)^n \operatorname{rank} C_n = \sum_{n \in \mathbb{Z}} (-1)^n \dim_{\mathbb{K}}(C_n \otimes \mathbb{K}) = \sum_{n \in \mathbb{Z}} (-1)^n \dim_{\mathbb{K}} H_n(C_* \otimes \mathbb{K})$$

for every choice of coefficient field \mathbb{K} . Now choosing $\mathbb{K} = \mathbb{Q}$ and applying (39.2) identifies the last expression with $\sum_{n \in \mathbb{Z}} (-1)^n \dim_{\mathbb{K}}(H_n(C_*) \otimes \mathbb{Q}) = \sum_{n \in \mathbb{Z}} (-1)^n \operatorname{rank} H_n(C_*)$ and thus completes the proof. \square

Notice what we gain if this result is combined with the isomorphism $H_*^{\text{CW}}(X; G) \cong H_*(X; G)$ for X a finite cell complex: in this case $H_*(X)$ is the homology of a finitely-generated chain complex of free abelian groups $C_*^{\text{CW}}(X)$, and this knowledge alone is sufficient to compute the Euler characteristic $\chi(X)$ *without needing to compute the homology!* All we actually need to know for the computation is the rank of each chain group $C_n^{\text{CW}}(X)$, which is the same as the number of n -cells in the complex. But since $\chi(X)$ is a topological invariant, we also learn from this that the alternating sum of these counts of cells does not depend on the choice of cell decomposition:

COROLLARY 39.11. *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. \square

Computing Euler characteristics of cell complexes is now quite easy.

EXAMPLE 39.12. 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 $k = 0, \dots, n$.

EXAMPLE 39.13. For the closed surface Σ_g of genus $g \geq 0$, we computed $H_*(\Sigma_g)$ in Example 36.3: the nontrivial homology groups were $H_0(\Sigma_g) \cong \mathbb{Z}$, $H_1(\Sigma_g) \cong \mathbb{Z}^{2g}$ and $H_2(\Sigma_g) \cong \mathbb{Z}$, thus

$$\chi(\Sigma_g) = 1 - 2g + 1 = 2 - 2g.$$

But one can also compute $\chi(\Sigma_g)$ without computing $H_*(\Sigma_g)$, just by observing that Σ_g has a cell decomposition with one 0-cell, one 2-cell and $2g$ cells of dimension 1; this is the same cell decomposition we used in Example 36.3, but there is no longer any need to compute the boundary map.

Here is an application of a more combinatorial nature. Recall that a **graph** (*Graph*) consists of a set V whose elements are called **vertices** (*Ecken* or *Punkte*), and a set E whose elements are called **edges** (*Kanten*), each of which is associated to a particular pair of vertices. Graphs are typically depicted by drawing a point for each vertex and drawing a curve for each edge such that its end points are the two vertices associated to that edge, and in this way every graph Γ naturally gives rise to a 1-dimensional CW-complex $|\Gamma|$ whose 0-cells are the vertices and 1-cells are the edges. The space $|\Gamma|$ is compact if and only if the graph Γ is finite, meaning both V and E are finite, and we say that Γ is **connected** if $|\Gamma|$ is a connected space. A finite connected graph is called a **tree** (*Baum*) if it contains no **cycles**, meaning there does not exist any finite sequence of distinct vertices $v_0, \dots, v_N \in V$ together with a finite sequence of distinct edges e_0, \dots, e_N such that the end points of e_j are v_j and v_{j+1} for $j = 0, \dots, N-1$ but the end points of e_N are v_N and v_0 . Now, since $|\Gamma|$ is a 1-dimensional CW-complex, we have $H_k(|\Gamma|) = 0$ for all k except 0 and 1. If Γ is connected, then $|\Gamma|$ is also path-connected and therefore $H_0(|\Gamma|) \cong \mathbb{Z}$. Since there are no 2-cells, $H_1(|\Gamma|)$ is isomorphic to the subgroup of 1-cycles in $C_1^{CW}(|\Gamma|)$, but it is not hard to prove that if Γ is a tree, then there are also no nontrivial 1-cycles in the chain complex, so $H_1(|\Gamma|) = 0$. This proves $\chi(|\Gamma|) = 1$, and combining it with Corollary 39.11, we then have:

THEOREM 39.14. *For any finite graph Γ with v vertices and e edges, if Γ is a tree, then $v - e = 1$. \square*

Let's conclude this discussion with an application to covering spaces.

EXERCISE 39.15. Suppose X is a compact cell complex and $\pi : Y \rightarrow X$ is a covering map of finite degree $d \in \mathbb{N}$. Show that Y admits a cell decomposition such that $Y^n = \pi^{-1}(X^n)$ for every n , and every individual n -cell $e_\alpha^n \subset X$ corresponds to exactly d cells in Y whose characteristic maps $\mathbb{D}^n \rightarrow Y$ are lifts of the characteristic map $\mathbb{D}^n \rightarrow X$ for e_α^n .

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 .

The exercise implies:

THEOREM 39.16. *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 39.17. *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 39.18. 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. The Lefschetz fixed point theorem

As another application of cellular homology, in this lecture I'd like to address the following general question:

QUESTION 40.1. *What topological conditions on a map $X \xrightarrow{f} X$ are sufficient to guarantee that f has a fixed point?*

We saw one example last semester: by the Brouwer fixed point theorem, no conditions at all are needed for f if X is a disk. We also saw in Lecture 33 that for $X = S^n$, every map f that does not have degree $(-1)^{n+1}$ must have a fixed point—this is a homotopy-invariant condition, but of course it is important to include the exception in this statement, as e.g. the antipodal map does not have any fixed points.

Our goal for today is a much more general homotopy-invariant criterion for the existence of fixed points. The rough idea is as follows: if $f : X \rightarrow X$ has no fixed points but X is a nice enough space to admit cell decompositions, then we would like to find a special cell decomposition of X such that after adjusting f by a homotopy, f becomes a cellular map sending each cell to *different* cells. In other words, the induced chain map $f_* : C_*^{\text{CW}}(X) \rightarrow C_*^{\text{CW}}(X)$ given by

$$f_* e_\alpha^n = \sum_{e_\beta^n \subset X} [e_\beta^n : e_\alpha^n] e_\beta^n$$

will then have the property that its *diagonal* terms all vanish:

$$[e_\alpha^n : e_\alpha^n] = 0 \quad \text{for all } n\text{-cells } e_\alpha^n \subset X.$$

In this situation, the chain map f_* is represented by a matrix that has zeroes along the diagonal, so its trace vanishes. At this point it is useful to introduce an algebraic result that has much in common with the previous lecture’s Proposition 39.10:

THEOREM 40.2 (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*

$$\sum_{n \in \mathbb{Z}} (-1)^n \text{tr} \left(C_n \xrightarrow{f} C_n \right) = \sum_{n \in \mathbb{Z}} (-1)^n \text{tr} \left(H_n(C_*) \xrightarrow{f_*} H_n(C_*) \right).$$

PROOF. For the boundary maps $\partial_n : C_n \rightarrow C_{n-1}$ for each $n \in \mathbb{Z}$, abbreviate

$$Z_n := \ker \partial_n \subset C_n, \quad B_n := \text{im } \partial_{n+1} \subset C_n, \quad H_n := Z_n/B_n.$$

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

$$\begin{array}{ccccccc} 0 & \longrightarrow & Z_n & \hookrightarrow & C_n & \xrightarrow{\partial_n} & B_{n-1} & \longrightarrow & 0 \\ & & \downarrow f_{Z_n} & & \downarrow f_{C_n} & & \downarrow f_{B_{n-1}} & & \\ 0 & \longrightarrow & Z_n & \hookrightarrow & C_n & \xrightarrow{\partial_n} & B_{n-1} & \longrightarrow & 0 \end{array}$$

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—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

$$\begin{array}{ccccccccc} 0 & \longrightarrow & B_n & \hookrightarrow & Z_n & \longrightarrow & H_n & \longrightarrow & 0 \\ & & \downarrow f_{B_n} & & \downarrow f_{Z_n} & & \downarrow f_{H_n} & & \\ 0 & \longrightarrow & B_n & \hookrightarrow & Z_n & \longrightarrow & H_n & \longrightarrow & 0 \end{array}$$

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

DEFINITION 40.3. For any space X and a field \mathbb{K} such that $H_*(X; \mathbb{K})$ is finite dimensional, the **Lefschetz number** (*Lefschetz-Zahl*) of a map $f : X \rightarrow X$ 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 40.4. We will not need to know this for our discussion, but it's interesting to note that while the definition above makes $L(f)$ a rational number, it is secretly always an integer. If X is a finite CW-complex and f a cellular map, then this follows easily from the Hopf trace formula, as $L_{\mathbb{Q}}(f)$ is then the same as the alternating sum of the traces of maps $f_* : C_n^{\text{CW}}(X) \otimes \mathbb{Q} \rightarrow C_n^{\text{CW}}(X) \otimes \mathbb{Q}$, each of which can be written in the canonical basis as a matrix with integer entries. Without these assumptions, it follows more generally from the universal coefficient theorem, which will give us a natural isomorphism $H_*(X; \mathbb{Q}) \cong H_*(X) \otimes \mathbb{Q}$, so that the maps $f_* : H_*(X; \mathbb{Q}) \rightarrow H_*(X; \mathbb{Q})$ can also be presented as matrices with integer entries. More precisely, every endomorphism $H_n(X) \rightarrow H_n(X)$ preserves the torsion subgroup $T_n \subset H_n(X)$ and thus descends to an endomorphism of the *free part* of $H_n(X)$,

$$H_n(X)/T_n \xrightarrow{f_*} H_n(X)/T_n,$$

which is a free abelian group. Thus f_* can again be presented as an integer matrix with respect to any basis of this free group, and the alternating sum of the traces of these matrices is the integer $L(f)$.

EXERCISE 40.5. Show that if X has finitely-generated homology and $f : X \rightarrow X$ is homotopic to the identity map, $L(f) = \chi(X)$.

Here is the main result of this lecture.

THEOREM 40.6 (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 33.4), and its proof requires some knowledge of the *flow* of a smooth vector field from differential geometry.

COROLLARY 40.7. *For any closed smooth manifold M with $\chi(M) \neq 0$, there is no continuous vector field on M that is nowhere zero.*

PROOF. If such a vector field exists, then we can approximate it with a smooth vector field X that is also nowhere zero. The flow of X for some small but nonzero time $t > 0$ is then a diffeomorphism $\varphi_X^t : M \rightarrow M$ with no fixed points, but is clearly also homotopic to the identity, thus by Exercise 40.5, $L(\varphi_X^t) = \chi(M) = 0$. \square

REMARK 40.8. Another easy corollary of the theorem is that it also holds for spaces somewhat more general than compact polyhedra: it holds in particular whenever X is a compact **Euclidean neighborhood retract**, meaning X admits a topological embedding $X \hookrightarrow \mathbb{R}^N$ for some $N \in \mathbb{N}$ such that some neighborhood $\mathcal{U} \subset \mathbb{R}^N$ of X admits a retraction to X . It is not so hard to prove (see [Hat02, Corollary A.9]) that all compact topological manifolds have this property, even those which do not admit triangulations. In this situation, even if X does not have a triangulation, we can triangulate \mathbb{R}^N finely enough so that all simplices touching $X \subset \mathbb{R}^N$ are contained in the neighborhood \mathcal{U} , and the retraction $r : \mathcal{U} \rightarrow X$ 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 40.9 below to prove $L_{\mathbb{K}}(i \circ f \circ r) = L_{\mathbb{K}}(f)$, so that Theorem 40.6 guarantees a fixed point for $i \circ f \circ r$. But $i \circ f \circ r(x) = x$ implies $x \in X$ and $f(x) = x$.

EXERCISE 40.9. 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 $\mathrm{tr}(\mathbf{AB}) = \mathrm{tr}(\mathbf{BA})$.

REMARK 40.10. Lefschetz’s original version of the fixed point theorem applied only to manifolds and was thus more restrictive, but it has the following nice feature that Theorem 40.6 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 40.7 to the statement that on a closed oriented manifold M , for any vector field that has at most finitely many zeroes, the algebraic count of these zeroes is $\chi(M)$; this is known as the *Poincaré-Hopf theorem*.

REMARK 40.11. It is easy to see that the compactness of X in Theorem 40.6 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.




FIGURE 19. A compact space X with $\chi(X) = 1$ admitting maps homotopic to the identity that have no fixed point.

REMARK 40.12. 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.)

To prove Theorem 40.6, we need to make precise the idea sketched at the beginning of this lecture: 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. Since we are working with polyhedra, it is natural to consider not just cellular but also *simplicial* maps. You may want to take a moment to review the definitions given in Lecture 29 for the combinatorial notion of a simplicial map and the induced continuous map on polyhedra, which maps each n -simplex linearly to a k -simplex for some $k \leq n$. In this context, the following question seems natural, though its immediate answer may disappoint you:

QUESTION 40.13. *Given two polyhedra X and Y , is every continuous map $f : X \rightarrow Y$ homotopic to a simplicial map?*

The answer is **no**: for example, if Y happens to have more vertices than X , then a simplicial map $f : X \rightarrow Y$ can never be surjective. We could for instance take X and Y to be the sphere S^n with two distinct triangulations such that Y has more vertices, in which case only maps $S^n \rightarrow S^n$ with degree zero have any hope of being homotopic to simplicial maps. The message of this answer is that we asked the wrong question. Suppose we eliminate this counterexample as follows: instead of fixing given triangulations of X and Y , we first make the triangulation of X finer by *subdividing*




FIGURE 20. Barycentric subdivision of a 2-simplex.

it until it has at least as many vertices as Y . If this extra step is allowed, then it is no longer obvious that any given map $f : X \rightarrow Y$ cannot be homotopic to a simplicial map. I remind you that when I say “subdivision,” I typically mean *barycentric* subdivision, as shown in Figure 20.

Here’s the main technical result we need. I will give only a sketch of the proof, but the main idea is not so hard to understand. (For a more detailed proof, see [Hat02, §2.C].)

THEOREM 40.14 (simplicial approximation). *If X and Y are compact polyhedra and $f : X \rightarrow Y$ is any continuous map, then after modifying the triangulation of X by finitely many barycentric subdivisions, f is homotopic to a simplicial map $g : X \rightarrow Y$ such that for every $x \in X$, $g(x)$ is contained in the smallest simplex of Y containing $f(x)$.*

SKETCH OF THE PROOF. For each vertex $v \in X$, define the so-called **open star** of v as the open neighborhood

$$\text{st } v \subset X$$

of v formed by the union of the interiors of all simplices in X that have v as a vertex. Figure 21 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 40.15 below). 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 21 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 , the exercise below implies that the set $\{w_{v_0}, \dots, w_{v_n}\}$ is also a simplex of Y . One can now check that g is indeed an “approximation” of f in the sense that $g(x)$ is contained in the smallest simplex of Y containing $f(x)$ for every $x \in X$. In light of this, a homotopy $h : I \times X \rightarrow Y$ 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

EXERCISE 40.15. 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 .

We can now prove the Lefschetz-Hopf theorem.

PROOF OF THEOREM 40.6. 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 21. 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 40.14 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} .

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

41. Coefficients and the Tor functor

In this lecture we'll begin a closer examination of precisely how the groups $H_*(X, A; G_1)$ and $H_*(X, A; G_2)$ are related to each other for different choices of abelian coefficient group G_1 and G_2 . Recall that $H_*(\cdot; G)$ is the composition of three functors:

$$(41.1) \quad \text{Top}_{\text{rel}} \xrightarrow{C_*} \text{Chain} \xrightarrow{\otimes G} \text{Chain} \xrightarrow{H_*} \text{Ab}_{\mathbb{Z}},$$

where the first sends each pair of spaces (X, A) to its singular chain complex $C_*(X, A)$ with integer coefficients, the second replaces the latter with $C_*(X, A; G) := C_*(X, A) \otimes G$, and the third computes the homology of the chain complex. If we were only concerned with *defining* topological invariants and not with *computing* them, then we could stop with $C_* : \text{Top}_{\text{rel}} \rightarrow \text{Chain}$, as the singular chain complex $C_*(X, A)$ is in itself a topological invariant, and it contains in principle all of the information that we could ever want to extract from any version of the singular homology of (X, A) . The problem is that, as an invariant in itself, $C_*(X, A)$ is horribly unwieldy and impractical: the group is absurdly large, and if you want to prove (X, A) and (Y, B) are not homeomorphic, you typically cannot do it by proving directly that $C_*(X, A)$ and $C_*(Y, B)$ are not isomorphic chain complexes. This is where replacing the complexes with their homology groups is useful: strictly speaking, we lose a lot of information when we do this, but it's worth it if the information that remains afterwards is manageable.

The situation is slightly different for the cellular homology functor $H_*^{\text{CW}}(\cdot; G) : \text{CW}_{\text{rel}} \rightarrow \text{Ab}_{\mathbb{Z}}$, which is also a composition of three functors

$$\text{CW}_{\text{rel}} \xrightarrow{C_*^{\text{CW}}} \text{Chain} \xrightarrow{\otimes G} \text{Chain} \xrightarrow{H_*} \text{Ab}_{\mathbb{Z}}.$$

For a CW-pair (X, A) , the chain complex $C_*^{\text{CW}}(X, A)$ is typically much more manageable, but it depends on the cell decomposition and is thus not a topological invariant. Passing to the homology $H_*^{\text{CW}}(X, A)$ is thus necessary in order to obtain something that depends only on the topology of (X, A) . One could say the same thing about simplicial homology, which of course is just a special case of cellular homology.

For this lecture we're going to focus on the purely algebraic aspects that are common to both singular and cellular homology, namely the two functors

$$\text{Chain} \xrightarrow{\otimes G} \text{Chain} \xrightarrow{H_*} \text{Ab}_{\mathbb{Z}}.$$

Each of them destroys some information in general, e.g. the homology of a chain complex can easily be trivial even when the complex itself is not, and applying $\otimes \mathbb{Q}$ to an abelian group with nontrivial torsion elements will always kill them. The case $G = \mathbb{Z}$ is special, because $\otimes \mathbb{Z}$ is actually the identity functor, so there seems to be less potential for losing information if we stick with integer coefficients. Of course we have also seen cases (e.g. the Klein bottle) where $H_n(X; \mathbb{Z})$ is trivial while $H_n(X; \mathbb{Z}_2)$ is not. The main result of this and the next lecture will show however that $H_n(X; \mathbb{Z}_2)$ is nevertheless determined by the collection of all the groups $H_k(X; \mathbb{Z})$; in fact it suffices to consider these groups for $k = n$ and $k = n - 1$. Results like this should not be interpreted to mean that homology with different coefficient groups is *unnecessary*—we’ve seen for instance that when X is a manifold with a non-orientable triangulation, it is in some ways more natural to consider $H_*(X; \mathbb{Z}_2)$ than $H_*(X; \mathbb{Z})$. But as a computational device, it is also often useful to know that the homology with integer coefficients determines everything else. Moreover, there is a kind of converse to this statement, allowing information about $H_*(X; \mathbb{Z})$ to be deduced from a collection of groups $H_*(X; \mathbb{K})$ for various choices of field coefficients \mathbb{K} .

The universal coefficient theorem. Given a chain complex C_* , we would now like to compare the homology of this complex with that of $C_* \otimes G$ for some coefficient group G . The first important observation is that for every $n \in \mathbb{Z}$, there is a natural homomorphism

$$(41.2) \quad H_n(C_*) \otimes G \xrightarrow{h} H_n(C_* \otimes G) : [c] \otimes g \mapsto [c \otimes g].$$

Indeed, this map is well defined since $\partial c = 0$ implies $(\partial \otimes \mathbb{1})(c \otimes g) = 0$ and $c = \partial a$ implies $c \otimes g = (\partial \otimes \mathbb{1})(a \otimes g)$. One might hope for this map to be an isomorphism, and we will show that this is true in many important cases, but not always. Here is the general result.

THEOREM 41.1 (universal coefficient theorem). *There exists a functor⁵⁶*

$$\text{Tor} : \text{Ab} \times \text{Ab} \rightarrow \text{Ab},$$

covariant in both variables, such that the following is true. For any chain complex C_ of free abelian groups, a fixed abelian group G and $n \in \mathbb{Z}$, 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 map in (41.2). Moreover, the sequence (but not its splitting) is natural in the sense that for any chain map of $\Phi : A_ \rightarrow B_*$ between two chain complexes of free abelian groups, 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 .*

⁵⁶While it is conventional to call Tor a “functor,” we will see from the definition that this is cheating a little bit, as one cannot simply feed a pair of groups G and H into Tor and extract a well-defined group $\text{Tor}(G, H)$. In reality, the definition of $\text{Tor}(G, H)$ requires some auxiliary choices beyond the groups G and H , but we will see that for any two sets of these choices, there is a canonical isomorphism between our two definitions of $\text{Tor}(G, H)$. In this sense, $\text{Tor}(G, H)$ is well defined in the same sense that “the one-point space” is well defined: one can take any set of one element and label it “{pt},” and any two spaces defined in this way are not technically the same, but there is a canonical homeomorphism between them.

In this lecture our main objective will be to explain what the mysterious extra term $\text{Tor}(H_{n-1}(C_*), G)$ means, and we will proceed to prove the theorem in the next lecture. The splitting of the exact sequence means that there is always an isomorphism

$$H_n(C_* \otimes G) \cong (H_n(C_*) \otimes G) \oplus \text{Tor}(H_{n-1}(C_*), G),$$

so that is the sense in which $H_n(C_*)$ and $H_{n-1}(C_*)$ determine $H_n(C_* \otimes G)$. It should be emphasized however that this isomorphism depends in general on non-canonical choices and does not fit into any nice commutative diagrams together with maps induced by chain maps—this is what is meant when we say that the splitting is “not natural”. The isomorphism is in any case not very useful unless one has a means of computing $\text{Tor}(H_{n-1}(C_*), G)$, but the following result makes this quite doable in most cases that arise in practice:

THEOREM 41.2. *The functor Tor has the following properties for all abelian groups A, B, G :*

- (1) $\text{Tor}(A, G) = 0$ whenever A is free;
- (2) $\text{Tor}(A \oplus B, G) \cong \text{Tor}(A, G) \oplus \text{Tor}(B, G)$;
- (3) For every $k \in \mathbb{N}$, $\text{Tor}(\mathbb{Z}_k, G)$ is isomorphic to the kernel of the map $G \rightarrow G : g \mapsto kg$.

The proof of this theorem will be an easy exercise once the definition of Tor is understood. The following sample application is also an extremely useful computational tool in itself.

COROLLARY 41.3. *If A is any finitely-generated⁵⁷ abelian group and G is an abelian group with trivial torsion, then $\text{Tor}(A, G) = 0$.*

PROOF. By the classification of finitely-generated abelian groups, $A \cong F \oplus \mathbb{Z}_{n_1} \oplus \dots \oplus \mathbb{Z}_{n_k}$ for some free abelian group F and $n_1, \dots, n_k \in \mathbb{N}$. Thus $\text{Tor}(F, G) = 0$, and

$$\text{Tor}(A, G) = \text{Tor}(F, G) \oplus \left(\bigoplus_{i=1}^k \text{Tor}(\mathbb{Z}_{n_i}, G) \right) = \bigoplus_{i=1}^k \ker(G \xrightarrow{n_i} G) = 0$$

since G is torsion free. □

The following easy consequence was already mentioned in the previous two lectures:

COROLLARY 41.4. *For any field \mathbb{K} of characteristic zero and any chain complex C_* of free abelian groups such that $H_n(C_*)$ is finitely generated for every n ,⁵⁸ the natural map $h : H_*(C_*) \otimes \mathbb{K} \rightarrow H_*(C_* \otimes \mathbb{K})$ is an isomorphism. In particular, if the singular homology of a pair (X, A) is finitely generated in every dimension, then there is a natural vector space isomorphism $H_*(X, A; \mathbb{K}) \cong H_*(X, A) \otimes \mathbb{K}$. □*

EXAMPLE 41.5. The characteristic zero assumption in the above corollary cannot be dropped in general: we know this already from the example of the Klein bottle K^2 , which has $H_2(K^2) = 0$ but $H_2(K^2; \mathbb{Z}_2) \cong \mathbb{Z}_2 \neq 0 = H_2(K^2) \otimes \mathbb{Z}_2$. Here is a sanity check: since $H_1(K^2) \cong \mathbb{Z} \oplus \mathbb{Z}_2$, the universal coefficient theorem in this case gives

$$H_2(K^2; \mathbb{Z}_2) \cong (H_2(K^2) \otimes \mathbb{Z}_2) \oplus \text{Tor}(H_1(K^2), \mathbb{Z}_2) \cong \text{Tor}(\mathbb{Z} \oplus \mathbb{Z}_2, \mathbb{Z}_2),$$

and by Theorem 41.2,

$$\text{Tor}(\mathbb{Z} \oplus \mathbb{Z}_2, \mathbb{Z}_2) = \text{Tor}(\mathbb{Z}, \mathbb{Z}_2) \oplus \text{Tor}(\mathbb{Z}_2, \mathbb{Z}_2) = \text{Tor}(\mathbb{Z}_2, \mathbb{Z}_2) \cong \ker(\mathbb{Z}_2 \xrightarrow{2} \mathbb{Z}_2) = \ker(\mathbb{Z}_2 \xrightarrow{0} \mathbb{Z}_2) = \mathbb{Z}_2.$$

⁵⁷The corollary is also true without assuming A to be finitely generated, the proof just takes a bit more effort; see [Hat02, Prop. 3A.5].

⁵⁸As with the footnote on Corollary 41.3, it is not actually necessary to assume $H_n(C_*)$ is finitely generated, but this assumption is satisfied in every example we are likely to care about.

Projective resolutions and Tor. The explanation of Tor begins with a thought-experiment. The simplest case in which we might hope to compare $H_*(C_*)$ with $H_*(C_* \otimes G)$ is when $H_*(C_*)$ is trivial, i.e. the complex C_* is an exact sequence. Let's simplify further and consider a *short* exact sequence

$$0 \longrightarrow A \xrightarrow{i} B \xrightarrow{j} C \longrightarrow 0$$

of (not necessarily free) abelian groups, and ask whether the associated sequence

$$0 \longrightarrow A \otimes G \xrightarrow{i \otimes \mathbb{1}} B \otimes G \xrightarrow{j \otimes \mathbb{1}} C \otimes G \longrightarrow 0$$

must also be exact. The answer is in general no: for example,

$$0 \longrightarrow \mathbb{Z} \xrightarrow{\cdot 2} \mathbb{Z} \xrightarrow{\text{pr}} \mathbb{Z}_2 \longrightarrow 0$$

is an exact sequence, but applying $\otimes \mathbb{Z}_2$ to it produces

$$0 \longrightarrow \mathbb{Z} \otimes \mathbb{Z}_2 \xrightarrow{\cdot 2} \mathbb{Z} \otimes \mathbb{Z}_2 \xrightarrow{\text{pr} \otimes \mathbb{1}} \mathbb{Z}_2 \otimes \mathbb{Z}_2 \longrightarrow 0,$$

which is not exact, because the first map $\mathbb{Z}_2 \xrightarrow{\cdot 2} \mathbb{Z}_2$ is trivial and thus not injective. The problem here is with the first nontrivial term, and this turns out to be the only problem in general:

PROPOSITION 41.6. *If $A \xrightarrow{i} B \xrightarrow{j} C \rightarrow 0$ is an exact sequence of abelian groups, then for every abelian group G , the sequence $A \otimes G \xrightarrow{i \otimes \mathbb{1}} B \otimes G \xrightarrow{j \otimes \mathbb{1}} C \otimes G \rightarrow 0$ is also exact.*

PROOF. Exactness of the first sequence means that j is surjective and $\text{im } i = \ker j \subset B$. Then for every $c \in C$ and $g \in G$, we can find $b \in B$ with $j(b) = c$ and write $(j \otimes \mathbb{1})(b \otimes g) = c \otimes g$, implying that $j \otimes \mathbb{1}$ is surjective. Clearly $(j \otimes \mathbb{1}) \circ (i \otimes \mathbb{1}) = (j \circ i) \otimes \mathbb{1} = 0$ since $j \circ i = 0$, so we only still need to show $\ker(j \otimes \mathbb{1}) \subset \text{im}(i \otimes \mathbb{1})$, which is now equivalent to showing that the map

$$(41.3) \quad (B \otimes G) / \text{im}(i \otimes \mathbb{1}) \rightarrow C \otimes G : [b \otimes g] \mapsto j(b) \otimes g$$

is injective (in which case it is an isomorphism). We can do this by constructing its inverse: define $\Phi : C \otimes G \rightarrow (B \otimes G) / \text{im}(i \otimes \mathbb{1})$ by $\Phi(c \otimes g) := b \otimes g$ for any $b \in j^{-1}(c)$. This is well defined because for any two choices $b, b' \in j^{-1}(c)$, we have $b' - b \in \ker j = \text{im } i$ and thus $b' - b = i(a)$ for some $a \in A$, thus $b' \otimes g - b \otimes g = i(a) \otimes g \in \text{im}(i \otimes \mathbb{1})$. It is easy to check that Φ is an inverse for the map (41.3). \square

The technical term for the phenomenon in Proposition 41.6 is that $\otimes G$ is a **right-exact functor**. We would drop the prefix “right-” from this terminology if an initial “ $0 \rightarrow$ ” could be added to both sequences, but the example above with $G = \mathbb{Z}_2$ shows that this does not work in general. The situation for *split* exact sequences is better:

EXERCISE 41.7. Show that if $0 \rightarrow A \xrightarrow{i} B \xrightarrow{j} C \rightarrow 0$ is a split exact sequence of abelian groups, then for every abelian group G , the sequence $0 \rightarrow A \otimes G \xrightarrow{i \otimes \mathbb{1}} B \otimes G \xrightarrow{j \otimes \mathbb{1}} C \otimes G \rightarrow 0$ is also split exact.

Hint: Up to isomorphism, you can assume $B = A \oplus C$ and write very simple formulas for i and j .

We will see that the Tor term appearing in the universal coefficient theorem is in some sense measuring the failure of $\otimes G$ to be an exact functor. To say more precisely what this means, it is natural at this point to generalize somewhat beyond abelian groups and discuss more general modules over a commutative ring R with unit. Recall that an **R -module homomorphism** $G \rightarrow H$ is then a group homomorphism $\Phi : G \rightarrow H$ that additionally respects the “scalar multiplication” by elements of R , i.e. it satisfies

$$\Phi(\lambda g) = \lambda \Phi(g) \quad \text{for all } \lambda \in R, g \in G.$$

If $R = \mathbb{Z}$, then nothing has changed and we are still talking about abelian groups with group homomorphisms; on the other end of the spectrum, one can choose R to be a field \mathbb{K} , so that R -modules are vector spaces over \mathbb{K} and R -module homomorphisms are \mathbb{K} -linear maps. In this more general context, there are obvious notions of chain complexes, chain maps, chain homotopies, homology groups and exact sequences, where all objects are assumed to be R -modules and all morphisms are R -module homomorphisms.

DEFINITION 41.8. An R -module F is called **projective** if for every surjective R -module homomorphism $\pi : G \rightarrow H$, every R -module homomorphism $\varphi : F \rightarrow H$ can be lifted to an R -module homomorphism $\tilde{\varphi} : F \rightarrow G$ so that the following diagram commutes:

$$\begin{array}{ccc} & & G \\ & \nearrow \tilde{\varphi} & \downarrow \pi \\ F & \xrightarrow{\varphi} & H \end{array}$$

EXAMPLE 41.9. Every free abelian group is a projective \mathbb{Z} -module. Indeed, if F has a basis $B \subset F$, then the required lift $\tilde{\varphi} : F \rightarrow G$ can be defined by choosing any $\tilde{\varphi}(b) \in \pi^{-1}(\varphi(b))$ for each $b \in B$ and extending $\tilde{\varphi}$ to the unique homomorphism with these values on the basis elements.

EXAMPLE 41.10. The group \mathbb{Z}_2 is not a projective \mathbb{Z} -module. For example, the lift in the diagram

$$\begin{array}{ccc} & & \mathbb{Z} \\ & \nearrow \tilde{\varphi} & \downarrow \pi \\ \mathbb{Z}_2 & \xrightarrow{1} & \mathbb{Z}_2 \end{array}$$

can never exist since $\text{Hom}(\mathbb{Z}_2, \mathbb{Z}) = 0$.

Example 41.9 extends easily to R -modules for any R : recall that an R -module A is called **free** if it admits a **basis**, meaning a subset $B \subset A$ such that every element $x \in A$ is uniquely expressible as $\sum_{b \in B} x_b b$ with at most finitely many nonzero coefficients $x_b \in R$. This is equivalent to saying that A is isomorphic to

$$F_R^{\text{mod}}(S) := \bigoplus_{s \in S} R,$$

the so-called **free R -module on S** . Clearly every free R -module is also projective.

DEFINITION 41.11. A **projective resolution** (F_*, f_*) of an R -module A is an exact sequence

$$\dots \longrightarrow F_2 \xrightarrow{f_2} F_1 \xrightarrow{f_1} F_0 \xrightarrow{f_0} A \longrightarrow 0$$

such that all of the F_i for $i = 0, 1, 2, \dots$ are projective R -modules.

PROPOSITION 41.12. *Every R -module A admits a projective resolution.*

PROOF. Pick any generating set $S_0 \subset A$, e.g. it would suffice to choose $S_0 = A$, though smaller subsets are usually possible. We then set $F_0 = F_R^{\text{mod}}(S_0)$ and define $f_0 : F_0 \rightarrow A$ as the unique R -module homomorphism that extends the inclusion $S_0 \hookrightarrow A$, noting that f_0 is surjective by construction. Next, pick S_1 to be a generating subset of $\ker f_0 \subset F_0$, and define $F_1 = F_R^{\text{mod}}(S_1)$ and $f_1 : F_1 \rightarrow \ker f_0$ analogously; this defines $f_1 : F_1 \rightarrow F_0$ such that $\text{im } f_1 = \ker f_0$. Now continue this process inductively: all of the modules F_i produced in this way are projective since they are free. □

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 41.13. *Given an R -module homomorphism $\varphi : A \rightarrow B$ and any projective resolutions (F_*, f_*) of A and (G_*, g_*) of B , there exists a chain map φ_* from F_* to G_* , i.e. a sequence of R -module homomorphisms $\varphi_i : F_i \rightarrow G_i$ such that the diagram*

$$(41.4) \quad \begin{array}{ccccccc} \dots & \longrightarrow & F_2 & \xrightarrow{f_2} & F_1 & \xrightarrow{f_1} & F_0 & \xrightarrow{f_0} & A \\ & & \downarrow \varphi_2 & & \downarrow \varphi_1 & & \downarrow \varphi_0 & & \downarrow \varphi \\ \dots & \longrightarrow & G_2 & \xrightarrow{g_2} & G_1 & \xrightarrow{g_1} & G_0 & \xrightarrow{g_0} & B \end{array}$$

commutes. Moreover, this chain map is unique up to chain homotopy.

COROLLARY 41.14. *Any two projective resolutions of the same R -module are chain homotopy equivalent.*

PROOF. If (F_*, f_*) and (G_*, g_*) are both projective resolutions of A , then applying Proposition 41.13 with $\varphi : A \rightarrow A$ as the identity map produces chain maps $(F_*, f_*) \rightarrow (G_*, g_*)$ and $(G_*, g_*) \rightarrow (F_*, f_*)$, which can be composed in either direction to produce chain maps from (F_*, f_*) to itself and from (G_*, g_*) to itself. Uniqueness then implies that both of the latter are chain homotopy equivalent to the identity. \square

PROOF OF PROPOSITION 41.13. For convenience denote $\varphi_{-1} := \varphi$, and assume for some integer $k \geq 0$ that the maps $\varphi_{-1}, \dots, \varphi_{k-1}$ in (41.4) have already been constructed so that all the relevant squares commute. We must then find a map $\varphi_k : F_k \rightarrow G_k$ such that $g_k \varphi_k = \varphi_{k-1} f_k$. Notice that

$$g_{k-1} \varphi_{k-1} f_k = \varphi_{k-2} f_{k-1} f_k = 0,$$

thus $\text{im}(\varphi_{k-1} f_k) \subset \ker g_{k-1} = \text{im } g_k$, and we can therefore define φ_k to be any solution to the lifting problem

$$\begin{array}{ccc} & & G_k \\ & \nearrow \varphi_k & \downarrow g_k \\ F_k & \xrightarrow{\varphi_{k-1} f_k} & \ker g_{k-1} \end{array}$$

A solution exists since F_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 $h_k : F_k \rightarrow G_{k+1}$ for $k \geq 0$ satisfying

$$\varphi_k - \psi_k = g_{k+1} h_k + h_{k-1} f_k$$

for every k . For this to make sense when $k = 0$, we need also a map $h_{-1} : A \rightarrow G_0$, which we define as $h_{-1} := 0$. (Note that this makes sense if we regard both rows of (41.4) as chain complexes extended to the right with a sequence of zero terms, since $\varphi_{-1} - \psi_{-1} = \varphi - \varphi = 0$ by definition.) 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 $h_k : F_k \rightarrow G_{k+1}$ such that

$$g_{k+1} h_k = \varphi_k - \psi_k - h_{k-1} f_k.$$

We observe that by commutativity and the chain homotopy relation for $k - 1$,

$$\begin{aligned} g_k(\varphi_k - \psi_k - h_{k-1} f_k) &= (\varphi_{k-1} - \psi_{k-1}) f_k - g_k h_{k-1} f_k \\ &= (g_k h_{k-1} + h_{k-2} f_{k-1} - g_k h_{k-1}) f_k = h_{k-2} f_{k-1} f_k = 0, \end{aligned}$$

so $\text{im}(\varphi_k - \psi_k - h_{k-1}f_k) \subset \ker g_k = \text{im } g_{k+1}$, and h_k can now be defined as any solution to the lifting problem

$$\begin{array}{ccc}
 & & G_{k+1} \\
 & \nearrow h_k & \downarrow g_{k+1} \\
 F_k & \xrightarrow{\varphi_k - \psi_k - h_{k-1}f_k} & \ker g_k
 \end{array}$$

The result now follows again by induction on k . □

Your first instinct when you see a chain map like $\varphi_* : (F_*, f_*) \rightarrow (G_*, g_*)$ as in the above result should be to look at the homomorphisms it induces between the homologies of the two chain complexes. Unfortunately, that is not interesting in this case, as (F_*, f_*) and (G_*, g_*) are both exact sequences, so their homologies vanish. But this is where the failure of $\otimes G$ to be an exact functor comes into play: for any pair of abelian groups A and G with a chosen projective resolution (F_*, f_*) of A , the exact sequence $\dots F_2 \xrightarrow{f_2} F_1 \xrightarrow{f_1} F_0 \xrightarrow{f_0} A \rightarrow 0$ induces a chain complex

$$\dots \rightarrow F_2 \otimes G \xrightarrow{f_2 \otimes \mathbb{1}} F_1 \otimes G \xrightarrow{f_1 \otimes \mathbb{1}} F_0 \otimes G \xrightarrow{f_0 \otimes \mathbb{1}} A \otimes G \rightarrow 0,$$

whose homology groups

$$H_n(F_* \otimes G) = \ker(f_n \otimes \mathbb{1}) / \text{im}(f_{n+1} \otimes \mathbb{1})$$

might not be trivial, but are in any case well defined up to canonical isomorphisms due to Proposition 41.13:

COROLLARY 41.15. *Given abelian groups G, A and B , one can associate to any homomorphism $\varphi : A \rightarrow B$ and chosen projective resolutions (F_*, f_*) of A and (F'_*, f'_*) of B a sequence of canonical homomorphisms*

$$\varphi_* : H_n(F_* \otimes G) \rightarrow H_n(F'_* \otimes G), \quad n \in \mathbb{Z},$$

which are functorial in the sense that composable homomorphisms give $(\psi \circ \varphi)_* = \psi_* \varphi_*$ and $A \xrightarrow{\mathbb{1}} A$ with the same projective resolution on both sides gives $\mathbb{1}_* = \mathbb{1}$. In particular, any two choices of projective resolutions (F_*, f_*) and (F'_*, f'_*) for the same group A come with canonical isomorphisms $H_n(F_* \otimes G) \cong H_n(F'_* \otimes G)$. □

DEFINITION 41.16. Given abelian groups A and G , we define

$$\text{Tor}(A, G) := H_1(F_* \otimes G),$$

where F_* is any choice of projective resolution for A . For any homomorphism $\varphi : A \rightarrow B$ of abelian groups, an induced homomorphism

$$\text{Tor}(A, G) \xrightarrow{\varphi_*} \text{Tor}(B, G)$$

is defined as the canonical map $\varphi_* : H_1(F_* \otimes G) \rightarrow H_1(F'_* \otimes G)$ given in Corollary 41.15 for any choice of projective resolution F'_* of B .

EXERCISE 41.17. Show that $\text{Tor}(A, G)$ is also a covariant functor in the second variable, i.e. homomorphisms $G \rightarrow H$ naturally induce homomorphisms $\text{Tor}(A, G) \rightarrow \text{Tor}(A, H)$.

There are good reasons to single out the group $H_1(F_* \otimes G)$ rather than considering $H_n(F_* \otimes G)$ for other values of $n \geq 0$; we'll get into this in the next lecture.

EXERCISE 41.18. Prove Theorem 41.2 using the following hints for each of the three stated properties:

- (1) If A is a free abelian group, construct a projective resolution (F_*, f_*) with $F_n = 0$ for all $n \geq 1$.

- (2) If (F_*, f_*) and (F'_*, f'_*) are projective resolutions of A and B respectively, construct a projective resolution of $A \oplus B$ using the groups $F_n \oplus F'_n$.
- (3) Construct a projective resolution of \mathbb{Z}_k with $F_1 = F_0 = \mathbb{Z}$ and $F_n = 0$ for all $n \geq 2$.

REMARK 41.19. A few useful properties of Tor beyond those listed in Theorem 41.2 are proved in [Hat02, Prop. 3A.5].

42. Universal coefficient theorem and Künneth formula

In the previous lecture, we defined the group $\text{Tor}(A, G)$ for any pair of abelian groups A and G as the first homology group of a tensored chain complex,

$$\text{Tor}(A, G) = H_1(F_* \otimes G),$$

where (F_*, f_*) is any choice of projective resolution $\dots \rightarrow F_2 \rightarrow F_1 \rightarrow F_0 \rightarrow A \rightarrow 0$. We have seen why this group is well defined up to canonical isomorphisms, but you may still be wondering what is special about $H_1(F_* \otimes G)$, as opposed to $H_n(F_* \otimes G)$ for other integers $n \geq 0$. The answer comes from the following fundamental algebraic result, which we will also need to make essential use of in proving the universal coefficient theorem. You can find a proof in many standard textbooks on algebra, such as [Lan02].

PROPOSITION 42.1. *Every subgroup of a free abelian group is also a free abelian group.* \square

REMARK 42.2. The proof of this result is fairly elementary when the group is finitely generated (see [Lan02, §I.7]), though if the group is not finitely generated, then the standard proof requires Zorn's lemma, a version of the axiom of choice. If the axiom of choice makes you uncomfortable, then you will have to bury that discomfort now, as many of the free abelian groups we want to deal with are in fact uncountably infinitely generated, e.g. the singular chain groups of any interesting topological space.

COROLLARY 42.3. *Every abelian group A admits a projective resolution of the form $\dots \rightarrow 0 \rightarrow F_1 \xrightarrow{f_1} F_0 \xrightarrow{f_0} A \rightarrow 0$. In particular, $H_n(F_* \otimes G) = 0$ for every $n \geq 2$ and every abelian group G .*

PROOF. Take F_0 to be any free abelian group that admits a surjective homomorphism $f_0 : F_0 \rightarrow A$, e.g. the free abelian group on any generating subset of A . Then Proposition 42.1 implies that $\ker f_0 \subset F_0$ is also free and therefore projective, hence

$$\dots \rightarrow 0 \rightarrow \ker f_0 \hookrightarrow F_0 \xrightarrow{f_0} A \rightarrow 0$$

is a projective resolution of A . \square

The group $H_0(F_* \otimes G)$ is also trivial, for slightly simpler reasons: if (F_*, f_*) is a projective resolution, then $F_1 \rightarrow F_0 \rightarrow A \rightarrow 0$ is an exact sequence, so by Proposition 41.6 (i.e. the fact that $\otimes G$ is a *right-exact* functor), so is $F_1 \otimes G \rightarrow F_0 \otimes G \rightarrow A \otimes G \rightarrow 0$, making the homology at the term $F_0 \otimes G$ zero. The upshot is that when we apply the functor $\otimes G$ to a projective resolution (F_*, f_*) of A , the only nontrivial homology group that can emerge is $H_1(F_* \otimes G)$. It serves as a measurement of the failure of $\otimes G$ to be an exact functor, i.e. if we could also say that projective resolutions of the form $0 \rightarrow F_1 \rightarrow F_0 \rightarrow A \rightarrow 0$ as in Corollary 42.3 always give rise to exact sequences $0 \rightarrow F_1 \otimes G \rightarrow F_0 \otimes G \rightarrow A \otimes G \rightarrow 0$, then $H_1(F_* \otimes G)$ would also vanish. This is why it gets a special name.

We can now prove the universal coefficient theorem. Here is the statement again in short form.

THEOREM 42.4. *For any chain complex C_* of free abelian groups, a fixed abelian group G and $n \in \mathbb{Z}$, there exists a natural 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 canonical map defined in (41.2), and the sequence splits (but not naturally).

PROOF. For the given boundary maps $\partial_n : C_n \rightarrow C_{n-1}$, abbreviate

$$Z_n := \ker \partial_n \subset C_n, \quad B_n := \text{im } \partial_{n+1} \subset Z_n \subset C_n,$$

and note that since C_n is free, Proposition 42.1 implies that Z_n and B_n are also free. Writing $\text{pr} : Z_n \rightarrow H_n(C_*) = Z_n/B_n$ for the natural quotient projection, the short exact sequence

$$0 \longrightarrow B_n \xrightarrow{i_n} Z_n \xrightarrow{\text{pr}} H_n(C_*) \longrightarrow 0$$

then becomes a projective resolution of $H_n(C_*)$ if we insert an infinite sequence of zero terms at the left. Tensoring with G now gives the chain complex

$$\dots \longrightarrow 0 \longrightarrow B_n \otimes G \xrightarrow{i_n \otimes \mathbb{1}} Z_n \otimes G \xrightarrow{\text{pr} \otimes \mathbb{1}} H_n(C_*) \otimes G \longrightarrow 0,$$

which is exact at the last two terms by Proposition 41.6, thus $\text{pr} \otimes \mathbb{1}$ is surjective and descends to an isomorphism

$$(42.1) \quad \text{coker}(i_n \otimes \mathbb{1}) = (Z_n \otimes G) / \text{im}(i_n \otimes \mathbb{1}) \xrightarrow{\text{pr} \otimes \mathbb{1}} H_n(C_*) \otimes G.$$

By the definition of Tor, we also have

$$(42.2) \quad \ker(i_n \otimes \mathbb{1}) = \text{Tor}(H_n(C_*), G).$$

Now consider the short exact sequence $0 \rightarrow Z_* \hookrightarrow C_* \xrightarrow{\partial} B_{*-1} \rightarrow 0$, which is actually a short exact sequence of chain complexes if we define $Z_* := \bigoplus_{n \in \mathbb{Z}} Z_n$ and $B_{*-1} := \bigoplus_{n \in \mathbb{Z}} B_{n-1}$ as chain complexes with trivial boundary maps, i.e. we have a commuting diagram

$$\begin{array}{ccccccc} & & \vdots & & \vdots & & \vdots \\ & & \downarrow & & \downarrow & & \downarrow \\ 0 & \longrightarrow & Z_n & \hookrightarrow & C_n & \xrightarrow{\partial_n} & B_{n-1} \longrightarrow 0 \\ & & \downarrow 0 & & \downarrow \partial_n & & \downarrow 0 \\ 0 & \longrightarrow & Z_{n-1} & \hookrightarrow & C_{n-1} & \xrightarrow{\partial_{n-1}} & B_{n-2} \longrightarrow 0 \\ & & \downarrow 0 & & \downarrow \partial_{n-1} & & \downarrow 0 \\ & & \vdots & & \vdots & & \vdots \end{array}$$

where each row is a short exact sequence. Since the groups B_n are free, each row also splits, so that by Exercise 41.7, applying $\otimes G$ to each of these rows produces exact sequences $0 \rightarrow Z_n \otimes G \rightarrow C_n \otimes G \rightarrow B_{n-1} \otimes G$, thus forming another short exact sequence of chain complexes

$$0 \longrightarrow Z_* \otimes G \longrightarrow C_* \otimes G \longrightarrow B_{*-1} \otimes G \longrightarrow 0.$$

By the usual diagram-chasing result, this gives rise to a long exact sequence of the homology groups of those complexes:

$$\dots \longrightarrow B_n \otimes G \xrightarrow{\Phi} Z_n \otimes G \longrightarrow H_n(C_* \otimes G) \longrightarrow B_{n-1} \otimes G \xrightarrow{\Phi} Z_{n-1} \otimes G \longrightarrow \dots,$$

where Φ denotes the connecting homomorphisms in this long exact sequence. If you look closely at the diagram chase required for constructing Φ , you'll find that it can be described explicitly: the map $\Phi : B_n \otimes G \rightarrow Z_n \otimes G$ for each $n \in \mathbb{Z}$ is just $i_n \otimes \mathbb{1}$, where $i_n : B_n \hookrightarrow Z_n$ is again the obvious inclusion. Now we use the standard trick for turning a long exact sequence into a short exact sequence centered on a certain term, in this case $H_n(C_* \otimes G)$: the map from this to $B_{n-1} \otimes G$ is

surjective onto the kernel of Φ , while the map from $Z_n \otimes G$ preceding this descends to an injection on the quotient of $Z_n \otimes G$ by $\text{im } \Phi$, giving a short exact sequence

$$0 \longrightarrow \text{coker}(i_n \otimes \mathbf{1}) \longrightarrow H_n(C_* \otimes G) \longrightarrow \ker(i_{n-1} \otimes \mathbf{1}) \longrightarrow 0.$$

Now plug in the isomorphisms (42.1) and (42.2) and our short exact sequence becomes the one in the statement of the theorem.

It is a bit tedious but straightforward to establish the naturality of this sequence, i.e. that any chain map from C_* to another chain complex C'_* of free abelian groups produces a commutative diagram with the exact sequences beginning with $H_n(C_*) \otimes G$ and $H_n(C'_*) \otimes G$ as its two rows. I will leave this as an exercise.

But we should still discuss why the sequence splits, as this is less straightforward. By Exercise 28.5, it will suffice to construct a left-inverse for the injective map $h : H_n(C_*) \otimes G \rightarrow H_n(C_* \otimes G)$. Recall first that the exact sequence $0 \rightarrow Z_n \hookrightarrow C_n \xrightarrow{\partial_n} B_{n-1} \rightarrow 0$ splits since B_{n-1} is free due to Proposition 42.1, thus there exists a left-inverse of the inclusion $Z_n \hookrightarrow C_n$, i.e. a projection homomorphism

$$p : C_n \rightarrow Z_n, \quad p|_{Z_n} = \mathbf{1}_{Z_n}.$$

The composition of p with the quotient projection $\text{pr} : Z_n \rightarrow Z_n/B_n = H_n(C_*)$ then defines a chain map

$$(C_*, \partial) \xrightarrow{\text{pr} \circ p} (H_*(C_*), 0)$$

since $\text{pr} \circ p \circ \partial = 0$, and this induces a chain map

$$(C_* \otimes G, \partial \otimes \mathbf{1}) \xrightarrow{(\text{pr} \circ p) \otimes \mathbf{1}} (H_*(C_*) \otimes G, 0),$$

which then descends to a homomorphism on the homologies that we shall call

$$\pi := ((\text{pr} \circ p) \otimes \mathbf{1})_* : H_n(C_* \otimes G) \rightarrow H_n(C_*) \otimes G$$

for each $n \in \mathbb{Z}$. One can now check that $\pi \circ h$ is the identity map on $H_n(C_*) \otimes G$. \square

REMARK 42.5. It should be clear from the construction of the left-inverse $\pi : H_n(C_* \otimes G) \rightarrow H_n(C_*) \otimes G$ that it has little chance of interacting nicely with morphisms of chain complexes, as it depended rather seriously on the completely arbitrary choice of a left-inverse $p : C_n \rightarrow Z_n$ for the inclusion $Z_n \hookrightarrow C_n$. The latter can be derived from a right-inverse of $\partial_n : C_n \rightarrow B_{n-1}$, which exists because Proposition 42.1 guarantees that B_{n-1} has a basis, and we can then choose an element in $\partial_n^{-1}(b) \subset C_n$ for every basis element $b \in B_{n-1}$, but this step requires in general an explicit invocation of the axiom of choice, and we already used this axiom in a less obvious way by quoting Proposition 42.1 (cf. Remark 42.2). This discussion does not prove that the splitting is not natural, but it should take away any reason you may have had to believe that it is. For actual counterexamples to naturality, see [Hat02, pp. 264–265] or [Bre93, Chapter V, Example 7.7].

EXERCISE 42.6. Prove that for any space X with finitely-many path-components⁵⁹ and any abelian group G , $H_1(X; G) \cong H_1(X) \otimes G$. *Hint:* $H_0(X; \mathbb{Z})$ is always free.

Our next major topic is computing the homology of product spaces. We will tackle the algebraic aspect of this problem in the present lecture, since it is very similar to the universal coefficient theorem—in fact, the result we need to prove generalizes Theorem 42.4 to a statement about the homology of the tensor product of two chain complexes.

⁵⁹The assumption of finitely-many path-components is not necessary in general; we need it only because we have not proved anything about $\text{Tor}(A, G)$ when A is not finitely generated, though one could (see [Hat02, Prop. 3A.5]).

To see why we might want such a result, 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

$$\Phi_\alpha \times \Phi_\beta : I^{k+\ell} \rightarrow X \times Y : (s, t) \mapsto (\Phi_\alpha(s), \Phi_\beta(t)).$$

The bilinear operation

$$C_k^{\text{CW}}(X) \times C_\ell^{\text{CW}}(Y) \xrightarrow{\times} C_{k+\ell}^{\text{CW}}(X \times Y)$$

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 **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 $y \in \partial N$; one then has to think somewhat more carefully about orientations to get the signs right (see Problem Set 9).⁶⁰

PROPOSITION 42.7. *For any pair of CW-complexes X and Y with a k -cell $e_\alpha^k \subset X$ and an ℓ -cell $e_\beta^\ell \subset Y$,*

$$\partial(e_\alpha^k \times e_\beta^\ell) = \partial e_\alpha^k \times e_\beta^\ell + (-1)^k e_\alpha^k \times \partial e_\beta^\ell \in C_{k+\ell-1}^{\text{CW}}(X \times Y).$$

□

The following purely algebraic definition should now hopefully seem quite natural.

DEFINITION 42.8. Given chain complexes (A_*, ∂^A) and (B_*, ∂^B) , the **tensor product chain complex** $(A_* \otimes B_*, \partial)$ is defined by

$$(42.3) \quad (A_* \otimes B_*)_n = \bigoplus_{k+\ell=n} A_k \otimes B_\ell,$$

where the direct sum is understood to be over the set of all pairs of integers k, ℓ that add up to n , and the boundary map is determined by the formula

$$\partial(a \otimes b) = \partial^A a \otimes b + (-1)^k a \otimes \partial^B b \quad \text{for } a \in A_k, b \in B_\ell.$$

You should take a moment to assure yourself that this really defines a chain complex: ∂^2 includes some terms that vanish because $(\partial^A)^2 = (\partial^B)^2 = 0$, but also cross terms $\partial^A a \otimes \partial^B b$ that disappear due to sign cancellations. Here is another easy thing to check: given chain maps $f : A_* \rightarrow A'_*$ and $g : B_* \rightarrow B'_*$, there is a chain map

$$(42.4) \quad f \otimes g : A_* \otimes B_* \rightarrow A'_* \otimes B'_* : a \otimes b \mapsto f(a) \otimes g(b).$$

We can now rephrase Proposition 42.7 as follows:

PROPOSITION 42.9. *The cellular cross product determines an isomorphism of chain complexes*

$$C_*^{\text{CW}}(X) \otimes C_*^{\text{CW}}(Y) \rightarrow C_*^{\text{CW}}(X \times Y) : a \otimes b \rightarrow a \times b.$$

□

⁶⁰An explicit proof of the formula in Prop. 42.7 can also be found in [Hat02, Prop. 3B.1].

REMARK 42.10. The sign changes in all these formulas take some getting used to, but there is an easy rule of thumb for remembering them. Every element $a \in A_k$ in a chain complex A_* can be regarded as *even* or *odd* depending on whether k is even or odd, and we can also naturally regard boundary maps $\partial : A_* \rightarrow A_*$ as having degree -1 since they send A_k to A_{k-1} . If we view ∂ is an *odd* object in this sense, then the rule is 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 arise for reasons having to do with 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

$$(42.5) \quad \bigoplus_{k+\ell=n} H_k(A_*) \otimes H_\ell(B_*) \rightarrow H_n(A_* \otimes B_*)$$

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 42.11 (algebraic Künneth formula). *Given two chain complexes C_*, C'_* of free abelian groups, the map (42.5) 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).

The statement becomes a bit more concise if we define the operation \otimes on the category of \mathbb{Z} -graded abelian groups via (42.3) and define Tor as a functor $\text{Ab}_{\mathbb{Z}} \times \text{Ab}_{\mathbb{Z}} \rightarrow \text{Ab}_{\mathbb{Z}}$ by

$$(\text{Tor}(C_*, C'_*))_n := \bigoplus_{k+\ell=n} \text{Tor}(C_k, C'_\ell).$$

Then the exact sequence in Theorem 42.11 becomes

$$0 \rightarrow H_*(C_*) \otimes H_*(C'_*) \rightarrow H_*(C_* \otimes C'_*) \rightarrow (\text{Tor}(H_*(C_*), H_*(C'_*)))_{*-1} \rightarrow 0,$$

where the subscript “ $* - 1$ ” on the last term indicates the downward degree shift. The splitting gives rise to an 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 in the next lecture how this can be applied to singular homology more directly, without needing to know that singular and cellular homologies are isomorphic.

The proof of the Künneth formula requires the following easy exercise, which also provides a useful criterion (beyond Theorem 41.2) for the vanishing of Tor groups.

EXERCISE 42.12. Show that if C_* is any chain complex and G is a free abelian group, the canonical map $h : H_*(C_*) \otimes G \rightarrow H_*(C_* \otimes G)$ is an isomorphism.

Hint: Proving $\text{Tor}(H_n(C_), G) = 0$ does not suffice, as the chain groups in C_* are not assumed to be free, so you cannot just feed this into the universal coefficient theorem. But if G is free, you can write it as a direct sum of copies of \mathbb{Z} , so try rewriting the chain complex $C_* \otimes G$ similarly.*

PROOF OF THEOREM 42.11. As in the proof of the universal coefficient theorem, we abbreviate the subgroups 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 C_* (and therefore also its subgroup $B_* \subset C_*$) 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 now a short exact sequence of tensor product chain complexes. Note that the individual chain groups of $B_{*-1} \otimes C'_*$ are

$$(B_{*-1} \otimes C'_*)_n = \bigoplus_{k+l=n-1} B_k \otimes C'_l,$$

thus $H_n(B_{*-1} \otimes C'_*) = H_{n-1}(B_* \otimes C'_*)$. The short exact sequence of chain complexes gives rise as usual to a long exact sequence of homology groups

$$\dots \rightarrow H_n(B_* \otimes C'_*) \xrightarrow{i_n} H_n(Z_* \otimes C'_*) \rightarrow H_n(C_* \otimes C'_*) \rightarrow H_{n-1}(B_* \otimes C'_*) \xrightarrow{i_{n-1}} H_{n-1}(Z_* \otimes C'_*) \rightarrow \dots,$$

where the maps labeled i_n, i_{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:

$$(42.6) \quad 0 \rightarrow \text{coker } i_n \rightarrow H_n(C_* \otimes C'_*) \rightarrow \ker i_{n-1} \rightarrow 0,$$

Since B_* is free, Exercise 42.12 gives a canonical isomorphism $H_*(B_* \otimes C'_*) = B_* \otimes H_*(C'_*)$, which more precisely means

$$H_n(B_* \otimes C'_*) = \bigoplus_{k+l=n} B_k \otimes H_l(C'_*) \quad \text{for all } n \in \mathbb{Z},$$

and for the same reasons,

$$H_n(Z_* \otimes C'_*) = \bigoplus_{k+l=n} Z_k \otimes H_l(C'_*) \quad \text{for all } n \in \mathbb{Z}.$$

Inspecting the diagram-chase then reveals that the maps $i_n : H_n(B_* \otimes C'_*) \rightarrow H_n(Z_* \otimes C'_*)$ are exactly what one would expect, namely

$$i_n = \bigoplus_{k+l=n} i_{k,\ell} : \bigoplus_{k+l=n} B_k \otimes H_\ell(C'_*) \rightarrow \bigoplus_{k+l=n} Z_k \otimes H_\ell(C'_*),$$

where each of the maps $i_{k,\ell} : B_k \otimes H_\ell(C'_*) \rightarrow Z_k \otimes H_\ell(C'_*)$ is the tensor product of the inclusion $B_k \hookrightarrow Z_k$ with the identity map $H_\ell(C'_*) \rightarrow H_\ell(C'_*)$, so the exact sequence (42.6) can now be written as

$$(42.7) \quad 0 \rightarrow \bigoplus_{k+l=n} \text{coker } i_{k,\ell} \rightarrow H_n(C_* \otimes C'_*) \rightarrow \bigoplus_{k+l=n-1} \ker i_{k,\ell} \rightarrow 0.$$

Finally, to understand the terms $\text{coker } i_{k,\ell}$ and $\ker i_{k,\ell}$ we can look at the short exact sequence

$$0 \rightarrow B_k \hookrightarrow Z_k \xrightarrow{\text{pr}} H_k(C_*) \rightarrow 0,$$

which defines a projective resolution of $H_k(C_*)$ and thus gives rise for each $\ell \in \mathbb{Z}$ to a chain complex

$$0 \longrightarrow B_k \otimes H_\ell(C'_*) \xrightarrow{i_{k,\ell}} Z_k \otimes H_\ell(C'_*) \longrightarrow H_k(C_*) \otimes H_\ell(C'_*) \longrightarrow 0.$$

This sequence is exact except possibly at the first nontrivial term, where the non-exactness is measured by

$$\ker i_{k,\ell} \cong \operatorname{Tor}(H_k(C_*), H_\ell(C'_*)),$$

and the exactness at all other terms implies that the map $Z_k \otimes H_\ell(C'_*) \rightarrow H_k(C_*) \otimes H_\ell(C'_*)$ descends to an isomorphism

$$\operatorname{coker} i_{k,\ell} \cong H_k(C_*) \otimes H_\ell(C'_*).$$

The sequence we were looking for is now obtained by plugging these isomorphisms into (42.7).

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

There is a natural generalization of the Tor functor defined on the category Mod^R of R -modules over any commutative ring R with unit

$$\operatorname{Tor}^R : \operatorname{Mod}^R \times \operatorname{Mod}^R \rightarrow \operatorname{Mod}^R,$$

such that $\operatorname{Tor}^{\mathbb{Z}} = \operatorname{Tor}$. Recall that the **tensor product of R -modules** is defined analogously to the tensor product of abelian groups, but with the additional relation

$$(42.8) \quad a \otimes (rb) = r(a \otimes b) = (ra) \otimes b$$

for any $r \in R$; concretely, for two R -modules A and B we define

$$A \otimes_R B := F_R^{\operatorname{mod}}(A \times B) / N$$

where $F_R^{\operatorname{mod}}(A \times B)$ is the free R -module on the set $A \times B$, and $N \subset F_R^{\operatorname{mod}}(A \times B)$ 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 $(a, b) \in F_R^{\operatorname{mod}}(a \times b)$ in the quotient by $a \otimes b \in A \otimes_R B$. It should be clear that if $R = \mathbb{Z}$, then this definition reproduces the usual tensor product of abelian groups $A \otimes B$, but in general, $A \otimes_R B$ and $A \otimes B$ are typically different sets. The R -module $\operatorname{Tor}^R(A, G)$ is now defined by applying the R -module tensor product functor $\otimes_R G$ to a projective resolution (F_*, f_*) of A and then taking the first homology group:

$$\operatorname{Tor}^R(A, G) = H_1(F_* \otimes_R G).$$

It is straightforward to show that this satisfies the natural generalizations of both Theorem 41.2 and Exercise 42.12; see Problem Set 9 #1. The only detail on which the discussion differs is that in general, the higher homology groups $H_n(F_* \otimes_R G)$ for $n \geq 2$ can also be nonzero, leading to a whole sequence of so-called *derived* functors $\operatorname{Tor}_n^R(A, G)$ with $\operatorname{Tor}^R(A, G) = \operatorname{Tor}_1^R(A, G)$. This ceases to be true however if we add an algebraic assumption about R :

PROPOSITION 42.13 (see [Lan02, §III.7]). *If R is a principal ideal domain, then every submodule of a free R -module is also free.* \square

There is no need at this point to look up what a principal ideal domain is if you've forgotten: the only cases that we actually intend to consider are when $R = \mathbb{Z}$ (so the result is just a restatement of Proposition 42.1) or when R is a field \mathbb{K} . In the latter case, R -modules are the same thing as vector spaces over \mathbb{K} , so Proposition 42.13 follows immediately from the fact that *every* vector

space is a free module over its field of scalars, i.e. it admits a basis. (One should not overstate how trivial this fact is—in the infinite-dimensional case, it also depends on Zorn’s lemma.)⁶¹

The fact that subgroups of free abelian groups are free was the main ingredient of our proof of both the universal coefficient theorem and the Künneth formula, so in light of Proposition 42.13, both results extend to statements about chain complexes of free modules over a principal ideal domain R . In particular, we obtain natural split exact sequences (with non-natural splittings)

$$0 \longrightarrow H_n(C_*) \otimes_R G \xrightarrow{h} H_n(C_* \otimes_R G) \longrightarrow \text{Tor}^R(H_{n-1}(C_*), G) \longrightarrow 0,$$

and

$$0 \rightarrow \bigoplus_{k+\ell=n} H_k(C_*) \otimes_R H_\ell(C'_*) \rightarrow H_n(C_* \otimes_R C'_*) \rightarrow \bigoplus_{k+\ell=n-1} \text{Tor}^R(H_k(C_*), H_\ell(C'_*)) \rightarrow 0$$

whenever C_* and C'_* are chain complexes of free R -modules and G is another R -module.

While there is no obvious advantage to generalizing the universal coefficient theorem in this way, we will see when we discuss its analogue for cohomology that it is very useful to have the freedom to choose $R \neq \mathbb{Z}$. Notice that if R is chosen to be a field \mathbb{K} , then since every vector space over \mathbb{K} is a free \mathbb{K} -module, the natural generalization of Theorem 41.2 implies

$$\text{Tor}^{\mathbb{K}}(A, B) = 0$$

for every pair of \mathbb{K} -vector spaces A and B . This puts the Künneth formula in an especially simple form, thus proving the following variant of Theorem 42.11:

COROLLARY 42.14. *For any field \mathbb{K} and any two chain complexes C_* and C'_* of \mathbb{K} -vector spaces, the canonical map*

$$\bigoplus_{k+\ell=n} H_k(C_*) \otimes_{\mathbb{K}} H_\ell(C'_*) \rightarrow H_n(C_* \otimes_{\mathbb{K}} C'_*)$$

is a \mathbb{K} -linear isomorphism for every $n \in \mathbb{Z}$. □

This result is one of the reasons why it is often easier to compute homology with field coefficients than over the integers.

43. The cross product on homology

We saw in the previous lecture that if X and Y are CW-complexes and we assign the product cell decomposition to $X \times Y$, there is an obvious chain map

$$(43.1) \quad C_*^{\text{CW}}(X) \otimes C_*^{\text{CW}}(Y) \rightarrow C_*^{\text{CW}}(X \times Y) : a \otimes b \mapsto a \times b,$$

which sends the generators $e_\alpha^k \otimes e_\beta^\ell$ to the product $(k + \ell)$ -cell $e_\alpha^k \times e_\beta^\ell$ in $X \times Y$. Letting this map descend to homology and then composing it with the canonical map $H_*^{\text{CW}}(X) \otimes H_*^{\text{CW}}(Y) \rightarrow H_*(C_*^{\text{CW}}(X) \otimes C_*^{\text{CW}}(Y))$ that is defined for any tensor product of chain complexes, we obtain the so-called **cross product** on cellular homology, which is a homomorphism

$$H_k^{\text{CW}}(X) \otimes H_\ell^{\text{CW}}(Y) \xrightarrow{\times} H_{k+\ell}^{\text{CW}}(X \times Y)$$

defined for every pair $k, \ell \in \mathbb{Z}$, or equivalently, a bilinear map $\times : H_k^{\text{CW}}(X) \otimes H_\ell^{\text{CW}}(Y) \rightarrow H_{k+\ell}^{\text{CW}}(X \times Y)$. One crucial feature of this discussion is that the map (43.1) is actually an *isomorphism* of

⁶¹Note also that when we talk algebraically about a basis of a vector space, we do not mean the same thing that is meant when discussing bases of separable Hilbert spaces in functional analysis: the latter is a countable set B such that every element x in the space can be written uniquely as a *convergent* sum $\sum_{b \in B} x_b b$, where it is possible for infinitely many of the coefficients x_b to be nonzero if they decay fast enough. In algebra, our vector spaces have no topologies and thus no notion of convergence, so for B to be a basis means that every x can be written uniquely as a linear combination $\sum_{b \in B} x_b b$ that converges for the trivial reason that at most finitely many of the coefficients are nonzero.

chain complexes—it is determined by a bijective correspondence between the generators of the two complexes—thus it induces an isomorphism between $H_*^{CW}(X \times Y)$ and the homology of the tensor product chain complex. We can then feed this into the algebraic Künneth formula (Theorem 42.11) to produce the cellular version of what is known as the **topological Künneth formula**, a split exact sequence

$$0 \rightarrow \bigoplus_{k+\ell=n} H_k^{CW}(X) \otimes H_\ell^{CW}(Y) \xrightarrow{\times} H_{k+\ell}^{CW}(X \times Y) \rightarrow \bigoplus_{k+\ell=n-1} \text{Tor}(H_k^{CW}(X), H_\ell^{CW}(Y)) \rightarrow 0.$$

Our main goal in this lecture will be to drop the assumption that X and Y are CW-complexes and generalize this whole discussion to singular homology.

Let us first mention how the above discussion can be generalized beyond integer coefficients. There is no obvious way to define a chain map as in (43.1) on the cellular chain complex with an arbitrary coefficient group G ,⁶² but if we take G to be a commutative ring R with unit, then $C_*^{CW}(X; R)$ can be interpreted as the free R -module generated by the cells of the CW-complex X , so (43.1) uniquely determines a chain map of R -modules

$$C_*^{CW}(X; R) \otimes_R C_*^{CW}(Y; R) \xrightarrow{\times} C_*^{CW}(X \times Y; R).$$

The induced R -module homomorphism on homology then gives rise to an R -bilinear cross product

$$H_k^{CW}(X; R) \otimes_R H_\ell^{CW}(Y; R) \xrightarrow{\times} H_{k+\ell}^{CW}(X \times Y; R).$$

If R is additionally a principal ideal domain, then the Künneth formula also holds, producing an exact sequence of R -modules

$$\begin{aligned} 0 \longrightarrow \bigoplus_{k+\ell=n} H_k^{CW}(X; R) \otimes_R H_\ell^{CW}(Y; R) &\xrightarrow{\times} H_{k+\ell}^{CW}(X \times Y; R) \\ &\longrightarrow \bigoplus_{k+\ell=n-1} \text{Tor}^R(H_k^{CW}(X; R), H_\ell^{CW}(Y; R)) \longrightarrow 0, \end{aligned}$$

with the pleasing feature that the Tor term vanishes whenever R is taken to be a field.

REMARK 43.1. There is an annoying point that we’ve been glossing over throughout 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 Problem Set 9 #5). 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 $H_*(X \times Y; G) \cong H_*^{CW}(X \times Y; G)$ holds as usual. With this in mind, we shall assume from now on that $X \times Y$ carries the product topology.

Since cellular and singular homologies are isomorphic, the cellular cross product determines a homomorphism

$$(43.2) \quad H_k(X; R) \otimes_R H_\ell(Y; R) \xrightarrow{\times} H_{k+\ell}(X \times Y; R)$$

whenever X and Y come with cell decompositions, but it is far from obvious at this stage whether \times by this definition is independent of the choice of cell decomposition. We shall deal with this as

⁶²Though, for what it’s worth, one *can* define a product of the form $C_*^{CW}(X; G_1) \otimes C_*^{CW}(Y; G_2) \rightarrow C_*^{CW}(X \times Y; G_1 \otimes G_2)$ for any pair of abelian groups G_1 and G_2 . The version with ring coefficients is essentially a composition of this with the ring multiplication map $R \otimes R \rightarrow R : r \otimes s \mapsto rs$.

usual by replacing the cellular cross product with a more general 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 (43.2) should exist. If you like to think about elements of $H_k(X)$ as represented by closed oriented k -dimensional submanifolds $M \subset X$ as in Lecture 29, then since the product of two closed oriented manifolds is also a closed oriented manifold, it would make sense to define

$$[M] \times [N] := [M \times N] \in H_{k+\ell}(X \times Y)$$

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 (43.2) be defined in general?

Since there is always a canonical map $H_*(X; R) \otimes_R H_*(Y; R) \rightarrow H_*(C_*(X; R) \otimes_R C_*(Y; R))$, we would obtain a map (43.2) if we had a chain map

$$C_*(X) \otimes C_*(Y) \xrightarrow{\times} C_*(X \times Y)$$

to play the role of (43.1) on the singular chain complex. We then need to decide what $\sigma \times \tau \in C_{k+\ell}(X \times Y)$ 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, so we cannot simply write

$$(43.3) \quad \sigma \times \tau : \Delta^k \times \Delta^\ell \rightarrow X \times Y : (s, t) \mapsto (\sigma(s), \tau(t))$$

and call this a generator of $C_{k+\ell}(X \times Y)$. But we've dealt with this kind of thing before using subdivision: a natural approach would be to fix a reasonable oriented triangulation of $\Delta^{k+\ell}$ for every pair of integers $k, \ell \geq 0$, and define $\sigma \times \tau \in C_{k+\ell}(X \times Y)$ as the sum of singular simplices obtained by restricting the map (43.3) to the $(k + \ell)$ -simplices in this triangulation. This is the correct geometric picture, but in practice, it is cumbersome to have to choose a triangulation and then show that what we've defined is independent of the choice. We shall instead employ an algebraic trick that accomplishes the same result: the trick is known as the method of **acyclic models**. One can find in various textbooks (e.g. [Vic94, Spa95]) a result called the *acyclic model theorem*, which is applicable to a wide variety of problems but difficult to digest, as it is typically expressed in heavily abstract category-theoretic language. We shall instead follow the approach of [Bre93] and demonstrate the method by example.

Before we continue, we need to recall something from the end of Lecture 28: the reduced singular homology $\tilde{H}_*(X; R)$ of a space X with coefficients in a ring R is also the homology of the so-called *augmented* chain complex $\tilde{C}_*(X; R)$, defined by appending an extra nonzero term $\tilde{C}_{-1}(X; R) := R$ to the end of the usual singular chain complex:

$$\dots \longrightarrow C_2(X; R) \xrightarrow{\partial_2} C_1(X; R) \xrightarrow{\partial_1} C_0(X; R) \xrightarrow{\epsilon} R \longrightarrow 0 \longrightarrow \dots$$

The **augmentation** is the homomorphism $\epsilon : C_0(X; R) \rightarrow R$, which is defined on each of the generators $\sigma : \Delta^0 \rightarrow X$ of $C_0(X)$ by $\epsilon(\sigma) = 1$ (see Exercise 28.19), and is thus surjective with $\epsilon \circ \partial_1 = 0$.

LEMMA 43.2. For any commutative ring R with unit, one can associate to any two spaces X and Y a chain map of R -modules

$$\Phi : C_*(X; R) \otimes_R C_*(Y; R) \rightarrow C_*(X \times Y; R)$$

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

$$\begin{array}{ccc} C_*(X; R) \otimes_R C_*(Y; R) & \xrightarrow{\Phi} & C_*(X \times Y; R) \\ \downarrow f_* \otimes g_* & & \downarrow (f \times g)_* \\ C_*(X'; R) \otimes C_*(Y'; R) & \xrightarrow{\Phi} & C_*(X' \times Y'; R) \end{array}$$

commutes. Moreover, Φ with these properties is unique up to chain homotopy.

PROOF. The choice of the ring R will make no difference at all in the proof, so we shall drop it from the notation; just keep in mind throughout the following that all homomorphisms are R -linear and all tensor products are over R . 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

$$(43.4) \quad \partial \Phi(i_k \otimes i_\ell) = \Phi(\partial(i_k \otimes i_\ell)) \in C_{n-1}(\Delta^k \times \Delta^\ell),$$

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

$$(43.5) \quad \partial \Phi \partial(i_k \otimes i_\ell) = \Phi \partial^2(i_k \otimes i_\ell) = 0,$$

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 $\epsilon : C_0(\Delta^k \times \Delta^\ell) \rightarrow R$ 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,

$$(43.6) \quad \epsilon \Phi \partial(i_k \otimes i_\ell) = 0 \quad \text{when} \quad n = 1.$$

Now comes the crucial point: $\Delta^k \times \Delta^\ell$ is contractible, so its reduced singular homology is trivial. In light of (43.5) and (43.6), 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 (43.4) has solutions, and we can define $\Phi(i_k \otimes i_\ell) \in C_n(\Delta^k \times \Delta^\ell)$ to be any element such that

$$(43.7) \quad \Phi(i_k \otimes i_\ell) \in \partial^{-1}(\Phi\partial(i_k \otimes i_\ell)).$$

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 (42.4)). 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 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 $\Phi : C_*(X; R) \otimes_R C_*(Y; R) \rightarrow C_*(X \times Y; R)$ from Lemma 43.2 induces a map on homology, which can then be composed with the canonical map from $H_*(X; R) \otimes_R H_*(Y; R)$ to define what we will henceforth call the **singular cross product**

$$H_*(X; R) \otimes_R H_*(Y; R) \longrightarrow H_*(C_*(X; R) \otimes_R C_*(Y; R)) \xrightarrow{\Phi_*} H_*(X \times Y; R)$$

The uniqueness of Φ up to chain homotopy implies that this definition does not depend on any choices, and naturality gives rise to a commutative diagram

$$\begin{array}{ccc} H_*(X; R) \otimes_R H_*(Y; R) & \xrightarrow{\times} & H_*(X \times Y; R) \\ \downarrow f_* \otimes g_* & & \downarrow (f \times g)_* \\ H_*(X'; R) \otimes_R H_*(Y'; R) & \xrightarrow{\times} & H_*(X' \times Y'; R) \end{array}$$

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 $\Phi : C_*(X; R) \otimes_R C_*(Y; R) \rightarrow C_*(X \times Y; R)$ was not only a chain map, it was 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. But 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 43.3. *In the setting of Lemma 43.2, if R is a principal ideal domain, then there also exist natural chain maps*

$$\begin{aligned} C_*(X \times Y; R) &\xrightarrow{\theta} C_*(X; R) \otimes_R C_*(Y; R), \\ C_*(X; R) \otimes_R C_*(Y; R) &\xrightarrow{\alpha} C_*(X; R) \otimes_R C_*(Y; R), \\ C_*(X \times Y; R) &\xrightarrow{\beta} C_*(X \times Y; R), \end{aligned}$$

which are uniquely determined up to chain homotopy by 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

$$\begin{array}{ccc} C_*(X \times Y; R) & \xrightarrow{\theta} & C_*(X; R) \otimes_R C_*(Y; R) \\ \downarrow (f \times g)_* & & \downarrow f_* \otimes g_* \\ C_*(X' \times Y'; R) & \xrightarrow{\theta} & C_*(X'; R) \otimes_R C_*(Y'; R) \end{array}$$

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 43.4 (Eilenberg-Zilber theorem). *The natural chain maps $\Phi : C_*(X; R) \otimes_R C_*(Y; R) \rightarrow C_*(X \times Y; R)$ and $\theta : C_*(X \times Y; R) \rightarrow C_*(X; R) \otimes_R C_*(Y; R)$ are both chain homotopy equivalences, and the induced maps on homology are inverse to each other.* \square

PROOF OF LEMMA 43.3. As before, we shall mostly omit R from the notation since the role it plays is minimal. The definition of the desired chain map on 0-chains is given in each case and clearly satisfies naturality, 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 and R is a principal ideal domain, the algebraic Künneth formula implies

$$H_m(C_*(\Delta^n; R) \otimes_R C_*(\Delta^n; R)) \cong \bigoplus_{k+\ell=m} H_k(\Delta^n; R) \otimes_R H_\ell(\Delta^n; R) \cong \begin{cases} R & \text{if } m = 0, \\ 0 & \text{otherwise,} \end{cases}$$

where all the Tor terms have vanished because every $H_k(\Delta^n; R)$ is a free R -module. 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

$$(43.8) \quad \theta(d_n) \in \partial^{-1}(\theta\partial d_n).$$

The case $n = 1$ is special since $H_0(\Delta^n; R) \otimes_R H_0(\Delta^n; R) = R$ 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.⁶³ In either case $\theta(d_n)$ can be defined so that (43.8) 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

$$\theta(\sigma) = \theta((\pi_X \circ \sigma) \times (\pi_Y \circ \sigma))_* d_n := ((\pi_X \circ \sigma)_* \otimes (\pi_Y \circ \sigma)_*) \theta(d_n).$$

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

⁶³Equivalently, 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.

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 43.5. 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 43.3. It's relaxing.

Corollary 43.4 implies that the natural map

$$\Phi_* : H_*(C_*(X; R) \otimes_R C_*(Y; R)) \rightarrow H_*(X \times Y; R)$$

used in the definition of the singular cross product is an isomorphism whenever R is a principal ideal domain, so we can now use it to replace the middle term in the algebraic Künneth formula, proving:

COROLLARY 43.6 (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*

$$0 \longrightarrow \bigoplus_{k+\ell=n} H_k(X; R) \otimes_R H_\ell(Y; R) \xrightarrow{\times} H_{k+\ell}(X \times Y; R) \longrightarrow \bigoplus_{k+\ell=n-1} \text{Tor}(H_k(X; R), H_\ell(Y; R)) \longrightarrow 0,$$

and the sequence splits (but not naturally). \square

In particular, we can always choose field coefficients to make the Tor terms vanish:

COROLLARY 43.7. *For any spaces X and Y and any field \mathbb{K} , the cross product on singular homology with coefficients in \mathbb{K} defines natural \mathbb{K} -vector space isomorphisms*

$$\times : \bigoplus_{k+\ell=n} H_k(X; \mathbb{K}) \otimes_{\mathbb{K}} H_\ell(Y; \mathbb{K}) \xrightarrow{\cong} H_n(X \times Y; \mathbb{K}).$$

for every integer $n \geq 0$. \square

EXERCISE 43.8. The goal of this exercise is to prove the associativity of the cross product:

$$(A \times B) \times C = A \times (B \times C) \in H_*(X \times Y \times Z; R)$$

for all $A \in H_*(X; R)$, $B \in H_*(Y; R)$ and $C \in H_*(Z; R)$. Here R may be any commutative ring with unit.

- (a) Use acyclic models to prove that for triples of spaces X, Y, Z , all natural chain maps

$$\Phi : C_*(X; R) \otimes_R C_*(Y; R) \otimes_R C_*(Z; R) \rightarrow C_*(X \times Y \times Z; R)$$

that act on 0-chains by $\Phi(x \otimes y \otimes z) = (x, y, z)$ are chain homotopic.

Remark: The statement implicitly assumes that there is a well-defined notion of the tensor product of three chain complexes, which of course is true since there is a canonical chain isomorphism between $(C_(X; R) \otimes_R C_*(Y; R)) \otimes_R C_*(Z; R)$ and $C_*(X; R) \otimes_R (C_*(Y; R) \otimes_R C_*(Z; R))$. Right?*

- (b) Given $A \in H_*(X; R)$, $B \in H_*(Y; R)$ and $C \in H_*(Z; R)$, show that the products $(A \times B) \times C$ and $A \times (B \times C) \in H_*(X \times Y \times Z; R)$ can each be expressed via natural chain maps as in part (a), and conclude that they are identical.

EXERCISE 43.9. Fix a commutative ring R with unit and recall that for any path-connected space Y , there is a canonical isomorphism $H_0(Y; R) = R$. Let $[\text{pt}] \in H_0(Y; R)$ denote the homology class corresponding to the unit $1 \in R$ under this identification, so it is a canonical generator of the R -module $H_0(Y; R)$ which can also be defined by $[\sigma \otimes 1] \in H_0(C_*(Y) \otimes R)$ for any singular

0-simplex $\sigma : \Delta^0 \rightarrow Y$. Show that for any space X , the cross product of any $A \in H_n(X; R)$ with $[\text{pt}] \in H_0(Y; R)$ is given by

$$A \times [\text{pt}] = i_* A$$

for any inclusion map of the form $i : X \hookrightarrow X \times Y : x \mapsto (x, \text{const})$. A similar formula holds for cross products with $[\text{pt}] \in H_0(X; R)$ if X is path-connected. In particular, this means that the unit in $H_0(\{\text{pt}\}; R) = R$ acts as a multiplicative identity element with respect to the cross product, under the obvious identifications $X \times \{\text{pt}\} = X = \{\text{pt}\} \times X$.

Hint: Remember that \times is induced by a natural chain map $\Phi : C_(X; R) \otimes_R C_*(Y; R) \rightarrow C_*(X \times Y; R)$, so if you have the right formula for $\Phi : C_n(X; R) \otimes_R C_0(Y; R) \rightarrow C_n(X \times Y; R)$, the relation will become obvious. In general, one can make many choices in defining Φ , but there is an obvious choice that one “should” make when one of the chains is 0-dimensional. Review the construction of Φ via acyclic models to show that this choice is always possible.*

The alert reader may notice that there is at least one important question we have not addressed yet: if X and Y are CW-complexes, are the singular and cellular cross products the same? The answer is of course yes, but we will not discuss it at length, since we don’t plan to carry out any serious applications of the cellular cross product—it is useful to have in mind for intuition and motivation, but the product on singular homology will play a much more important role in further developments. One other (and closely related) question we have not addressed is how to define the cross product on *relative* singular homology. We will come back to this when we introduce the cohomology cup product.

44. Čech homology and inverse limits

By now I have said almost everything I can reasonably say about singular homology without bringing cohomology into the picture, so that will be the next major topic. But before that, I’d like to address a question that may have been nagging at you since we introduced the Eilenberg-Steenrod axioms many weeks ago: what other axiomatic homology theories are there?

Let me name a few theories that are not examples: first, cellular and simplicial homology are not axiomatic homology theories in the sense of Eilenberg-Steenrod [ES52], as they are not functors defined on Top_{rel} or any subcategory of Top_{rel} . Both require auxiliary choices beyond a pair of spaces (X, A) for their definitions, e.g. one cannot define the group $H_*^{\text{CW}}(X, A)$ without having chosen a cell decomposition for (X, A) , thus H_*^{CW} is a functor $\text{CW}_{\text{rel}} \rightarrow \text{Ab}_{\mathbb{Z}}$, and the fact that $H_*^{\text{CW}}(X, A)$ up to isomorphism depends only on the topology of (X, A) is a deep theorem, but not an intrinsic feature of the definition. You may have heard of other theories that produce topological invariants in spite of auxiliary choices: for example, *Morse homology* (see [AD14] or [Sch93]) associates to any triple (X, f, g) consisting of a closed oriented smooth manifold X with a generic function $f : X \rightarrow \mathbb{R}$ and Riemannian metric g a graded abelian group $H_*^{\text{Morse}}(X, f, g)$ that turns out (by a different deep theorem that has nothing to do with the axioms) to be isomorphic to $H_*(X)$. There is also a smooth variant of singular homology defined when X is a smooth manifold, where all singular simplices are required to be smooth maps. This is often used for proving de Rham’s theorem, which relates $H_*(X; \mathbb{R})$ to the *de Rham cohomology* $H_{\text{dR}}^*(X)$ of X , defined in terms of the exterior derivative on smooth differential forms; but there is a canonical isomorphism from smooth singular homology to $H_*(X)$, so it really is just another formulation of the same theory. Yet another variant is the *cubic* version of singular homology, which is used in place of $H_*(X)$ in a few books such as [Mas91]: here the point is to replace the standard n -simplex Δ^n with the standard n -cube I^n , which makes defining product structures easier but other things harder. In any case, cubic singular homology is indeed an axiomatic homology theory, but actually it is always isomorphic to $H_*(X)$, so it is not actually a different theory—which is why some authors feel free to use it as a substitute.

I want to describe a theory that satisfies the Eilenberg-Steenrod axioms and thus captures the same topological information as singular homology on nice spaces, but does not match it on all spaces and is based on a totally different idea. To explain the definition, we will also need to introduce inverse limits, the contravariant version of direct limits. For reasons that we will see, the Čech homology theory that I'm going to describe is not especially popular, partly because it requires some extra conditions in order to make it satisfy all the axioms. On the other hand, the closely related theory of Čech cohomology satisfies the analogous set of axioms without restriction, and is widely used in several branches of mathematics, especially in algebraic geometry. We will come back to that briefly after discussing the axioms for cohomology theories in a few lectures, but the present lecture is intended as a sketch of Čech homology in particular, and since it is only a sketch, we will leave several details unproved but give suitable references wherever possible. The main source for most of this material is [ES52].

The idea behind Čech homology is to measure the topology of a space X in terms of the combinatorial data formed by the overlaps of open sets in an arbitrarily fine open covering of X . The starting point is the observation that for any given open covering, the overlaps can be encoded in the form of a simplicial complex. (For the abstract combinatorial notion of a simplicial complex, you may want to review Lecture 29.)

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 44.1. For each open covering $\mathfrak{U} \in \mathcal{O}(X)$ of a space X , the **nerve** of \mathfrak{U} is the simplicial complex $\mathcal{N}(\mathfrak{U})$ whose set of vertices is \mathfrak{U} , and whose simplices are the finite subsets $\sigma \subset \mathfrak{U}$ such that

$$\bigcap_{\mathcal{U} \in \sigma} \mathcal{U} \neq \emptyset.$$

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 simplicial homology of the nerves of open coverings. To give the definition in its standard form, I need to add a detail about simplicial homology that we did not have occasion to discuss before: recall from Lecture 29 that if $K = (V, S)$ is a simplicial complex, the generators of the simplicial n -chain group $C_n^\Delta(K)$ are oriented n -simplices

$$[v_0, \dots, v_n] \in C_n^\Delta(K),$$

of which there are two equivalence classes for each individual n -simplex $\{v_0, \dots, v_n\} \in S$, with equivalence defined via even permutations of the order of vertices, and in $C_n^\Delta(K)$ they are also subject to an equivalence relation such that the two orientations of $\{v_0, \dots, v_n\}$ define linearly-dependent generators of opposite signs:

$$[v_0, v_1, v_2, \dots, v_n] = -[v_1, v_0, v_2, \dots, v_n] \in C_n^\Delta(K).$$

There is an alternative way to define a chain complex out of this data, and it is algebraically simpler, though less geometrically motivated. Let

$$C_n^o(K)$$

denote the free abelian group generated by the set of all ordered $(n+1)$ -tuples (v_0, \dots, v_n) of vertices of K such that there exists a simplex $\sigma \in S$ containing all of v_0, \dots, v_n . Note that we do not assume the vertices v_0, \dots, v_n are all distinct, though if they are, then it means $\{v_0, \dots, v_n\} \subset S$ and the ordered tuple is called an **ordered n -simplex** of K . There is an obvious surjective homomorphism

$$(44.1) \quad C_n^o(K) \rightarrow C_n^\Delta(K) : (v_0, \dots, v_n) \mapsto \begin{cases} 0 & \text{if } v_i = v_j \text{ for some } i \neq j, \\ [v_0, \dots, v_n] & \text{otherwise.} \end{cases}$$

This becomes a chain map if we define $\partial : C_n^o(K) \rightarrow C_{n-1}^o(K)$ for each $n \geq 1$ by

$$\partial(v_0, \dots, v_n) = \sum_{k=0}^n (-1)^k (v_0, \dots, v_{k-1}, v_{k+1}, \dots, v_n),$$

producing a chain complex $C_*^o(K) = \bigoplus_{n \in \mathbb{Z}} C_n^o(K)$, where as usual we set $C_n^o(K) := 0$ for $n < 0$. The following can be proved using the method of acyclic models (see [Spa95, §4.3] or [ES52, Theorem VI.6.9]):

PROPOSITION 44.2. *For any simplicial complex K , the canonical chain map $C_*^o(K) \rightarrow C_*^\Delta(K)$ defined by (44.1) is a chain homotopy equivalence. \square*

We shall denote the homology of the complex $C_*^o(K)$ with coefficients in an abelian group G by

$$H_*^o(K; G) := H_*(C_*^o(K) \otimes G),$$

and similarly denote by $H_*^o(K, L; G)$ the homology of the quotient complex for a simplicial pair (K, L) . Proposition 44.2 implies that there is always a canonical isomorphism $H_*^o(K, L; G) \cong H_*^\Delta(K, L; G)$; moreover, it is easy to check that this isomorphism is *natural* in the sense that it identifies the homomorphisms on H_*^o and H_*^Δ that are induced by simplicial maps. This means that, in practice, $H_*^o(K, L; G)$ is exactly the same invariant as our usual simplicial homology $H_*^\Delta(K, L; G)$, but the definition of $C_*^o(K)$ makes certain matters of bookkeeping slightly easier.

Figure 22 shows some examples of open covers \mathcal{U} of S^1 and the polyhedra $|\mathcal{N}(\mathcal{U})|$ that arise from their nerves. We see that in one case, $|\mathcal{N}(\mathcal{U})|$ is homeomorphic to S^1 ; this is not a coincidence, and we'll come back to it shortly. In general, however, $|\mathcal{N}(\mathcal{U})|$ need not be homeomorphic nor even homotopy equivalent to the space that is being covered, and in fact, the nerve of an open cover of X can easily be an infinite-dimensional simplicial complex even when X is something as tame as a compact polyhedron. 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 is to consider an open covering together with all of its possible refinements.

A **refinement** of an open cover $\mathcal{U} \in \mathcal{O}(X)$ is another open covering $\mathcal{U}' \in \mathcal{O}(X)$ such that every $\mathcal{U}' \in \mathcal{U}'$ is a subset of some $\mathcal{U} \in \mathcal{U}$. For pairs (X, A) , we say similarly that a refinement of




FIGURE 22. 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.

$(\mathfrak{U}, \mathfrak{U}_A) \in \mathcal{O}(X, A)$ 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

$$F : \mathfrak{U}' \rightarrow \mathfrak{U}, \quad F(\mathfrak{U}'_A) \subset \mathfrak{U}_A$$

such that for every $\mathcal{U} \in \mathfrak{U}'$, $\mathcal{U} \subset F(\mathcal{U})$. It follows that if $\sigma \subset \mathfrak{U}'$ is a simplex of $\mathcal{N}(\mathfrak{U}')$, then

$$\bigcap_{\mathcal{U} \in \sigma} F(\mathcal{U}) \supset \bigcap_{\mathcal{U} \in \sigma} \mathcal{U} \neq \emptyset,$$

hence $F(\sigma) \subset \mathfrak{U}$ is a simplex of $\mathcal{N}(\mathfrak{U})$, and similarly, F maps simplices of $\mathcal{N}(\mathfrak{U}'_A)$ to simplices of $\mathcal{N}(\mathfrak{U}_A)$. In other words, F is a simplicial map from $\mathcal{N}(\mathfrak{U}')$ to $\mathcal{N}(\mathfrak{U})$, and in the relative case, a map of simplicial pairs:

$$F : \mathcal{N}(\mathfrak{U}', \mathfrak{U}'_A) \rightarrow \mathcal{N}(\mathfrak{U}, \mathfrak{U}_A).$$

It therefore induces a chain map

$$(44.2) \quad F_* : C_*^o(\mathcal{N}(\mathfrak{U}', \mathfrak{U}'_A)) \rightarrow C_*^o(\mathcal{N}(\mathfrak{U}, \mathfrak{U}_A)).$$

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 44.3 ([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 induced chain map (44.2) on ordered simplicial chain complexes is independent of choices up to chain homotopy equivalence. \square*

It follows that for any choice of coefficient group G , we can associate to any refinement $\beta := (\mathfrak{U}', \mathfrak{U}'_A) \in \mathcal{O}(X, A)$ of an open covering $\alpha := (\mathfrak{U}, \mathfrak{U}_A) \in \mathcal{O}(X, A)$ a natural homomorphism of simplicial homology groups

$$\varphi_{\alpha\beta} : H_*^o(\mathcal{N}(\mathfrak{U}', \mathfrak{U}'_A); G) \rightarrow H_*^o(\mathcal{N}(\mathfrak{U}, \mathfrak{U}_A); G).$$

One can view this as defining something very similar to a direct system: indeed, let us define a pre-order on $\mathcal{O}(X, A)$ by writing

$$(\mathfrak{U}', \mathfrak{U}'_A) > (\mathfrak{U}, \mathfrak{U}_A)$$

whenever $(\mathfrak{U}', \mathfrak{U}'_A)$ is a refinement of $(\mathfrak{U}, \mathfrak{U}_A)$. (Note that it is not a partial order, as two open coverings can easily be refinements of each other without being identical.) Since any two open coverings have a common refinement, this makes $(\mathcal{O}(X, A), <)$ a directed set, and the result above associates a morphism of \mathbb{Z} -graded abelian groups to any pair $\alpha, \beta \in \mathcal{O}(X, A)$ with $\beta > \alpha$. The only trouble is that this morphism goes the wrong way: if the collection of graded abelian groups

$$\{H_*^o(\mathcal{N}(\alpha); G)\}_{\alpha \in \mathcal{O}(X, A)}$$

were to be viewed as a direct system, then we would have to have an associated morphism $\varphi_{\beta\alpha} : H_*^o(\mathcal{N}(\alpha); G) \rightarrow H_*^o(\mathcal{N}(\beta); G)$ for every $\beta > \alpha$, but instead we have $\varphi_{\alpha\beta} : H_*^o(\mathcal{N}(\beta); G) \rightarrow H_*^o(\mathcal{N}(\alpha); G)$. There is also a name for this.

DEFINITION 44.4. Given a category \mathcal{C} and a directed set $(I, <)$, an **inverse system** (*projektives System*) $\{X_\alpha, \varphi_{\alpha\beta}\}$ in \mathcal{C} over $(I, <)$ associates to each $\alpha \in I$ an object X_α of \mathcal{C} , along with morphisms

$$\varphi_{\alpha\beta} \in \text{Mor}(X_\beta, X_\alpha) \quad \text{for each } \alpha < \beta$$

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 $\alpha, \beta, \gamma \in I$ with $\alpha < \beta < \gamma$.

REMARK 44.5. In terms of the category \mathcal{I} corresponding to the directed set $(I, <)$ as in Remark 38.2, an inverse system in \mathcal{C} over $(I, <)$ is the same thing as a contravariant functor $\mathcal{I} \rightarrow \mathcal{C}$ (recall Definition 26.13).

Convergence of inverse systems is defined analogously to direct systems, the main difference being that most arrows go the other way.

DEFINITION 44.6. For an inverse system $\{X_\alpha, \varphi_{\alpha\beta}\}$ in \mathcal{C} over $(I, <)$, a **target** $\{Y, f_\alpha\}$ of the system consists of an object Y of \mathcal{C} together with associated morphisms $f_\alpha \in \text{Mor}(Y, X_\alpha)$ for each $\alpha \in I$ such that the diagram

$$\begin{array}{ccc} X_\alpha & \xleftarrow{\varphi_{\alpha\beta}} & X_\beta \\ & \swarrow f_\alpha & \searrow f_\beta \\ & Y & \end{array}$$

commutes for every pair $\alpha, \beta \in I$ with $\alpha < \beta$.

DEFINITION 44.7. A target $\{X_\infty, \varphi_\alpha\}$ of the inverse system $\{X_\alpha, \varphi_{\alpha\beta}\}$ is called an **inverse limit**⁶⁴ (*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 $f_\infty \in \text{Mor}(Y, X_\infty)$ such that the diagram

$$\begin{array}{ccc} X_\alpha & \xleftarrow{\varphi_\alpha} & X_\infty \\ & \searrow f_\alpha & \uparrow f_\infty \\ & & Y \end{array}$$

commutes for every $\alpha \in I$.

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 $\alpha < \beta < \gamma < \dots \in I$, and the defining feature of $\varprojlim \{X_\alpha\}$ is that the morphism indicated by the dashed arrow must exist and be unique whenever all the other morphisms in the diagram are given.

As with direct limits, there is no guarantee from these definitions that an inverse limit must exist, but for the categories we are most interested in, its existence can be established by describing it more concretely. One should not confuse the statement that an inverse limit exists with any claim that it is nonempty—the empty set is also a topological space and can appear as the limit of an inverse system in **Top** (see Example 44.12 below).

EXERCISE 44.8. If $\{X_\alpha, \varphi_{\alpha\beta}\}$ is an inverse system in **Top** over $(I, <)$, show that its inverse limit is the space

$$\varprojlim \{X_\alpha\} = \left\{ \{x_\alpha\} \in \prod_{\alpha \in I} X_\alpha \mid x_\alpha = \varphi_{\alpha\beta}(x_\beta) \text{ for all } \alpha, \beta \in I \text{ with } \alpha < \beta \right\},$$

with the associated morphisms $\varphi_\alpha : \varprojlim \{X_\beta\} \rightarrow X_\alpha$ defined via the obvious projections $\prod_{\beta \in I} X_\beta \rightarrow X_\alpha$ for each $\alpha \in I$. Conclude from this that the topology on $\varprojlim \{X_\alpha\}$ is the weakest for which the maps $\varphi_\alpha : \varprojlim \{X_\beta\} \rightarrow X_\alpha$ are all continuous.

REMARK 44.9. The exercise extends in an obvious way to describe inverse limits in the category **Set** of sets (with morphisms defined as arbitrary maps).

EXERCISE 44.10. Consider an inverse system $\{X_\alpha, \varphi_{\alpha\beta}\}$ in **Top** for which the spaces X_α are all subspaces of some fixed topological space X , $\beta > \alpha$ holds if and only if $X_\beta \subset X_\alpha$ and the maps $\varphi_{\alpha\beta} : X_\beta \rightarrow X_\alpha$ are all inclusions. Show that $\varprojlim \{X_\alpha\} = \bigcap_{\alpha \in I} X_\alpha$, with the associated morphisms $\varphi_\alpha : \varprojlim \{X_\beta\} \rightarrow X_\alpha$ given by the obvious inclusions.

Comment: The obvious analogue of this exercise involving direct limits and unions is only sometimes true, e.g. it works for viewing any CW-complex as the direct limit of its skeleta, but Problem Set 7 #5 shows an example in which the direct limit and the union are the same set with different topologies. In this sense, inverse systems in the category **Top** are somewhat better behaved than direct systems.

⁶⁴Inverse limits are also sometimes called **projective limits** or just plain **limits** (as opposed to *colimits*).

EXERCISE 44.11. Prove that for any inverse system $\{X_\alpha, \varphi_{\alpha\beta}\}$ of topological spaces such that every X_α is nonempty, compact and Hausdorff, $\varprojlim\{X_\alpha\} \neq \emptyset$.

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 $\{x^\beta \in \prod_\alpha X_\alpha\}_{\beta \in I}$ then defines a net in $\prod_\alpha X_\alpha$, which therefore has a cluster point. Prove that the cluster point belongs to $\varprojlim\{X_\alpha\}$. (For a slightly different argument that does not use nets, see [ES52, Theorem VIII.3.6]; it does still require Tychonoff's theorem.)

EXAMPLE 44.12. Combining the previous two exercises produces the well-known fact that in any Hausdorff space, the intersection of any collection of nonempty compact subsets that all have nonempty pairwise intersections is nonempty. It is easy to see that the compactness condition cannot be dropped from this statement: for instance, taking the collection of intervals $\{(0, 1/n]\}_{n \in \mathbb{N}}$ as an inverse system in the sense of Exercise 44.10, the inverse limit is

$$\varprojlim\{(0, 1/n]\}_{n \in \mathbb{N}} = \bigcap_{n \in \mathbb{N}} (0, 1/n] = \emptyset.$$

EXERCISE 44.13. If $\{G_\alpha, \varphi_{\alpha\beta}\}$ is an inverse system in \mathbf{Ab} over $(I, <)$, show that its inverse limit is a group of the form

$$\varprojlim\{G_\alpha\} = \left\{ \{g_\alpha\} \in \prod_{\alpha \in I} G_\alpha \mid g_\alpha = \varphi_{\alpha\beta}(g_\beta) \text{ for all } \alpha, \beta \in I \text{ with } \alpha < \beta \right\},$$

with the associated homomorphisms $\varphi_\alpha : \varprojlim\{G_\beta\} \rightarrow G_\alpha$ defined via the projections $\prod_{\beta \in I} G_\beta \rightarrow G_\alpha$ all $\alpha \in I$.

EXERCISE 44.14. Prove the obvious analogues of the result in Exercise 44.13 for inverse systems in the categories $\mathbf{Ab}_\mathbb{Z}$ of \mathbb{Z} -graded abelian groups and Chain of chain complexes.

We now have enough concepts in place to define the Čech homology groups.

DEFINITION 44.15. The Čech homology $\check{H}_*(X, A; G)$ of a pair of spaces (X, A) with coefficients in an abelian group G is defined as the \mathbb{Z} -graded abelian group

$$\check{H}_*(X, A; G) := \varprojlim \{H_*^o(\mathcal{N}(\mathfrak{U}, \mathfrak{U}_A); G)\}_{(\mathfrak{U}, \mathfrak{U}_A) \in \mathcal{O}(X, A)}.$$

It is slightly harder than for singular homology to see why this should define a functor $\mathbf{Top}_{\text{rel}} \rightarrow \mathbf{Ab}_\mathbb{Z}$, 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 $f^{-1}(\mathcal{U})$ for $\mathcal{U} \in \mathfrak{U}$, and whenever $\beta \in \mathcal{O}(Y, B)$ is a refinement of α , $f^*\beta \in \mathcal{O}(X, A)$ is clearly also a refinement of $f^*\alpha$. The obvious correspondence 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

$$f_* : H_*^o(\mathcal{N}(f^*\alpha); G) \rightarrow H_*^o(\mathcal{N}(\alpha); G)$$

for every $\alpha \in \mathcal{O}(Y, B)$. Using the universal property of the inverse limit, one can derive from this a morphism

$$f_* : \check{H}_*(X, A; G) \rightarrow \check{H}_*(Y, B; G)$$

between the corresponding inverse limits, and prove that it satisfies the usual conditions for $\check{H}_*(\cdot; G)$ 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 $\check{H}_*(X, A; G)$ to be the same as $H_*(X, A; G)$. To this end, consider the case where X is the polyhedron of a finite




FIGURE 23. The open stars of two neighboring vertices v_0 and v_1 in a simplicial complex.

simplicial complex $K = (V, S)$. We saw in Lecture 40 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 23). These sets define a distinguished open covering of X ,

$$\mathfrak{U}_K := \{\text{st } v \mid v \in V\},$$

and recall from Exercise 40.15 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 the complex K itself. Now if $\mathfrak{U} \in \mathcal{O}(X)$ is another open covering, since X is compact, we can always find a refinement of \mathfrak{U} in the form $\mathfrak{U}_{K'}$ by applying barycentric subdivision to the simplices of K enough times, producing a new simplicial complex K' with more and smaller simplices but a homeomorphic polyhedron $|K'| = X$, and since barycentric subdivision induces chain homotopy equivalences, one can show that the induced map

$$H_*(\mathcal{N}(\mathfrak{U}_{K'}); G) \rightarrow H_*(\mathcal{N}(\mathfrak{U}_K); G)$$

resulting from the fact that $\mathfrak{U}_{K'} > \mathfrak{U}_K$ is always an isomorphism. This is enough information to compute the inverse limit (see Problem Set 10 #4(g)) and prove:

THEOREM 44.16. *For any compact polyhedron $X = |K|$ with underlying simplicial complex K , $\check{H}_*(X; G) \cong H_*^\Delta(K; G)$ for every coefficient group G . \square*

It follows in particular that Čech homology is isomorphic to singular homology on compact polyhedra. Notice by the way that if we had not already constructed one example of an axiomatic homology theory, we could still use this argument to prove that simplicial homology is independent of the triangulation of a compact polyhedron—it follows now from the fact Čech homology is a topological invariant.

It is not *always* true however that $\check{H}_*(X; G) \cong H_*(X; G)$.

LEMMA 44.17. *If X is a connected space, then for every open cover \mathfrak{U} of X , the nerve $\mathcal{N}(\mathfrak{U})$ is connected.*

PROOF. If $\mathcal{N}(\mathfrak{U})$ is not connected then it can be decomposed as a disjoint union of two nonempty subcomplexes $\mathcal{N}(\mathfrak{U}) \cong K_0 \sqcup K_1$. 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 44.18. *For any connected space X and any coefficient group G , $\check{H}_0(X; G) \cong G$.*

PROOF. Lemma 44.17 implies that for every $\mathfrak{U} \in \mathcal{O}(X)$, $H_0^o(\mathcal{N}(\mathfrak{U}); G) \cong H_0^\Delta(\mathcal{N}(\mathfrak{U}); G) \cong G$. It is similarly easy to show that the canonical map $H_0^o(\mathcal{N}(\mathfrak{U}'); G) \rightarrow H_0^o(\mathcal{N}(\mathfrak{U}); G)$ 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 dimension 0, which splits over a direct sum of the *path-components* (not connected components) of each space. So, for instance, Figure 19 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 44.16. At least $\check{H}_*(\cdot; G)$ and $H_*(\cdot; G)$ will match on all CW-pairs if $\check{H}_*(\cdot; G)$ satisfies the Eilenberg-Steenrod axioms. So does it? The answer is a bit surprising.

THEOREM 44.19. *For every abelian group G , $\check{H}_*(\cdot; G) : \text{Top}_{\text{rel}} \rightarrow \text{Ab}_{\mathbb{Z}}$ satisfies all of the Eilenberg-Steenrod axioms except for exactness, but it does not satisfy exactness in general.*

An actual counterexample to the exactness axiom is explained in [ES52, §X.4]. It would take at least a few lectures to either explain that counterexample or prove that the rest of the axioms are satisfied, so we'll mostly skip it since this lecture is meant to be only a brief digression away from the main topic of the course. But it's worth taking a closer look at how one would naturally try to prove the exactness axiom, and why it fails in general. It also succeeds in some cases, so the negative statement in Theorem 44.19 is not the end of the story.

The problem with exactness is traceable to a problem with the behavior of exact sequences under inverse limits. If (X, A) is a pair of spaces and $(\mathfrak{U}, \mathfrak{U}_A) \in \mathcal{O}(X, A)$, then $\mathcal{N}(\mathfrak{U}_A) \subset \mathcal{N}(\mathfrak{U})$ is a subcomplex and the obvious short exact sequence of ordered simplicial chain complexes

$$0 \rightarrow C_*^o(\mathcal{N}(\mathfrak{U}_A); G) \rightarrow C_*^o(\mathcal{N}(\mathfrak{U}); G) \rightarrow C_*^o(\mathcal{N}(\mathfrak{U}, \mathfrak{U}_A); G) = C_*^o(\mathcal{N}(\mathfrak{U}); G) / C_*^o(\mathcal{N}(\mathfrak{U}_A); G) \rightarrow 0$$

gives rise to a long exact sequence of simplicial homology groups

$$(44.3) \quad \dots \rightarrow H_n^o(\mathcal{N}(\mathfrak{U}_A); G) \rightarrow H_n^o(\mathcal{N}(\mathfrak{U}); G) \rightarrow H_n^o(\mathcal{N}(\mathfrak{U}, \mathfrak{U}_A); G) \rightarrow H_{n-1}^o(\mathcal{N}(\mathfrak{U}_A); G) \rightarrow \dots$$

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

$$\begin{array}{cccccccc} \dots & \rightarrow & H_n^o(\mathcal{N}(\mathfrak{U}_A); G) & \rightarrow & H_n^o(\mathcal{N}(\mathfrak{U}); G) & \rightarrow & H_n^o(\mathcal{N}(\mathfrak{U}, \mathfrak{U}_A); G) & \rightarrow & H_{n-1}^o(\mathcal{N}(\mathfrak{U}_A); G) & \rightarrow & \dots \\ & & \uparrow & & \uparrow & & \uparrow & & \uparrow & & \\ \dots & \rightarrow & H_n^o(\mathcal{N}(\mathfrak{U}'_A); G) & \rightarrow & H_n^o(\mathcal{N}(\mathfrak{U}'); G) & \rightarrow & H_n^o(\mathcal{N}(\mathfrak{U}', \mathfrak{U}'_A); G) & \rightarrow & H_{n-1}^o(\mathcal{N}(\mathfrak{U}'_A); G) & \rightarrow & \dots \end{array}$$

An exact sequence is the same thing as a chain complex with trivial homology groups, so we can therefore regard this collection of exact sequences as an inverse system over $(\mathcal{O}(X, A), <)$ in the

category Chain of chain complexes. By Exercise 44.14, this system will have an inverse limit, which will be a chain complex

$$(44.4) \quad \dots \rightarrow \check{H}_n(A; G) \rightarrow \check{H}_n(X; G) \rightarrow \check{H}_n(X, A; G) \rightarrow \check{H}_{n-1}(A; G) \rightarrow \dots$$

But it is not obvious whether this sequence of maps is exact. If this had been a direct limit, we could now appeal to Proposition 38.18, which would produce a natural isomorphism between the homology of the direct limit and the direct limit of the homology groups in the system; the latter are all zero since the sequences are all exact, so this would imply that the limit sequence is also exact. But we don't have an analogue of Proposition 38.18 for inverse limits. As a matter of fact, the result we would like to prove at this point is false:

EXAMPLE 44.20. 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

$$\begin{array}{ccccccccccc} \dots & \longrightarrow & 0 & \longrightarrow & A_1 & \xrightarrow{\cdot 2} & B_1 & \xrightarrow{\text{pr}} & C_1 & \longrightarrow & 0 & \longrightarrow & \dots \\ & & \uparrow & & \cdot 3 \uparrow & & \cdot 3 \uparrow & & 1 \uparrow & & \uparrow & & \\ \dots & \longrightarrow & 0 & \longrightarrow & A_2 & \xrightarrow{\cdot 2} & B_2 & \xrightarrow{\text{pr}} & C_2 & \longrightarrow & 0 & \longrightarrow & \dots \\ & & \uparrow & & \cdot 3 \uparrow & & \cdot 3 \uparrow & & 1 \uparrow & & \uparrow & & \\ \dots & \longrightarrow & 0 & \longrightarrow & A_3 & \xrightarrow{\cdot 2} & B_3 & \xrightarrow{\text{pr}} & C_3 & \longrightarrow & 0 & \longrightarrow & \dots \\ & & \uparrow & & \uparrow & & \uparrow & & \uparrow & & \uparrow & & \\ & & \vdots & & \vdots & & \vdots & & \vdots & & \vdots & & \end{array}$$

commutes, meaning the $\varphi_{n-1,n}$ are all chain maps, and we can compose them to define further chain maps $\varphi_{m,n}$ for every $m < n$ and interpret this collection of data as an inverse system of chain complexes. By Exercises 44.13 and 44.14, the individual terms of the inverse limit complex are as follows: first,

$$\varprojlim \{A_n\} = \left\{ (a_1, a_2, a_3, \dots) \in \prod_{n \in \mathbb{N}} A_n \mid a_{n-1} = 3a_n \text{ for all } n \geq 2 \right\} = 0,$$

and $\varprojlim \{B_n\}$ similarly vanishes since no integer is divisible by arbitrarily large powers of 3. On the other hand,

$$\varprojlim \{C_n\} = \left\{ (c_1, c_2, c_3, \dots) \in \prod_{n \in \mathbb{N}} C_n \mid c_{n-1} = c_n \text{ for all } n \geq 2 \right\} \cong \mathbb{Z}_2,$$

so the full inverse limit chain complex 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 $(I, <)$, so have a commuting diagram for every $\beta > \alpha$ in the form

$$(44.5) \quad \begin{array}{ccccccc} \cdots & \longrightarrow & C_{n+1}^\alpha & \xrightarrow{\partial^\alpha} & C_n^\alpha & \xrightarrow{\partial^\alpha} & C_{n-1}^\alpha & \xrightarrow{\partial^\alpha} & \cdots \\ & & \varphi_{\alpha\beta} \uparrow & & \varphi_{\alpha\beta} \uparrow & & \varphi_{\alpha\beta} \uparrow & & \\ \cdots & \longrightarrow & C_{n+1}^\beta & \xrightarrow{\partial^\beta} & C_n^\beta & \xrightarrow{\partial^\beta} & C_{n-1}^\beta & \xrightarrow{\partial^\beta} & \cdots \end{array}$$

and assume moreover that the rows of these diagrams are always exact. The inverse limit is

$$C_*^\infty := \varprojlim \{C_*^\alpha\} = \left\{ \{x_\alpha\} \in \prod_{\alpha \in I} C_*^\alpha \mid \varphi_{\alpha\beta}(x_\beta) = x_\alpha \text{ for all } \beta > \alpha \right\},$$

where the chain complex boundary map can be written as

$$\partial^\infty := \prod_{\alpha \in I} \partial^\alpha \Big|_{C_*^\infty} : C_*^\infty \rightarrow C_*^\infty,$$

the restriction to the subgroup $C_*^\infty \subset \prod_{\alpha \in I} C_*^\alpha$ 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 $x = \{x_\alpha\}_{\alpha \in I} \in C_*^\infty$ with $\partial^\infty \{x_\alpha\} = 0$, we have $\partial^\alpha x_\alpha = 0$ for all $\alpha \in I$ and thus $x_\alpha = \partial^\alpha y_\alpha$ for some $y_\alpha \in C_{n+1}^\alpha$. The trouble is that these elements y_α are not generally unique, and if they are chosen arbitrarily, then they need not satisfy

$$(44.6) \quad \varphi_{\alpha\beta}(y_\beta) = y_\alpha \quad \text{for all } \beta > \alpha,$$

without which $\{y_\alpha\}_{\alpha \in I}$ will not be an element of C_{n+1}^∞ .

To get a firmer handle on this problem, define for each $\alpha \in I$ the nonempty subset

$$K_\alpha := (\partial^\alpha)^{-1}(x_\alpha) \subset C_{n+1}^\alpha.$$

The chain map relation and the condition $\varphi_{\alpha\beta}(x_\beta) = x_\alpha$ then imply

$$\varphi_{\alpha\beta}(K_\beta) \subset K_\alpha \quad \text{for all } \beta > \alpha,$$

which makes the collection of sets $\{K_\alpha\}_{\alpha \in I}$ with maps $K_\beta \xrightarrow{\varphi_{\alpha\beta}} K_\alpha$ into an inverse system in Set over $(I, <)$. By Exercise 44.8 and Remark 44.9, $\varprojlim \{K_\alpha\}$ is then the set of all elements $\{y_\alpha\} \in \prod_{\alpha \in I} K_\alpha$ such that (44.6) 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 44.20 implies that the answer must sometimes be no, and indeed, we know from Example 44.12 that an inverse limit of nonempty sets or topological spaces can easily be the empty set.

To make progress, we need to add more assumptions. Suppose first of all that the individual groups C_n^α for each $n \in \mathbb{Z}$ and $\alpha \in I$ are finite. Then the sets K_α are also finite, and if we assign them the discrete topology, we can view them all as nonempty compact Hausdorff spaces. In this case there is a positive result we can use: Exercise 44.11 implies that $\varprojlim \{K_\alpha\}$ will then always be nonempty, which fills the gap at the end of our proof that the limit sequence is exact!

I would like to point out that this trick for the case of finite groups is fairly abstract: hidden inside Exercise 44.11 is Tychonoff's theorem on the compactness of arbitrary products of compact spaces (cf. Lecture 6 from last semester), which depends on Zorn's lemma and thus the axiom of choice. As a consequence, we are guaranteed the existence of some $y \in (\partial^\infty)^{-1}(x)$ whenever $\partial^\infty x = 0$, but we cannot even begin to suggest how one might find y in practice. In the classic book of Eilenberg and Steenrod (see [ES52, Theorem 5.7 and Lemma 5.8 in Chapter VIII]), there

is a linear-algebraic variation on this trick that also uses Zorn's lemma, and similarly solves the problem whenever the groups C_n^α are all assumed to be finite-dimensional vector spaces over a field \mathbb{K} , with ∂^α and $\varphi_{\alpha\beta}$ as \mathbb{K} -linear maps. These two scenarios are relevant to Čech homology under certain assumptions: in particular, suppose the coefficient group G is either finite or a finite-dimensional vector space over a field, and (X, A) is a *compact pair*, meaning X is a compact Hausdorff space and $A \subset X$ is closed. In this case, our open coverings of (X, A) always have finite refinements, whose nerves are then finite simplicial pairs, and the groups in the sequence (44.3) 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 44.21. *If G is either a finite abelian group or a finite-dimensional vector space over a field, then the restriction of Čech homology to the category Cpct_{rel} of compact pairs defines an axiomatic homology theory*

$$\check{H}_*(\cdot; G) : \text{Cpct}_{\text{rel}} \rightarrow \text{Ab}_{\mathbb{Z}}.$$

□

The restriction to compact pairs means that some details of the theory we have developed for axiomatic homology need to be handled with a bit more care: for instance, only the weaker form of the excision axiom (see Remark 30.4) makes sense in this category, so some of the excision tricks we used, e.g. for computing the homology of spheres and the isomorphism $h_*(X, A) \cong \tilde{h}_*(X/A)$ for good pairs, need to be modified a bit when $h_* = \check{H}_*(\cdot; G)$. But this can be done, with the result that if G satisfies the conditions in the theorem above, then there is always a natural isomorphism

$$\check{H}_*(X, A; G) \cong H_*(X, A; G)$$

when (X, A) is a compact CW-pair.

Čech homology also has one nice property that singular homology does not: it is *continuous* with respect to inverse limits of spaces. The statement can be formulated for any axiomatic homology theory h_* as follows: suppose $\{(X_\alpha, A_\alpha), \varphi_{\alpha\beta}\}$ is an inverse system of pairs of spaces over some directed set $(I, <)$. The associated morphisms $\varphi_\alpha : \varprojlim\{(X_\beta, A_\beta)\} \rightarrow (X_\alpha, A_\alpha)$ then induce homomorphisms

$$\Phi_\alpha := (\varphi_\alpha)_* : h_*(\varprojlim\{(X_\beta, A_\beta)\}) \rightarrow h_*(X_\alpha, A_\alpha),$$

which make $\{h_*(\varprojlim\{(X_\beta, A_\beta)\}), \Phi_\alpha\}$ a target of the inverse system of \mathbb{Z} -graded abelian groups

$$\{h_*(X_\alpha, A_\alpha), (\varphi_{\alpha\beta})_*\}.$$

By the universal property of inverse limits, there is then a canonical limit morphism

$$\Phi_\infty : h_*(\varprojlim\{(X_\alpha, A_\alpha)\}) \rightarrow \varprojlim\{h_*(X_\alpha, A_\alpha)\}.$$

The following result is often quoted as the selling point of the Čech theory in comparison with singular homology. One can show in fact that every compact pair is the inverse limit of some inverse system of compact pairs that are homotopy equivalent to CW-pairs, thus the theorem can be used to understand the topology of very “wild” spaces for which singular homology cannot be expected to give a reasonable answer.

THEOREM 44.22 (continuity in Čech homology; see [ES52, Chapter X]). *For any inverse system of compact pairs $\{(X_\alpha, A_\alpha), \varphi_{\alpha\beta}\}$ and any abelian coefficient group G , the canonical map*

$$\check{H}_*(\varprojlim\{(X_\alpha, A_\alpha)\}; G) \rightarrow \varprojlim\{\check{H}_*(X_\alpha, A_\alpha; G)\}$$

is an isomorphism.

□

EXERCISE 44.23. 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 Exercise 44.10.

45. Singular cohomology

Motivation. Singular cohomology assigns to each topological space X and each abelian group G a \mathbb{Z} -graded abelian group denoted by

$$H^*(X; G) = \bigoplus_{n \in \mathbb{Z}} H^n(X; G).$$

It is closely related to singular homology, and in many (though not all) cases is isomorphic to it, but it has a slightly different structure. The most obvious difference is that as a functor from \mathbf{Top} to $\mathbf{Ab}_{\mathbb{Z}}$, it is *contravariant*, meaning that continuous maps $f : X \rightarrow Y$ induce homomorphisms

$$f^* : H^n(Y; G) \rightarrow H^n(X; G),$$

going the opposite direction from homology. You may at this stage rightfully question what is to be gained from this cosmetic difference: as we will see, the most significant advantage is that if we choose the coefficient group G to be a ring R , then $H^*(X; R)$ has a natural product structure, called the **cup product**

$$H^k(X; R) \otimes_R H^\ell(X; R) \xrightarrow{\cup} H^{k+\ell}(X; R).$$

It is closely related to the homology cross product, but the latter is something that we use to relate the homologies of two spaces X and Y to that of their product $X \times Y$, whereas \cup produces extra algebraic structure on $H^*(X; R)$ itself. This can be extremely useful in computations. Moreover, we will see that in the special case where X is a closed oriented n -manifold, \cup gives rise to a product structure on homology that has deep geometric meaning, the **intersection product**

$$H_{n-k}(X) \otimes H_{n-\ell}(X) \rightarrow H_{n-k-\ell}(X) : [M] \otimes [N] \mapsto [M] \cdot [N] := [M \cap N].$$

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

$$H^k(X) \xrightarrow{\cong} H_{n-k}(X)$$

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.

THEOREM 45.1. *For every even $n \geq 0$, every continuous map $f : \mathbb{C}P^n \rightarrow \mathbb{C}P^n$ has a fixed point.*

PROOF (MODULO TECHNICAL DETAILS). The complex projective spaces $\mathbb{C}\mathbb{P}^n$ have not previously appeared in these notes (though they did appear on Problem Set 7), so let's quickly recall the definition: $\mathbb{C}\mathbb{P}^n$ is 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. We saw in Problem Set 7 #2 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:

$$H^k(\mathbb{C}\mathbb{P}^n) \cong H_k(\mathbb{C}\mathbb{P}^n) \cong \begin{cases} \mathbb{Z} & \text{for } k = 0, 2, 4, \dots, 2n, \\ 0 & \text{for all other } k. \end{cases}$$

We will also see that there is a universal coefficient theorem that expresses $H^k(X; G)$ up to isomorphism in terms of $H_k(X)$, $H_{k-1}(X)$ and G , and implies moreover 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. Thus for a map $f : \mathbb{C}\mathbb{P}^n \rightarrow \mathbb{C}\mathbb{P}^n$, we can write

$$L(f) = \sum_{k \in \mathbb{Z}} (-1)^k \operatorname{tr} (H^k(\mathbb{C}\mathbb{P}^n) \xrightarrow{f^*} H^k(\mathbb{C}\mathbb{P}^n)) = \sum_{k=0}^n \operatorname{tr} (H^{2k}(\mathbb{C}\mathbb{P}^n) \xrightarrow{f^*} H^{2k}(\mathbb{C}\mathbb{P}^n)).$$

Now we take advantage of the cup product on $H^*(\mathbb{C}\mathbb{P}^n)$, which has the following properties:

- It is *natural*, i.e. for all $\alpha, \beta \in H^*(\mathbb{C}\mathbb{P}^n)$, $f^*(\alpha \cup \beta) = f^*\alpha \cup f^*\beta$. (This is a general property of the cup product with respect to continuous maps between arbitrary spaces.)
- If $\alpha \in H^2(\mathbb{C}\mathbb{P}^n) \cong \mathbb{Z}$ is a generator, then for each $k = 0, 1, \dots, n$,

$$\alpha^k := \underbrace{\alpha \cup \dots \cup \alpha}_k \in H^{2k}(\mathbb{C}\mathbb{P}^n) \cong \mathbb{Z}$$

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 $\alpha \in H^2(\mathbb{C}\mathbb{P}^n)$, 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 $H^2(\mathbb{C}\mathbb{P}^n) \cong \mathbb{Z}$. It follows via the two properties above that for each $k = 0, \dots, n$, the generator $\alpha^k \in H^{2k}(\mathbb{C}\mathbb{P}^n)$ 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

$$L(f) = 1 + m + \dots + m^k \in \mathbb{Z}.$$

This is clearly not equal to 0 if $m = 1$. On the other hand, if $m \neq 1$, then we can rewrite this as

$$L(f) = \frac{1 - m^{n+1}}{1 - m},$$

which is zero if and only if $m^{n+1} = 1$. Since m is an integer and we have already excluded the case $m = 1$, this can only happen if $m = -1$, and then only if n is odd. The result thus follows from the Lefschetz fixed point theorem. \square

The functor $\text{Hom}(\cdot, G)$ and cochains. Let's talk about algebra. Given a chain complex (C_*, ∂) of abelian groups, we obtain its homology by applying the functor $H_* : \text{Chain} \rightarrow \text{Ab}_{\mathbb{Z}}$, which discards some of the information in (C_*, ∂) in the hope of obtaining something computable. For a little more flexibility, we can also choose an abelian coefficient group G and “pre-process” the chain complex via the functor $\otimes G$, producing the composition of functors

$$(45.1) \quad \text{Chain} \xrightarrow{\otimes G} \text{Chain} \xrightarrow{H_*} \text{Ab}_{\mathbb{Z}},$$

which sends the chain complex C_* to $H_*(C_* \otimes G)$.

The idea of cohomology is to pre-process the chain complex in a different way: instead of applying $\otimes G$, we apply the functor $\text{Hom}(\cdot, G)$ and thus *dualize it*. You are certainly already familiar with the notion of the dual space of a vector space; more generally, the **dual** of a module A over a commutative ring R is defined as the module of R -module homomorphisms to R ,

$$\text{Hom}_R(A, R) := \{\lambda \in \text{Hom}(A, R) \mid \lambda(ra) = r\lambda(a) \text{ for all } r \in R, a \in A\},$$

which reproduces the definition familiar from linear algebra if R is a field. Restricting to the case $R = \mathbb{Z}$ defines the dual of an abelian group A to be

$$A^* := \text{Hom}(A, \mathbb{Z}).$$

More generally, we can fix an arbitrary abelian group G and consider the functor

$$\text{Ab} \rightarrow \text{Ab} : A \mapsto \text{Hom}(A, G).$$

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 26.13): 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.

We next define what $\text{Hom}(C_*, G)$ should mean when C_* is a chain complex with boundary map $\partial : C_* \rightarrow C_{*-1}$. Since C_* is a \mathbb{Z} -graded abelian group, we would like $\text{Hom}(C_*, G)$ to be another \mathbb{Z} -graded abelian group: 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 subgroup of elements with degree n in $\text{Hom}(C_*, G)$.⁶⁵ 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 45.3), we're going to introduce an extra sign and define

$$(45.2) \quad \delta : \text{Hom}(C_*, G) \rightarrow \text{Hom}(C_*, G) : \alpha \mapsto (-1)^{|\alpha|+1} \partial^* \alpha,$$

⁶⁵If 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 group 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.

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 abelian group A_* 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 $A_* \rightarrow B_*$ of degree -1 that satisfies the usual chain homotopy relation

$$\varphi - \psi = h\delta_A + \delta_B h.$$

The **homology** of a cochain complex (A_*, δ) is the \mathbb{Z} -graded abelian group

$$H_*(A_*, \delta) = \ker \delta / \text{im } \delta,$$

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 \mathbb{Z} -graded abelian group

$$H^*(C_*, \partial; G) := H_*(\text{Hom}(C_*, G), \delta).$$

Much like (45.1), the functor that replaces a chain complex with its cohomology can be expressed as the composition of two functors:

$$\text{Chain} \xrightarrow{\text{Hom}(\cdot, G)} \text{Cochain} \xrightarrow{H_*} \text{Ab}_{\mathbb{Z}}.$$

Here **Cochain** denotes the category whose objects are cochain complexes, with morphisms defined as chain maps, and $H_* : \text{Cochain} \rightarrow \text{Ab}_{\mathbb{Z}}$ is a covariant functor sending each cochain complex to its homology and chain maps to the induced homomorphisms on homology. The functor

$$\text{Hom}(\cdot, G) : \text{Chain} \rightarrow \text{Cochain}$$

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

$$\varphi^* : (\text{Hom}(B_*, G), \delta_B) \rightarrow (\text{Hom}(A_*, G), \delta_A),$$

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 functor $H^*(\cdot; G) : \text{Chain} \rightarrow \text{Ab}_{\mathbb{Z}}$ is also contravariant: it associates to each chain map $\varphi : (A_*, \partial_A) \rightarrow (B_*, \partial_B)$ the homomorphism $H_*(\text{Hom}(B_*, G), \delta_B) \rightarrow H_*(\text{Hom}(A_*, G), \delta_A)$ induced by the chain map φ^* , and we shall also denote the induced morphism of \mathbb{Z} -graded abelian groups by

$$\varphi^* : H^*(B_*, \partial_B; G) \rightarrow H^*(A_*, \partial_A; G).$$

Two further algebraic observations are worth recording before we go back to topology.

PROPOSITION 45.2. *If $\varphi, \psi : (A_*, \partial_A) \rightarrow (B_*, \partial_B)$ are chain maps between chain complexes and $h : A_* \rightarrow B_{*+1}$ is a chain homotopy between φ and ψ , then the map $\eta : \text{Hom}(B_*, G) \rightarrow \text{Hom}(A_{*-1}, G)$ 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

$$\text{Hom}(\cdot, G) : \text{Chain}^h \rightarrow \text{Cochain}^h,$$

where Cochain^h is the category with cochain complexes as objects and chain homotopy classes of chain maps as morphisms. As a consequence, $H^*(\cdot; G)$ likewise descends to a functor

$$H^*(\cdot; G) : \text{Chain}^h \rightarrow \text{Ab}_{\mathbb{Z}}.$$

The second observation is that for any chain complex (C_*, ∂) , the canonical pairing

$$(45.3) \quad \text{Hom}(C_n, G) \times C_n \rightarrow G : (\alpha, c) \mapsto \alpha(c)$$

descends to homology to give a well-defined pairing

$$(45.4) \quad H^n(C_*, \partial; G) \times H_n(C_*, \partial) \rightarrow G : ([\alpha], [c]) \mapsto \langle [\alpha], [c] \rangle := \alpha(c).$$

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 (45.4) as the **evaluation** of cohomology classes on homology classes.

REMARK 45.3. The reason for the sign in (45.2) can be understood in terms of the “chain-level” evaluation map (45.3). 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

$$(45.5) \quad \text{Hom}(C_*, G) \otimes C_* \rightarrow G$$

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 (45.4) is well defined as a corollary of the fact that (45.5) is a chain map. For this we need to make sense of $\text{Hom}(C_*, G) \otimes C_*$ as a tensor product chain complex, even though $\text{Hom}(C_*, G)$ 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 (45.5) vanishes on all elements of degree nonzero in the tensor product chain complex, so it becomes natural to understand the right hand side as a chain complex that has the group G in degree 0 and the trivial group in all other degrees. With this convention in place, the boundary map on the right hand side is zero, so the chain map condition demands that for all $\alpha \in \text{Hom}(C_k, G)$ 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 (45.2) 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].

REMARK 45.4. The entirety of this discussion admits a straightforward generalization in which (C_*, ∂) is a chain complex of modules over a commutative ring R with unit, G is also an R -module, and the functor $\text{Hom}(\cdot, G)$ is replaced by $\text{Hom}_R(\cdot, G)$, which transforms a chain complex of R -modules into a cochain complex of R -modules. The homology $H_*(C_*, \partial)$ and cohomology $H^*(C_*, \partial; G)$ then both also have natural R -module structures, and the evaluation (45.4) becomes an R -module homomorphism

$$H^*(C_*, \partial; G) \otimes_R H_*(C_*, \partial) \rightarrow G,$$

which just means that (45.4) is R -bilinear. This situation arises naturally if we start with (C_*, ∂) as a chain complex of abelian groups but then introduce R as a coefficient ring by replacing it with $(C_* \otimes R, \partial \otimes \mathbf{1})$, which has a natural R -module structure. The following exercise shows that if we now take the cohomology of $(C_* \otimes R, \partial \otimes \mathbf{1})$ with coefficients in an R -module G by applying $\text{Hom}_R(\cdot, G)$ and then H_* , we obtain exactly the same result as the cohomology of (C_*, ∂) with coefficients in G , except that instead of just the pairing of abelian groups that is defined in (45.4), we have a bilinear R -module pairing

$$H^*(C_*, \partial; G) \times H_*(C_* \otimes R, \partial \otimes \mathbf{1}) \rightarrow G.$$

EXERCISE 45.5. Assume R is a commutative ring with unit, A is an abelian group and G is an R -module.

- (a) Show that $A \otimes G$ and $\text{Hom}(A, G)$ each have natural R -module structures defined via the relations

$$\begin{aligned} r(a \otimes g) &:= a \otimes (rg) && \text{for } r \in R, a \in A, g \in G, \\ (r\Phi)(a) &:= r(\Phi(a)) && \text{for } r \in R, a \in A, \Phi \in \text{Hom}(A, G), \end{aligned}$$

and that $\otimes G$ and $\text{Hom}(\cdot, R)$ can each be understood as functors (covariant and contravariant respectively) from the category of abelian groups to the category of R -modules.

- (b) Show that there is a canonical R -module isomorphism between $\text{Hom}(A, G)$ and $\text{Hom}_R(A \otimes R, G)$, which defines a natural transformation between the functors

$$\text{Ab} \rightarrow \text{Mod}^R : A \mapsto \text{Hom}(A, G) \quad \text{and} \quad \text{Ab} \rightarrow \text{Mod}^R : A \mapsto \text{Hom}_R(A \otimes R, G).$$

The singular cochain complex. The **singular cohomology** of a pair (X, A) with coefficients in an abelian group G is now defined by applying the algebraic processing described above to the singular chain complex: that is,

$$H^*(X, A; G) := H^*(C_*(X, A); G) = H_*(\text{Hom}(C_*(X, A), G)).$$

For the case $G = \mathbb{Z}$, we sometimes abbreviate

$$H^*(X, A) := H^*(X, A; \mathbb{Z}),$$

though we will occasionally also use this shorter notation for cohomology with arbitrary coefficients when there is no danger of confusion. It is standard to abbreviate the cochain complex $\text{Hom}(C_*(X, A), G)$ by

$$C^*(X, A; G) := \text{Hom}(C_*(X, A), G)$$

and refer to elements of $C^*(X, A; G)$ as **singular cochains**. Elements of $\ker \delta \subset C^*(X, A; G)$ and $\text{im } \delta \subset C^*(X, A; G)$ are likewise called (singular) **cocycles** and **coboundaries** respectively. Each element $\varphi \in C^n(X; G)$ is a homomorphism $\varphi : C_n(X) \rightarrow G$, and since $C_n(X)$ is a free abelian group, all such homomorphisms can be described uniquely via their values on the generators, i.e. the singular n -simplices in X . We thus have a canonical identification

$$C^n(X; G) = G^{\mathcal{K}_n(X)} = \prod_{\sigma \in \mathcal{K}_n(X)} G = \{\text{functions } \varphi : \mathcal{K}_n(X) \rightarrow G\},$$

where $\mathcal{K}_n(X)$ again denotes the set of all singular n -simplices in X . We will often use this identification to regard cochains $\varphi \in C^n(X; G)$ simply as functions $\varphi : \mathcal{K}_n(X) \rightarrow G$. With this understood, we plug in (45.2) and the usual formula for the boundary operator $\partial : C_{n+1}(X) \rightarrow C_n(X)$ to find a corresponding formula for the **coboundary operator** $\delta : C^n(X; G) \rightarrow C^{n+1}(X; G)$, in the form

$$(45.6) \quad (\delta\varphi)(\sigma) = (-1)^{n+1} \sum_{k=0}^{n+1} (-1)^k \varphi(\sigma|_{\partial_{(k)}\Delta^{n+1}}) \quad \text{for } \varphi : \mathcal{K}_n(X) \rightarrow G, \quad \sigma \in \mathcal{K}_{n+1}(X).$$

In the relative case, we can think of a homomorphism $\varphi : C_n(X, A) = C_n(X)/C_n(A) \rightarrow G$ as equivalent to a homomorphism $\varphi : C_n(X) \rightarrow G$ that vanishes on the subgroup $C_n(A) \subset C_n(X)$, so this is the same thing as a function $\mathcal{K}_n(X) \rightarrow G$ that vanishes on the subset $\mathcal{K}_n(A) \subset \mathcal{K}_n(X)$:

$$C^n(X, A; G) = \left\{ \varphi : \mathcal{K}_n(X) \rightarrow G \mid \varphi|_{\mathcal{K}_n(A)} = 0 \right\}.$$

The formula (45.6) then gives the correct homomorphism $\delta : C^n(X, A; G) \rightarrow C^{n+1}(X, A; G)$ by restriction.

As a functor, $H^*(\cdot; G) : \text{Top}_{\text{rel}} \rightarrow \text{Ab}_{\mathbb{Z}}$ is the composition of three functors,

$$\text{Top}_{\text{rel}} \xrightarrow{C_*} \text{Chain} \xrightarrow{\text{Hom}(\cdot, G)} \text{Cochain} \xrightarrow{H_*} \text{Ab}_{\mathbb{Z}},$$

one of which is contravariant, thus $H^*(\cdot; G)$ is also contravariant. Concretely, this means that continuous maps of pairs $f : (X, A) \rightarrow (Y, B)$ induce “pullback” homomorphisms

$$f^* : H^n(Y, B; G) \rightarrow H^n(X, A; G)$$

for every $n \in \mathbb{Z}$. These maps are induced by the chain map $f^* : C^*(Y, B; G) \rightarrow C^*(X, A; G)$ defined by

$$(f^*\varphi)(c) := \varphi(f_*c) \quad \text{for } \varphi \in C^n(Y, B; G), \quad c \in C_n(X, A).$$

By the previous algebraic discussion, there is a natural pairing

$$H^*(X, A; G) \otimes H_*(X, A) \rightarrow G : [\varphi] \otimes [c] \mapsto \langle [\varphi], [c] \rangle := \varphi(c),$$

which we call the **evaluation** of the cohomology class $[\varphi]$ on the homology class $[c]$, and it satisfies

$$(45.7) \quad \langle f^*[\varphi], [c] \rangle = \langle [\varphi], f_*[c] \rangle \quad \text{for } [\varphi] \in H^*(Y, B; G), \quad [c] \in H_*(X, A), \quad (X, A) \xrightarrow{f} (Y, B).$$

More generally, if G is a module over a commutative ring R with unit, one can define an R -bilinear pairing

$$H^*(X, A; G) \otimes_R H_*(X, A; R) \rightarrow G$$

as in Remark 45.4.

Let us conclude this lecture with two straightforward but revealing computations of $H^n(X; G)$ for particular values of n . We start with the case $n = 0$.

For any space X , $C^{-1}(X; G) = 0$, thus $H^0(X; G)$ is simply the kernel of the map $C^0(X; G) \xrightarrow{\delta} C^1(X; G)$, 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 $\gamma : I \rightarrow X$, (45.6) 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 45.6. *For any space X and abelian group G , there is a canonical isomorphism*

$$H^0(X; G) \cong \prod_{\pi_0(X)} G.$$

□

REMARK 45.7. This proves that $H^0(X; G) \cong H_0(X; G)$ if X has only finitely-many path-components, but otherwise $H^0(X; G)$ is larger than $H_0(X; G)$. Indeed, for any collection of groups $\{G_\alpha\}_{\alpha \in I}$, the direct sum $\bigoplus_{\alpha \in I} G_\alpha$ can be identified with the subgroup of the direct product $\prod_{\alpha \in I} G_\alpha$ consisting of tuples $\{g_\alpha\}_{\alpha \in I}$ that have at most finitely-many nonzero coordinates. For example, if the index set I is \mathbb{N} and $G_\alpha = \mathbb{Z}_2$ for every $\alpha \in I$, then $\bigoplus_{\alpha \in I} G_\alpha$ is countably infinite but $\prod_{\alpha \in I} G_\alpha$ is uncountable.

The second computation relates $H^1(X; G)$ to $\pi_1(X)$; we shall give a brief sketch and leave the details as exercises. Assume X is a path-connected space, and identify Δ^1 with $I = [0, 1]$ as usual so that singular 1-cochains $\varphi \in C^1(X; G)$ can be interpreted as functions from the set of paths $\{\gamma : I \rightarrow X\}$ to G .

EXERCISE 45.8. Show that a singular 1-cochain $\varphi \in C^1(X; G)$ is a cocycle if and only if it satisfies both of the following:

- (i) For all paths $\gamma : I \rightarrow X$, $\varphi(\gamma) \in G$ 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 45.9. Show that a singular 1-cochain $\varphi \in C^1(X; G)$ is a coboundary if and only if there exists a function⁶⁶ $\psi : X \rightarrow G$ such that for all paths $\gamma : I \rightarrow X$, $\varphi(\gamma) = \psi(\gamma(1)) - \psi(\gamma(0))$.

EXERCISE 45.10. 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

$$\Psi_\varphi([\gamma]) = \varphi(\gamma) \quad \text{for } x \xrightarrow{\gamma} x.$$

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 45.9. For surjectivity, it might help to observe that since $H_1(X)$ is the abelianization of $\pi_1(X, x)$ and G is abelian, $\text{Hom}(\pi_1(X, x), G) = \text{Hom}(H_1(X), G)$, so the map $\Psi : H^1(X; G) \rightarrow \text{Hom}(\pi_1(X, x), G)$ can then be identified with

$$H^1(X; G) \rightarrow \text{Hom}(H_1(X), G) : [\varphi] \mapsto \langle [\varphi], \cdot \rangle.$$

⁶⁶Note that since G is not assumed to have any topology in this discussion, there is no continuity assumption on the function $\psi : X \rightarrow G$.

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 $C_1(X) \rightarrow G$. Use the fact that $0 \rightarrow Z_1 \hookrightarrow C_1(X) \xrightarrow{\partial} B_0 \rightarrow 0$ is a split exact sequence. (Why?)

46. 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 46.1. An **axiomatic cohomology theory** h^* is a contravariant functor

$$\mathbf{Top}_{\text{rel}} \rightarrow \mathbf{Ab}_{\mathbb{Z}} : (X, A) \mapsto h^*(X, A) = \bigoplus_{n \in \mathbb{Z}} h^n(X, A),$$

together with a natural transformation δ^* from the functor $\mathbf{Top}_{\text{rel}} \rightarrow \mathbf{Ab} : (X, A) \mapsto h^n(A)$ to the functor $\mathbf{Top}_{\text{rel}} \rightarrow \mathbf{Ab} : (X, A) \mapsto h^{n+1}(X, A)$ for each $n \in \mathbb{Z}$ such that the following axioms are satisfied:

- (EXACTNESS) For all pairs (X, A) with inclusion maps $i : A \hookrightarrow X$ and $j : (X, \emptyset) \hookrightarrow (X, A)$, 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.

- (HOMOTOPY) For any two homotopic maps $f, g : (X, A) \rightarrow (Y, B)$, the induced morphisms $f^*, g^* : h^*(Y, B) \rightarrow h^*(X, A)$ are identical.
- (EXCISION) For any pair (X, A) and any subset $B \subset X$ with closure in the interior of A , the inclusion $(X \setminus B, A \setminus B) \hookrightarrow (X, A)$ induces an isomorphism

$$h^*(X, A) \xrightarrow{\cong} h^*(X \setminus B, A \setminus B).$$

- (DIMENSION) For any space $\{\text{pt}\}$ containing only one point, $h^n(\{\text{pt}\}) = 0$ for all $n \neq 0$.
- (ADDITIVITY) For any collection of spaces $\{X_\alpha\}_{\alpha \in J}$ with inclusion maps $i_\alpha : X_\alpha \hookrightarrow \coprod_{\beta \in J} X_\beta$, the induced homomorphisms $i_\alpha^* : h^*\left(\coprod_{\beta \in J} X_\beta\right) \rightarrow h^*(X_\alpha)$ determine an isomorphism

$$\prod_{\alpha \in J} i_\alpha^* : h^*\left(\coprod_{\beta \in J} X_\beta\right) \xrightarrow{\cong} \prod_{\alpha \in J} h^*(X_\alpha).$$

The group $h^0(\{\text{pt}\})$ is called the **coefficient group** of the theory.

THEOREM 46.2. For any abelian group G , the singular cohomology $H^*(\cdot; G)$ is an axiomatic cohomology theory with coefficient group G .

PROOF. Exactness follows from the fact that if we dualize the usual short exact sequence of singular chain complexes $0 \rightarrow C_*(A) \xrightarrow{i_*} C_*(X) \xrightarrow{j_*} C_*(X, A) \rightarrow 0$, then the resulting sequence of chain maps

$$(46.1) \quad 0 \longleftarrow C^*(A; G) \xleftarrow{i^*} C^*(X; G) \xleftarrow{j^*} C^*(X, A; G) \longleftarrow 0$$

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

$$j^* : \left\{ \varphi : \mathcal{K}_n(X) \rightarrow G \mid \varphi|_{\mathcal{K}_n(A)} = 0 \right\} \hookrightarrow \{ \varphi : \mathcal{K}_n(X) \rightarrow G \},$$

and i^* becomes the restriction map

$$i^* : \{\varphi : \mathcal{K}_n(X) \rightarrow G\} \rightarrow \{\varphi : \mathcal{K}_n(A) \rightarrow G\} : \varphi \mapsto \varphi|_{\mathcal{K}_n(A)},$$

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 in the next lecture that not every short exact sequence of abelian groups remains exact after it is dualized. But this one does. As a result, (46.1) is what we may sensibly call a short exact sequence of *cochain* complexes, which is the same thing as a short exact sequence of chain complexes except that the coboundary operator raises degrees instead of lowering them. The usual diagram-chasing argument therefore produces from this a long exact sequence of the homology groups of the complexes, with a connecting homomorphism that raises the degree by 1.

The main reason for the homotopy axiom is Proposition 45.2 in the previous lecture, which implies that if the two chain maps $f_*, g_* : C_*(X, A) \rightarrow C_*(Y, B)$ are chain homotopic, then so are the two chain maps $f^*, g^* : C^*(Y, B; G) \rightarrow C^*(X, A; G)$.

For excision, recall from Theorem 28.2 that if $B \subset \bar{B} \subset \mathring{A} \subset A \subset X$, then the inclusion $i : (X \setminus B, A \setminus B) \hookrightarrow (X, A)$ induces a chain homotopy equivalence $i_* : C_*(X \setminus B, A \setminus B) \rightarrow C_*(X, A)$, meaning in particular that there is a chain map $\rho_* : C_*(X, A) \rightarrow C_*(X \setminus B, A \setminus B)$ such that $\rho_* i_*$ and $i_* \rho_*$ are each chain homotopy equivalent to the identity. Dualizing both i_* and ρ_* then produces chain maps $i^* : C^*(X, A; G) \rightarrow C^*(X \setminus B, A \setminus B; G)$ and $\rho^* : C^*(X \setminus B, A \setminus B; G) \rightarrow C^*(X, A; G)$ such that by Proposition 45.2, $i^* \rho^*$ 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 $\sigma_n \in \mathcal{K}_n(\{\text{pt}\})$ for each $n \geq 0$, giving canonical isomorphisms

$$C^n(\{\text{pt}\}; G) \xrightarrow{\cong} G : \varphi \mapsto \varphi(\sigma_n).$$

The map $\delta : C^n(\{\text{pt}\}; G) \rightarrow C^{n+1}(\{\text{pt}\}; G)$ 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 $H^n(\{\text{pt}\}; G) = 0$. For $n > 0$ odd, we instead have $\ker \delta_n = 0$ and thus $H^n(\{\text{pt}\}; G) = 0$. The only special case is $n = 0$, for which $H^0(\{\text{pt}\}; G) = \ker \delta_0 = G$.

The additivity axiom is a straightforward consequence of the fact that since no individual singular simplex can have image in more than one component of a disjoint union, the chain complex $C_*(\coprod_{\beta} X_{\beta})$ splits naturally into a direct sum of chain complexes $\bigoplus_{\beta} C_*(X_{\beta})$. Dualizing then changes the direct sum to a direct product as we saw in the computation of $H^0(X; G)$ in the previous lecture. We leave the details as an exercise. \square

EXERCISE 46.3. Describe a cohomological version of the “braid” diagram in Lecture 30 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.

Reduced cohomology. Every cohomology theory h^* also has a **reduced** version, which is again defined in terms of the unique map

$$\epsilon : X \rightarrow \{\text{pt}\}.$$

Choosing any embedding $i : \{\text{pt}\} \rightarrow X$, the fact that $\epsilon \circ i$ is the identity map implies that

$$(\epsilon \circ i)^* = i^* \epsilon^* : h^*(\{\text{pt}\}) \rightarrow h^*(\{\text{pt}\})$$

is also the identity, so $\epsilon^* : h^*(\{\text{pt}\}) \rightarrow h^*(X)$ 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

$$0 \longrightarrow h^*(\{\text{pt}\}) \xrightarrow{\epsilon^*} h^*(X) \longrightarrow \tilde{h}^*(X) \longrightarrow 0,$$

implying 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 $\epsilon^* : h^*(\{\text{pt}\}) \rightarrow h^*(X)$ is thus an isomorphism, so its cokernel is trivial:

THEOREM 46.4. *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 46.5. *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

EXERCISE 46.6. Show that $\tilde{H}^*(X; G)$ is also the cohomology of the *augmented* chain complex

$$\dots \longrightarrow C_2 \xrightarrow{\partial_2} C_1 \xrightarrow{\partial_1} C_0 \xrightarrow{\epsilon} G \longrightarrow 0 \longrightarrow 0 \longrightarrow \dots$$

described in Remark 28.18.

EXERCISE 46.7. Adapt the proof of Theorem 30.13 to prove that for any axiomatic cohomology theory h^* and any space X , there is an isomorphism $\tilde{h}^n(X) \rightarrow \tilde{h}^{n+1}(SX)$ for every $n \in \mathbb{Z}$.

EXERCISE 46.8. For any axiomatic cohomology theory h^* and two spaces X and Y with maps $\epsilon_X : X \rightarrow \{\text{pt}\}$ and $\epsilon_Y : Y \rightarrow \{\text{pt}\}$, 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

$$(\epsilon_X^*, \epsilon_Y^*) : h^*(\{\text{pt}\}) \rightarrow h^*(X) \times h^*(Y).$$

Then apply this in the case $X = Y = \{\text{pt}\}$ to identify $\tilde{h}^0(\{\text{pt}\} \amalg \{\text{pt}\})$ with the cokernel of the diagonal map $G \rightarrow G \times G$, where $G = h^0(\{\text{pt}\})$. 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 46.9. Combine the previous two exercises to prove by induction on $n \in \mathbb{N}$ that for any axiomatic cohomology theory h^* with coefficient group G ,

$$h^k(S^n) \cong \begin{cases} G & \text{if } k = 0 \text{ or } k = n, \\ 0 & \text{otherwise.} \end{cases}$$

EXERCISE 46.10. Adapt the proof of Theorem 30.21 to prove that for any axiomatic cohomology theory h^* and any good pair (X, A) , there is a natural isomorphism

$$h^*(X, A) \cong \tilde{h}^*(X/A).$$

REMARK 46.11. You may by now be getting the impression that cohomology is always isomorphic to homology, especially in light of the computation above for S^n . There is a grain of truth in this, but the whole story is more complicated: e.g. we will see in the next lecture that $H^*(X; G)$ is fully determined by $H_*(X)$, but it is not always the same, especially if $H_*(X)$ 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 $H_*(\cdot; G)$, not every axiomatic cohomology theory h^* has a corresponding axiomatic homology theory h_* (cf. Remark 46.20 at the end of this lecture).

The Mayer-Vietoris sequence. One can use a diagram-chase as in Exercise 31.3 to derive from the axioms a Mayer-Vietoris sequence for any axiomatic cohomology theory, but for singular cohomology it also can be seen more directly. Indeed, suppose $A, B \subset X$ are subsets whose interiors cover X , let

$$a : A \cap B \hookrightarrow A, \quad b : A \cap B \hookrightarrow B$$

denote the obvious continuous inclusions of spaces, and

$$\alpha : C_*(A) \hookrightarrow C_*(A) + C_*(B), \quad \beta : C_*(B) \hookrightarrow C_*(A) + C_*(B)$$

the obvious inclusions of subgroups of $C_*(X)$. The Mayer-Vietoris sequence in singular homology was derived in Lecture 31 from a short exact sequence of chain complexes in the form

$$0 \longrightarrow C_*(A \cap B) \xrightarrow{(a_*, -b_*)} C_*(A) \oplus C_*(B) \xrightarrow{\alpha \oplus \beta} C_*(A) + C_*(B) \longrightarrow 0.$$

Applying to this the functor $\text{Hom}(\cdot; G)$ and using the natural isomorphism

$$\begin{aligned} \text{Hom}(C_*(A), G) \oplus \text{Hom}(C_*(B), G) &\cong \text{Hom}(C_*(A) \oplus C_*(B), G) \\ (\varphi, \psi) &\mapsto \varphi \oplus \psi \end{aligned}$$

transforms it into to the sequence

$$0 \longleftarrow C^*(A \cap B; G) \xleftarrow{a^* \oplus (-b^*)} C^*(A; G) \oplus C^*(B; G) \xleftarrow{(\alpha^*, \beta^*)} C^*(A + B; G) \longleftarrow 0,$$

where we are abbreviating

$$C^*(A + B; G) := \text{Hom}(C_*(A) + C_*(B), G).$$

The dual maps

$$\begin{aligned} a^* : C^*(A; G) &\rightarrow C^*(A \cap B; G), & b^* : C^*(B; G) &\rightarrow C^*(A \cap B; G), \\ \alpha^* : C^*(A + B; G) &\rightarrow C^*(A; G), & \beta^* : C^*(A + B; G) &\rightarrow C^*(B; G) \end{aligned}$$

are all canonical restriction maps, e.g. α^* replaces a homomorphism $\varphi : C_*(A) + C_*(B) \rightarrow G$ with its restriction to the subgroup $C_*(A)$. It is now an easy exercise to check that the dualized sequence is also exact. To make use of this, we need to identify the homology of the cochain complex $C^*(A + B; G)$ with something more familiar. By Lemma 28.1, the condition $X = \mathring{A} \cup \mathring{B}$ guarantees that the inclusion $j : C_*(A) + C_*(B) \hookrightarrow C_*(X)$ is a chain homotopy equivalence, so by Proposition 45.2, the dual map

$$j^* : C^*(X; G) \rightarrow C^*(A + B; G)$$

is also a chain homotopy equivalence and therefore induces an isomorphism

$$H^*(X; G) \xrightarrow{\cong} H_*(C^*(A + B; G)).$$

Combining this with the usual diagram-chasing result gives:

THEOREM 46.12 (Mayer-Vietoris sequence for cohomology). *If $A, B \subset X$ are subsets such that $X = \mathring{A} \cup \mathring{B}$ and*

$$i_A : A \cap B \hookrightarrow A, \quad i_B : A \cap B \hookrightarrow B, \quad j_A : A \hookrightarrow X, \quad j_B : B \hookrightarrow X,$$

denote the obvious inclusions, then there exist connecting homomorphisms $\delta^ : H^n(A \cap B; G) \rightarrow H^{n+1}(X; G)$ for every $n \in \mathbb{Z}$ such that the sequence*

$$\begin{aligned} \dots \longleftarrow H^{n+1}(X; G) \xleftarrow{\delta^*} H^n(A \cap B; G) \xleftarrow{i_A^* \oplus (-i_B^*)} H^n(A; G) \oplus H^n(B; G) \\ \xleftarrow{(j_A^*, j_B^*)} H^n(X; G) \xleftarrow{\delta^*} H^{n-1}(A \cap B; G) \longleftarrow \dots \end{aligned}$$

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

EXERCISE 46.13. Adapt the diagram-chasing arguments in Lecture 31 to show that every axiomatic cohomology theory h^* admits a Mayer-Vietoris sequence under the same hypotheses on $X = A \cup B$, and that it also works if h^* is replaced by \tilde{h}^* .

Cellular cohomology. The cellular cohomology of a CW-pair (X, A) with coefficients in G is defined as the cohomology of the cellular chain complex, or equivalently,

$$H_{\text{CW}}^*(X, A; G) := H_*(C_{\text{CW}}^*(X, A; G)),$$

where we define the **cellular cochain complex**

$$C_{\text{CW}}^*(X, A; G) := \text{Hom}(C_{\text{CW}}^{\text{CW}}(X, A), G).$$

This gives a contravariant functor $H_{\text{CW}}^* : \text{CW}_{\text{rel}} \rightarrow \text{Ab}_{\mathbb{Z}}$ that is 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

$$\varphi_\alpha^n \in C_{\text{CW}}^n(X; \mathbb{Z}), \quad \varphi_\alpha^n(e_\beta^n) := \begin{cases} 1 & \text{if } \beta = \alpha, \\ 0 & \text{otherwise.} \end{cases}$$

These generators form a basis of $C_{\text{CW}}^n(X; \mathbb{Z})$ 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 , thus it suffices to write down a formula for $\delta\varphi_\alpha^n \in C_{\text{CW}}^{n+1}(X; \mathbb{Z})$. We have

$$(\delta\varphi_\alpha^n)(e_\beta^{n+1}) = (-1)^{n+1} \varphi_\alpha^n(\partial e_\beta^{n+1}) = (-1)^{n+1} \sum_{e_\gamma^n \subset X} \varphi_\alpha^n([e_\gamma^n : e_\beta^{n+1}]e_\gamma^n) = (-1)^{n+1} [e_\alpha^n : e_\beta^{n+1}],$$

and thus,

$$\delta\varphi_\alpha^n = (-1)^{n+1} \sum_{e_\beta^{n+1} \subset X} [e_\alpha^n : e_\beta^{n+1}] \varphi_\beta^{n+1}.$$

As with homology, cellular cohomology provides a powerful tool for computing arbitrary axiomatic cohomology theories on spaces that have cell decompositions:

THEOREM 46.14. *For any axiomatic cohomology theory h^* with coefficient group G and every CW-pair (X, A) , there exists an isomorphism $H_{CW}^*(X, A; G) \rightarrow h^*(X, A)$, which is natural in the sense that every cellular map $f : (X, A) \rightarrow (Y, B)$ gives rise to a commutative diagram*

$$\begin{array}{ccc} H_{CW}^*(X, A; G) & \xrightarrow{\cong} & h^*(X, A) \\ f^* \uparrow & & f^* \uparrow \\ H_{CW}^*(Y, B; G) & \xrightarrow{\cong} & h^*(Y, B) \end{array}$$

For finite-dimensional complexes, this theorem can be proved in a way that closely parallels the corresponding argument for cellular homology carried out in Lectures 36 and 37. One starts by deriving from the axioms 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 , then using 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 $C_{CW}^n(X; G)$ 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 diagram

$$\begin{array}{ccccccc} & & & & h^n(X^{n-1}) & \xlongequal{\quad} & 0 \\ & & & & \uparrow & & \\ h^{n+1}(X^n) & \leftarrow & h^{n+1}(X^{n+1}) & \xleftarrow{j_{n+1}^*} & h^{n+1}(X^{n+1}, X^n) & \xleftarrow{\delta_n^*} & h^n(X^n) & \xleftarrow{i_n^*} & h^n(X^{n+1}) & \leftarrow & h^n(X^{n+1}, X^n) \\ & & & \swarrow \gamma_n & \uparrow j_n^* & & \uparrow & & & & \parallel & \\ & & & & h^n(X^n, X^{n-1}) & & & & & & & 0 \\ & & & & \uparrow \delta_{n-1}^* & & & & & & & \\ & & & & h^{n-1}(X^{n-1}) & & & & & & & \end{array}$$

in which the diagonal arrow defines maps γ_n so that the sequence

$$h^0(X^0) \xrightarrow{\gamma_0} h^1(X^1, X^0) \xrightarrow{\gamma_1} h^2(X^2, X^1) \longrightarrow \dots$$

becomes a cochain complex. One can check that γ_n is equivalent to the cellular coboundary map $C_{CW}^n(X; G) \xrightarrow{\delta} C_{CW}^{n+1}(X; G)$ under the natural isomorphisms $h^n(X^n, X^{n+1}) \cong C_{CW}^n(X; G)$. The diagram then allows us to deduce that the map $i_n^* : h^n(X^{n+1}) \rightarrow h^n(X^n)$ is injective and j_n^* descends to an isomorphism

$$\ker \gamma_n / \text{im } \gamma_{n-1} \xrightarrow{j_n^*} \text{im } i_n^* \cong h^n(X^{n+1}) \cong h^n(X^{n+2}) \cong \dots,$$

thus giving an isomorphism $H_{CW}^n(X; G) \cong h^n(X)$ if $X = X^N$ for some $N \in \mathbb{N}$ sufficiently large. To handle CW-pairs (X, A) with $A \neq \emptyset$, one carries out this same argument with $h^n(X^n, X^{n-1})$

replaced by $h^n(X^n \cup A, X^{n-1} \cup A)$ and the long exact sequence of (X^n, X^{n-1}) replaced by the sequence of the triple $(X^n \cup A, X^{n-1} \cup A, A)$.

EXERCISE 46.15. Work out the further details of the proof of Theorem 46.14 for finite-dimensional CW-pairs.

Some additional arguments are needed if the CW-pair (X, A) is infinite dimensional. In Lecture 38, we handled this for singular homology by viewing any CW-complex as the direct limit of its skeleta and establishing sufficient conditions for singular homology to behave continuously under direct limits. You may recall that this issue was a little bit subtle: some topological condition is required on a direct system of spaces in order for the singular chain complex functor to behave continuously under direct limits, but fortunately, all CW-complexes satisfy the condition (cf. Prop. 39.2). For cohomology, the situation is worse: since $H^*(\cdot; G)$ is contravariant, the cohomology groups of the skeleta $\{H^*(X^n; G)\}_{n \geq 0}$ define an *inverse* system of graded abelian groups, and we've seen in Lecture 44 that in contrast to direct systems, the algebraic functor taking chain complexes to their homology groups does not always behave continuously under inverse limits. This does not mean that the situation is hopeless, but it does make things more complicated: it means that one must introduce another “derived” functor (analogous to Tor) to account for the nonexactness of the functor that takes each inverse system to its inverse limit. This is done in [Mil62] in the more general context of axiomatic cohomology theories, but we will not discuss it any further here since we do not need that level of generality. For our purposes, it will suffice if we can prove the general case of the isomorphism $H_{CW}^*(X, A; G) \cong H^*(X, A; G)$, and in the next lecture we will see a cheap way of deriving this from facts already proven about $H_*(X, A; G)$ via a universal coefficient theorem.

Other cohomology theories. Finally, I would like to give brief sketches of two axiomatic cohomology theories other than singular cohomology. They will demonstrate in particular that the two properties of $H^*(\cdot; G)$ we discussed at the end of the previous lecture,

$$H^0(X; G) \cong \prod_{\pi_0(X)} G \quad \text{and} \quad H^1(X; G) \cong \text{Hom}(\pi_1(X), G),$$

do not follow from the axioms, but are distinctive to the singular theory.

We begin with Čech cohomology. Recall from Lecture 44 that for any open cover $(\mathfrak{U}, \mathfrak{U}_A) \in \mathcal{O}(X, A)$ of a pair of spaces (X, A) , one can define its *nerve* $\mathcal{N}(\mathfrak{U}, \mathfrak{U}_A)$, which consists of a simplicial complex $\mathcal{N}(\mathfrak{U})$ and subcomplex $\mathcal{N}(\mathfrak{U}_A) \subset \mathcal{N}(\mathfrak{U})$. The *ordered simplicial cohomology* of the nerve with coefficients in G is defined in the obvious way as the cohomology of the ordered simplicial chain complex

$$H_o^*(\mathcal{N}(\mathfrak{U}, \mathfrak{U}_A); G) := H^*(C_*^o(\mathfrak{U}, \mathfrak{U}_A); G).$$

Recall moreover that the set $\mathcal{O}(X, A)$ of all open coverings of (X, A) is a directed set with respect to refinement, and any refinement $(\mathfrak{U}', \mathfrak{U}'_A)$ of $(\mathfrak{U}, \mathfrak{U}_A)$ gives rise to a simplicial map $F : \mathcal{N}(\mathfrak{U}', \mathfrak{U}'_A) \rightarrow \mathcal{N}(\mathfrak{U}, \mathfrak{U}_A)$ whose induced chain map

$$F_* : C_*^o(\mathcal{N}(\mathfrak{U}', \mathfrak{U}'_A)) \rightarrow C_*^o(\mathcal{N}(\mathfrak{U}, \mathfrak{U}_A))$$

is unique up to chain homotopy. It follows via Proposition 45.2 that dualizing this map produces a map of cochain complexes that is also unique up to chain homotopy, producing a canonically defined morphism

$$F^* : H_o^*(\mathcal{N}(\mathfrak{U}, \mathfrak{U}_A)) \rightarrow H_o^*(\mathcal{N}(\mathfrak{U}', \mathfrak{U}'_A)).$$

Notice what has happened as a result of dualization: the collection of simplicial homology groups $\{H_o^*(\mathcal{N}(\mathfrak{U}, \mathfrak{U}_A); G)\}_{(\mathfrak{U}, \mathfrak{U}_A) \in \mathcal{O}(X, A)}$ in Lecture 44 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

$$\check{H}^*(X, A; G) := \varinjlim \{H_o^*(\mathcal{N}(\mathfrak{U}, \mathfrak{U}_A); G)\}_{(\mathfrak{U}, \mathfrak{U}_A) \in \mathcal{O}(X, A)}.$$

There is a huge technical advantage in the fact that $\check{H}^*(X, A; G)$ is defined via a direct limit instead of an inverse limit: exactness of sequences is preserved under direct limits (cf. Prop. 39.1), and one can use this to prove that unlike $\check{H}_*(\cdot; G)$, the cohomology $\check{H}^*(\cdot; G)$ satisfies the exactness axiom without any restrictions. It also satisfies all the other axioms:

THEOREM 46.16 (see [ES52, Spa95]). *For any abelian group G , the Čech cohomology $\check{H}(\cdot; G) : \text{Top}_{\text{rel}} \rightarrow \text{Ab}_{\mathbb{Z}}$ 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 44.17 that whenever $\mathfrak{U} \in \mathcal{O}(X)$ is an open covering of a connected space X , the nerve $\mathcal{N}(\mathfrak{U})$ is also connected, thus $H^0(\mathcal{N}(\mathfrak{U}); G) \cong G$. One can deduce from this that if X is connected, $\check{H}^0(X; G) \cong G$ and the reduced Čech cohomology of X in degree zero vanishes. Exercise 46.7 then implies $\check{H}^1(SX; G) = 0$. But if X has more than one path-component, then $\check{H}^0(X; G)$ and $H^1(SX; G)$ are both nontrivial; the latter is isomorphic to $\text{Hom}(\pi_1(SX), G)$ since the suspension SX is always path-connected, thus SX is an example of a space for which $\check{H}^1(SX; G) \not\cong \text{Hom}(\pi_1(SX), G)$.

The **Alexander-Spanier cohomology** is yet another theory that satisfies all of the Eilenberg-Steenrod axioms but is based on a different idea of how to detect topological information. Let us describe the absolute version.

For integers $n \geq 0$, let $\bar{C}^n(X; G)$ denote the additive abelian group of equivalence classes of functions

$$\varphi : X^{n+1} = \underbrace{X \times \dots \times X}_{n+1} \rightarrow G,$$

where we say $\varphi \sim \psi$ whenever φ and ψ are identical on some neighborhood of the **diagonal**

$$\Delta := \{(x, \dots, x) \in X^{n+1} \mid x \in X\}.$$

The group operation on $\bar{C}^n(X; G)$ is defined via pointwise addition, so for two equivalence classes $[\varphi], [\psi] \in \bar{C}^n(X; G)$, $[\varphi] + [\psi] \in \bar{C}^n(X; G)$ is represented by the function $\varphi + \psi : X^{n+1} \rightarrow G$ defined by

$$(\varphi + \psi)(x_0, \dots, x_n) := \varphi(x_0, \dots, x_n) + \psi(x_0, \dots, x_n).$$

You should take a moment to assure yourself that the equivalence class of $\varphi + \psi$ is independent of the choice of representatives $\varphi \in [\varphi]$ and $\psi \in [\psi]$. Note that since the group G is not assumed to have a topology, there is no continuity condition on the functions $X^{n+1} \rightarrow G$ representing elements of $\bar{C}^n(X; G)$. Instead, this group detects the topology of X via the notion of “neighborhoods of $\Delta \subset X^{n+1}$ ” that is used to define the equivalence relation.

To make the collection of groups $\bar{C}^n(X; G)$ 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 group of $(n+1)$ -functions to the group of $(n+2)$ -functions such that $\delta^2 = 0$, and it preserves the subgroup of functions that vanish near the diagonal, thus it descends to a coboundary homomorphism

$$\delta : \bar{C}^n(X; G) \rightarrow \bar{C}^{n+1}(X; G).$$

Extending this to all $n \in \mathbb{Z}$ by defining $\bar{C}^n(X; G) = 0$ for $n < 0$, we obtain a cochain complex $(\bar{C}^*(X; G), \delta)$, and its homology is the Alexander-Spanier cohomology of X , denoted by

$$\bar{H}^*(X; G) := H_*(\bar{C}^*(X; G)).$$

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

$$f^* : \bar{C}^*(Y; G) \rightarrow \bar{C}^*(X; G) : \varphi \mapsto \varphi \circ (f \times \dots \times f),$$

thus inducing homomorphisms $f^* : \bar{H}^*(Y; G) \rightarrow \bar{H}^*(X; G)$. With some more effort, one can also define relative groups $\bar{H}^*(X, A; G)$ and prove that \bar{H}^* satisfies all of the Eilenberg-Steenrod axioms for a cohomology theory.⁶⁷ I recommend thinking through the following two exercises in order to gain some intuition on how $\bar{H}^*(X; G)$ measures the topology of X .

EXERCISE 46.17. Show that for any space X , $\bar{H}^0(X; G)$ is a direct product of copies of G , one for each connected component of X .

EXERCISE 46.18. Show by explicit computation that $\bar{H}^1(\mathbb{R}; G) = 0$ and $\bar{H}^1(S^1; G) \cong G$.
Hint: Consider the map $\bar{H}^1(X; G) \rightarrow G : [\varphi] \mapsto \sum_{k=1}^N \varphi(\gamma(t_k), \gamma(t_{k-1}))$ for some loop $\gamma : [0, 1] \rightarrow X$ and partition $0 = t_0 < t_1 < \dots < t_N = 1$. It is well defined because φ represents a cocycle. (Why?)

Both \check{H}^* and \bar{H}^* satisfy an “extra” axiom that singular cohomology does not, the so-called **continuity** axiom (cf. Theorem 44.22). Since both are contravariant functors, any inverse system of spaces $\{X_\alpha, \varphi_{\beta\alpha}\}$ gives rise to direct systems of cohomology groups, e.g. in the Alexander-Spanier theory, we obtain the system $\{\bar{H}^*(X_\alpha; G), \varphi_{\beta\alpha}^*\}$. It turns out that whenever the spaces X_α are all compact and Hausdorff, there is an isomorphism

$$\bar{H}^*(\varinjlim \{X_\alpha\}; G) \cong \varinjlim \{\bar{H}^*(X_\alpha; G)\},$$

and the same is true for Čech cohomology. It is not hard to find examples (e.g. involving spaces that are connected but not path-connected) for which this is not true in singular cohomology, cf. Exercise 44.23.

REMARK 46.19. One can show that every compact Hausdorff space is an inverse limit of some inverse system of compact Hausdorff spaces homotopy equivalent to CW-complexes. It follows that up to isomorphism, there is only one cohomology theory on compact Hausdorff spaces that satisfies all of the Eilenberg-Steenrod axioms plus continuity. In particular, $\check{H}^*(X; G) \cong \bar{H}^*(X; G)$ whenever X is compact and Hausdorff, though both may be different from $H^*(X; G)$. (This result can be generalized beyond compact spaces using sheaf cohomology; details are carried out in [Spa95, Chapter 6].)

REMARK 46.20. It is interesting to note that $\bar{C}^*(X; G)$ is not in any obvious way the dual complex of a chain complex, thus it is far from obvious at this stage what the definition of “Alexander-Spanier homology” might be. A corresponding homology theory was defined in an appendix of [Spa48], but its definition is much more complicated, requiring inverse limits, and as a result it suffers from the same drawbacks as Čech homology, i.e. it fails to satisfy the exactness axiom of Eilenberg-Steenrod.

⁶⁷For a good exposition of the details, see [Spa95, §6.4–6.5].

47. Universal coefficients and the Ext functor

The goal of this lecture is to understand how $H^*(X; G)$ is determined in general by $H_*(X)$, and in particular, under what circumstances the natural homomorphism

$$(47.1) \quad h : H^n(X; G) \rightarrow \text{Hom}(H_n(X), G) : \varphi \mapsto \langle \varphi, \cdot \rangle$$

is an isomorphism. The answer closely parallels the universal coefficient theorem for homology that we discussed in Lectures 41 and 42.

We shall assume throughout this lecture that R is a commutative ring with unit, G is an R -module and C_* is a chain complex of free R -modules. Recall that in the homological story, the failure of the natural map $H_*(C_*) \otimes_R G \rightarrow H_*(C_* \otimes_R G)$ to be an isomorphism in some cases is closely related to the fact that $\otimes_R G$ is not an *exact functor*, i.e. it does not preserve exactness when applied to short exact sequences, though the failure of exactness is limited to the first nontrivial term in the sequence. We should therefore begin the cohomological discussion by investigating the analogous question for the functor $\text{Hom}_R(\cdot, G) : \text{Mod}^R \rightarrow \text{Mod}^R$. For instance, applying $\text{Hom}(\cdot, \mathbb{Z}_2)$ to the short exact sequence

$$0 \rightarrow \mathbb{Z} \xrightarrow{\cdot 2} \mathbb{Z} \xrightarrow{\text{pr}} \mathbb{Z}_2 \rightarrow 0$$

gives

$$0 \leftarrow \mathbb{Z}_2 \xleftarrow{0} \mathbb{Z}_2 \xleftarrow{1} \mathbb{Z}_2 \leftarrow 0,$$

which is exact everywhere except at the leftmost nontrivial term, since the zero map $\mathbb{Z}_2 \rightarrow \mathbb{Z}_2$ is not surjective. It is not hard to show that exactness at the other terms must always hold:

EXERCISE 47.1. Show that if $A \xrightarrow{i} B \xrightarrow{j} C \rightarrow 0$ is an exact sequence of R -modules, then for every R -module G , the sequence $0 \rightarrow \text{Hom}_R(C, G) \xrightarrow{j^*} \text{Hom}_R(B, G) \xrightarrow{i^*} \text{Hom}_R(A, G)$ is also exact.

Given a short exact sequence $0 \rightarrow A \xrightarrow{i} B \xrightarrow{j} C \rightarrow 0$, the failure of the dualized sequence $0 \rightarrow \text{Hom}_R(C, G) \xrightarrow{j^*} \text{Hom}_R(B, G) \xrightarrow{i^*} \text{Hom}_R(A, G) \rightarrow 0$ to be exact in general at the last term has an easy explanation: e.g. if we view the injective map $i : A \hookrightarrow B$ as the inclusion of a submodule $A \subset B$, then exactness at $\text{Hom}_R(A, G)$ would mean that

$$i^* : \text{Hom}_R(B, G) \rightarrow \text{Hom}_R(A, G) : \beta \mapsto \beta \circ i = \beta|_A$$

is surjective, which is true if and only if every homomorphism $A \rightarrow G$ can be extended to a homomorphism $B \rightarrow G$. But it is easy to come up with examples where this extension problem cannot be solved, such as extending

$$2\mathbb{Z} \rightarrow \mathbb{Z} : m \mapsto \frac{1}{2}m$$

over the larger group \mathbb{Z} that contains $2\mathbb{Z}$. Notice, incidentally, that the subgroup $2\mathbb{Z} \subset \mathbb{Z}$ also appears in our favorite example of a short exact sequence that does not split, namely $0 \rightarrow 2\mathbb{Z} \hookrightarrow \mathbb{Z} \rightarrow \mathbb{Z}_2 \rightarrow 0$. It should not surprise you to learn that the extension problem can always be solved in the split case:

EXERCISE 47.2 (cf. Exercise 41.7). Show that if $0 \rightarrow A \xrightarrow{i} B \xrightarrow{j} C \rightarrow 0$ is a split exact sequence of R -modules, then for every R -module G , the sequence $0 \rightarrow \text{Hom}_R(C, G) \xrightarrow{j^*} \text{Hom}_R(B, G) \xrightarrow{i^*} \text{Hom}_R(A, G) \rightarrow 0$ is also split exact.

In category-theoretic terms, Exercise 47.1 shows that $\text{Hom}_R(\cdot, G) : \text{Mod}^R \rightarrow \text{Mod}^R$ is what is called a **left-exact functor**, but like the functor $\otimes_R G$ studied in Lecture 41, it is not generally exact. The failure of $\otimes_R G$ to be exact was measured by a derived functor called Tor^R , and we will now introduce another derived functor that plays the analogous role for $\text{Hom}_R(\cdot, G)$.

Suppose A and G are R -modules, and (F_*, f_*) is a projective resolution of A . We showed in Lecture 41 that projective resolutions always exist and they are unique up to chain homotopy equivalence. We can then treat the exact sequence

$$\dots \longrightarrow F_2 \xrightarrow{f_2} F_1 \xrightarrow{f_1} F_0 \xrightarrow{f_0} A \longrightarrow 0$$

as a chain complex of R -modules and apply $\text{Hom}_R(\cdot, G)$ to it, producing a cochain complex

$$\dots \longleftarrow \text{Hom}_R(F_2, G) \xleftarrow{f_2^*} \text{Hom}_R(F_1, G) \xleftarrow{f_1^*} \text{Hom}_R(F_0, G) \xleftarrow{f_0^*} \text{Hom}_R(A, G) \longleftarrow 0.$$

By Proposition 45.2 and the uniqueness of projective resolutions, this dual complex is also unique up to chain homotopy equivalence. In fact, for any R -module homomorphism $\varphi : A \rightarrow A'$ with choices of projective resolutions (F_*, f_*) for A and (F'_*, g'_*) for A' , Proposition 41.13 gives a chain map from (F_*, f_*) to (F'_*, g'_*) which can be dualized, giving a chain map between cochain complexes

$$(47.2) \quad \begin{array}{ccccccc} \dots & \longleftarrow & \text{Hom}_R(F_2, G) & \xleftarrow{f_2^*} & \text{Hom}_R(F_1, G) & \xleftarrow{f_1^*} & \text{Hom}_R(F_0, G) & \xleftarrow{f_0^*} & \text{Hom}_R(A, G) \\ & & \varphi_2^* \uparrow & & \varphi_1^* \uparrow & & \varphi_0^* \uparrow & & \varphi^* \uparrow \\ \dots & \longleftarrow & \text{Hom}_R(F'_2, G) & \xleftarrow{g_2^*} & \text{Hom}_R(F'_1, G) & \xleftarrow{g_1^*} & \text{Hom}_R(F'_0, G) & \xleftarrow{g_0^*} & \text{Hom}_R(A', G) \end{array}$$

that is also unique up to chain homotopy. It follows that we can associate to A the (co-)homology groups

$$H^n(F_*; G) = H_n(\text{Hom}_R(F_*, G)) = \ker f_{n+1}^* / \text{im } f_n^*,$$

which depend on A but are independent of the choice of projective resolution (F_*, f_*) up to canonical isomorphisms, and moreover, they behave as contravariant functors since homomorphisms $\varphi : A \rightarrow A'$ induce maps

$$\varphi_n^* : H^n(F'_*; G) \rightarrow H^n(F_*; G)$$

via the chain map in (47.2). The left-exactness of $\text{Hom}_R(\cdot, G)$ implies that $H_0(\text{Hom}_R(F_*, G)) = 0$, and if R is a principal ideal domain, then one can always choose a projective resolution (F_*, f_*) with $F_n = 0$ for all $n \geq 2$, proving $H_n(\text{Hom}_R(F_*, G)) = 0$ for $n \geq 2$. The remaining homology group in the sequence is the one that measures the failure of $\text{Hom}_R(\cdot, G)$ to be an exact functor, thus it gets a name:

DEFINITION 47.3. For any pair of R -modules A and G and a choice of projective resolution (F_*, f_*) for A , we define

$$\text{Ext}^R(A, G) := H^1(F_*; G) = H_1(\text{Hom}_R(F_*, G)).$$

In the case $R = \mathbb{Z}$, we abbreviate

$$\text{Ext}(A, G) := \text{Ext}^{\mathbb{Z}}(A, G).$$

The discussion above implies that any two choices of projective resolution for A give rise to canonically isomorphic definitions of $\text{Ext}^R(A, G)$, thus we shall not distinguish between them in the notation. Moreover, homomorphisms $\varphi : A \rightarrow A'$ induce natural homomorphisms

$$\varphi^* : \text{Ext}^R(A', G) \rightarrow \text{Ext}^R(A, G),$$

making $\text{Ext}^R(\cdot, G) : \text{Mod}^R \rightarrow \text{Mod}^R$ a contravariant functor.

EXERCISE 47.4 (cf. Exercise 41.17). Show that for any fixed R -module A , $\text{Ext}^R(A, \cdot) : \text{Mod}^R \rightarrow \text{Mod}^R$ is naturally a covariant functor.

As with Tor , Ext^R is not difficult to compute for simple classes of R -modules, which include all finitely-generated abelian groups.

EXERCISE 47.5. Make intelligent choices of projective resolutions in order to prove the following properties of Ext^R :

- (1) $\text{Ext}^R(A, G) = 0$ whenever A is a free R -module;
- (2) $\text{Ext}^R(A \oplus B, G) \cong \text{Ext}^R(A, G) \oplus \text{Ext}^R(B, G)$;
- (3) $\text{Ext}(\mathbb{Z}_k, G)$ is isomorphic for each $k \in \mathbb{N}$ and every abelian group G to the cokernel of the map $G \xrightarrow{\cdot k} G$, i.e. the quotient group G/kG .

The following easy application of this exercise is worth recording for future use:

PROPOSITION 47.6. *For any finitely generated abelian group A , $\text{Ext}(A, \mathbb{Z})$ is isomorphic to the torsion subgroup of A .*

PROOF. By the classification of finitely generated abelian groups, A is isomorphic to the direct sum of some free abelian group F with the torsion subgroup $T \subset A$, and T in turn is a finite direct sum of finite cyclic groups $\mathbb{Z}_{k_1}, \dots, \mathbb{Z}_{k_N}$. Exercise 47.5 then gives

$$\text{Ext}(A, \mathbb{Z}) \cong \text{Ext}(F, \mathbb{Z}) \oplus \left(\bigoplus_{j=1}^N \text{Ext}(\mathbb{Z}_{k_j}, \mathbb{Z}) \right) \cong \bigoplus_{j=1}^N \mathbb{Z}_{k_j} \cong T.$$

□

With this bit of machinery in place, we can now state the universal coefficient theorem for cohomology.

THEOREM 47.7 (universal coefficient theorem). *For any chain complex C_* of free modules over a principal ideal domain R , a fixed R -module G and $n \in \mathbb{Z}$, the map $h : H^n(C_*; G) \rightarrow \text{Hom}(H_n(C_*), G)$ defined as in (47.1) fits into a split exact sequence*

$$0 \longrightarrow \text{Ext}^R(H_{n-1}(C_*), G) \longrightarrow H^n(C_*; G) \xrightarrow{h} \text{Hom}_R(H_n(C_*), G) \longrightarrow 0.$$

Moreover, the sequence (though not its splitting) is natural in the sense that for any chain map $\varphi : A_* \rightarrow B_*$ between two chain complexes of free R -modules, the exact sequences for both fit into a commutative diagram

$$\begin{array}{ccccccc} 0 & \longrightarrow & \text{Ext}^R(H_{n-1}(A_*), G) & \longrightarrow & H^n(A_*; G) & \xrightarrow{h} & \text{Hom}_R(H_n(A_*), G) \longrightarrow 0 \\ & & \varphi_* \uparrow & & \varphi_* \uparrow & & \varphi_* \uparrow \\ 0 & \longrightarrow & \text{Ext}^R(H_{n-1}(B_*), G) & \longrightarrow & H^n(B_*; G) & \xrightarrow{h} & \text{Hom}_R(H_n(B_*), G) \longrightarrow 0, \end{array}$$

where the map $\varphi_* : \text{Ext}^R(H_{n-1}(B_*), G) \rightarrow \text{Ext}^R(H_{n-1}(A_*), G)$ arises from the functoriality of Ext^R , and $\varphi_* : \text{Hom}_R(H_{n-1}(B_*), G) \rightarrow \text{Hom}_R(H_{n-1}(A_*), G)$ is defined by dualizing $\varphi_* : H_{n-1}(A_*) \rightarrow H_{n-1}(B_*)$.

Before going into the proof, let us discuss some consequences. When applied to the singular chain complex of a pair (X, A) , the splitting of the sequence gives an isomorphism

$$H^n(X, A; G) \cong \text{Hom}(H_n(X, A), G) \oplus \text{Ext}(H_{n-1}(X, A), G),$$

revealing that $H^n(X, A; G)$ is determined up to isomorphism by $H_n(X, A)$, $H_{n-1}(X, A)$ and G . More generally, if G is a module over a principal ideal domain R , then we can view $H_*(X, A; R)$ as the homology of a chain complex of free R -modules $C_*(X, A) \otimes R$, and $H^*(X, A; G)$ is the cohomology of the same chain complex due to Exercise 45.5, so the theorem also gives an R -module isomorphism

$$H^n(X, A; G) \cong \text{Hom}_R(H_n(X, A; R), G) \oplus \text{Ext}^R(H_{n-1}(X, A; R), G),$$

which reduces to the same statement again if $R = \mathbb{Z}$. All of this applies equally well to cellular or simplicial chain complexes, since these are also freely generated.

There is a particularly appealing corollary whenever R and G are both chosen to be a field \mathbb{K} . All modules in the picture are then vector spaces over \mathbb{K} , which are automatically free, so the vector space $\text{Ext}^{\mathbb{K}}(H_{n-1}(X, A; \mathbb{K}), \mathbb{K})$ is trivial by Exercise 47.5, and we conclude that cohomology is just the dual vector space of homology:

COROLLARY 47.8. *For any field \mathbb{K} and any pair of spaces (X, A) , the natural map*

$$H^n(X, A; \mathbb{K}) \rightarrow \text{Hom}_{\mathbb{K}}(H_n(X, A; \mathbb{K}), \mathbb{K}) : [\varphi] \mapsto \langle [\varphi], \cdot \rangle$$

is an isomorphism. □

Here is another situation in which the Ext term vanishes automatically: since $H_0(X)$ is the free abelian group generated by $\pi_0(X)$, Exercise 47.5 implies $\text{Ext}(H_0(X), G) = 0$ for every X and G .

COROLLARY 47.9. *The natural map $H^1(X; G) \rightarrow \text{Hom}(H_1(X), G)$ is an isomorphism for all spaces X and abelian groups G .* □

REMARK 47.10. Note that if X is path-connected, then $\text{Hom}(H_1(X), G)$ is canonically isomorphic to $\text{Hom}(\pi_1(X), G)$; indeed, since G is abelian, every homomorphism $\pi_1(X) \rightarrow G$ vanishes on the commutator subgroup and thus descends to the abelianization, which is $H_1(X)$. Corollary 47.9 thus confirms the result of Exercise 45.10.

If we take $R = \mathbb{Z}$, then the Ext term need not vanish in general, but one still obtains something revealing whenever the homology groups are finitely generated. Assume C_* is a chain complex of free abelian groups, abbreviate $H_n := H_n(C_*)$ and $H^n := H^n(C_*) = H^n(C_*; \mathbb{Z})$, and let $T_n \subset H_n$ and $T^n \subset H^n$ denote their respective torsion subgroups. If H_n and H^n are finitely generated, we can define their **free parts** as the quotients

$$H_n^{\text{free}}(C_*) = H_n^{\text{free}} := H_n/T_n, \quad H_{\text{free}}^n(C_*) = H_{\text{free}}^n := H^n/T^n,$$

and these are finitely-generated free abelian groups such that there exist isomorphisms

$$H_n \cong H_n^{\text{free}} \oplus T_n, \quad H^n \cong H_{\text{free}}^n \oplus T^n.$$

Applying the universal coefficient theorem with $G = \mathbb{Z}$ now produces the formula

$$H^n \cong \text{Hom}(H_n, \mathbb{Z}) \oplus \text{Ext}(H_{n-1}, \mathbb{Z}).$$

Since all homomorphisms $H_n \rightarrow \mathbb{Z}$ kill torsion elements and $\text{Hom}(\mathbb{Z}^m, \mathbb{Z}) \cong \mathbb{Z}^m$ for each $m \in \mathbb{N}$, we have $\text{Hom}(H_n, \mathbb{Z}) \cong \text{Hom}(H_n^{\text{free}} \oplus T_n, \mathbb{Z}) \cong \text{Hom}(H_n^{\text{free}}, \mathbb{Z}) \cong H_n^{\text{free}}$, and if H_{n-1} is also finitely generated, Proposition 47.6 gives an isomorphism $\text{Ext}(H_{n-1}, \mathbb{Z}) \cong T_{n-1}$, resulting in the formula

$$(47.3) \quad H^n \cong H_n^{\text{free}} \oplus T_{n-1}.$$

This implies $H_{\text{free}}^n \cong H_n^{\text{free}}$ and $T^n \cong T_{n-1}$. The first isomorphism can also be understood as follows. According to Theorem 47.7, the natural map $h : H^n \rightarrow \text{Hom}(H_n, \mathbb{Z})$ descends to an isomorphism

$$(47.4) \quad H^n/T \xrightarrow{h} \text{Hom}(H_n, \mathbb{Z}),$$

where $T \subset H^n$ is the image of the injective map $\text{Ext}(H_{n-1}, \mathbb{Z}) \rightarrow H^n$ in the long exact sequence. Composing homomorphisms $H_n \rightarrow \mathbb{Z}$ with the quotient projection $H_n \rightarrow H_n^{\text{free}}$ gives a natural isomorphism $\text{Hom}(H_n^{\text{free}}, \mathbb{Z}) = \text{Hom}(H_n, \mathbb{Z})$ since $\text{Hom}(T_n, \mathbb{Z}) = 0$, and $\text{Hom}(H_n^{\text{free}}, \mathbb{Z}) \cong H_n^{\text{free}}$ is a finitely generated free abelian group. Since $\text{Ext}(H_{n-1}, \mathbb{Z}) \cong T_{n-1}$ is a torsion group, the subgroup $T \subset H^n$ is necessarily contained in the torsion subgroup, but it must in fact be all of it since the isomorphism (47.4) implies that H^n/T is free. Combining all this with the naturality in Theorem 47.7 gives:

COROLLARY 47.11. For any chain complex C_* of free abelian groups and any $n \in \mathbb{Z}$ such that $H_n(C_*)$ and $H_{n-1}(C_*)$ are both finitely generated, the natural map $h : H^n(C_*) \rightarrow \text{Hom}(H_n(C_*), \mathbb{Z}) = \text{Hom}(H_n^{\text{free}}(C_*), \mathbb{Z})$ descends to a natural isomorphism

$$H_{\text{free}}^n(C_*) \xrightarrow{h} \text{Hom}(H_n^{\text{free}}(C_*), \mathbb{Z}),$$

where naturality means that for any two chain complexes A_*, B_* satisfying the above conditions and a chain map $\varphi_* : A_* \rightarrow B_*$, there is a commutative diagram

$$\begin{array}{ccc} H_{\text{free}}^n(A_*) & \xrightarrow{\cong} & \text{Hom}(H_n^{\text{free}}(A_*), \mathbb{Z}) \\ \varphi_* \uparrow & & \varphi_* \uparrow \\ H_{\text{free}}^n(B_*) & \xrightarrow{\cong} & \text{Hom}(H_n^{\text{free}}(B_*), \mathbb{Z}). \end{array}$$

In particular, $H_{\text{free}}^n(C_*) \cong H_n^{\text{free}}(C_*)$. Moreover, the torsion of $H^n(C_*)$ is isomorphic to the torsion of $H_{n-1}(C_*)$. □

EXERCISE 47.12. Use cellular cohomology to compute $H^*(\mathbb{R}\mathbb{P}^2)$, and compare the result with $H_*(\mathbb{R}\mathbb{P}^2)$ in light of Corollary 47.11.

Recall the numerical invariants defined in Lecture 39, namely the Betti numbers $b_n(X) = \text{rank } H_n(X)$ and the Euler characteristic $\chi(X) = \sum_{n=0}^{\infty} (-1)^n b_n(X)$. The universal coefficient theorem now gives us the freedom to compute these in terms of cohomology instead of homology:

COROLLARY 47.13. For any space X such that $H_*(X)$ is finitely generated,

$$b_n(X) = \text{rank } H^n(X) = \dim_{\mathbb{K}} H^n(X; \mathbb{K})$$

for every integer $n \geq 0$ and any field \mathbb{K} of characteristic zero. Moreover, if X is a compact CW-complex, then the formula

$$\chi(X) = \sum_{n=0}^{\infty} (-1)^n \dim_{\mathbb{K}} H^n(X; \mathbb{K})$$

holds for any field \mathbb{K} .

PROOF. Since $b_n(X)$ depends only on the free part of $H_n(X)$, it matches $\text{rank } H^n(X)$ due to Corollary 47.11. For any field \mathbb{K} of characteristic zero, the universal coefficient theorem for homology then implies $\text{rank } H_n(X) = \dim_{\mathbb{K}} H_n(X; \mathbb{K})$, and the latter matches $\dim_{\mathbb{K}} H^n(X; \mathbb{K})$ by Corollary 47.8. Finally if X is a finite cell complex, then we already know by applying Proposition 39.10 to the cellular chain complex that $\chi(X) = \sum_{n=0}^{\infty} (-1)^n \dim_{\mathbb{K}} H_n(X; \mathbb{K})$ holds for every field \mathbb{K} , and Corollary 47.8 enables us to replace $H_n(X; \mathbb{K})$ by $H^n(X; \mathbb{K})$ in this expression. □

We can also use these kinds of tricks to compute the Lefschetz number of a map $f : X \rightarrow X$. Here is one of the technical results that was used without proof when we argued two lectures ago that $L(f) \neq 0$ for every map $f : \mathbb{C}\mathbb{P}^n \rightarrow \mathbb{C}\mathbb{P}^n$ if n is even (cf. Theorem 45.1).

COROLLARY 47.14. Assume X is a space such that $H_*(X)$ is finitely generated. Then for any map $f : X \rightarrow X$ and any field \mathbb{K} ,

$$L_{\mathbb{K}}(f) = \sum_{n=0}^{\infty} (-1)^n \text{tr} \left(H^n(X; \mathbb{K}) \xrightarrow{f^*} H^n(X; \mathbb{K}) \right),$$

while for the special case $\mathbb{K} = \mathbb{Q}$,

$$L(f) = \sum_{n=0}^{\infty} (-1)^n \text{tr} \left(H_{\text{free}}^n(X) \xrightarrow{f^*} H_{\text{free}}^n(X) \right),$$

where $H_{\text{free}}^n(X)$ denotes the free part of $H^n(X; \mathbb{Z})$.

PROOF. Under the isomorphism of $H^n(X; \mathbb{K})$ with the dual space of $H_n(X; \mathbb{K})$ given by Corollary 47.8, the map $f^* : H^n(X; \mathbb{K}) \rightarrow H^n(X; \mathbb{K})$ is the transpose of f_* , as we have

$$\langle f^*[\varphi], [c] \rangle = (f^*\varphi)(c) = \varphi(f_*c) = \langle [\varphi], f_*[c] \rangle$$

for all $[\varphi] \in H^n(X; \mathbb{K})$ and $[c] \in H_n(X; \mathbb{K})$. This implies the formula above for $L_{\mathbb{K}}(f)$ since every linear map has the same trace as its transpose. The formula for $L(f)$ follows similarly from the naturality statement in Corollary 47.11, which gives a commutative diagram

$$\begin{array}{ccc} H_{\text{free}}^n(X) & \xrightarrow{\cong} & \text{Hom}(H_n^{\text{free}}(X), \mathbb{Z}) \\ \varphi^* \uparrow & & \varphi^* \uparrow \\ H_{\text{free}}^n(X) & \xrightarrow{\cong} & \text{Hom}(H_n^{\text{free}}(X), \mathbb{Z}) \end{array}$$

and thus identifies $f^* : H_{\text{free}}^n(X) \rightarrow H_{\text{free}}^n(X)$ with the transpose of $f_* : H_n^{\text{free}}(X) \rightarrow H_n^{\text{free}}(X)$. \square

We still have a loose end to tie up regarding cellular cohomology: the argument of the previous lecture only proves $H_{\text{CW}}^*(X; G) \cong H^*(X; G)$ when the CW-complex X is finite dimensional. Without this assumption, what it proves is that

$$H_{\text{CW}}^n(X; G) \cong H^n(X^N; G)$$

for every $n \geq 0$ and $N \geq n + 1$. Thus it will suffice to prove:

LEMMA 47.15. *For any CW-complex X , the inclusion $X^{n+2} \hookrightarrow X$ induces an isomorphism $H^n(X; G) \rightarrow H^n(X^{n+2}; G)$ for every integer $n \geq 0$.*

PROOF. The direct limit approach in Lecture 38 proves that the inclusion $X^{n+2} \hookrightarrow X$ induces isomorphisms $H_k(X^{n+2}) \xrightarrow{\cong} H_k(X)$ for every $k \leq n + 1$, and the long exact sequence

$$\dots \rightarrow H_k(X^{n+2}) \xrightarrow{\cong} H_k(X) \rightarrow H_k(X, X^{n+2}) \rightarrow H_{k-1}(X^{n+2}) \xrightarrow{\cong} H_{k-1}(X) \rightarrow \dots$$

then implies $H_k(X, X^{n+2}) = 0$. Plugging the relative singular chain complex $C_*(X, X^{n+2})$ into the universal coefficient theorem now gives a short exact sequence

$$0 \rightarrow \text{Ext}(H_{k-1}(X, X^{n+2}), G) \rightarrow H^k(X, X^{n+2}; G) \rightarrow \text{Hom}(H_k(X, X^{n+2}), G) \rightarrow 0,$$

in which the first and last terms both vanish, therefore so does $H^k(X, X^{n+2}; G)$. The long exact sequence of (X, X^{n+2}) in cohomology then has a segment of the form

$$0 = H^n(X, X^{n+2}; G) \rightarrow H^n(X; G) \rightarrow H^n(X^{n+2}; G) \rightarrow H^{n+1}(X, X^{n+2}; G) = 0,$$

implying that $H^n(X; G) \rightarrow H^n(X^{n+2}; G)$ is an isomorphism. \square

COROLLARY 47.16. *The isomorphism $H_{\text{CW}}^*(X; G) \cong H^*(X; G)$ holds for all (not just finite-dimensional) CW-complexes X .* \square

EXERCISE 47.17. Extend this discussion to prove that for all CW-pairs (X, A) , the obvious inclusions of pairs induce isomorphisms

$$\begin{aligned} H_n(X^{n+1} \cup A, A; G) &\xrightarrow{\cong} H_n(X, A; G), \\ H^n(X, A; G) &\xrightarrow{\cong} H^n(X^{n+2} \cup A, A; G) \end{aligned}$$

for all $n \geq 0$, and conclude from this that the isomorphisms $H_*^{\text{CW}}(X, A; G) \cong H_*(X, A; G)$ and $H_{\text{CW}}^*(X, A; G) \cong H^*(X, A; G)$ hold for all (possibly infinite-dimensional) CW-pairs (X, A) .

Hint: For homology, you need to extend the direct limit discussion in Lecture 38 to accommodate direct limits in Top_{rel} . You can then derive the cohomological statement from this by plugging the

singular chain complex of $(X, X^{n+2} \cup A)$ into the universal coefficient theorem and using the exact sequence of the triple $(X, X^{n+1} \cup A, A)$.

Finally, let's prove the universal coefficient theorem. The proof is at every step parallel to the theorem we proved for homology with coefficients in Lecture 42, but instead of applying the functor $\otimes G$, we apply $\text{Hom}(\cdot, G)$ and reverse arrows as needed.

PROOF OF THEOREM 47.7. We use the usual abbreviations $Z_n := \ker \partial_n \subset C_n$ and $B_n := \text{im } \partial_{n+1} \subset C_n$ for the boundary map $\partial_n : C_n \rightarrow C_{n-1}$ in the chain complex C_* . Since C_n for each $n \in \mathbb{Z}$ is assumed to be a free R -module and R is a principal ideal domain, Z_n and B_n are also free R -modules, which makes the exact sequence

$$\dots \longrightarrow 0 \longrightarrow 0 \longrightarrow B_n \xrightarrow{i_n} Z_n \xrightarrow{\text{pr}} H_n(C_*) \longrightarrow 0$$

a projective resolution of $H_n(C_*)$. Applying $\text{Hom}_R(\cdot, G)$ then gives a cochain complex

$$\dots \longleftarrow 0 \longleftarrow \text{Hom}_R(B_n, G) \xleftarrow{i_n^*} \text{Hom}_R(Z_n, G) \xleftarrow{\text{pr}^*} \text{Hom}_R(H_n(C_*), G) \longleftarrow 0,$$

which is exact at the two rightmost nontrivial terms by Exercise 47.1, so pr^* is injective and defines an isomorphism

$$(47.5) \quad \text{Hom}_R(H_n(C_*), G) \xrightarrow{\text{pr}^*} \ker i_n^* \subset \text{Hom}_R(Z_n, G).$$

The definition of Ext meanwhile gives

$$(47.6) \quad \text{coker } i_n^* = \text{Hom}_R(B_n, G) / \text{im } i_n^* = \text{Ext}^R(H_n(C_*), G).$$

Now regard $Z_* := \bigoplus_{n \in \mathbb{Z}} Z_n$ and $B_* := \bigoplus_{n \in \mathbb{Z}} B_n$ as chain complexes with trivial boundary operators, so that $0 \rightarrow Z_* \hookrightarrow C_* \xrightarrow{\partial} B_{*-1} \rightarrow 0$ is a short exact sequence of chain complexes. It also splits since B_{*-1} is free, so by Exercise 47.2, we can apply $\text{Hom}_R(\cdot, G)$ to it and obtain a short exact sequence of cochain complexes

$$0 \longleftarrow \text{Hom}_R(Z_*, G) \longleftarrow \text{Hom}_R(C_*, G) \longleftarrow \text{Hom}_R(B_{*-1}, G) \longleftarrow 0.$$

Here $\text{Hom}_R(Z_*, G)$ and $\text{Hom}_R(B_{*-1}, G)$ are each regarded as cochain complexes with trivial coboundary map, so they do not change if we replace them with their respective homologies. The usual diagram-chasing result thus produces a long exact sequence

$$\dots \longleftarrow \text{Hom}_R(B_n, G) \xleftarrow{i_n^*} \text{Hom}_R(Z_n, G) \longleftarrow H^n(C_*; G) \longleftarrow \text{Hom}_R(B_{n-1}, G) \xleftarrow{i_{n-1}^*} \text{Hom}_R(Z_{n-1}, G) \longleftarrow \dots,$$

where the maps i_n^* and i_{n-1}^* are actually connecting homomorphisms in the long exact sequence. (One must check this by reviewing the definition of the connecting homomorphisms via the diagram-chase.) Finally, we turn this into a short exact sequence with $H^n(C_*; G)$ at the center by replacing $\text{Hom}_R(Z_n, G)$ with $\ker i_n^*$ and $\text{Hom}_R(B_{n-1}, G)$ with $\text{coker } i_{n-1}^*$, giving

$$0 \longleftarrow \ker i_n^* \longleftarrow H^n(C_*; G) \longleftarrow \text{coker } i_{n-1}^* \longleftarrow 0.$$

In light of the isomorphisms (47.5) and (47.6), this is exactly the sequence we were aiming for.

I will leave it as an exercise to check the naturality of the sequence. The splitting follows easily from the chain map

$$(C_*, \partial) \xrightarrow{\text{pr} \circ p} (H_*(C_*), 0)$$

that we derived at the end of the proof of Theorem 41.1: dualizing this map gives a chain map of cochain complexes

$$(\text{pr} \circ p)^* : \text{Hom}_R(H_*(C_*), G) \rightarrow \text{Hom}_R(C_*, G),$$

where the first complex has a trivial coboundary map, so this induces a homomorphism

$$\mathrm{Hom}_R(H_n(C_*), G) \rightarrow H^n(C_*; G)$$

for each $n \in \mathbb{Z}$. One can check that the latter is a right-inverse of the canonical map $h : H^n(C_*; G) \rightarrow \mathrm{Hom}_R(H_n(C_*), G)$. \square

48. Products on cohomology, part 1

In this lecture we shall define the cup product

$$\cup : H^k(X; R) \otimes_R H^\ell(X; R) \rightarrow H^{k+\ell}(X; R)$$

and thus turn $H^*(X; R)$ into a ring. The definition is based on the cross product which we defined for homology in Lecture 43: we shall first have to adapt that definition to define a cohomological version of the cross product, and then use the contravariance of $H^*(\cdot; R)$ to derive the cup product from this and prove its basic properties.

We assume throughout this lecture that R is a commutative ring with unit; it will not need to be a principal ideal domain.

The cross product in singular homology was based on the existence and uniqueness (up to chain homotopy) of natural chain maps

$$(48.1) \quad \Phi_{(X,Y)} : C_*(X; R) \otimes_R C_*(Y; R) \rightarrow C_*(X \times Y; R).$$

Indeed, $\times : H_*(X; R) \otimes_R H_*(Y; R) \rightarrow H_*(X \times Y; R)$ was defined as the composition of the map induced by $\Phi_{(X,Y)}$ on homology with the canonical map $H_*(X; R) \otimes_R H_*(Y; R) \rightarrow H_*(C_*(X; R) \otimes_R C_*(Y; R))$, the latter being a purely algebraic construct that can be defined for any pair of chain complexes of R -modules. Exercise 43.8 outlined an argument via acyclic models to show that \times is associative, and we saw in Exercise 43.9 that the canonical generator of $H_0(\{\mathrm{pt}\}; R)$ acts as a multiplicative identity element. The question of commutativity is a bit subtler: in the first place, the relation $A \times B = B \times A$ would not make sense in general since for $A \in H_*(X; R)$ and $B \in H_*(Y; R)$, $A \times B$ and $B \times A$ do not belong to the same group, i.e. $H_*(X \times Y; R)$ is strictly speaking not the same thing as $H_*(Y \times X; R)$. Of course we can choose the obvious homeomorphism

$$\tau : X \times Y \rightarrow Y \times X : (x, y) \mapsto (y, x),$$

and use it to identify these two groups via the isomorphism

$$H_*(X \times Y; R) \xrightarrow{\tau_*} H_*(Y \times X; R),$$

but even then things are not so simple: instead of strict commutativity, we run into the usual convention that requires a sign change whenever the order of two elements with odd degree is interchanged.

PROPOSITION 48.1. *For any $A \in H_k(X; R)$ and $B \in H_\ell(Y; R)$, $\tau_*(A \times B) = (-1)^{k\ell}(B \times A)$.*

PROOF. Define $\Psi : C_*(X; R) \otimes_R C_*(Y; R) \rightarrow C_*(Y; R) \otimes_R C_*(X; R)$ by

$$\Psi(a \otimes b) := (-1)^{k\ell}(b \otimes a) \quad \text{for } a \in C_k(X; R), b \in C_\ell(Y; R).$$

This is a chain map, since for $a \in C_k(X; R)$ and $b \in C_\ell(Y; R)$,

$$\Psi \partial(a \otimes b) = \Psi(\partial a \otimes b + (-1)^k a \otimes \partial b) = (-1)^{(k-1)\ell} b \otimes \partial a + (-1)^k (-1)^{k(\ell-1)} \partial b \otimes a,$$

which is the same as

$$\partial \Psi(a \otimes b) = (-1)^{k\ell} \partial(b \otimes a) = (-1)^{k\ell}(\partial b \otimes a + (-1)^\ell b \otimes \partial a).$$

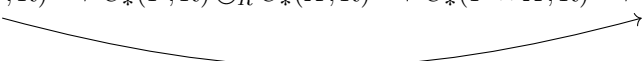
It also satisfies a naturality property: if $f : X \rightarrow X'$ and $g : Y \rightarrow Y'$ are continuous maps, then we have a commutative diagram

$$\begin{array}{ccc} C_*(X; R) \otimes_R C_*(Y; R) & \xrightarrow{\Psi} & C_*(Y; R) \otimes_R C_*(X; R) \\ \downarrow f_* \otimes g_* & & \downarrow g_* \otimes f_* \\ C_*(X'; R) \otimes_R C_*(Y', R) & \xrightarrow{\Psi} & C_*(Y; R) \otimes_R C_*(X; R). \end{array}$$

Similarly, $\tau_*^{-1} : C_*(Y \times X; R) \rightarrow C_*(X \times Y; R)$ is a natural chain map in the sense that for any f and g as above, the diagram

$$\begin{array}{ccc} C_*(Y \times X; R) & \xrightarrow{\tau_*^{-1}} & C_*(X \times Y; R) \\ \downarrow (g \times f)_* & & \downarrow (f \times g)_* \\ C_*(Y' \times X'; R) & \xrightarrow{\tau_*^{-1}} & C_*(X' \times Y'; R) \end{array}$$

also commutes. We can therefore compose three natural chain maps

$$C_*(X; R) \otimes_R C_*(Y; R) \xrightarrow{\Psi} C_*(Y; R) \otimes_R C_*(X; R) \xrightarrow{\Phi_{(Y,X)}} C_*(Y \times X; R) \xrightarrow{\tau_*^{-1}} C_*(X \times Y; R),$$


obtaining a natural chain map $C_*(X; R) \otimes_R C_*(Y; R) \rightarrow C_*(X \times Y; R)$ that acts on 0-chains in the obvious way $x \otimes y \mapsto (x, y)$. By the uniqueness statement for natural chain maps that we proved in Lemma 43.2, it follows that $\tau_*^{-1} \circ \Phi_{(Y,X)} \circ \Psi$ is chain homotopic to $\Phi_{(X,Y)} : C_*(X; R) \otimes_R C_*(Y; R) \rightarrow C_*(X \times Y; R)$, so these two chain maps induce the same map on homology, which proves $A \times B = (-1)^{k\ell} \tau_*^{-1}(B \times A)$ for $A \in H_k(X; R)$ and $B \in H_\ell(Y; R)$. \square

EXERCISE 48.2. Adapt Proposition 48.1 to prove the analogous statement about the cellular cross product $\times : H_k^{\text{CW}}(X; R) \otimes_R H_\ell^{\text{CW}}(Y; R) \rightarrow H_{k+\ell}^{\text{CW}}(X \times Y; R)$, where X and Y are assumed to be CW-complexes and $X \times Y$ and $Y \times X$ carry the resulting product cell decompositions.

Remark: The tricky part here is that one must compute the relevant incidence numbers for the cellular map $\tau : X \times Y \rightarrow Y \times X$; this is where the strange sign change will come from.

We would now like to define a similar product for the singular cohomology of two spaces: for each pair of integers $k, \ell \geq 0$, this should define a map

$$(48.2) \quad H^k(X; R) \otimes_R H^\ell(Y; R) \xrightarrow{\times} H^{k+\ell}(X \times Y; R)$$

that is dual to the homology cross product in the sense of the natural pairing of $H^*(\cdot; R)$ with $H_*(\cdot; R)$. The first step is an easy algebraic observation: cohomology groups are just the homology groups of dualized chain complexes, which are in this case cochain complexes of R -modules, so in light of the canonical map

$$H_*(A_*) \otimes_R H_*(B_*) \rightarrow H_*(A_* \otimes_R B_*) : [a] \otimes [b] \mapsto [a \otimes b]$$

that exists for any pair of chain complexes of R -modules, plugging in $A_* = C^*(X; R)$ and $B_* = C^*(Y; R)$ gives a canonical map

$$(48.3) \quad H^k(X; R) \otimes_R H^\ell(Y; R) \rightarrow H_{k+\ell}(C^*(X; R) \otimes_R C^*(Y; R)).$$

The complex at the right of this expression is to be understood as a *tensor product cochain complex*, meaning its degree n group is the direct sum of all $C^k(X; R) \otimes_R C^\ell(Y; R)$ for $k + \ell = n$, and its coboundary map is determined by the usual “graded Leibnitz rule”

$$\delta(\varphi \otimes \psi) = \delta\varphi \otimes \psi + (-1)^k \varphi \otimes \delta\psi \quad \text{for } \varphi \in C^k(X; R), \psi \in C^\ell(Y; R).$$

The obvious way we should try to get from here to (48.2) is by finding a natural chain map $C^*(X; R) \otimes_R C^*(Y; R) \rightarrow C^*(X \times Y; R)$ and composing the induced map on homology with (48.3). This can be done as follows. Recall from Lemma 43.3 that the natural chain map $\Phi : C_*(X) \otimes C_*(Y) \rightarrow C_*(X \times Y)$ has a natural chain homotopy inverse

$$\theta : C_*(X \times Y) \rightarrow C_*(X) \otimes C_*(Y),$$

which was also unique up to chain homotopy. (Note that this is a statement about chain complexes of abelian groups, i.e. \mathbb{Z} -modules, so it requires the fact that \mathbb{Z} is a principal ideal domain, but it has nothing to do with the ring R .) We can dualize this to produce a chain map of cochain complexes

$$\theta^* : \text{Hom}(C_*(X) \otimes C_*(Y), R) \rightarrow C^*(X \times Y; R),$$

and the desired chain map $C^*(X; R) \otimes_R C^*(Y; R) \rightarrow C^*(X \times Y; R)$ can then be defined as the composition

$$(48.4) \quad C^*(X; R) \otimes_R C^*(Y; R) \xrightarrow{F} \text{Hom}(C_*(X) \otimes C_*(Y), R) \xrightarrow{\theta^*} C^*(X \times Y; R),$$

$\searrow \hspace{10em} \nearrow$
 $\hspace{10em} \times$

where the map at the left arises canonically from the following purely algebraic construct:

EXERCISE 48.3. Show that for any two chain complexes of abelian groups A_* , B_* , the canonical map

$$F : \text{Hom}(A_*, R) \otimes_R \text{Hom}(B_*, R) \rightarrow \text{Hom}(A_* \otimes B_*, R)$$

defined for $\alpha \in \text{Hom}(A_k, R)$ and $\beta \in \text{Hom}(B_\ell, R)$ by

$$F(\alpha \otimes \beta)(a \otimes b) = \begin{cases} (-1)^{k\ell} \alpha(a)\beta(b) & \text{if } a \in A_k \text{ and } b \in B_\ell, \\ 0 & \text{otherwise} \end{cases}$$

is a chain map of cochain complexes.

Hint: To evaluate each of $\delta F(\alpha \otimes \beta)$ and $F\delta(\alpha \otimes \beta)$ on some element $a \otimes b$, you need to distinguish two cases, depending on the individual degrees of a and b . Getting all the signs right is a bit tricky. *Remark:* The sign convention in the definition of F would need to change if we had defined coboundary operators simply by $\delta = \partial^*$, instead of with the extra sign that appears in (45.2).

In the following, we shall remove the symbol F from the notation and simply regard elements of $C^*(X; R) \otimes_R C^*(Y; R)$ as homomorphisms $C_*(X) \otimes C_*(Y) \rightarrow R$ via the formula that results from Exercise 48.3, namely

$$(\varphi \otimes \psi)(a \otimes b) := \begin{cases} (-1)^{|\psi||a|} \varphi(a)\psi(b) & \text{if } |\varphi| = |a| \text{ and } |\psi| = |b|, \\ 0 & \text{otherwise.} \end{cases}$$

We shall denote the image of $\varphi \otimes \psi \in C^k(X; R) \otimes_R C^\ell(Y; R)$ under the map (48.4) by $\varphi \times \psi \in C^{k+\ell}(X \times Y; R)$. This can now be written as a homomorphism $C_{k+\ell}(X \times Y) \rightarrow R$ in the form

$$\varphi \times \psi = (\varphi \otimes \psi) \circ \theta : C_{k+\ell}(X \times Y) \rightarrow R,$$

and the fact that (48.4) is a chain map means that it satisfies the Leibniz rule

$$\delta(\varphi \times \psi) = \delta\varphi \times \psi + (-1)^{|\varphi|} \varphi \times \delta\psi.$$

It follows that \times descends to the singular cohomology groups, and we can define an R -bilinear product

$$H^k(X; R) \otimes_R H^\ell(Y; R) \rightarrow H^{k+\ell}(X \times Y; R) : [\varphi] \otimes [\psi] \mapsto [\varphi] \times [\psi] := [\varphi \times \psi].$$

Note that the cochain-level cross product $\varphi \times \psi$ depends in general on an arbitrary choice, namely the chain map $\theta : C_*(X \times Y) \rightarrow C_*(X) \otimes C_*(Y)$, but since the latter is unique up to chain homotopy, its dual θ^* is similarly unique up to chain homotopy, and so therefore is the resulting chain map $C^*(X; R) \otimes_R C^*(Y; R) \rightarrow C^*(X \times Y; R)$, proving that $[\varphi] \times [\psi]$ does not depend on our choices.

REMARK 48.4. After Lecture 43, you may have been tempted to try constructing the chain map $C^*(X; R) \otimes_R C^*(Y; R) \rightarrow C^*(X \times Y; R)$ via the method of acyclic models. This idea runs into the following difficulty: a cochain group such as $C^k(X; \mathbb{Z}) = \text{Hom}(C_k(X), \mathbb{Z})$ is not generally free, it is a direct *product* rather than a direct sum, thus it does not admit a basis. It is therefore quite unclear what models one might try to use for defining a map on $C^*(X; R) \otimes_R C^*(Y; R)$. The solution is to do the same thing we have done all along in our treatment of singular cohomology: reuse results that we've already proved about homology, but dualize them where appropriate.

PROPOSITION 48.5. *The cross product on singular cohomology has the following properties:*

- (1) *It is associative:* $(\varphi \times \psi) \times \eta = \varphi \times (\psi \times \eta) \in H^*(X \times Y \times Z; R)$ for any $\varphi \in H^*(X; R)$, $\psi \in H^*(Y; R)$ and $\eta \in H^*(Z; R)$.
- (2) *It is graded commutative:* $\varphi \times \psi = (-1)^{k\ell} \tau^*(\psi \times \varphi)$ for any $\varphi \in H^k(X; R)$, $\psi \in H^\ell(Y; R)$ and the homeomorphism $\tau : X \times Y \rightarrow Y \times X : (x, y) \mapsto (y, x)$.
- (3) *It is dual to the homology cross product: we have*

$$(48.5) \quad \langle \varphi \times \psi, A \times B \rangle = (-1)^{k\ell} \langle \varphi, A \rangle \langle \psi, B \rangle$$

for all $\varphi \in H^k(X; R)$, $A \in H_k(X)$, $\psi \in H^\ell(Y; R)$ and $B \in H_\ell(Y)$.⁶⁸

PROOF. We shall leave the first two properties as exercises; they can be solved as in Exercise 43.8 and Proposition 48.1 using the uniqueness up to chain homotopy of certain natural chain maps. For associativity, for instance, it is possible to express each of $(\varphi \times \psi) \times \eta$ and $\varphi \times (\psi \times \eta)$ as the composition of $\varphi \otimes \psi \otimes \eta$ with some natural chain map.

Let's quickly check the third property: given cycles $a \in C_k(X)$, $b \in C_\ell(Y)$ and cocycles $\varphi \in C^k(X; R)$, $\psi \in C^\ell(Y; R)$, we have by definition

$$\begin{aligned} \langle [\varphi] \times [\psi], [a] \times [b] \rangle &= \langle [\varphi \times \psi], [\Phi(a \otimes b)] \rangle = (\varphi \times \psi)(\Phi(a \otimes b)) \\ &= (\varphi \otimes \psi) \circ \theta \circ \Phi(a \otimes b) = (\varphi \otimes \psi) \circ (\mathbf{1} + \partial h + h \partial)(a \otimes b) = (\varphi \otimes \psi)(a \otimes b) \\ &= (-1)^{k\ell} \varphi(a) \psi(b) = (-1)^{k\ell} \langle [\varphi], [a] \rangle \langle [\psi], [b] \rangle, \end{aligned}$$

where in the second line we've used the existence of a chain homotopy h between $\theta \circ \Phi$ and the identity map, plus the fact that $\partial(a \otimes b)$ and $(\varphi \otimes \psi) \circ \partial$ both vanish since $\partial a = \partial b = 0$ and $\delta \varphi = \delta \psi = 0$. □

EXERCISE 48.6. The cross product on cellular cohomology is defined in the same way as above, but with the simplifying feature that since the chain map $\Phi : C_*^{\text{CW}}(X) \otimes C_*^{\text{CW}}(Y) \rightarrow C_*^{\text{CW}}(X \times Y)$ is an isomorphism, the required map $\theta : C_*^{\text{CW}}(X \times Y) \rightarrow C_*^{\text{CW}}(X) \otimes C_*^{\text{CW}}(Y)$ is simply its inverse, and is unique. Prove that the resulting product $H_{\text{CW}}^*(X; R) \otimes_R H_{\text{CW}}^*(Y; R) \rightarrow H_{\text{CW}}^*(X \times Y; R)$ satisfies the same (or analogous) properties as in Proposition 48.5.

EXERCISE 48.7. Prove that the naturality formula

$$(f \times g)^*(\varphi \times \psi) = f^* \varphi \times g^* \psi$$

⁶⁸The relation (48.5) uses R as the coefficients for cohomology and *integer* coefficients for homology, but since R is a ring, there is also an R -bilinear cross product $\times : H_*(X; R) \otimes_R H_*(Y; R) \rightarrow H_*(X \times Y; R)$ and natural evaluation pairings of $H^*(\cdot; R)$ with $H_*(\cdot; R)$. It is straightforward to check as a consequence of (48.5) that the same relation holds for the R -bilinear homology cross product and evaluation pairings.

holds for all $\varphi \in H^*(X; R)$ and $\psi \in H^*(Y; R)$ given continuous maps $f : X' \rightarrow X$ and $g : Y' \rightarrow Y$.

EXERCISE 48.8. The goal of this problem is to show that the canonical generator $1 \in R = H^0(\{\text{pt}\}; R)$ acts as a multiplicative identity element for the cross product.

- (a) Suppose Ψ associates to every space X a chain map $\Psi : C_*(X) \rightarrow C_*(X)$, and call this a **natural chain map** $C_*(X) \rightarrow C_*(X)$ if it acts as the identity map on 0-chains and satisfies $\Psi \circ f_* = f_* \circ \Psi$ for every continuous map $f : X \rightarrow Y$. Use the method of acyclic models to show that any two choices of natural chain maps in this sense are chain homotopic for all X .
- (b) Identify the chain complex $C_*(X \times \{\text{pt}\})$ with $C_*(X)$ via the obvious canonical isomorphism between them, and consider the following two maps:

$$\begin{aligned} C_*(X \times \{\text{pt}\}) &\xrightarrow{\theta} C_*(X) \otimes C_*(\{\text{pt}\}) \xrightarrow{1 \otimes \epsilon} C_*(X) \otimes \mathbb{Z} = C_*(X), \\ C_*(X \times \{\text{pt}\}) &\xrightarrow{(\pi_X)_*} C_*(X), \end{aligned}$$

where $\pi_X : X \times \{\text{pt}\} \rightarrow X$ is the canonical projection, θ is any natural chain homotopy inverse for the natural chain map $\Phi : C_*(X) \otimes C_*(\{\text{pt}\}) \rightarrow C_*(X \times \{\text{pt}\})$ as used in the construction of the cross product, and $\epsilon : C_*(\{\text{pt}\}) \rightarrow \mathbb{Z}$ is the *augmentation* map, which vanishes on $C_n(\{\text{pt}\})$ for $n \neq 0$ and sends the generator $\sigma \in C_0(\{\text{pt}\})$ to 1. Verify that both of these define natural chain maps, hence by part (a), they are chain homotopic.

- (c) Deduce from part (b) that for any space X , the cross product of $\varphi \in H^*(X; R)$ with $1 \in R \subset H^0(\{\text{pt}\}; R)$ satisfies $\varphi \times 1 = \pi_X^* \varphi$.

Thus far the development of cohomology looks quite similar to that of homology, i.e. every theorem or construction for homology has had a cohomological analogue. But we can now introduce something in cohomology that has no homological analogue: it is possible due to the fact that cohomology is *contravariant*, so in particular, the diagonal map $d : X \rightarrow X \times X$ induces a natural map $H^*(X \times X; R) \rightarrow H^*(X; R)$.

DEFINITION 48.9. The **cup product** on $H^*(X; R)$ is an R -bilinear map $\cup : H^k(X; R) \oplus H^\ell(X; R) \rightarrow H^{k+\ell}(X; R)$ for each $k, \ell \geq 0$ defined by

$$\varphi \cup \psi := d^*(\varphi \times \psi),$$

where $d : X \rightarrow X \times X$ is the diagonal map $x \mapsto (x, x)$.

On the cochain level, one can write $[\varphi] \cup [\psi] = [\varphi \cup \psi]$ if

$$\cup : C^k(X; R) \otimes_R C^\ell(X; R) \rightarrow C^{k+\ell}(X; R) : \varphi \otimes \psi \mapsto \varphi \cup \psi$$

is defined by

$$(48.6) \quad \varphi \cup \psi := d^*(\varphi \times \psi) = (\varphi \otimes \psi) \circ \theta \circ d_*.$$

This really is just the composition of our previous chain map $C^*(X; R) \otimes_R C^*(X; R) \rightarrow C^*(X \times X; R)$ with the chain map $d^* : C^*(X \times X; R) \rightarrow C^*(X; R)$ induced by the diagonal map, thus it is also a chain map and therefore satisfies the Leibniz rule

$$\delta(\varphi \cup \psi) = \delta\varphi \cup \psi + (-1)^{|\varphi|} \varphi \cup \delta\psi.$$

The presence of the mysterious object $\theta : C_*(X \times X) \rightarrow C_*(X) \otimes C_*(X)$ in (48.6) prevents it from being a useful formula on its own; we should also keep in mind that it is not uniquely defined since θ itself is not unique. We will see however that one can exploit this freedom to produce useful formulas.

DEFINITION 48.10. A **diagonal approximation** is an assignment to every space X of a chain map

$$\Psi_X : C_*(X) \rightarrow C_*(X) \otimes C_*(X)$$

that is defined on 0-chains $x \in C_0(X)$ by $\Psi_X(x) = x \otimes x$ and is natural in the sense that every continuous map $f : X \rightarrow Y$ gives rise to a commutative diagram

$$\begin{array}{ccc} C_*(X) & \xrightarrow{\Psi_X} & C_*(X) \otimes C_*(X) \\ \downarrow f_* & & \downarrow f_* \otimes f_* \\ C_*(Y) & \xrightarrow{\Psi_Y} & C_*(Y) \otimes C_*(Y). \end{array}$$

EXERCISE 48.11. Show that the map $\theta \circ d_*$ appearing in (48.6) is a diagonal approximation.

EXERCISE 48.12. Prove via acyclic models that all diagonal approximations are chain homotopic, and deduce that for any two cocycles $\varphi, \psi \in C^*(X; R)$, the cohomology class of $(\varphi \otimes \psi) \circ \Psi \in C^*(X; R)$ is independent of the choice of diagonal approximation Ψ .

Exercise 48.12 reveals an alternative but equivalent definition of the cup product that we could have taken if we had wanted to leave the cross product out of the discussion: \cup on cohomology is induced by any cochain-level product $\cup : C^k(X; R) \otimes_R C^\ell(X; R) \rightarrow C^{k+\ell}(X; R)$ of the form

$$\varphi \cup \psi := (\varphi \otimes \psi) \circ \Psi,$$

where Ψ is an arbitrary choice of diagonal approximation.

The most popular diagonal approximation in the literature is called the **Alexander-Whitney** diagonal approximation and is defined as follows. Number the vertices of the standard n -simplex $\Delta^n \subset \mathbb{R}^{n+1}$ as $0, \dots, n$, and given any integers $0 \leq j_0 < j_1 < \dots < j_k \leq n$, let

$$[j_0, \dots, j_k] \subset \Delta^n$$

denote the k -simplex spanned by the vertices j_0, \dots, j_k , which is identified naturally with the standard k -simplex. For instance, in this notation, the 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 $\Psi : C_*(X) \rightarrow C_*(X) \otimes C_*(X)$ on each singular n -simplex $\sigma : \Delta^n \rightarrow X$ by

$$\Psi(\sigma) := \sum_{k+\ell=n} (\sigma|_{[0, \dots, k]}) \otimes (\sigma|_{[k, \dots, n]})$$

EXERCISE 48.13. Verify that Ψ as defined above is 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 $n = k + \ell$,⁶⁹

$$(\varphi \cup \psi)(\sigma) = (-1)^{k\ell} \varphi(\sigma|_{[0, \dots, k]}) \psi(\sigma|_{[k, \dots, n]}).$$

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 there are more elegant ways to carry out all the computations we need.

⁶⁹The 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 (45.2) for the definition of coboundary maps.

To that end, let us now establish some properties of the cup product that will be essential in further developments. To understand the second property below, we need to be aware that there is always a canonical inclusion

$$R \subset H^0(X; R)$$

that makes the coefficient ring a submodule of $H^0(X; R)$: namely, each $r \in R$ is identified with the cohomology class of the cocycle $\varphi_r : C_0(X) \rightarrow R$ that has the value $\varphi_r(\sigma) = r$ on every generator $\sigma : \Delta^0 \rightarrow X$.

THEOREM 48.14. *The cup product $\cup : H^*(X; R) \otimes_R H^*(X; R) \rightarrow H^*(X; R)$ has the following properties.*

- (1) *It is natural: for all continuous maps $f : Y \rightarrow X$ and $\varphi, \psi \in H^*(X; R)$, $f^*(\varphi \cup \psi) = f^*\varphi \cup f^*\psi$.*
- (2) *It has a unit: under the canonical inclusion $R \subset H^0(X; R)$, $1 \in R$ satisfies $1 \cup \varphi = \varphi \cup 1 = \varphi$ for all $\varphi \in H^*(X; R)$.*
- (3) *It is associative: $(\varphi \cup \psi) \cup \eta = \varphi \cup (\psi \cup \eta)$ for all $\varphi, \psi, \eta \in H^*(X; R)$.*
- (4) *It is graded commutative: $\varphi \cup \psi = (-1)^{k\ell} \psi \cup \varphi$ for all $\varphi \in H^k(X; R)$ and $\psi \in H^\ell(X; R)$.*
- (5) *It is related to the cross product by*

$$\varphi \times \psi = \pi_X^* \varphi \cup \pi_Y^* \psi \quad \text{for } \varphi \in H^*(X; R), \psi \in H^*(Y; R),$$

where $\pi_X : X \times Y \rightarrow X$ and $\pi_Y : X \times Y \rightarrow Y$ are the natural projections.

PROOF. Most of these properties are relatively straightforward exercises using some combination of the cross product properties proved earlier and acyclic model arguments. Let's go quickly through the list:

Naturality is an easy consequence of the similar formula $(f \times g)^*(\varphi \times \psi) = f^*\varphi \times g^*\psi$ from Exercise 48.7. The unit property follows similarly from Exercise 48.8, while associativity and graded commutativity are easy consequences of the corresponding properties for the cross product.

For the last property, we observe that for any cocycles $\varphi \in C^*(X; R)$, $\psi \in C^*(Y; R)$ and any choice of diagonal approximation $\Psi : C_*(X \times Y) \rightarrow C_*(X \times Y) \otimes C_*(X \times Y)$, the resulting expression for the cocycle $\pi_X^* \varphi \cup \pi_Y^* \psi : C_*(X \times Y) \rightarrow R$ is the composition

$$C_*(X \times Y) \xrightarrow{\Psi} C_*(X \times Y) \otimes C_*(X \times Y) \xrightarrow{(\pi_X)_* \otimes (\pi_Y)_*} C_*(X) \otimes C_*(Y) \xrightarrow{\varphi \otimes \psi} R.$$

As can easily be checked, this is the composition of $\varphi \otimes \psi$ with a natural chain map $C_*(X \times Y) \rightarrow C_*(X) \otimes C_*(Y)$ that is defined in the canonical way on 0-chains, thus the chain map is chain homotopic to θ and we end up with the usual formula for $\varphi \times \psi$. □

49. Products on cohomology, part 2

The cohomology rings of tori. Let's start today with a computation. Theorem 48.14 in the previous lecture states that for any space X and commutative ring R with unit $1 \in R$, $H^*(X; R) = \bigoplus_{n=0}^\infty H^n(X; R)$ is a graded commutative ring with unit $1 \in R \subset H^0(X; R)$, where the ring structure is defined via the cup product $H^k(X; R) \times H^\ell(X; R) \rightarrow H^{k+\ell}(X; R) : (\varphi, \psi) \mapsto \varphi \cup \psi$. Recall that "graded commutative" means something slightly different than "commutative": homogeneous elements commute up to a sign that depends on their degrees,

$$\varphi \cup \psi = (-1)^{k\ell} \psi \cup \varphi \quad \text{for } \varphi \in H^k(X; R), \psi \in H^\ell(X; R).$$

In order to access the full range of applications of cohomology, one often needs to compute not just the individual groups $H^k(X; R)$ for $k \geq 0$, but also the ring structure of $H^*(X; R)$. We saw an example of this with $X = \mathbb{C}P^n$ in Theorem 45.1. At this point we do not yet have enough machinery to compute the ring $H^*(\mathbb{C}P^n)$, but we can compute $H^*(\mathbb{T}^n)$ for every $n \geq 0$.

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 (see Problem Set 9 #3). 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, $H_*(S^1)$ is a finitely generated free abelian group. Let's call its canonical generators

$$[\text{pt}] \in H_0(S^1), \quad [S^1] \in H_1(S^1),$$

i.e. $[\text{pt}]$ 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 $H_1(S^1) \cong \pi_1(S^1)$. Now suppose we assume for a given integer $n \geq 2$ that $H_*(\mathbb{T}^q)$ is finitely generated and free for every $q \leq n - 1$. Then the Künneth formula gives for every $m \geq 0$ an isomorphism

$$\times : \bigoplus_{k+\ell=m} H_k(\mathbb{T}^{n-1}) \otimes H_\ell(S^1) \xrightarrow{\cong} H_m(\mathbb{T}^n)$$

since all Tor terms vanish, while $H_\ell(S^1)$ on the left hand side is only nontrivial for $\ell = 0, 1$ and is then \mathbb{Z} , giving

$$H_m(\mathbb{T}^n) \cong H_{m-1}(\mathbb{T}^{n-1}) \oplus H_m(\mathbb{T}^{n-1}).$$

This proves that $H_m(\mathbb{T}^n)$ is also a finitely-generated free abelian group, and its rank is an entry in Pascal's triangle,

$$\text{rank } H_m(\mathbb{T}^n) = \binom{n}{m}.$$

Moreover, the cross product provides a canonical set of generators of $H_m(\mathbb{T}^n)$: for each choice of integers $1 \leq j_1 < \dots < j_m \leq n$, we define

$$e_{j_1, \dots, j_m} := A_{j_1} \times \dots \times A_{j_m} \in H_m(\mathbb{T}^n)$$

by setting $A_{j_i} := [S^1]$ for each $i = 1, \dots, m$ and $A_j := [\text{pt}]$ for all other $j = 1, \dots, n$. So far so good.

It will be useful to have an alternative description of the degree 1 generators $e_j \in H_1(\mathbb{T}^n)$. Pick a base point $t_0 \in S^1$ and consider the embedding

$$(49.1) \quad i_j : S^1 \hookrightarrow \mathbb{T}^n : x \mapsto (\underbrace{t_0 \times \dots \times t_0}_{j-1}, x, \underbrace{t_0 \times \dots \times t_0}_{n-j}).$$

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 $(i_j)_* : H_*(S^1) \rightarrow H_*(\mathbb{T}^n)$ is independent of this choice.

LEMMA 49.1. For each $j = 1, \dots, n$, $(i_j)_*[S^1] = e_j$.

PROOF. Consider first the case $j = n$. Under the obvious identification of S^1 with $\{\text{pt}\} \times S^1$, we can then write $i_n = c \times \text{Id} : \{\text{pt}\} \times S^1 \rightarrow \mathbb{T}^{n-1} \times S^1$, where $c : \{\text{pt}\} \rightarrow \mathbb{T}^{n-1}$ denotes the constant map with value (t_0, \dots, t_0) . The naturality of the cross product now gives a commutative diagram

$$\begin{array}{ccc} H_0(\{\text{pt}\}) \otimes H_1(S^1) & \xrightarrow{\times} & H_1(S^1) \\ \downarrow c_* \otimes \mathbb{1} & & \downarrow (i_n)_* \\ H_0(\mathbb{T}^{n-1}) \otimes H_1(S^1) & \xrightarrow{\times} & H_1(\mathbb{T}^n). \end{array}$$

In light of Exercise 43.9, this proves $(i_n)_*[S^1] = (i_n)_*([\text{pt}] \times [S^1]) = c_*[\text{pt}] \times [S^1] = [\text{pt}] \times [S^1] = e_n$ since $c_* : H_0(\{\text{pt}\}) \rightarrow H_0(\mathbb{T}^{n-1})$ is an isomorphism relating the canonical generators $[\text{pt}]$. The general case $j \in \{1, \dots, n\}$ follows from this same argument after permuting the coordinates. \square

The computation of $H^m(\mathbb{T}^n; R)$ for each m is now an easy application of the universal coefficient theorem: the Ext terms vanish since $H_*(\mathbb{T}^n)$ is always free, implying that the canonical map

$$H^m(\mathbb{T}^n; R) \rightarrow \text{Hom}(H_m(\mathbb{T}^n), R) : \varphi \mapsto \langle \varphi, \cdot \rangle$$

is an isomorphism. Since $H_m(\mathbb{T}^n) \cong \mathbb{Z}^{\binom{n}{m}}$, this means $H^m(\mathbb{T}^n; R) \cong R^{\binom{n}{m}}$, and we can write down a canonical set of generators as follows. For $n = 1$, define

$$\lambda \in H^1(S^1; R)$$

to be the unique cohomology class such that

$$\langle \lambda, [S^1] \rangle = 1.$$

Now for each choice of integers $1 \leq j_1 < \dots < j_m \leq n$, define

$$\lambda_{j_1, \dots, j_m} := \alpha_1 \times \dots \times \alpha_n \in H^m(\mathbb{T}^n; R),$$

where we choose $\alpha_{j_i} := \lambda$ for each $i = 1, \dots, m$ and $\alpha_j = 1 \in H^0(S^1; R)$ for all other $j = 1, \dots, n$. By (48.5), 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 $H^*(\mathbb{T}^n; R)$ as a free R -module.

To describe $H^*(\mathbb{T}^n; R)$ as a ring, we now need to compute each product of the form $\lambda_{j_1, \dots, j_m} \cup \lambda_{k_1, \dots, k_q} \in H^{m+q}(\mathbb{T}^n; R)$. We start with an observation about the 1-dimensional classes $\lambda_j \in H^1(\mathbb{T}^n; R)$. 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 (49.1) above by

$$\pi_j \circ i_k = \begin{cases} \text{Id} : S^1 \rightarrow S^1 & \text{if } j = k, \\ \text{constant} & \text{if } j \neq k. \end{cases}$$

LEMMA 49.2. *We have $\pi_j^* \lambda = \lambda_j$ for each $j = 1, \dots, n$.*

PROOF. In light of Lemma 49.1 and the isomorphism $H^1(\mathbb{T}^n; R) \cong \text{Hom}(H_1(\mathbb{T}^n), R)$, the classes $\pi_j^* \lambda \in H^1(\mathbb{T}^n; R)$ are characterized by

$$\langle \pi_j^* \lambda, e_k \rangle = \langle \lambda, (\pi_j)_*(i_k)_*[S^1] \rangle = \langle \lambda, (\pi_j \circ i_k)_*[S^1] \rangle = \begin{cases} 1 & \text{if } j = k, \\ 0 & \text{if } j \neq k. \end{cases}$$

At the same time, (48.5) implies

$$\langle \lambda_j, e_k \rangle = \begin{cases} 1 & \text{if } j = k, \\ 0 & \text{if } j \neq k, \end{cases}$$

and the result follows. \square

We are now in a position to compute $\lambda_{j_1} \cup \dots \cup \lambda_{j_m} \in H^m(\mathbb{T}^n; R)$ for any set of integers $1 \leq j_1 < \dots < j_m \leq n$. Indeed, writing $\alpha_{j_i} = \lambda \in H^1(S^1; R)$ for each $i = 1, \dots, m$ and $\alpha_j = 1 \in H^0(S^1; R)$ for all other j , we have $\pi_{j_i}^* \alpha_{j_i} = \lambda_j \in H^1(\mathbb{T}^n; R)$ in the first case and $\pi_j^* \alpha_j = \pi_j^* 1 = 1 \in H^0(\mathbb{T}^n; R)$ in the second case, thus by Theorem 48.14(5),

$$\lambda_{j_1} \cup \dots \cup \lambda_{j_m} = \pi_1^* \alpha_1 \cup \dots \cup \pi_n^* \alpha_n = \alpha_1 \times \dots \times \alpha_n = \lambda_{j_1, \dots, j_m}.$$

This means that all of our basis elements for $H^*(\mathbb{T}^n; R)$ can be obtained as cup products of the degree 1 elements $\lambda_1, \dots, \lambda_n$, and moreover, this relation fully determines all cup products in $H^*(\mathbb{T}^n; R)$; indeed, graded commutativity implies

$$\lambda_i \cup \lambda_j = -\lambda_j \cup \lambda_i$$

for all i and j , so in particular $\lambda_i \cup \lambda_i$ always vanishes, and all other products of basis elements $\lambda_{j_1, \dots, j_m}$ can be derived from this via associativity. We've proved:

THEOREM 49.3. *For any commutative ring R with unit and any $n \in \mathbb{N}$, the ring $H^*(\mathbb{T}^n; R)$ is isomorphic to the exterior algebra $\Lambda_R[\lambda_1, \dots, \lambda_n]$ on n generators of degree 1, where the generators $\lambda_j \in H^1(\mathbb{T}^n; R)$ can be defined in terms of the projections $\pi_j : \mathbb{T}^n \rightarrow S^1$ and the canonical generator $\lambda \in H^1(S^1; R)$ by $\lambda_j = \pi_j^* \lambda$. \square*

We will be able to compute more examples of cohomology rings after we discuss Poincaré duality, which provides the most useful geometric interpretation of the cup product in terms of intersections.

The relative case. Our whole discussion of products so far has focused on absolute homology and cohomology, so you may be wondering how it extends to pairs of spaces (X, A) with $A \neq \emptyset$. As usual, some valuable intuition for this comes from cellular homology. If (X, A) and (Y, B) are CW-pairs, then $x \otimes y \in C_*^{\text{CW}}(X) \otimes C_*^{\text{CW}}(Y)$ represents a trivial element of $C_*^{\text{CW}}(X, A) \otimes C_*^{\text{CW}}(Y, B)$ whenever either $x \in C_*^{\text{CW}}(A)$ or $y \in C_*^{\text{CW}}(B)$. Thus the chain-level cellular cross product defined in (43.1) becomes an isomorphism of relative 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 by

$$(49.2) \quad (X, A) \times (Y, B) := (X \times Y, (A \times Y) \cup (X \times B)).$$

You should take a moment to convince yourself that this does in fact define a CW-pair. The most general version of the cross product on relative cellular homology thus takes the form

$$H_k^{\text{CW}}(X, A; R) \otimes_R H_\ell^{\text{CW}}(Y, B; R) \xrightarrow{\times} H_{k+\ell}^{\text{CW}}((X, A) \times (Y, B); R),$$

and the Künneth formula (in the case where R is a principal ideal domain) then becomes

$$0 \longrightarrow \bigoplus_{k+\ell=n} H_k^{\text{CW}}(X, A; R) \otimes_R H_\ell^{\text{CW}}(Y, B; R) \xrightarrow{\times} H_n^{\text{CW}}((X, A) \times (Y, B); R) \\ \longrightarrow \bigoplus_{k+\ell=n-1} \text{Tor}^R(H_k^{\text{CW}}(X, A; R), H_\ell^{\text{CW}}(Y, B; R)) \longrightarrow 0.$$

Adapting this discussion for singular homology is slightly nontrivial, and it does not 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 (49.2) as a definition of the product of two objects in the category Top_{rel} .⁷⁰ Applying the naturality of the chain map $\Phi : C_*(X; R) \otimes_R C_*(Y; R) \rightarrow C_*(X \times Y; R)$ to the inclusions $A \hookrightarrow X$ and $B \hookrightarrow Y$, we see that Φ maps $C_*(A; R) \otimes_R C_*(Y; R)$ into $C_*(A \times Y; R)$ and $C_*(X; R) \otimes_R C_*(B; R)$ into $C_*(X \times B; R)$, thus it descends to a natural chain map

$$\Phi : C_*(X, A; R) \otimes_R C_*(Y, B; R) \rightarrow C_*((X, A) \times (Y, B); R),$$

⁷⁰This definition of $(X, A) \times (Y, B)$ is the right one for talking about the cross product and Künneth's formula, but it is not so good for other purposes: in particular it suffers from the fact that the obvious projection maps from $X \times Y$ to X or Y do not generally define morphisms from $(X, A) \times (Y, B)$ to (X, A) or (Y, B) under this definition. (Think about it.) There is a more obvious alternative definition of the product for an arbitrary collection of pairs of spaces which does not have this problem with projection maps. That is the right definition to use if, say, one wants an explicit description of inverse limits in Top_{rel} , in the spirit of Exercise 44.8.

so that the cross product on relative homology is well defined:

$$\times : H_k(X, A; R) \otimes_R H_\ell(Y, B; R) \rightarrow H_{k+\ell}((X, A) \times (Y, B); R).$$

We run into a complication, however, if we either want to define the cross product on relative cohomology or prove a relative Künneth formula: both require the chain homotopy inverse $\theta : C_*(X \times Y; R) \rightarrow C_*(X; R) \otimes_R C_*(Y; R)$, and this does not always descend to a map

$$C_*((X, A) \times (Y, B); R) \rightarrow C_*(X, A; R) \otimes_R C_*(Y, B; R).$$

The point is that if we are given a chain in the subspace $(A \times Y) \cup (X \times B)$, there is generally no reason to expect that θ will send it into $(C_*(A; R) \otimes_R C_*(Y; R)) + (C_*(X; R) \otimes_R C_*(B; R))$. What we can immediately say instead is that Φ and θ descend to chain homotopy inverses between the two quotient complexes

(49.3)

$$\frac{C_*(X; R) \otimes_R C_*(Y; R)}{(C_*(A; R) \otimes_R C_*(Y; R)) + (C_*(X; R) \otimes_R C_*(B; R))} \xrightarrow{\Phi} \frac{C_*(X \times Y; R)}{C_*(A \times Y; R) + C_*(X \times B; R)} \xleftarrow{\theta}$$

The complex at the left is just $C_*(X, A; R) \otimes_R C_*(Y, B; R)$, which is what we want, but the one at the right is not the same as $C_*((X, A) \times (Y, B); R)$. However, the identity map does descend to a natural chain map

$$\frac{C_*(X \times Y; R)}{C_*(A \times Y; R) + C_*(X \times B; R)} \longrightarrow \frac{C_*(X \times Y; R)}{C_*((A \times Y) \cup (X \times B); R)} = C_*((X, A) \times (Y, B); R),$$

and it will happen sometimes that this map is a chain homotopy equivalence, so that Φ descends after all to an isomorphism from the homology of $C_*(X, A; R) \otimes_R C_*(Y, B; R)$ to $H_*((X, A) \times (Y, B); R)$. The situation should remind you of the technical hurdles we had to overcome in order to prove excision or define the Mayer-Vietoris sequence, and it can be dealt with in a similar way.

DEFINITION 49.4. Given a space X , two subspaces $X_1, X_2 \subset X$ are called an **excisive couple** if the natural chain map defined by the inclusion

$$C_*(X_1) + C_*(X_2) \hookrightarrow C_*(X_1 \cup X_2)$$

induces isomorphisms on the homology groups.⁷¹

LEMMA 49.5. *Two subspaces $X_1, X_2 \subset X$ form an excisive couple if and only if the canonical chain map*

$$\frac{C_*(X)}{C_*(X_1) + C_*(X_2)} \longrightarrow C_*(X, X_1 \cup X_2)$$

descends to an isomorphism on the homology groups.

PROOF. Assume either that X_1 and X_2 form an excisive couple or that the map of quotient complexes induces isomorphisms on homology. There is a commutative diagram

$$\begin{array}{ccccccc} 0 & \longrightarrow & C_*(X_1) + C_*(X_2) & \hookrightarrow & C_*(X) & \longrightarrow & \frac{C_*(X)}{C_*(X_1) + C_*(X_2)} \longrightarrow 0 \\ & & \downarrow & & \downarrow \mathbb{1} & & \downarrow \\ 0 & \longrightarrow & C_*(X_1 \cup X_2) & \hookrightarrow & C_*(X) & \longrightarrow & C_*(X, X_1 \cup X_2) \longrightarrow 0, \end{array}$$

⁷¹Notice what this definition does not say: it would seem natural to ask for this inclusion to be a chain homotopy equivalence, but we are requiring something slightly weaker. We will in fact want to use the weaker conditions in some situations where the stronger one might not hold; an example is the third case of Lemma 49.7.

where both rows are short exact sequences of chain complexes and all arrows represent chain maps induced by either inclusions or quotient projections. Transforming both rows into long exact sequences of homology groups then produces a diagram in which two out of every three vertical maps are isomorphisms, so the five-lemma implies that the third one is as well. \square

EXERCISE 49.6. Show that if $X_1, X_2 \subset X$ are an excisive couple, then the relevant induced maps on homology or cohomology with an arbitrary coefficient group are also isomorphisms. *Hint: Use the naturality of the universal coefficient theorems.*

The lemma implies that all important results regarding products $(X, A) \times (Y, B)$ in homology or cohomology will hold as long as the two subsets $A \times Y$ and $X \times B$ in $X \times Y$ form an excisive couple.

LEMMA 49.7. *Given two pairs or spaces (X, A) and (Y, B) , the subsets $A \times Y$ and $X \times B$ in $X \times Y$ form an excisive couple whenever any of the following conditions holds:*

- (1) $A \subset X$ and $B \subset Y$ are both open subsets;
- (2) $A = \emptyset$ or $B = \emptyset$;
- (3) (X, A) and (Y, B) are both CW-pairs.

PROOF. The first case follows by barycentric subdivision, as we showed when we proved the excision axiom (see Lemma 24.11 from last semester). The second case is trivial. The third can be proven by replacing singular with cellular chain complexes and appealing to the isomorphism of singular with cellular homology: the necessary condition is obvious on the cellular chain complex since $C_*^{\text{CW}}(A) + C_*^{\text{CW}}(B)$ and $C_*^{\text{CW}}(A \cup B)$ are exactly the same. \square

Now using the quotient complex at the right hand side of (49.3) as a stand-in for $C_*((X, A) \times (Y, B); R)$, we obtain a relative version of the Eilenberg-Zilber theorem and therefore a relative Künneth formula:

THEOREM 49.8. *If R is a principal ideal domain and (X, A) and (Y, B) are pairs such that the subsets $A \times Y$ and $X \times B$ in $X \times Y$ form an excisive couple, then there is a natural short exact sequence*

$$0 \longrightarrow \bigoplus_{k+\ell=n} H_k(X, A; R) \otimes_R H_\ell(Y, B; R) \xrightarrow{\times} H_{k+\ell}((X, A) \times (Y, B); R) \\ \longrightarrow \bigoplus_{k+\ell=n-1} \text{Tor}^R(H_k(X, A; R), H_\ell(Y, B; R)) \longrightarrow 0,$$

and the sequence splits. \square

EXERCISE 49.9. Using Exercise 49.6 to identify $H^*((X, A) \times (Y, B); R)$ with the cohomology of the quotient complex on the right hand side of (49.3), write down a chain-level definition of the cross product on relative singular cohomology,

$$\times : H^k(X, A; R) \otimes_R H^\ell(Y, B; R) \rightarrow H^{k+\ell}((X, A) \times (Y, B); R),$$

under the assumption that $A \times Y, X \times B \subset X \times Y$ form an excisive couple.

Here is an interesting application of the relative Künneth formula. If (X, x_0) and (Y, y_0) are two pointed spaces, their **smash product** $X \wedge Y$ is defined as the quotient space

$$X \wedge Y := (X \times Y) / ((\{x_0\} \times Y) \cup (X \times \{y_0\})).$$

Strictly speaking, this construction depends on the choice of base points, but we shall suppress this in the notation. Notice that the subset being quotiented out is homeomorphic to the wedge sum $X \vee Y$, so it is sensible to write

$$X \wedge Y = (X \times Y) / (X \vee Y).$$

It is now straightforward to check that for any base-point preserving continuous maps $f : (X, x_0) \rightarrow (X', x'_0)$ 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 49.10. 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 49.7, the Künneth formula is valid for the pairs $(X, \{x_0\})$ and $(Y, \{y_0\})$. Since $(X, \{x_0\}) \times (Y, \{y_0\}) = (X \times Y, X \vee Y)$, the Künneth formula now takes the form

$$\begin{aligned} 0 \rightarrow \bigoplus_{k+\ell=n} H_k(X, \{x_0\}) \otimes H_\ell(Y, \{y_0\}) &\xrightarrow{\times} H_n(X \times Y, X \vee Y) \\ &\longrightarrow \bigoplus_{k+\ell=n-1} \text{Tor}(H_k(X, \{x_0\}), H_\ell(Y, \{y_0\})) \rightarrow 0, \end{aligned}$$

or under the natural isomorphisms $H_*(X, A) = \tilde{H}_*(X/A)$ for good pairs,

$$0 \rightarrow \bigoplus_{k+\ell=n} \tilde{H}_k(X) \otimes \tilde{H}_\ell(Y) \xrightarrow{\times} \tilde{H}_n(X \wedge Y) \longrightarrow \bigoplus_{k+\ell=n-1} \text{Tor}(\tilde{H}_k(X), \tilde{H}_\ell(Y)) \rightarrow 0.$$

EXERCISE 49.11. Show that for the cross product on reduced homology as described above and the identification of $S^k \wedge S^\ell$ with $S^{k+\ell}$ as indicated in Example 49.10, 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 49.12. 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.

With significantly more effort, one can proceed from Exercise 49.12 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].

50. Relative cup and cap products, orientations

There are two more things about products that we should discuss before moving on to the next major topic.

The relative cup product. Recall that the cup product can be defined in terms of any map on cochains $C^k(X; R) \otimes_R C^\ell(X; R) \rightarrow C^{k+\ell}(X; R) : \varphi \otimes \psi \mapsto \varphi \cup \psi$ of the form

$$(50.1) \quad \varphi \cup \psi = (\varphi \otimes \psi) \circ \Psi : C_*(X) \rightarrow R,$$

where $\Psi : C_*(X) \rightarrow C_*(X) \otimes C_*(X)$ is any choice of natural chain map that acts as the diagonal map on 0-chains, i.e. Ψ is a *diagonal approximation*. This formulation will be convenient in the following because it avoids any reference to the cross product, whose definition in the relative case we have seen involves some subtleties. We claim that whenever $A, B \subset X$ are two subspaces that form an excisive couple (see Definition 49.4), there is a well-defined relative cup product

$$\cup : H^k(X, A; R) \otimes_R H^\ell(X, B; R) \rightarrow H^{k+\ell}(X, A \cup B; R).$$

Indeed, under this assumption, Exercise 49.6 identifies $H^*(X, A \cup B; R)$ 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 $\varphi \in C^*(X, A; R)$ and $\psi \in C^*(X, B; R)$, 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 48.14, though we need to be a bit careful about its relation to the cross product. While the diagonal map $d : X \rightarrow X \times X$ always gives a well-defined map of pairs

$$(X, A \cup B) \xrightarrow{d} (X, A) \times (X, B),$$

the formula $\varphi \cup \psi = d^*(\varphi \times \psi)$ might *not* make sense under the assumption above, because $A \times X$ and $X \times B$ might not be an excisive couple in $X \times X$, in which case the cross product in this expression is not well defined. If both $A, B \subset X$ and $A \times X, X \times B \subset X \times X$ are excisive couples, then both definitions of $\varphi \cup \psi \in H^*(X, A \cup B; R)$ do make sense, and they match.

As a special case, the product

$$\cup : H^k(X, A; R) \otimes_R H^\ell(X, A; R) \rightarrow H^{k+\ell}(X, A; R)$$

is well defined for *every* pair (X, A) , as $A, A \subset X$ always trivially forms an excisive couple. This is true even though the cross product $\times : H^k(X, A; R) \otimes_R H^\ell(X, A; R) \rightarrow H^{k+\ell}((X, A) \times (X, A); R)$ might not always make sense.

The cap product. The cap product is another pairing that intertwines cohomology with homology to produce a homology class, and its main property is that it is in some sense *dual* to the cup product. For this discussion, we use a fixed commutative ring R with unit as the coefficient ring for both homology and cohomology, though in order to prevent notational clutter, we will also work with the chain groups $C_n(X)$ over \mathbb{Z} and make use of the fact that since $C_n(X; R) = C_n(X) \otimes R$, there is a canonical isomorphism

$$\text{Hom}(C_n(X), G) = \text{Hom}_R(C_n(X; R), G) \quad \text{for any } R\text{-module } G,$$

cf. Exercise 45.5. We'll use the notation $\langle \cdot, \cdot \rangle$ to denote both of the natural pairings

$$\begin{aligned} C^n(X; R) \otimes C_n(X) &\rightarrow R : \varphi \otimes c \mapsto \langle \varphi, c \rangle := \varphi(c), \\ C^n(X; R) \otimes_R C_n(X; R) &\rightarrow R : \varphi \otimes (c \otimes r) \mapsto \langle \varphi, c \otimes r \rangle := r\varphi(c), \end{aligned}$$

so for fixed $\varphi \in C^n(X; R)$, $\langle \varphi, \cdot \rangle$ can be interpreted either as a group homomorphism $C_n(X) \rightarrow R$ or as the equivalent R -module homomorphism $C_n(X; R) \rightarrow R$. We shall also adopt the convention of defining

$$\langle \varphi, c \rangle := 0 \quad \text{if } |\varphi| \neq |c|,$$

so that the pairing makes sense for all $\varphi \in C^*(X; R)$ and c in $C_*(X)$ or $C_*(X; R)$.

PROPOSITION 50.1. *Given a diagonal approximation $\Psi : C_*(X) \rightarrow C_*(X) \otimes C_*(X)$ and the associated cup product of cochains defined in (50.1), there exists a unique R -module homomorphism*

$$(50.2) \quad C^k(X; R) \otimes_R C_\ell(X; R) \rightarrow C_{\ell-k}(X; R) : \varphi \otimes c \mapsto \varphi \cap c$$

for each pair of integers k, ℓ such that the relation

$$\langle \psi \cup \varphi, c \rangle = \langle \psi, \varphi \cap c \rangle$$

is satisfied for all $\psi, \varphi \in C^*(X; R)$ and $c \in C_*(X; R)$. Moreover, \cap satisfies

$$(50.3) \quad \partial(\varphi \cap c) = \delta\varphi \cap c + (-1)^k \varphi \cap \partial c \quad \text{for all } \varphi \in C^k(X; R), c \in C_\ell(X; R).$$

REMARK 50.2. The degrees appearing in (50.2) become easy to remember if you regard $C^*(X; R)$ in the spirit of Remark 45.3 as a *chain* (not cochain) complex by reversing the degrees and writing $C^*(X; R)_k := C^{-k}(X; R)$. We can then regard $C^*(X; R) \otimes_R C_*(X; R)$ as a tensor product chain complex, and the Leibniz rule (50.3) becomes the statement that $\cap : C^*(X; R) \otimes_R C_*(X; R) \rightarrow C_*(X; R)$ is a chain map.

PROOF OF PROPOSITION 50.1. The uniqueness of \cap is easy to see, because if \cap and \cap' are two such maps that both satisfy $\langle \psi \cup \varphi, c \rangle = \langle \psi, \varphi \cap c \rangle = \langle \psi, \varphi \cap' c \rangle$ for all ψ, φ, c , then for every $\varphi \in C^k(X; R)$ and $c \in C_\ell(X; R)$ we have

$$\langle \psi, \varphi \cap c - \varphi \cap' c \rangle = 0 \quad \text{for all } \psi \in C^{\ell-k}(X; R).$$

Since $C_{\ell-k}(X; R)$ is a free R -module, this cannot happen unless $\varphi \cap c - \varphi \cap' c = 0$; see Exercise 50.3 below.

For existence, we can write down a formula for \cap in terms of the diagonal approximation $\Psi : C_*(X) \rightarrow C_*(X) \otimes C_*(X)$, though the formula will take some effort to digest. Fix $\varphi \in C^k(X; R)$ and regard this as a group homomorphism $C_*(X) \rightarrow R$ that is trivial on $C_\ell(X)$ for every $\ell \neq k$. We can then form the composed homomorphism

$$\begin{array}{ccc} C_*(X) & \xrightarrow{\Psi} & C_*(X) \otimes C_*(X) \xrightarrow{\mathbb{1} \otimes \varphi} C_*(X) \otimes R = C_*(X; R). \\ & \searrow \varphi \cap (\cdot) & \nearrow \end{array}$$

Feeding an ℓ -chain $c \in C_\ell(X)$ into this composition produces at first a finite sum

$$\Psi(c) \in \bigoplus_{p+q=\ell} C_p(X) \otimes C_q(X),$$

and the only term in this sum that does not vanish when fed into $\mathbb{1} \otimes \varphi$ is the one with $q = k$ and $p = \ell - k$, thus $\varphi \cap (\cdot)$ maps $C_\ell(X)$ into $C_{\ell-k}(X; R)$. In light of the identification $\text{Hom}(C_*(X), C_*(X; R)) = \text{Hom}_R(C_*(X) \otimes R, C_*(X; R))$, $\varphi \cap (\cdot)$ has a canonical R -linear extension to an R -module homomorphism $C_*(X; R) \rightarrow C_*(X; R)$, and moreover, fixing $c \in C_\ell(X; R)$ instead produces an R -linear map $C^*(X; R) \rightarrow C_*(X; R) : \varphi \mapsto \varphi \cap c$, hence \cap can now be interpreted as either a group homomorphism $C^*(X; R) \otimes C_*(X) \rightarrow C_*(X; R)$ or an R -module homomorphism

$C^*(X; R) \otimes_R C_*(X; R) \rightarrow C_*(X; R)$. The former is easier to write down in an explicit formula, namely

$$(50.4) \quad \varphi \cap c = (\mathbb{1} \otimes \varphi) \circ \Psi(c).$$

Note that since φ may have either odd or even degree, it is understood in this expression that the usual sign convention must be obeyed when evaluating $\mathbb{1} \otimes \varphi$ on a product chain: in particular, if $\varphi \in C^k(X; R)$, $a \in C_{\ell-k}(X)$, $b \in C_k(X)$ and $\psi \in C^{\ell-k}(X; R)$, then

$$\langle \psi, (\mathbb{1} \otimes \varphi)(a \otimes b) \rangle = (-1)^{k(\ell-k)} \langle \psi, a \otimes \varphi(b) \rangle = (-1)^{k(\ell-k)} \psi(a) \varphi(b) = (\psi \otimes \varphi)(a \otimes b),$$

hence the map $C_*(X) \otimes C_*(X) \rightarrow R$ defined by $\langle \psi, (\mathbb{1} \otimes \varphi)(\cdot) \rangle$ is the same as $\psi \otimes \varphi$. As a consequence, for every $c \in C_\ell(X)$ we have

$$\langle \psi, \varphi \cap c \rangle = \langle \psi, (\mathbb{1} \otimes \varphi)(\Psi(c)) \rangle = (\psi \otimes \varphi)(\Psi(c)) = \langle \psi \cup \varphi, c \rangle,$$

and the same relation therefore holds if $c \in C_\ell(X; R)$ and $\langle \cdot, \cdot \rangle$ is the R -bilinear pairing. This proves existence.

The Leibniz rule (50.3) can now be deduced via $\langle \psi \cup \varphi, c \rangle = \langle \psi, \varphi \cap c \rangle$ from the corresponding Leibniz rule for \cup , but this is a slightly annoying computation in which getting all the signs right is tricky,⁷² so let's instead describe a more "highbrow" version of the same argument. The relation between \cup and \cap can be interpreted as saying that the diagram

$$\begin{array}{ccc} C^*(X; R) \otimes_R C^*(X; R) \otimes_R C_*(X; R) & \xrightarrow{\mathbb{1} \otimes \cap} & C^*(X; R) \otimes_R C_*(X; R) \\ \downarrow \cup \otimes \mathbb{1} & & \downarrow \langle \cdot, \cdot \rangle \\ C^*(X; R) \otimes_R C_*(X; R) & \xrightarrow{\langle \cdot, \cdot \rangle} & R \end{array}$$

commutes. Notice that if we view $C^*(X; R)$ as a chain complex with $C^*(X; R)_n = C^{-n}(X; R)$ as in Remark 45.3, then all maps in this diagram other than $\mathbb{1} \otimes \cap$ are already known to be chain maps. The composition $\langle \cdot, \cdot \rangle \circ (\mathbb{1} \otimes \cap)$ is therefore also a chain map, implying that if we take any pair of elements $\varphi \in C^*(X; R)$ and $x \in C^*(X; R) \otimes_R C_*(X; R)$, we will have

$$((\mathbb{1} \otimes \cap) \circ \partial - \partial \circ (\mathbb{1} \otimes \cap))(\varphi \otimes x) \in \ker \langle \cdot, \cdot \rangle.$$

Applying the usual graded Leibniz rule for $\partial(\varphi \otimes x)$ gives

$$((\mathbb{1} \otimes \cap) \circ \partial - \partial \circ (\mathbb{1} \otimes \cap))(\varphi \otimes x) = (-1)^{|\varphi|} \varphi \otimes ((\cap \circ \partial - \partial \circ \cap)(x)),$$

so after plugging this into $\langle \cdot, \cdot \rangle$, we deduce that

$$\langle \varphi, (\cap \circ \partial - \partial \circ \cap)x \rangle = 0$$

holds for all $\varphi \in C^*(X; R)$ and $x \in C^*(X; R) \otimes_R C_*(X; R)$. It follows via Exercise 50.3 that $(\cap \circ \partial - \partial \circ \cap)x = 0$ for all x , i.e. \cap is a chain map. \square

EXERCISE 50.3. Show that if A is a free R -module and $a \in A$ is a nontrivial element, then $\varphi(a) \neq 0$ for some $\varphi \in \text{Hom}_R(A, R)$.

Hint: This is not true in general without the freeness assumption, e.g. it is clearly false for the \mathbb{Z} -module \mathbb{Z}_2 , since $\text{Hom}(\mathbb{Z}_2, \mathbb{Z}) = 0$. Use a basis of A in your proof.

EXERCISE 50.4. Prove that $\cap : C^k(X; R) \otimes_R C_\ell(X; R) \rightarrow C_{\ell-k}(X; R)$ has the following naturality property: for any map $f : X \rightarrow Y$ with a cochain $\varphi \in C^k(Y; R)$ and chain $c \in C_\ell(X; R)$,

$$f_*(f^*\varphi \cap c) = \varphi \cap f_*c.$$

⁷²I got it right on the third try.

Since \cap is a chain map, it descends to a pairing

$$H^k(X; R) \otimes_R H^\ell(X; R) \rightarrow H_{\ell-k}(X; R) : [\varphi] \otimes [c] \mapsto [\varphi] \cap [c] := [\varphi \cap c],$$

which we call the **cap product**. One can use the formula (50.4) and the fact that diagonal approximations are unique up to chain homotopy to prove that different choices of diagonal approximation define the same cap product on the level of homology.

Since the diagonal approximation in the definition of \cap can be chosen freely, we might as well make our “favorite” choice to write down an explicit formula: recall from Lecture 48 the Alexander-Whitney diagonal approximation, defined on singular n -simplices $\sigma : \Delta^n \rightarrow X$ by

$$\Psi(\sigma) = \sum_{k+\ell=n} (\sigma|_{[0,\dots,k]}) \otimes (\sigma|_{[k,\dots,n]}),$$

which leads to the chain-level cup product formula

$$(\varphi \cup \psi)(\sigma) = (-1)^{|\varphi||\psi|} \varphi(\sigma|_{[0,\dots,k]}) \psi(\sigma|_{[k,\dots,n]}).$$

This formula has the convenient feature that it satisfies

$$(50.5) \quad 1 \cup \varphi = \varphi = \varphi \cup 1$$

and

$$(50.6) \quad (\varphi \cup \psi) \cup \eta = \varphi \cup (\psi \cup \eta)$$

for all cochains φ, ψ, η , where $1 \in C^0(X; R)$ is the cochain sending all singular 0-simplices to $1 \in R$,⁷³ which represents the unit element $1 \in H^0(X; R)$. Under other choices of diagonal approximation, these formulas may not hold except “up to chain homotopy” (so that they do of course hold after descending to cohomology). The corresponding explicit formula for \cap in this setting is

$$\varphi \cap \sigma = (-1)^{k(\ell-k)} \varphi(\sigma|_{[\ell-k,\dots,\ell]}) \sigma|_{[0,\dots,\ell-k]} \in C_{\ell-k}(X; R) \quad \text{for } \varphi \in C^k(X; R) \text{ and } \sigma : \Delta^\ell \rightarrow X.$$

We will not use this formula for anything, but now you’ve seen it. The relations (50.5) and (50.6), however, are concretely useful in the following exercise.

EXERCISE 50.5. Use the relation $\langle \psi \cup \varphi, c \rangle = \langle \psi, \varphi \cap c \rangle$ to prove that if \cap is defined using the Alexander-Whitney diagonal approximation, then it satisfies

$$1 \cap c = c \quad \text{for all } c \in C_*(X; R)$$

and

$$(\varphi \cup \psi) \cap c = \varphi \cap (\psi \cap c) \quad \text{for all } \varphi, \psi \in C^*(X; R), c \in C_*(X; R).$$

Letting all these properties descend to the level of homology and cohomology, here is a summary of what we have proved so far about the cap product:

THEOREM 50.6. *The cap product $\cap : H^*(X; R) \otimes_R H_*(X; R) \rightarrow H_*(X; R)$ has the following properties.*

- (1) $f_*(f^*\varphi \cap A) = \varphi \cap f_*A$ for all continuous maps $f : Y \rightarrow X$, $\varphi \in H^*(Y; R)$ and $A \in H_*(X; R)$.
- (2) $1 \cap A = A$ for all $A \in H_*(X; R)$.
- (3) $(\varphi \cup \psi) \cap A = \varphi \cap (\psi \cap A)$ for all $\varphi, \psi \in H^*(X; R)$ and $A \in H_*(X; R)$.
- (4) $\langle 1, \varphi \cap A \rangle = \langle \varphi, A \rangle$ for any $\varphi \in H^*(X; R)$ and $A \in H_*(X; R)$ with the same degree.

□

⁷³In different contexts, we have also previously referred to the cochain $1 : C_0(X) \rightarrow R$ as the *augmentation* of the singular chain complex.

REMARK 50.7. The identity

$$(50.7) \quad \langle \varphi \cup \psi, A \rangle = \langle \varphi, \psi \cap A \rangle \quad \text{for all } \psi \in H^k(X; R), A \in H_\ell(X; R), \varphi \in H^{\ell-k}(X; R)$$

was left out of the above theorem only because it is redundant: it follows from the third and fourth identities by writing

$$\langle \varphi \cup \psi, A \rangle = \langle 1, (\varphi \cup \psi) \cap A \rangle = \langle 1, \varphi \cap (\psi \cap A) \rangle = \langle \varphi, \psi \cap A \rangle.$$

Conversely, the fourth identity is a special case of this one.

The relative cap product takes the form

$$(50.8) \quad \cap : H^*(X, A; R) \otimes_R H_*(X, A \cup B; R) \rightarrow H_*(X, B; R)$$

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

$$C^*(X; R) \otimes C_*(X) \rightarrow C_*(X; R) : \varphi \otimes c \mapsto \varphi \cap c = (\mathbf{1} \otimes \varphi) \circ \Psi(c)$$

always descends to a well-defined map on the relative complexes

$$C^*(X, A; R) \otimes \frac{C_*(X)}{C_*(A) + C_*(B)} \rightarrow C_*(X, B; R),$$

as $\varphi \in C^*(X, A; R)$ 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 $\varphi \cap c \in C_*(B; R)$. Now if $A, B \subset X$ are an excisive couple, the homology of $C_*(X)/(C_*(A) + C_*(B))$ with coefficients in R has a natural identification with $H_*(X, A \cup B; R)$, thus making sense of (50.8).

Topological manifolds and orientations. At this point we've covered the essential properties of products, so the next few lectures will focus on a new topic: the global topology of finite-dimensional topological manifolds.⁷⁴

There is a basic fact about manifolds that was briefly mentioned in the context of the Lefschetz fixed point theorem 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 $r : P \rightarrow M$, and since P necessarily has finitely generated homology, it follows that the same is true for M :

THEOREM 50.8. *For every compact manifold M , $H_*(M)$ is finitely generated.* \square

I would now like to discuss what it means for a topological manifold to be *orientable*. We discussed this somewhat in Lecture 29 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 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.

⁷⁴I 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. This will only change when we discuss intersection theory.

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 33). Recall that if $\dim M = n$, then for every $x \in M$ 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

$$(50.9) \quad \begin{aligned} H_k(M, M \setminus \{x\}; G) &\cong H_k(\mathcal{U}_x, \mathcal{U}_x \setminus \{x\}; G) \cong H_k(\mathbb{R}^n, \mathbb{R}^n \setminus \{0\}; G) \cong H_k(\mathbb{D}^n, \partial \mathbb{D}^n; G) \\ &\cong \tilde{H}_{k-1}(S^{n-1}; G) \cong \begin{cases} G & \text{if } k = n, \\ 0 & \text{otherwise.} \end{cases} \end{aligned}$$

We call $H_n(M, M \setminus \{x\}; G)$ the **local homology** group of M at x , and a **local orientation** of M at x is defined to be a choice of generator

$$[M]_x \in H_n(M, M \setminus \{x\}; \mathbb{Z}) \cong \mathbb{Z}.$$

At every point there are clearly two possible choices of local orientations. The question now is: if we have chosen a local orientation of M at every point $x \in M$, what should it mean to say that these orientations *vary continuously* with x ? The answer emerges from the following observation: suppose $\varphi : \mathcal{U} \xrightarrow{\cong} \mathbb{R}^n$ is a chart defined on some open set $\mathcal{U} \subset M$ and $A \subset \mathcal{U}$ denotes the subset $\varphi^{-1}(\mathbb{D}^n)$. Then for any point $x \in A$, the obvious inclusions of pairs fit into a commutative diagram

$$\begin{array}{ccccc} H_n(M, M \setminus \{x\}) & \xleftarrow{\cong} & H_n(\mathcal{U}, \mathcal{U} \setminus \{x\}) & \xleftarrow[\varphi_*^{-1}]{\cong} & H_n(\mathbb{R}^n, \mathbb{R}^n \setminus \{\varphi(x)\}) \\ & & \uparrow j_{x,A} & & \uparrow \cong \\ H_n(M, M \setminus A) & \xleftarrow{\cong} & H_n(\mathcal{U}, \mathcal{U} \setminus A) & \xleftarrow[\varphi_*^{-1}]{\cong} & H_n(\mathbb{R}^n, \mathbb{R}^n \setminus \mathbb{D}^n), \end{array}$$

in which the map at the right is an isomorphism due to a combination of homotopy equivalence and the five-lemma, proving that the map

$$j_{x,A} : H_n(M, M \setminus A) \rightarrow H_n(M, M \setminus \{x\})$$

is an isomorphism. We shall say in this situation that $A \subset M$ is a **disk-like neighborhood** of $x \in M$.

DEFINITION 50.9. An **orientation** of an n -dimensional topological manifold M is a choice of local orientations $[M]_x$ for every $x \in M$ satisfying the following consistency condition: every $x \in M$ is contained in some disk-like neighborhood $A \subset M$ such that for all $y \in A$,

$$j_{y,A} \circ j_{x,A}^{-1} [M]_x = [M]_y.$$

A manifold equipped with an orientation will be called an **oriented manifold** (*orientierte Mannigfaltigkeit*). In light of (50.9), you can imagine an orientation as a choice for every $x \in M$ of a favorite generator $[S_x] \in \tilde{H}_{n-1}(S_x) \cong \mathbb{Z}$ 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 $\tilde{H}_{n-1}(S_x) \rightarrow \tilde{H}_{n-1}(S_y)$ sending $[S_x]$ to $[S_y]$. You should take a moment to contemplate why this description matches Definition 50.9 in the case $M = \mathbb{R}^n$.

We now have enough language in place to state one of the most important theorems in this course, whose proof will be spread over the next few lectures. Recall that a manifold is called **closed** (*geschlossen*) if it is compact and has no boundary.

THEOREM 50.10. *For every closed oriented topological n -manifold, there exists a distinguished class $[M] \in H_n(M)$ such that the map induced by the inclusion $i^x : (M, \emptyset) \hookrightarrow (M, M \setminus \{x\})$ for every $x \in M$ satisfies*

$$i_*^x [M] = [M]_x,$$

and this map is an isomorphism if M is connected. Moreover, for every $k = 0, 1, \dots, n$, the map

$$H^k(M) \rightarrow H_{n-k}(M) : \varphi \mapsto \varphi \cap [M]$$

is an isomorphism.

The distinguished class $[M] \in H_n(M)$ is called the **fundamental class** of M , and the isomorphism $H^k(M) \rightarrow H_{n-k}(M)$ defined by capping with $[M]$ is the **Poincaré duality** map. We will see that it also works with other choices of coefficients, which sometimes also allow the orientation condition to be relaxed, e.g. since $\tilde{H}_{n-1}(S^{n-1}; \mathbb{Z}_2) \cong \mathbb{Z}_2$ allows only one choice of generator, every manifold is orientable and admits a unique orientation in the sense of \mathbb{Z}_2 coefficients, and the \mathbb{Z}_2 version of Theorem 50.10 is then true for every closed manifold. We'll prove the existence of the fundamental class next time.

51. The orientation bundle

The previous lecture's discussion of orientations can be put in a slightly more general framework. Assume M is a topological n -manifold (without boundary) and G is an abelian group; we will be most interested in cases where G is a ring, but there is no need to impose any such restriction just yet. It will be convenient to introduce the shorthand notation

$$H_*(M | A; G) := H_*(M, M \setminus A; G)$$

for each $A \subset M$. The chain complex underlying $H_*(M | A; G)$ does not see any chains that fail to intersect A , and by subdivision, we can also restrict our attention to arbitrarily "small" chains, which means that $H_*(M | A; G)$ really only depends on the topology of arbitrarily small neighborhoods of A in M . For this reason we sometimes call it the "homology of M restricted to A ". Note that if $B \subset A \subset M$, the identity map on M defines a natural inclusion of pairs $(M, M \setminus A) \hookrightarrow (M, M \setminus B)$, which therefore induces natural homomorphisms

$$j_{B,A} : H_*(M | A; G) \rightarrow H_*(M | B; G).$$

Whenever A is a single point $\{x\} \subset M$, we shall abbreviate

$$\Theta_x^G := H_n(M | x; G) := H_n(M | \{x\}; G) = H_n(M, M \setminus \{x\}; G),$$

which associates to every point $x \in M$ a group Θ_x^G isomorphic to G . For $x \in A \subset M$, we denote the associated homomorphism by

$$j_{x,A} : H_*(M | A; G) \rightarrow \Theta_x^G.$$

DEFINITION 51.1. The **orientation bundle** of M with coefficients in G is the set⁷⁵

$$\Theta^G := \bigcup_{x \in M} \Theta_x^G,$$

endowed with the topology generated by the collection of subsets

$$\mathcal{B} := \{ \mathcal{U}_c \subset \Theta^G \mid \mathcal{U} \subset M \text{ open and } c \in H_n(M | \bar{\mathcal{U}}; G) \},$$

where for $\mathcal{U} \subset M$ and $c \in H_n(M | \bar{\mathcal{U}}; G)$ we define

$$\mathcal{U}_c := \{ j_{x, \bar{\mathcal{U}}}(c) \in \Theta_x^G \mid x \in \mathcal{U} \}.$$

⁷⁵We are regarding Θ_x^G and Θ_y^G as disjoint sets whenever $x \neq y$, so $\Theta^G = \bigcup_{x \in M} \Theta_x^G$ is set-theoretically their disjoint union. I am avoiding writing it as $\coprod_{x \in M} \Theta_x^G$ since this notation normally carries implications about the topology of the union, and those implications would be inconsistent with the topology we actually want to define on Θ^G .


PROPOSITION 51.2. *The collection of subsets $\mathcal{B} = \{\mathcal{U}_c\}$ appearing in Definition 51.1 is the base of a topology on Θ^G for which the natural projection map*

$$p : \Theta^G \rightarrow M$$

sending Θ_x^G 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^G , 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 $c \in \Theta_x^G$, 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 in the previous lecture that $j_{x,\bar{\mathcal{U}}} : H_n(M|\bar{\mathcal{U}};G) \rightarrow \Theta_x^G$ 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 $a \in H_n(M|\bar{\mathcal{U}};G)$ and $b \in H_n(M|\bar{\mathcal{V}};G)$. Then $\mathcal{U}_a \cap \mathcal{V}_b \subset \bigcup_{x \in \mathcal{U} \cap \mathcal{V}} \Theta_x^G$, 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 $H_n(M|A;G)$:



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 $c \in \Theta_x^G$ 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 $p : \Theta^G \rightarrow M$ 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 $j_{x,A} : H_n(M|A;G) \rightarrow H_n(M|x;G)$ factors through $H_n(M|\bar{\mathcal{U}};G)$ for any smaller open neighborhood $\mathcal{U} \subset \mathring{A}$ of x , implying that $j_{x,\bar{\mathcal{U}}} : H_n(M|\bar{\mathcal{U}};G) \rightarrow H_n(M|x;G)$ is also an isomorphism. One can use this to show that for each $x \in \mathring{A}$, assigning the discrete topology to Θ_x^G makes the map

$$(p, j_{x,A}) : p^{-1}(A) \rightarrow A \times \Theta_x^G$$

a homeomorphism. Using it to identify $p^{-1}(A)$ with $A \times \Theta_x^G$ turns $p^{-1}(A) \xrightarrow{p} A$ into the trivial covering map $A \times \Theta_x^G \rightarrow A : (a, c) \mapsto a$. \square

REMARK 51.3. 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^G are groups—this makes $p : \Theta^G \rightarrow M$ a *sheaf* of abelian groups, or if we choose G to be a ring R so that each homology group is an R -module, a sheaf of R -modules. For readers who may know what this

means and find it interesting: $p : \Theta^G \rightarrow M$ is the *completion* of the *presheaf* that associates to each open subset $\mathcal{U} \subset M$ the abelian group $H_n(M | \bar{\mathcal{U}}; G)$.

EXERCISE 51.4. Given a point $x \in X$, let I denote the set of all open neighborhoods of x , and write $\mathcal{U} < \mathcal{V}$ whenever $\mathcal{V} \subset \mathcal{U}$. This makes $(I, <)$ into a directed set, and whenever $\mathcal{U} < \mathcal{V}$ there is an associated homomorphism $j_{\bar{\mathcal{V}}, \bar{\mathcal{U}}} : H_n(M | \bar{\mathcal{U}}; G) \rightarrow H_n(M | \bar{\mathcal{V}}; G)$, so that the collection of abelian groups $\{H_n(M | \bar{\mathcal{U}}; G)\}_{\mathcal{U} \in I}$ forms a direct system. Find a canonical isomorphism

$$\varinjlim \{H_n(M | \bar{\mathcal{U}}; G)\} \xrightarrow{\cong} H_n(M | x; G).$$

DEFINITION 51.5. For each subset $A \subset M$, we denote

$$\Theta^G|_A := p^{-1}(A) \subset \Theta^G$$

and call the covering map $\Theta^G|_A \xrightarrow{p} A$ the **restriction** of the orientation bundle to A . A **section** (*Schnitt*) of Θ^G along A is by definition a continuous map $s : A \rightarrow \Theta^G$ such that $p \circ s = \text{Id}_A$, i.e. it continuously associates to each $x \in A$ an element $s(x) \in \Theta_x^G$. The set of all sections of Θ^G along A will be denoted by $\Gamma(\Theta^G|_A)$, with the special case $A = M$ denoted simply by $\Gamma(\Theta^G)$. We say a section $s \in \Gamma(\Theta^G|_A)$ 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

$$\Gamma_c(\Theta^G|_A) \subset \Gamma(\Theta^G|_A).$$

EXERCISE 51.6. Show that $\Gamma(\Theta^G|_A)$ and $\Gamma_c(\Theta^G|_A)$ are both naturally abelian groups, where addition of sections is defined pointwise, i.e.

$$(s_1 + s_2)(x) := s_1(x) + s_2(x) \in \Theta_x^G.$$

If $G = R$ is a commutative ring with unit, show similarly that $\Gamma(\Theta^G|_A)$ and $\Gamma_c(\Theta^G|_A)$ are modules over R .

For the special case $G = \mathbb{Z}$, we shall denote the orientation bundle by

$$\Theta := \Theta^{\mathbb{Z}}, \quad \Theta_x := \Theta_x^{\mathbb{Z}} = H_n(M | x; \mathbb{Z}) \cong \mathbb{Z}.$$

The definition of orientations in the previous lecture can now be recouched in the following terms.

DEFINITION 51.7. An **orientation** of M along a subset $A \subset M$ is a section $s \in \Gamma(\Theta|_A)$ such that $s(x) = [M]_x$ generates $\Theta_x \cong \mathbb{Z}$ for every $x \in A$. We say M is **orientable** if it admits an orientation.

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** .

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 Θ_x , that is, the two local orientations of M at x . Let us denote

$$\widetilde{M} := \{c \in \Theta \mid c \text{ is a local orientation}\},$$

in other words, \widetilde{M} is the union for all $x \in M$ of the two generators of $\Theta_x \cong \mathbb{Z}$. Assigning to $\widetilde{M} \subset \Theta$ the subspace topology, it is easy to see that the restriction of $p : \Theta \rightarrow M$ defines a two-to-one covering map

$$\pi := p|_{\widetilde{M}} : \widetilde{M} \rightarrow M : [M]_x \mapsto x.$$

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 $\Theta|_A$ for which $-s + s = 0$. The images

of these two sections are disjoint, but by the definition of the topology on Θ , they are both also open subsets of $\pi^{-1}(A) \subset \widetilde{M}$, implying that $\pi^{-1}(A)$ is disconnected. Conversely:

EXERCISE 51.8. If $A \subset M$ is connected and $\pi^{-1}(A) \subset \widetilde{M}$ has more than one connected component, show that each component intersects Θ_x for every $x \in A$. (Hint: Show that the set of $x \in A$ for which Θ_x intersects the component is both open and closed.) Conclude that $\pi^{-1}(A) \subset \widetilde{M}$ therefore has exactly two components, each of which is the image of a section of Θ along A .

Combining the exercise with the previous remarks proves:

PROPOSITION 51.9. *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 51.10. 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 51.11. *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

EXERCISE 51.12. Use the universal coefficient theorem to show that for every abelian group G , there is a natural isomorphism $\Phi_x : \Theta_x \otimes G \rightarrow \Theta_x^G$ for every $x \in M$ such that if $s \in \Gamma(\Theta|_A)$ is a section and $g \in G$, then $s'(x) := \Phi_x(s(x) \otimes g)$ defines a section $s' \in \Gamma(\Theta^G|_A)$. Deduce that if M is orientable along A , then it is also R -orientable along A for every choice of R .

We would now like to formulate a relationship between the group of sections $\Gamma(\Theta^G|_A)$ and the homology group $H_n(M|A; G)$.

LEMMA 51.13. *For every closed subset $A \subset M$, there exists a homomorphism*

$$J_A : H_n(M|A; G) \rightarrow \Gamma_c(\Theta^G|_A) : c \mapsto s_c$$

defined by $s_c(x) := j_{x,A}(c)$ for $x \in A$.

PROOF. We need to show two things about the map $s_c : A \rightarrow \Theta^G|_A$, 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 $[c] \in H_n(M|A; G)$ represented by a relative cycle $c \in C_n(M; G)$ with $\partial c \in C_{n-1}(M \setminus A; G)$, we can write c as a finite linear combination $\sum_i m_i \sigma_i$ of singular n -simplices $\sigma : \Delta^n \rightarrow M$ with coefficients $m_i \in G$. Since Δ^n is compact and the sum is finite, there exists a compact subset $K \subset M$ that contains the images of all the σ_i , so for any $x \in A$ with $x \notin K$, c is an n -chain in $M \setminus \{x\}$, implying that its image under the chain map induced by $(M, M \setminus A) \hookrightarrow (M, M \setminus \{x\})$ is trivial and thus $s_c(x) = j_{x,A}[c] = 0$. The support of s_c is therefore contained in the compact subset $A \cap K \subset A$.

For continuity, we start with the observation that if $A \subset X$ happens to have the property that $j_{x,A}$ is an isomorphism for every $x \in A$, then the same argument as in the proof of Proposition 51.2 identifies $\Theta^G|_A$ with $A \times \Theta_x^G$ so that s_c looks like a "constant section" $x \mapsto (x, g)$ for some $g \in \Theta_x^G$ and is thus obviously continuous. We can reduce the situation to this case as follows. For

$[c] \in H_n(M, M \setminus A; G)$ represented by a relative cycle $c \in C_n(M; G)$ as above, $\partial c \in C_{n-1}(M \setminus A; G)$ is a chain in some compact subset $K \subset M \setminus A$, where $M \setminus A$ is open since A is closed. This implies that every $x \in A$ admits a disk-like neighborhood $\bar{U} \subset M$ disjoint from K , so that ∂c also defines an $(n-1)$ -chain in $M \setminus \bar{U}$ and c can therefore be regarded as a relative n -cycle in $(M, M \setminus \bar{U})$, representing a class

$$[c] \in H_n(M \mid \bar{U}; G).$$

There is then a well-defined and necessarily continuous section of $\Theta^G|_{\mathcal{U}}$ defined by the same formula $x \mapsto j_{x, \bar{U}}([c])$, and our original $s_c : A \rightarrow \Theta^G|_A$ near $x \in A$ is the restriction of this section to $A \cap \mathcal{U}$, which is therefore continuous. \square

Here is the main theorem about the orientation bundle.

THEOREM 51.14. *If M is a topological n -manifold without boundary, then for every closed subset $A \subset M$, the map $J_A : H_n(M \mid A; G) \rightarrow \Gamma_c(\Theta^G|_A)$ is an isomorphism, and $H_k(M \mid A; G) = 0$ for all $k > n$.*

We shall postpone the proof of this theorem until the beginning of the next lecture, and instead focus for now on its corollaries.

COROLLARY 51.15. *Assume M is a connected topological n -manifold without boundary and G is an abelian group. Then:*

- (1) $H_k(M; G) = 0$ for all $k > n$.
- (2) If M is noncompact, then additionally $H_n(M; G) = 0$.

If M is also compact, then the following additional statements hold:

- (3) If M is orientable, then $H_n(M; G) \cong G$.
- (4) If M is not orientable, then $H_n(M; G) \cong \{g \in G \mid 2g = 0\}$.
- (5) For any commutative ring R with unit, any R -orientation $s \in \Gamma(\Theta^R)$ of M determines a generator

$$[M] \in H_n(M; R) \cong R \quad \text{such that} \quad j_{x, M}[M] = s(x) \text{ for all } x \in M.$$

PROOF. We work through the claims one by one:

- (1) Follows from $H_k(M \mid A; G) = 0$ with $A = M$.
- (2) If M is noncompact then $\Gamma_c(\Theta^G) = 0$ since any section with compact support must equal zero somewhere; indeed, continuity then implies that the subset $\{x \in M \mid s(x) = 0\}$ is both open and closed, so it is all of M .
- (3) Taking $A = M$ gives an isomorphism $H_n(M; G) \cong \Gamma(\Theta^G)$, where the compact support condition is irrelevant since M is compact. Then given an orientation $s \in \Gamma(\Theta)$ and the natural isomorphisms $\Phi_x : \Theta_x \otimes G \rightarrow \Theta_x^G$ from the universal coefficient theorem (cf. Exercise 51.12), we obtain an isomorphism

$$G \rightarrow \Gamma(\Theta^G) : g \mapsto s_g \quad \text{where} \quad s_g(x) := \Phi_x(s(x) \otimes g).$$

- (4) For the subgroup $G_0 = \{g \in G \mid 2g = 0\}$, we can again use the isomorphisms $\Phi_x : \Theta_x \otimes G \rightarrow \Theta_x^G$ 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 $\pm[M]_x \in \Theta_x$ 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: in particular, since \tilde{M} is connected, given any $x \in M$, there is no section taking the value $[M]_x \otimes g$ at x for some generator $[M]_x \in \Theta_x$ and $g \in G$ unless $g = -g$.

- (5) This is immediate from the isomorphism $J_M : H_n(M; R) \xrightarrow{\cong} \Gamma(\Theta^R)$ and the definition of an R -orientation. □

DEFINITION 51.16. The generator $[M] \in H_n(M; R) \cong R$ associated to any R -orientation of a closed connected n -manifold M in the above corollary is called the **fundamental class** of M (over R). More generally, the fundamental class of any finite disjoint union of closed, connected and R -oriented n -manifolds M_1, \dots, M_N is defined as

$$[M_1 \amalg \dots \amalg M_N] := i_*^1[M_1] + \dots + i_*^N[M_N]$$

for the inclusions $i_j : M_j \hookrightarrow \coprod_{i=1}^N M_i, j = 1, \dots, N$.

It is not hard to show that for this extension of the fundamental class to a non-connected but closed manifold M with R -orientation $s \in \Gamma(\Theta^R)$, $[M] \in H_n(M; R)$ is uniquely characterized by the condition

$$j_{x,M}[M] = s(x) \quad \text{for all } x \in M.$$

Indeed, uniqueness follows from Theorem 51.14, as the difference between any two classes in $H_n(M; R)$ satisfying this condition is then a class $c \in H_n(M; R)$ such that $J_A(c) = 0 \in \Gamma(\Theta^R)$, implying $c = 0$. In particular, the results of Lecture 34 show that the classes $[M] \in H_n(M; \mathbb{Z})$ or $[M] \in H_n(M; \mathbb{Z}_2)$ that we've previously constructed for closed triangulated n -manifolds by summing the n -simplices in the triangulation match our new definition of fundamental classes. The latter are defined for all closed topological manifolds with orientations and do not require any triangulation.

52. Existence of the fundamental class

Before proving Theorem 51.14, let's discuss one further extension of its range of applications. The previous lecture focused entirely on manifolds without boundary, but there is an easy extension of the fundamental class to compact manifolds M with $\partial M \neq \emptyset$. The construction in Lecture 29 via triangulations lends some intuition on what to expect: $[M]$ will not be an element of $H_n(M; R)$ in general, but a *relative* class in $H_n(M, \partial M; R)$. We first need to be clear on what an *orientation* of M should be if $\partial M \neq \emptyset$. The local homology groups $H_n(M | x; G)$ are the same as usual if $x \notin \partial M$; indeed, the interior $\overset{\circ}{M} \subset M$ is a manifold without boundary that contains x , and excision implies that the inclusion $(\overset{\circ}{M}, \overset{\circ}{M} \setminus \{x\}) \hookrightarrow (M, M \setminus \{x\})$ induces an isomorphism

$$G \cong H_n(\overset{\circ}{M} | x; G) \xrightarrow{\cong} H_n(M | x; G).$$

If $x \in \partial M$ however, then $H_n(M | x; \mathbb{Z})$ is trivial, so the notion of a local orientation at a boundary point does not make sense. (Exercise!) The solution is simply to ignore the boundary points.

DEFINITION 52.1. For M an n -manifold with boundary, an **R -orientation of M** is defined to be an R -orientation of its interior $\overset{\circ}{M}$.

We will need the following basic observation from point-set topology: if ∂M is compact, then it has a so-called **collar neighborhood** (*Kragenumgebung*) in M , meaning a neighborhood $\mathcal{U} \subset M$ of ∂M that is homeomorphic to $(-1, 0] \times \partial M$ via a homeomorphism sending ∂M to $\{0\} \times \partial M$. This is not completely obvious, but the proof is not hard (see e.g. [Hat02, Proposition 3.42]). It follows that M is homotopy equivalent to its interior, hence the latter has finitely generated homology if M is compact.

Since $\overset{\circ}{M}$ is a manifold without boundary, Theorem 51.14 gives an isomorphism

$$J_A : H_n(\overset{\circ}{M}, \overset{\circ}{M} \setminus A; G) \rightarrow \Gamma_c(\Theta^G|_A)$$

for any closed subset $A \subset \mathring{M}$ and abelian group G . Now set G to be a commutative ring R with unit, and assume M has an R -orientation $s \in \Gamma(\Theta^R|_{\mathring{M}})$, which fixes a generator $[M]_x := s(x) \in H_n(M | x; R) \cong R$ for every $x \in \mathring{M}$. We will refer to a relative homology class

$$[M] \in H_n(M, \partial M; R)$$

as a **relative fundamental class** for M if the natural map $i^x : H_n(M, \partial M; R) \rightarrow H_n(M | x; R)$ defined via the inclusion $(M, \partial M) \hookrightarrow (M, M \setminus \{x\})$ for every $x \in \mathring{M}$ satisfies

$$i_*^x[M] = [M]_x.$$

If $\partial M = \emptyset$, this matches our previous characterization of the fundamental class of a closed manifold.

THEOREM 52.2. *If M is a compact manifold with boundary carrying an R -orientation $s \in \Gamma(\Theta^R|_{\mathring{M}})$, then there exists a unique relative fundamental class $[M] \in H_n(M, \partial M; R)$, which is a generator of $H_n(M, \partial M; R) \cong R$ if M is connected.*

PROOF. We shall assume M is connected, as the extension to finite disjoint unions of connected manifolds will follow the same as in the case without boundary.

Identify a neighborhood of ∂M in M with $(-1, 0] \times \partial M$ and for $\epsilon > 0$ small, let $A_\epsilon \subset M$ denote the complement of $(-\epsilon, 0] \times \partial M \subset M$, which is a compact set homotopy equivalent to M . Now if $x \in A_\epsilon$, consider the commuting diagram

$$\begin{array}{ccccc} H_n(M, \partial M; R) & \xrightarrow{\cong} & H_n(M, M \setminus A_\epsilon; R) & \xleftarrow{\cong} & H_n(\mathring{M}, \mathring{M} \setminus A_\epsilon; R) \\ & \searrow^{i_x^*} & \downarrow j_{x, A_\epsilon} & & \downarrow j_{x, A_\epsilon} \\ & & H_n(M, M \setminus \{x\}; R) & \xleftarrow{\cong} & H_n(\mathring{M}, \mathring{M} \setminus \{x\}; R), \end{array}$$

where several maps are labeled as isomorphisms due to homotopy invariance. Given an R -orientation $s \in \Gamma(\Theta^R|_{\mathring{M}})$ and the resulting generator $[M]_x := s(x) \in H_n(M | x; R)$ for each $x \in A_\epsilon$, Theorem 51.14 provides a unique $[M]_\epsilon \in H_n(\mathring{M} | A_\epsilon; R)$ such that $j_{x, A_\epsilon}[M]_\epsilon = [M]_x$ for every $x \in A_\epsilon$, where uniqueness follows from the fact that if two such classes existed, applying J_{A_ϵ} to their difference would give a section that vanishes somewhere and therefore (since M is connected) everywhere. 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 $i_*^x[M] = [M]_x \in H_n(M | x; R)$ for all $x \in A_\epsilon$. We leave it as an exercise to check that this definition of $[M] \in H_n(M, \partial M; R)$ does not depend on the choice of $\epsilon > 0$. (Hint: the isomorphism of Theorem 51.14 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 $H_n(M | A_\epsilon \cap A_\delta; R)$, which is an isomorphism.) Since any $x \in \mathring{M}$ is in A_ϵ for sufficiently small $\epsilon > 0$, the result follows. \square

We still need to prove Theorem 51.14, that the natural map

$$J_A : H_n(M | A; G) \rightarrow \Gamma_c(\Theta^G|_A) : c \mapsto s_c$$

is an isomorphism and $H_k(M | A; G) = 0$ for all $k > n$ whenever M is a topological n -manifold without boundary and $A \subset X$ is a closed subset. 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 (but informal) bit of terminology, we shall call a subset $A \subset M$ in an n -manifold M **convex** if A is contained in a Euclidean neighborhood $\mathcal{U} \subset M$ with a chart $\varphi : \mathcal{U} \xrightarrow{\cong} \mathbb{R}^n$ such that $\varphi(A) \subset \mathbb{R}^n$ is convex.

EXERCISE 52.3. Show that if M is a topological n -manifold, every compact subset $A \subset M$ can be written as $A = \bigcap_{i=1}^{\infty} A_i$, where each $A_i \subset M$ is a finite union of compact convex subsets.
Hint: Show that for any $\epsilon > 0$, A can be covered by a finite union of compact convex sets that each have diameter less than ϵ with respect to some fixed metric on M .

The advantage of convexity is that intersections of convex sets in \mathbb{R}^n are always convex, however, this is not true for convex subsets of a manifold M since two such sets might not lie entirely in the same Euclidean neighborhood. This detail makes an extra step necessary in the following proof, where we will alter our notion of “small” subsets in M to allow *all* compact subsets of Euclidean neighborhoods and not just those that are convex.

PROOF OF THEOREM 51.14. In the following we omit the coefficient group G from the notation wherever possible since it will play no significant role, so e.g. $H_*(M|A)$ and Θ_x should be understood as abbreviations for $H_*(M|A; G)$ and Θ_x^G respectively.

Step 1: We claim that the theorem is true whenever $A \subset M$ is a compact convex subset. Indeed, A is in this case a disk-like neighborhood, so $j_{x,A} : H_n(M, M \setminus A) \rightarrow \Theta_x$ is an isomorphism for every $x \in A$, and since two sections of Θ along a connected subset must be identical whenever they match at one point, this makes $J_A : H_n(M|A) \rightarrow \Gamma(\Theta|_A)$ an isomorphism. For $k > n$, we have the usual computation

$$H_k(M, M \setminus A) \cong H_k(\mathbb{R}^n, \mathbb{R}^n \setminus \mathbb{D}^n) \cong \tilde{H}_{k-1}(S^{n-1}) = 0.$$

Step 2: For the first of three “inductive” steps, we show that if $A, B \subset M$ are two subsets such that the theorem holds for A, B and $A \cap B$, then it also holds for $A \cup B$. The tool required for this is the relative Mayer-Vietoris sequence from Lecture 32. Since A and B are both closed, the complements of these and $A \cap B$ and $A \cup B$ are all open, so we can apply Mayer-Vietoris for the pairs $(M, M \setminus A)$ and $(M, M \setminus B)$, giving a long exact sequence

$$\begin{aligned} \dots \rightarrow H_{k+1}(M, M \setminus A) \oplus H_{k+1}(M, M \setminus B) &\rightarrow H_{k+1}(M, (M \setminus A) \cup (M \setminus B)) \rightarrow H_k(M, (M \setminus A) \cap (M \setminus B)) \\ &\rightarrow H_k(M, M \setminus A) \oplus H_k(M, M \setminus B) \rightarrow H_k(M, (M \setminus A) \cup (M \setminus B)) \rightarrow \dots \end{aligned}$$

We can of course abbreviate this by writing $H_*(M, (M \setminus A) \cap (M \setminus B)) = H_*(M|A \cup B)$ and $H_*(M, (M \setminus A) \cup (M \setminus B)) = H_*(M|A \cap B)$. 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 second map is surjective, as it might not be possible to extend a given section along $A \cap B$ to a section along A or B , but it should be evident that the sequence is exact at all other terms. It can also be checked that the maps in these sequences commute with the natural maps from homology groups to groups of sections, producing a commutative diagram

$$\begin{array}{ccccccc} H_{n+1}(M|A) \oplus H_{n+1}(M|B) & \rightarrow & H_{n+1}(M|A \cap B) & \rightarrow & H_n(M|A \cup B) & \rightarrow & H_n(M|A) \oplus H_n(M|B) \rightarrow H_n(M|A \cap B) \\ \downarrow \cong & & \downarrow \cong & & \downarrow J_{A \cup B} & & \cong \downarrow J_A \oplus J_B & & \cong \downarrow J_{A \cap B} \\ 0 & \longrightarrow & 0 & \longrightarrow & \Gamma_c(\Theta|_{A \cup B}) & \longrightarrow & \Gamma_c(\Theta|_A) \oplus \Gamma_c(\Theta|_B) & \longrightarrow & \Gamma_c(\Theta|_{A \cap B}) \end{array}$$

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. 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 abelian groups, 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 that $H_*(M | A_\infty)$ is in fact the direct limit $\varinjlim \{H_*(M | A_i)\}_{i=1}^\infty$. For this, we need to show that if H is another \mathbb{Z} -graded abelian group with a sequence of morphisms $\Phi_i : H_*(M | A_i) \rightarrow H$ making the diagram

$$\begin{array}{ccccccc} H_*(M | A_1) & \longrightarrow & H_*(M | A_2) & \longrightarrow & H_*(M | A_3) & \longrightarrow & \dots \longrightarrow H_*(M | A_\infty) \\ & & & & \searrow \Phi_1 & & \searrow \Phi_3 \quad \searrow \Phi_\infty \\ & & & & & & \downarrow \\ & & & & & & H \end{array}$$

then the map Φ_∞ indicated by the dashed arrow exists and is unique. Indeed, we can define $\Phi_\infty[c]$ for any given class $[c] \in H_k(M | A_\infty)$ as follows: represent $[c]$ by a relative cycle $c \in C_k(M)$, which means $\partial c \in C_{k-1}(M \setminus A_\infty)$, and note that $M \setminus A_\infty$ contains a compact subset K that contains the images of all singular simplices appearing in ∂c . Since A_∞ is compact, we can then find an open neighborhood $\mathcal{U} \subset M$ of A_∞ that is disjoint from K , and we also have $A_N \subset \mathcal{U}$ for $N \in \mathbb{N}$ sufficiently large. It follows that $\partial c \in C_{k-1}(M \setminus A_N)$, so c is also a relative cycle in $(M, M \setminus A_N)$ and thus defines a class $[c] \in H_k(M | A_N)$. We now define $\Phi_\infty[c]$ as $\Phi_N[c]$ after reinterpreting $[c]$ in this way.

By restricting sections to smaller domains, we also have a sequence of restriction homomorphisms

$$\Gamma(\Theta|_{A_1}) \rightarrow \Gamma(\Theta|_{A_2}) \rightarrow \Gamma(\Theta|_{A_3}) \rightarrow \dots \rightarrow \Gamma(\Theta|_{A_\infty}),$$

and we can use a similar trick to identify $\varinjlim \{\Gamma(\Theta|_{A_i})\}_{i=1}^\infty$ with $\Gamma(\Theta|_{A_\infty})$. Indeed, the problem now is to show that any sequence of homomorphisms $\varphi_i : \Gamma(\Theta|_{A_i}) \rightarrow H$ as in the diagram

$$\begin{array}{ccccccc} \Gamma(\Theta|_{A_1}) & \longrightarrow & \Gamma(\Theta|_{A_2}) & \longrightarrow & \Gamma(\Theta|_{A_3}) & \longrightarrow & \dots \longrightarrow \Gamma(\Theta|_{A_\infty}) \\ & & & & \searrow \varphi_1 & & \searrow \varphi_3 \quad \searrow \varphi_\infty \\ & & & & & & \downarrow \\ & & & & & & H \end{array}$$

gives rise to a unique map $\varphi_\infty : \Gamma(\Theta|_{A_\infty}) \rightarrow H$. The key here is the observation that since $p : \Theta \rightarrow M$ is a covering map, A_∞ has an open neighborhood $\mathcal{U} \subset M$ such that every section $A \rightarrow \Theta$ has a unique extension over \mathcal{U} , which is therefore defined on A_N for $N \in \mathbb{N}$ sufficiently large. The desired map φ_∞ is thus defined on any $s \in \Gamma(\Theta|_{A_\infty})$ by extending s to A_N and then applying φ_N .

With these preliminaries in place, we can combine both direct systems into a commuting diagram

$$\begin{array}{ccccccc} H_n(M | A_1) & \longrightarrow & H_n(M | A_2) & \longrightarrow & H_n(M | A_3) & \longrightarrow & \dots \longrightarrow H_n(M | A_\infty) \\ \cong \downarrow J_{A_1} & & \cong \downarrow J_{A_2} & & \cong \downarrow J_{A_3} & & \downarrow J_{A_\infty} \\ \Gamma(\Theta|_{A_1}) & \longrightarrow & \Gamma(\Theta|_{A_2}) & \longrightarrow & \Gamma(\Theta|_{A_3}) & \longrightarrow & \dots \longrightarrow \Gamma(\Theta|_{A_\infty}), \end{array}$$

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_∞} , 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_∞} .

Step 4: Everything we've proved so far applies only to *compact* subsets $A \subset M$, but the third inductive step introduces noncompact subsets by allowing infinite disjoint unions. Let us call a collection of compact subsets $\{A_\alpha \subset M\}_{\alpha \in I}$ **separated** if they admit a collection of open neighborhoods $\{A_\alpha \subset \mathcal{U}_\alpha \subset M\}_{\alpha \in I}$ such that $\mathcal{U}_\alpha \cap \mathcal{U}_\beta = \emptyset$ 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 $A := \bigcup_{\alpha \in I} A_\alpha$. The point of the separation condition is that if we write $\mathcal{U} := \bigcup_{\alpha \in I} \mathcal{U}_\alpha$, then $(\mathcal{U}, \mathcal{U} \setminus A) \cong \coprod_{\alpha} (\mathcal{U}_\alpha, \mathcal{U}_\alpha \setminus A_\alpha)$, 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

$$\begin{array}{ccccccc} H_n(M|A) & \xleftarrow{\cong} & H_n(\mathcal{U}|A) & \xleftarrow{\cong} & \bigoplus_{\alpha} H_n(\mathcal{U}_\alpha|A_\alpha) & \xrightarrow{\cong} & \bigoplus_{\alpha} H_n(M|A_\alpha) \\ \downarrow J_A & & \downarrow J_A & & \downarrow \bigoplus_{\alpha} J_{A_\alpha} & & \downarrow \bigoplus_{\alpha} J_{A_\alpha} \\ \Gamma_c(\Theta|A) & \xlongequal{\quad} & \Gamma_c(\Theta|A) & \xleftarrow{\cong} & \bigoplus_{\alpha} \Gamma_c(\Theta|A_\alpha) & \xlongequal{\quad} & \bigoplus_{\alpha} \Gamma_c(\Theta|A_\alpha), \end{array}$$

where the isomorphism $\bigoplus_{\alpha} \Gamma_c(\Theta|A_\alpha) \rightarrow \Gamma_c(\Theta|A)$ sends each $\sum_{\alpha} s_{\alpha} \in \bigoplus_{\alpha} \Gamma_c(\Theta|A_\alpha)$ to the unique section $s \in \Gamma_c(\Theta|A)$ such that $s|_{A_\alpha} = s_{\alpha}$ for every $\alpha \in I$. (Note that s necessarily has compact support since only finitely many of the summands in $\sum_{\alpha} s_{\alpha}$ can be nonzero.) 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 Exercise 52.3, any such set is the intersection of a nested sequence of sets that are each finite unions of 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 convex, so combining steps 1 and 2 proves that the theorem holds for all the finite unions of convex sets, and step 3 then establishes it for A .

Step 6: We extend the theorem to arbitrary compact subsets $A \subset M$. In light of Exercise 52.3, this now follows directly from steps 5, 2 and 3, as A is the intersection of a nested sequence of compact sets that are each finite unions of sets contained in Euclidean neighborhoods. (The fact that those sets can be assumed convex is no longer relevant, but since any intersection between them is contained in a Euclidean neighborhood, step 5 now replaces step 1.)

Step 7: The extension of the theorem to an arbitrary closed $A \subset M$ can now be achieved as follows. I need to appeal to a slightly nontrivial point-set topological fact about manifolds: every finite-dimensional topological manifold M has a one-point compactification M^* that is *metrizable*. Recall that the one-point compactification of any space X is defined as the union of X with one extra point $X^* := X \cup \{\infty\}$, where a subset of X^* is considered open if it is either an open set in X or takes the form $(X \setminus K) \cup \{\infty\}$ 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. Problem Set 4 #6 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 can now conclude from step 2 that the theorem also holds for $A = B \cup C$. \square

53. Poincaré duality

The classical perspective on Poincaré duality is demonstrated by Figure 24. 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

$$C_*^\Delta(M) \rightarrow C_*^{\text{CW}^*}(M),$$

therefore giving rise to an isomorphism $H_k(M) \cong H^{n-k}(M)$ for each $k = 0, \dots, n$.

REMARK 53.1. 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




FIGURE 24. A triangulation of a surface and its dual cell decomposition.

isomorphism $C_*^\Delta(M) \rightarrow C_{CW}^{n-*}(M)$. 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 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 natural 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 53.2 (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*

$$H^k(M; R) \xrightarrow{\text{PD}} H_{n-k}(M; R) : \varphi \mapsto \varphi \cap [M]$$

is an isomorphism for every $k \in \mathbb{Z}$.

Before getting into the proof, let's pick some low-hanging fruit and state a few corollaries. Recall that by the universal coefficient theorem, the Betti numbers of a space can be expressed as ranks of either the homology or the cohomology groups, which are the same in corresponding degrees. Poincaré duality thus gives a nontrivial relation between them:

COROLLARY 53.3. For every closed orientable n -manifold M ,

$$b_k(M) = b_{n-k}(M)$$

for all $k \in \mathbb{Z}$. Moreover, without any orientability assumption, the same relation also holds for the so-called “ \mathbb{Z}_2 Betti numbers,” i.e.

$$\dim_{\mathbb{Z}_2} H_k(M; \mathbb{Z}_2) = \dim_{\mathbb{Z}_2} H_{n-k}(M; \mathbb{Z}_2)$$

for all $k \in \mathbb{Z}$. □

COROLLARY 53.4. Every closed odd-dimensional manifold M satisfies $\chi(M) = 0$.

PROOF. In the oriented case, this follows because $b_k(M)$ and $b_{n-k}(M)$ cancel each other in the alternating sum that defines $\chi(M)$.

If M is not orientable but has a cell decomposition, then here are two possible arguments. First, Proposition 39.10 identifies $\chi(M)$ with the alternating sum of the numbers of cells, which does not change if we switch to \mathbb{Z}_2 -coefficients, so $\chi(M) = \sum_{k \in \mathbb{Z}} (-1)^k \dim_{\mathbb{Z}_2} H_k(M; \mathbb{Z}_2) = 0$ due to the \mathbb{Z}_2 case of Corollary 53.3. Alternatively, one could argue in terms of the orientation double cover $\pi : \widetilde{M} \rightarrow M$, which satisfies $\chi(\widetilde{M}) = 2\chi(M)$ by Theorem 39.16 but also $\chi(\widetilde{M}) = 0$ since \widetilde{M} is in this case also a closed n -manifold and is *always* orientable. (In fact, \widetilde{M} is the space of local orientations of M and thus has a “tautological” orientation.)

Both arguments actually suffice even if M does not have a cell decomposition, due to an important (but difficult) theorem of [KS69] stating that every compact topological manifold is homotopy equivalent to a finite CW-complex. In the case of the covering argument, one can use the lifting theorem to produce a commutative diagram that relates $\pi : \widetilde{M} \rightarrow M$ via homotopy equivalences to a similar double cover of CW-complexes, thus proving that $\chi(\widetilde{M}) = 2\chi(M)$ holds even for topological manifolds M that are not CW-complexes.⁷⁶ □

Poincaré duality also provides considerable information about the ring structure of $H^*(M; R)$ as a consequence of the relation $\langle \psi \cup \varphi, [M] \rangle = \langle \psi, \varphi \cap [M] \rangle$. For each $k = 0, \dots, n$, consider the quadratic form

$$\begin{aligned} H^k(M; R) \otimes_R H^{n-k}(M; R) &\xrightarrow{Q} R \\ \varphi \otimes \psi &\mapsto Q(\varphi, \psi) := \langle \varphi \cup \psi, [M] \rangle. \end{aligned}$$

For reasons that we will discuss in the next lecture, this is called the **intersection form** on M . In the case $R = \mathbb{Z}$, $Q(\varphi, \psi)$ vanishes whenever either φ or ψ is torsion, thus it descends to a quadratic form on the free part $H_{\text{free}}^*(M) := H^*(M)/\text{torsion}$,

$$Q : H_{\text{free}}^k(M) \otimes H_{\text{free}}^{n-k}(M) \rightarrow \mathbb{Z}.$$

For a general pair of abelian groups A and B , a bilinear map $Q : A \times B \rightarrow G$ (or equivalently a group 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. There is an obvious analogue of this definition for R -modules and R -module homomorphisms.

⁷⁶Thanks to Levent Kotan for suggesting the argument via the orientation double cover.

COROLLARY 53.5. For any closed n -manifold M with a \mathbb{K} -orientation and corresponding fundamental class $[M] \in H_n(M; \mathbb{K})$ for some field \mathbb{K} , the intersection form

$$Q : H^k(M; \mathbb{K}) \otimes_{\mathbb{K}} H^{n-k}(M; \mathbb{K}) \rightarrow \mathbb{K}$$

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 quadratic form $H_{\text{free}}^k(M; \mathbb{Z}) \otimes H_{\text{free}}^{n-k}(M; \mathbb{Z}) \rightarrow \mathbb{Z}$.

PROOF. With integer coefficients, we saw in Lecture 47 that the canonical map $h : H_{\text{free}}^{n-k}(M) \rightarrow \text{Hom}(H_{\text{free}}^{n-k}(M), \mathbb{Z}) : \varphi \mapsto \langle \varphi, \cdot \rangle$ is an isomorphism. Since the duality map $\text{PD} : H^k(M) \rightarrow H_{n-k}(M)$ 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

$$\begin{array}{ccc} H_{\text{free}}^{n-k}(M) & \xrightarrow[\cong]{h} & \text{Hom}(H_{n-k}^{\text{free}}(M), \mathbb{Z}) & \xrightarrow[\cong]{\text{PD}^*} & \text{Hom}(H_{\text{free}}^k(M), \mathbb{Z}) \\ & \searrow & \Phi & \nearrow & \\ & & \cong & & \end{array}$$

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 53.6. If M is a closed oriented n -manifold and $\varphi \in H^k(M; \mathbb{Z})$ is a primitive⁷⁷ non-torsion 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 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

EXERCISE 53.7. We can now compute the ring structure of $H^*(\mathbb{C}\mathbb{P}^n)$. 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 $\alpha_k \in H^{2k}(\mathbb{C}\mathbb{P}^n) \cong \mathbb{Z}$ denote the generator that evaluates to 1 on the generator of $H_{2k}(\mathbb{C}\mathbb{P}^n)$ represented by the $2k$ -cell.

- (a) Use Corollary 53.6 to prove $\alpha_k \cup \alpha_{n-k} = \pm \alpha_n$ for every k .
- (b) Generalize part (a) to show that $\alpha_k \cup \alpha_\ell = \pm \alpha_{k+\ell}$ for every $k, \ell \in \mathbb{N}$ with $k + \ell \leq n$.

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 $H^*(\mathbb{C}\mathbb{P}^n)$ is generated by the single element $\alpha := \alpha_1 \in H^2(\mathbb{C}\mathbb{P}^n)$, subject only to the relation $\alpha^{n+1} = 0$ since $H^k(\mathbb{C}\mathbb{P}^n) = 0$ for all $k > 2n$. We conclude that there is an

⁷⁷i.e. φ is not $m\psi$ for any $\psi \in H^k(M; \mathbb{Z})$ and an integer $m \geq 2$

isomorphism of \mathbb{Z} -graded rings⁷⁸

$$H^*(\mathbb{C}\mathbb{P}^n) \cong \mathbb{Z}[\alpha]/(\alpha^{n+1}), \quad |\alpha| = 2,$$

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 $\mathbb{C}\mathbb{P}^n \hookrightarrow \mathbb{C}\mathbb{P}^\infty$ to find a graded ring isomorphism $H^*(\mathbb{C}\mathbb{P}^\infty) \cong \mathbb{Z}[\alpha]$, where again $|\alpha| = 2$.

REMARK 53.8. The computation in Exercise 53.7 fills in the last remaining gap in our proof from Lecture 45 (see Theorem 45.1) that all maps $f : \mathbb{C}\mathbb{P}^n \rightarrow \mathbb{C}\mathbb{P}^n$ have fixed points when n is even.

EXERCISE 53.9. Compute each of the following cohomology rings:

- (a) $H^*(\mathbb{R}\mathbb{P}^n; \mathbb{Z}_2) \cong \mathbb{Z}_2[\alpha]/(\alpha^{n+1})$ with $|\alpha| = 1$.
 (b) $H^*(\mathbb{R}\mathbb{P}^\infty; \mathbb{Z}_2) \cong \mathbb{Z}_2[\alpha]$ with $|\alpha| = 1$.

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.

Suppose M is an n -manifold without boundary, not necessarily compact, but endowed with an R -orientation $s \in \Gamma(\Theta^R)$. This section does not have compact support if M is noncompact, but if we choose a compact subset $K \subset M$, then $s|_K \in \Gamma(\Theta^R|_K)$ trivially does have compact support and therefore corresponds under Theorem 51.14 to a distinguished homology class

$$[M]_K := J_K^{-1}(s) \in H_n(M|K; R).$$

Recall from Lecture 50 that there is a relative cap product pairing

$$\cap : H^k(M, M \setminus K; R) \otimes_R H_n(M, M \setminus K; R) \rightarrow H_n(M; R),$$

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

$$\text{PD}_K : H^k(M|K; R) \rightarrow H_n(M; R) : \varphi \mapsto \varphi \cap [M]_K.$$

Now consider what happens to this map if we replace K by a larger compact subset $K' \subset M$ that contains K : first, since $[M]_K \in H_n(M|K; R)$ and $[M]_{K'} \in H_n(M|K'; R)$ are determined by the same globally-defined section $s \in \Gamma(\Theta^R)$, 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 of the cap product (i.e. Theorem 50.6(1)) then implies that for all $\varphi \in H^k(M|K; R)$,

$$i_*(i^*\varphi \cap [M]_{K'}) = \text{PD}_{K'}(i^*\varphi) = \varphi \cap [M]_K = \text{PD}_K(\varphi),$$

⁷⁸A \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 $ab \in R_{k+\ell}$.

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

$$(53.1) \quad \begin{array}{ccc} H^k(M | K; R) & \xrightarrow{i^*} & H^k(M | K'; R) \\ & \searrow \text{PD}_K & \downarrow \text{PD}_{K'} \\ & & H_{n-k}(M; R), \end{array}$$

which means that we can view the maps $\text{PD}_K : H^k(M | K; R) \rightarrow H_{n-k}(M; R)$ as defining a target of a direct system of abelian groups $\{H^k(M | K; R)\}_K$ 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

$$\text{PD} : \varinjlim \{H^k(M | K; R)\}_K \rightarrow H_{n-k}(M; R).$$

DEFINITION 53.10. For any space X , we define the **compactly supported cohomology** of X as the direct limit

$$H_c^*(X; G) := \varinjlim \{H^*(X | K; G)\}_K,$$

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.

EXERCISE 53.11. Show that if M is compact, there is a natural isomorphism $H_c^*(M; R) \cong H^*(M; R)$ which identifies the map $\text{PD} : H_c^k(M; R) \rightarrow H_{n-k}(M; R)$ defined above with the usual map $\varphi \mapsto \varphi \cap [M]$.

EXERCISE 53.12. In the following, suppose G is any abelian group.

- (a) Prove that $H_c^n(\mathbb{R}^n; G) \cong G$ and $H_c^k(\mathbb{R}^n; G) = 0$ for all $k \neq n$.
- (b) Construct a canonical isomorphism between $H_c^*(X; G)$ and the homology of the subcomplex $C_c^*(X; G) \subset C^*(X; G)$ consisting of every cochain $\varphi : C_k(X) \rightarrow G$ 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** 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. Give an example showing that this need not be true if X and Y are only homotopy equivalent.
- (e) In contrast to part (c), show that $H_c^*(\cdot; G)$ does *not* define a functor on the usual category of topological spaces with morphisms defined to be continuous (but not necessarily proper) maps.

Hint: Think about maps between \mathbb{R}^n and the one-point space.

Here is the noncompact version of Poincaré duality, which has the compact version as a corollary in light of Exercise 53.11.

THEOREM 53.13. 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.

The proof will proceed by a form of induction similar to the construction of the fundamental class in the previous lecture. We start with a purely local result to begin the induction.

NOTATION. Since the coefficient ring R plays no significant role in the proof of Theorem 53.13, we shall omit it from the notation wherever possible for the rest of this lecture.

LEMMA 53.14. For either choice of orientation of \mathbb{R}^n , the map $\text{PD} : H_c^k(\mathbb{R}^n) \rightarrow H_{n-k}(\mathbb{R}^n)$ is an isomorphism for every $k \in \mathbb{Z}$.

PROOF. There is an obvious cofinal family⁷⁹ 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 $H^n(\mathbb{R}^n|\mathbb{D}^n) \cong \text{Hom}_R(H_n(\mathbb{R}^n|\mathbb{D}^n), R)$ by evaluation of cohomology classes on homology classes, so that $\langle \varphi, \cdot \rangle$ generates $\text{Hom}_R(H_n(\mathbb{R}^n|\mathbb{D}^n), R)$ 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 53.12 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 $H^*(X \setminus K) \rightarrow H^*(X \setminus L)$ induced by the inclusion $(X, X \setminus L) \hookrightarrow (X, X \setminus K)$ produces a map $H^*(A|K) \rightarrow H^*(X|L)$:

$$\begin{array}{ccc} & H^*(X|K) & \\ & \swarrow \cong & \searrow \\ H^*(A|K) & \xrightarrow{\quad} & H^*(X|L). \end{array}$$

⁷⁹In a directed set $(I, <)$, a subset $S \subset I$ is called a **cofinal family** if for every $\alpha \in I$, there exists a $\beta \in S$ such that $\alpha < \beta$.

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

$$H^*(A) \rightarrow H_c^*(X).$$

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 53.15. *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 (53.1), 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). 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

$$\begin{array}{ccc} & (M, M \setminus K) & \\ \swarrow & & \searrow \\ (M, M \setminus (K \cup L)) & & (M, M \setminus (K \cap L)) \\ \searrow & & \swarrow \\ & (M, M \setminus L) & \end{array}$$

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:

$$(53.2) \quad \begin{array}{ccccccc} \dots & \longrightarrow & H^k(M|K \cap L) & \longrightarrow & H^k(M|K) \oplus H^k(M|L) & \longrightarrow & H^k(M|K \cup L) & \longrightarrow & H^{k+1}(M|K \cap L) & \longrightarrow & \dots \\ & & \downarrow \cong & & \downarrow \cong & & \parallel & & \downarrow \cong & & \\ \dots & \longrightarrow & H^k(A \cap B|K \cap L) & \longrightarrow & H^k(A|K) \oplus H^k(B|L) & \longrightarrow & H^k(M|K \cup L) & \longrightarrow & H^{k+1}(A \cap B|K \cap L) & \longrightarrow & \dots \\ & & \downarrow \text{PD}_{K \cap L} & & \downarrow \text{PD}_K \oplus \text{PD}_L & & \downarrow \text{PD}_{K \cup L} & & \downarrow \text{PD}_{K \cap L} & & \\ \dots & \longrightarrow & H_{n-k}(A \cap B) & \longrightarrow & H_{n-k}(A) \oplus H_{n-k}(B) & \longrightarrow & H_{n-k}(M) & \longrightarrow & H_{n-k-1}(A \cap B) & \longrightarrow & \dots \end{array}$$

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 53.15 imply that this diagram commutes, with the possible exception of the bottom right square involving connecting homomorphisms. It turns out that this square also commutes, and the proof is not especially deep, but it is a tedious chain-level calculation involving barycentric subdivision, so we will skip it and simply refer to [Hat02, pp. 246–247]. The result is:

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 53.16. *The diagram in (53.2) 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*

$$(53.3) \quad \begin{array}{ccccccc} \dots & \rightarrow & H_c^k(A \cap B) & \longrightarrow & H_c^k(A) \oplus H_c^k(B) & \longrightarrow & H_c^k(M) \longrightarrow H_c^{k+1}(A \cap B) \rightarrow \dots \\ & & \downarrow \text{PD} & & \downarrow \text{PD} \oplus \text{PD} & & \downarrow \text{PD} & & \downarrow \text{PD} \\ \dots & \rightarrow & H_{n-k}(A \cap B) & \rightarrow & H_{n-k}(A) \oplus H_{n-k}(B) & \rightarrow & H_{n-k}(M) \rightarrow H_{n-k-1}(A \cap B) \rightarrow \dots \end{array}$$

in which both rows are exact.

SKETCH OF THE PROOF. Aside from the tedious verification that (53.2) commutes, the claim that the top row of (53.3) is exact is slightly nontrivial: this follows from the general fact that direct limits of exact sequences are always exact. Indeed, we proved in Proposition 39.1 that the functor $H_* : \text{Chain} \rightarrow \text{Ab}_{\mathbb{Z}}$ is continuous under direct limits, and an exact sequence is nothing other than a chain complex with trivial homology. (Recall from Lecture 44 however that the analogous statement for inverse limits is false, so this detail should not be taken for granted.) \square

Applying the five-lemma now gives:

COROLLARY 53.17. *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 53.14 and 53.16 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 53.18. *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 groups $H_{n-k}(\mathcal{U}_i)$ 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 38.20 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 abelian group A such that the diagram

$$\begin{array}{ccccccc}
 H_c^*(\mathcal{U}_1) & \longrightarrow & H_c^*(\mathcal{U}_2) & \longrightarrow & H_c^*(\mathcal{U}_3) & \longrightarrow & \dots \longrightarrow H_c^*(M) \\
 & & & & & \searrow f_3 & \downarrow f_\infty \\
 & & & & & & A \\
 & & & \nearrow f_1 & & & \\
 & & & & & &
 \end{array}$$

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 53.15, we now obtain a commutative diagram

$$\begin{array}{ccccccc}
 H_c^k(\mathcal{U}_1) & \longrightarrow & H_c^k(\mathcal{U}_2) & \longrightarrow & H_c^k(\mathcal{U}_3) & \longrightarrow & \dots \\
 \downarrow \text{PD} & & \downarrow \text{PD} & & \downarrow \text{PD} & & \\
 H_{n-k}(\mathcal{U}_1) & \longrightarrow & H_{n-k}(\mathcal{U}_2) & \longrightarrow & H_{n-k}(\mathcal{U}_3) & \longrightarrow & \dots
 \end{array}
 \tag{53.4}$$

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:

$$\varinjlim \{H_c^k(\mathcal{U}_i)\}_{i=1}^\infty = H_c^k(M) \xrightarrow[\cong]{\text{PD}} H_{n-k}(M) = \varinjlim \{H_{n-k}(\mathcal{U}_i)\}_{i=1}^\infty.$$

□

PROOF OF THEOREM 53.13. Lemmas 53.14 and 53.16 prove the theorem for all finite unions of convex open sets \mathbb{R}^n , and feeding this into Lemma 53.18 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 53.16 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 53.18. □

EXERCISE 53.19. Assume M is a compact R -oriented n -manifold with boundary and $[M] \in H_n(M, \partial M; R)$ is the resulting relative fundamental class. The relative cap product with $[M]$ then gives rise to two natural maps

$$\text{PD} : H^k(M, \partial M; R) \rightarrow H_{n-k}(M; R),
 \tag{53.5}$$

$$\text{PD} : H^k(M; R) \rightarrow H_{n-k}(M, \partial M; R),
 \tag{53.6}$$

both defined by $\text{PD}(\varphi) = \varphi \cap [M]$. The theorem that both are isomorphisms is sometimes called **Lefschetz duality**.

- (a) Find a cofinal family of compact subsets $A \subset \overset{\circ}{M}$ such that the natural maps in the diagram

$$H^*(\overset{\circ}{M} | A; R) \longleftarrow H^*(M | A; R) \longrightarrow H^*(M, \partial M; R)$$

are isomorphisms. Use this to find a natural isomorphism (cf. Exercise 38.9)

$$H_c^*(M; R) \cong H^*(M, \partial M, R),$$

and deduce via Theorem 53.13 that (53.5) 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

$$\begin{array}{cccccccc} \dots & \rightarrow & H^k(M, \partial M; R) & \xrightarrow{j^*} & H^k(M; R) & \xrightarrow{i^*} & H^k(\partial M; R) & \xrightarrow{\delta^*} & H^{k+1}(M, \partial M; R) & \rightarrow & \dots \\ & & \downarrow \cdot \cap [M] & & \downarrow \cdot \cap [M] & & \downarrow \cdot \cap [\partial M] & & \downarrow \cdot \cap [M] & & \\ \dots & \rightarrow & H_{n-k}(M; R) & \xrightarrow{j_*} & H_{n-k}(M, \partial M; R) & \xrightarrow{\partial_*} & H_{n-k-1}(\partial M; R) & \xrightarrow{i_*} & H_{n-k-1}(M; R) & \rightarrow & \dots \end{array}$$

where $i : \partial M \hookrightarrow M$ and $j : (M, \emptyset) \hookrightarrow (M, \partial M)$ denote the usual inclusions.

Hint: Work directly with chains and cochains. It helps to know that if $c \in C_n(M; R)$ is a relative n -cycle representing $[M] \in H_n(M, \partial M; R)$, then the $(n - 1)$ -cycle $\partial c \in C_{n-1}(\partial M; R)$ represents $[\partial M] \in H_{n-1}(\partial M; R)$; see Problem Set 12 #3.

- (c) Deduce from the diagram in part (b) that the map in (53.6) is also an isomorphism.
 (d) If M has a triangulation, interpret the isomorphisms (53.5) and (53.6) in terms of the dual cell decomposition.

REMARK 53.20. Here is the promised addendum to Remark 53.1. When M is compact and has an oriented triangulation, $C_\Delta^k(M)$ has an obvious identification with the free abelian group generated by all the k -simplices in the triangulation: indeed, if we fix an orientation on each k -simplex and call $\mathcal{K}_k(M)$ the resulting set of oriented k -simplices so that $C_k^\Delta(M) = \bigoplus_{\sigma \in \mathcal{K}_k(M)} \mathbb{Z}$, then the dual elements $\varphi_\sigma : C_k^\Delta(M) \rightarrow \mathbb{Z}$ 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 $C_\Delta^k(M)$. In this case, we obtain a chain isomorphism

$$C_\Delta^k(M) \rightarrow C_{n-k}^{\text{CW}}(M)$$

by sending each of the k -cochains φ_σ to the $(n - k)$ -cell dual to σ , and the isomorphism $H^k(M) \cong H_{n-k}(M)$ follows. The trouble if M is not compact is that $C_k^\Delta(M)$ is now an infinitely-generated free abelian group, so its dual $C_\Delta^k(M)$ is not isomorphic to it, but is actually much larger: the cochains φ_σ do not form a basis for $C_\Delta^k(M)$ since they only span the subgroup of homomorphisms $C_\Delta^k(M) \rightarrow \mathbb{Z}$ that are nonzero on finitely many simplices. As a consequence, $C_\Delta^k(M)$ and $C_{n-k}^{\text{CW}}(M)$ 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 $C_\Delta^k(M)$, whose homology is the simplicial version of $H_c^*(M)$.

54. The intersection product

Today's topic is an addendum to Poincaré duality: I want to describe the natural product structure on *homology* that arises from the combination of Poincaré duality with the cup product. Unlike the cup product, the product on $H_*(X)$ will not be defined for arbitrary spaces X , but makes sense only when Poincaré duality holds, i.e. when X is a closed manifold with an orientation over the chosen coefficient ring. This is a bit restrictive, but the restriction pays off: in fact, if we restrict further and assume X is a *smooth* manifold, then the intersection product provides the nicest possible geometric interpretation of the cup product, namely as something that measures (in homological terms) the intersection between submanifolds.

Our standing assumptions throughout this lecture are as follows: M is a closed, connected and R -oriented smooth manifold of dimension n , where the coefficient ring R will always be either \mathbb{Z} , \mathbb{Z}_2 , \mathbb{Q} or \mathbb{R} . One can allow more general choices for R , but these are the main ones of interest. The assumption of an R -orientation actually just means that M is oriented if R is \mathbb{Z} , \mathbb{Q} or \mathbb{R} , and in the case $R = \mathbb{Z}_2$ it is a vacuous assumption. The smoothness assumption can also be relaxed somewhat, but it is the quickest way to achieve the conditions that we will actually need at various points: namely that suitable intersections $A \cap B$ between two submanifolds $A, B \subset M$ are also submanifolds (of the “correct” dimension), and that neighborhoods of submanifolds can be identified with certain vector bundles (the tubular neighborhood theorem). For this reason, the present lecture will assume some knowledge of the basic theory of smooth manifolds and their tangent spaces.

DEFINITION 54.1. The **intersection product** on M associates to each pair of integers $k, \ell = 0, \dots, n$ a bilinear map

$$H_{n-k}(M; R) \otimes_R H_{n-\ell}(M; R) \rightarrow H_{n-(k+\ell)}(M; R) : A \otimes B \mapsto A \cdot B$$

uniquely defined by the condition

$$\text{PD}(\varphi) \cdot \text{PD}(\psi) = \text{PD}(\psi \cup \varphi),$$

for $\varphi \in H^k(M; R)$ and $\psi \in H^\ell(M; R)$, where $\text{PD} : H^m(M; R) \rightarrow H_{n-m}(M; R)$ denotes the Poincaré duality isomorphism. In the case $k + \ell = n$, we use the canonical isomorphism $H_0(M; R) \xrightarrow{\cong} R : c \mapsto \langle 1, c \rangle$ to regard $A \cdot B$ as a number in R , the **intersection number** between A and B ,

$$H_{n-k}(M; R) \otimes_R H_k(M; R) \rightarrow R : A \otimes B \mapsto A \cdot B.$$

The intersection number is equivalent to what we called the *intersection form* $Q : H^k(M; R) \otimes_R H^{n-k}(M; R) \rightarrow R$ in the previous lecture: if $\varphi \in H^k(M; R)$ and $\psi \in H^{n-k}(M; R)$ have Poincaré dual classes $A := \text{PD}(\varphi) \in H_{n-k}(M; R)$ and $B := \text{PD}(\psi) \in H_k(M; R)$, then a precise relation is given by

$$\begin{aligned} A \cdot B &= \langle 1, A \cdot B \rangle = \langle 1, \text{PD}(\psi \cup \varphi) \rangle \\ (54.1) \quad &= \langle 1, (\psi \cup \varphi) \cap [M] \rangle = \langle 1, \psi \cap (\varphi \cap [M]) \rangle = \langle \psi, \varphi \cap [M] \rangle \\ &= \langle \psi \cup \varphi, [M] \rangle = Q(\psi, \varphi). \end{aligned}$$

Corollary 53.5 therefore implies:

COROLLARY 54.2. *If M is a closed, connected and oriented n -manifold, the intersection number defines a nonsingular bilinear form*

$$H_{n-k}^{\text{free}}(M) \otimes H_k^{\text{free}}(M) \rightarrow \mathbb{Z} : A \otimes B \mapsto A \cdot B,$$

for every $k = 0, \dots, n$, so in particular, for every primitive non-torsion element $A \in H_{n-k}(M)$ there exists a class $B \in H_k(M)$ with $A \cdot B = 1$. If M is instead assumed to be orientable over a field \mathbb{K} , then the intersection number similarly defines a nonsingular \mathbb{K} -bilinear form

$$H_{n-k}(M; \mathbb{K}) \otimes_{\mathbb{K}} H_k(M; \mathbb{K}) \rightarrow \mathbb{K}.$$

□

We can also extract from the end of the second line of (54.1) the following useful formula: for every $\varphi \in H^k(M; R)$ and $A \in H_k(M; R)$,

$$\langle \varphi, A \rangle = A \cdot \text{PD}(\varphi).$$

Once we have understood how to interpret the intersection product geometrically, this formula will yield some intuitive insight into the isomorphism $\text{PD} : H^k(M; R) \rightarrow H_{n-k}(M; R)$, i.e. it transforms

the natural evaluation of cohomology classes on homology classes into the operation of counting intersections between homology classes. In many situations, this can be used to compute Poincaré dual classes explicitly.

To understand what $A \cdot B$ means in terms of “counting intersections,” we consider the following scenario: suppose M has a smooth structure and A and B are closed *smooth submanifolds* of M with dimensions $n - k$ and $n - \ell$ respectively, i.e. their *codimensions* are k and ℓ . We also assume that A and B carry R -orientations, in which case they have well-defined fundamental classes $[A] \in H_{n-k}(A; R)$ and $[B] \in H_{n-\ell}(B; R)$, and we shall use the same notation for the classes in $H_*(M; R)$ obtained by feeding these into the maps induced by the inclusions $A, B \hookrightarrow M$, that is,

$$[A] \in H_{n-k}(M; R), \quad [B] \in H_{n-\ell}(M; R).$$

Note that in practice, the orientation assumption just means that A and B are both oriented if the coefficient ring R is chosen from among \mathbb{Z} , \mathbb{Q} and \mathbb{R} , whereas there is no orientation assumption at all if we use $R = \mathbb{Z}_2$.

We now need a few basic notions from differential topology. The first is the *smooth* version of an orientation, which is simpler than what we have defined for topological manifolds. For our purposes, it is best to express this in terms of tangent spaces: since M has a smooth structure, there is a tangent space $T_x M$ associated to every point $x \in M$, which is a real vector space of dimension n . In general, an **orientation** of an n -dimensional vector space V is defined to be an equivalence class of bases of V , where two bases are equivalent if and only if one can be deformed to the other through a continuous family of bases. There are always two equivalence classes, due to the fact that the group $\text{GL}(n, \mathbb{R})$ has two connected components, distinguished by the sign of the determinant: we call the bases in the preferred equivalence class **positively oriented** and all others **negatively oriented**. An orientation of the smooth manifold M can then be defined as a choice of orientation for every tangent space $T_x M$ that varies continuously with respect to x . To make this precise, we would need to define the appropriate topology on the *tangent bundle* $TM = \bigcup_{x \in M} T_x M$ and discuss what it means for a vector field to be continuous, but it should at least be intuitively clear what is meant, and since this is only meant as a survey, we'll leave it at that. One can check that this notion of orientation is equivalent to the various other notions of orientation that we've seen before, i.e. in terms of local homology groups or orientation-preserving coordinate transformations.

An important fact about tangent spaces is that if $A \subset M$ is a smooth submanifold, then each of its tangent spaces $T_x A$ is naturally a linear subspace of $T_x M$. This is enough background information to define the important notion of *transversality*.

DEFINITION 54.3. We say that two smooth submanifolds $A, B \subset M$ are **transverse** and write “ $A \pitchfork B$ ” if for every $x \in A \cap B$,

$$T_x A + T_x B = T_x M.$$

It is an easy exercise in linear algebra to show that if the condition $T_x A + T_x B = T_x M$ holds where $T_x A \subset T_x M$ and $T_x B \subset T_x M$ have codimensions k and ℓ respectively, then $T_x A \cap T_x B \subset T_x M$ is a subspace with codimension $k + \ell$. The following nonlinear version of this observation is a standard application of the implicit function theorem.

PROPOSITION 54.4. *If $A, B \subset M$ are closed smooth submanifolds of codimensions k and ℓ respectively and are transverse to each other, then $A \cap B \subset M$ is also a closed smooth submanifold, with codimension $k + \ell$.* \square

We can also add orientations to this picture. Suppose $V, W \subset \mathbb{R}^n$ are two oriented linear subspaces of codimensions k and ℓ such that $V + W = \mathbb{R}^n$, so $V \cap W$ has codimension $k + \ell$. Define the orientation of \mathbb{R}^n so that the standard basis (e_1, \dots, e_n) is considered positively oriented. We

can then define an orientation of $V \cap W$ as follows. If either of V or W is contained in the other, then their intersection is the smaller subspace and thus already has an orientation, so assume this is not the case. Then given a basis $(X_1, \dots, X_{n-(k+\ell)})$ of $V \cap W$, it is always possible to choose additional vectors $Y_1, \dots, Y_\ell \in V$ and $Z_1, \dots, Z_k \in W$ such that the ordered tuples

$$(X_1, \dots, X_{n-(k+\ell)}, Y_1, \dots, Y_\ell) \text{ in } V, \quad (X_1, \dots, X_{n-(k+\ell)}, Z_1, \dots, Z_k) \text{ in } W$$

both form positively oriented bases, and

$$(X_1, \dots, X_{n-(k+\ell)}, Y_1, \dots, Y_\ell, Z_1, \dots, Z_k)$$

is then a basis of \mathbb{R}^n . We define the orientation of $V \cap W$ such that $(X_1, \dots, X_{n-(k+\ell)})$ is positively oriented if and only if this basis of \mathbb{R}^n is positively oriented. One can check that the only choices this definition depends on are the orientations of V , W and \mathbb{R}^n , plus the choice to write V in front of W instead of vice versa. Doing the same thing with the tangent spaces at all points $x \in A \cap B$ of a transverse intersection gives:

PROPOSITION 54.5. *In the setting of Proposition 54.4, if A , B and M are all oriented, then $A \cap B$ inherits from this data a natural orientation. \square*

REMARK 54.6. It should be emphasized that according to the definition above, the orientation of $A \cap B$ may be different from that of $B \cap A$, i.e. they are the same submanifold, but the choice to write A in front of B or vice versa may determine different orientations. It is an easy exercise to show that the orientations differ if and only if the codimensions of A and B are both odd. (This should remind you of a graded commutativity relation—there is good reason for that!)

We should add a word about the case $k + \ell = n$, which does not quite fit into the above discussion since we have not properly defined what an orientation of a 0-dimensional vector space should mean. The definition is consistent with the notion that there should always be exactly two choices of orientation: if $\dim V = 0$, we define an **orientation** of V to mean a choice of sign ± 1 . To see that this is a sensible definition, consider the situation where $V, W \subset \mathbb{R}^n$ are transverse linear subspaces of dimensions k and $n - k$, which means $V \oplus W = \mathbb{R}^n$, and $V \cap W$ is thus a 0-dimensional subspace. There are exactly two possibilities: choosing positively oriented bases (X_1, \dots, X_k) of V and (Y_1, \dots, Y_{n-k}) of W , the basis

$$(X_1, \dots, X_k, Y_1, \dots, Y_{n-k})$$

of \mathbb{R}^n is either positively or negatively oriented, and we define the orientation of $V \cap W$ to be $+1$ or -1 accordingly. Applying this idea to tangent spaces, the intersection $A \cap B$ between two transverse oriented closed submanifolds $A, B \subset M$ of complementary dimensions $\dim A + \dim B = \dim M$ is simply a finite set of points $x \in A \cap B$ with attached signs

$$\epsilon(x) = \pm 1,$$

determined as described above from the orientations of the complementary tangent spaces $T_x A \oplus T_x B = T_x M$. You should take a moment to convince yourself that this notion of the orientation of a 0-manifold is consistent with the definition we already had for *topological* 0-manifolds: indeed, if M is a discrete set, then each of the local homology groups $H_0(M, M \setminus \{x\}; \mathbb{Z})$ is *canonically* isomorphic to $H_0(\{\text{pt}\}; \mathbb{Z}) = \mathbb{Z}$, so a local orientation is a choice of generator of the group \mathbb{Z} , i.e. either $+1$ or -1 .

We can now state the main theorem of this lecture.

THEOREM 54.7. *Assume M is a closed, connected, smooth and R -oriented manifold of dimension n , and $A, B \subset M$ are closed, smooth, R -oriented submanifolds of codimensions k and ℓ respectively, such that $A \pitchfork B$. Then for the induced R -orientation on $A \cap B$ from Proposition 54.5,*

$$[A] \cdot [B] = [A \cap B].$$




FIGURE 25. The surface and 1-dimensional submanifolds discussed in Example 54.9.

As usual, the case $k + \ell = n$ deserves special comment. If A and B are oriented in this case, $A \cap B$ is a finite set of points x with attached signs $\epsilon(x) = \pm 1$, and since M is connected, the canonical isomorphism $H_0(M; \mathbb{Z}) = \mathbb{Z}$ identifies $[A \cap B]$ with the integer

$$[A] \cdot [B] = \sum_{x \in A \cap B} \epsilon(x) \in \mathbb{Z}.$$

The right hand side of this expression is sometimes called the *algebraic* (or *signed*) count of transverse intersections between A and B .

EXERCISE 54.8. Draw some examples of pairs of transversely intersecting closed oriented 1-dimensional submanifolds $A, B \subset S^2$, and convince yourself that the signed count of intersections between them will always be 0. (Indeed, this must be true for at least two reasons: first, since $H_1(S^2) = 0$, both submanifolds represent the trivial homology class and it follows that $[A] \cdot [B] = 0$. Alternatively, if you believe the implication of Theorem 54.7 that the signed count of intersections between A and B only depends on their homology classes, then you can easily adjust A or B by a homotopy (i.e. contracting A so that it lies in an arbitrarily small neighborhood) so that $A \cap B = \emptyset$.)

EXAMPLE 54.9. Figure 25 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 Problem Set 12 #3(c)),

$$\begin{array}{ccccc} H_1(\beta) & \xrightarrow{0} & H_1(\Sigma_+) & \longrightarrow & H_1(\Sigma), \\ & & \searrow i_* & \nearrow & \\ & & & & \end{array}$$

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

$$[\gamma] \cdot [\delta] = \pm 1.$$

This proves that both of the classes $[\gamma], [\delta] \in H_1(\Sigma)$ are not only nontrivial but also primitive.

The nonseparating loops in Example 54.9 admit the following interesting generalization. If M is an n -manifold, a submanifold $\Sigma \subset M$ is called a **hypersurface** if $\dim \Sigma = n - 1$. Assuming M is connected, we say that $\Sigma \subset M$ **separates** M if $M \setminus \Sigma$ is disconnected.

THEOREM 54.10. *Suppose M is a closed, connected and R -oriented smooth n -manifold containing a closed, connected and R -oriented smooth hypersurface $\Sigma \subset M$. Then the homology class $[\Sigma] \in H_{n-1}(M; R)$ is trivial if and only if Σ separates M .*

PROOF. If Σ separates M then we can write $M = M_+ \cup_{\Sigma} M_-$ where M_{\pm} are two compact R -oriented n -manifolds with boundary $\partial M_{\pm} = \Sigma$, so the same argument as in Example 54.9 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 $\mathbb{R}^{n-1} \times \{0\} \subset \mathbb{R}^n$,⁸¹ but on opposite sides of Σ , i.e. the n th coordinates of z_+ 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 25) a smooth loop $\gamma \subset M$ that intersects Σ exactly once and transversely. It follows that

$$[\Sigma] \cdot [\gamma] = \pm 1,$$

hence $[\Sigma] \in H_{n-1}(M; R)$ and $[\gamma] \in H_1(M; R)$ are both nontrivial. □

If you've been wondering why non-orientable surfaces like $\mathbb{R}P^2$ 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 54.11. *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 54.10) 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 orientable over \mathbb{Z}_2), the intersection number of Σ with γ is then

$$[\Sigma] \cdot [\gamma] = 1 \in \mathbb{Z}_2,$$

implying $[\Sigma] \neq 0 \in H_{n-1}(S^n; \mathbb{Z}_2)$ and $[\gamma] \neq 0 \in H_1(S^n; \mathbb{Z}_2)$. This contradicts our computation of $H_*(S^n; \mathbb{Z}_2)$. □

REMARK 54.12. The fact that S^n is orientable is not the decisive factor in Corollary 54.11, as there is no obstruction in general to embedding closed non-orientable hypersurfaces into closed orientable manifolds. An easy example is $\mathbb{R}P^2 \hookrightarrow \mathbb{R}P^3$.

NOTATION. For the rest of this lecture, all homology and cohomology groups are to be understood with coefficient ring R , but we will omit R in the notation wherever possible.

To prove Theorem 54.7, we need another basic result from differential topology, which gives a model for the neighborhood of any smooth submanifold $A \subset M$. Assume $\dim M = n$ and $\dim A = n - k$, so A has codimension k . The **normal bundle** of A is defined as the union of k -dimensional quotient vector spaces

$$N^M A := \bigcup_{x \in A} N_x^M A, \quad \text{where} \quad N_x^M A := T_x M / T_x A.$$

⁸¹One 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 $k = \dim \Sigma$.

When there is no ambiguity about the ambient manifold, we will usually abbreviate

$$NA := N^M A, \quad N_x A := N_x^M A.$$

This is an example of a *vector bundle*; its **fibers** are the individual vector spaces $N_x A$, which are also the preimages of points under the natural projection map

$$\pi : NA \rightarrow A : N_x A \mapsto x \text{ for all } x \in A.$$

One can define a natural topology and smooth structure on NA so that it becomes a smooth n -dimensional manifold and $\pi : NA \rightarrow A$ is a smooth map. The determining feature of this topology and smooth structure is that the map $\pi : NA \rightarrow A$ is **locally trivializable**, meaning every point $x \in A$ has a neighborhood $\mathcal{U} \subset A$ admitting a diffeomorphism

$$\Phi_{\mathcal{U}} : \pi^{-1}(\mathcal{U}) \rightarrow \mathcal{U} \times \mathbb{R}^k$$

that restricts to a linear isomorphism $N_y A \rightarrow \{y\} \times \mathbb{R}^k$ for each $y \in \mathcal{U}$. A local trivialization thus identifies the map $\pi^{-1}(\mathcal{U}) \xrightarrow{\pi} \mathcal{U}$ with the obvious projection map $\mathcal{U} \times \mathbb{R}^k \rightarrow \mathcal{U}$. The existence of local trivializations is a straightforward consequence of the existence of slice charts near every point of A . If NA and A are both endowed with orientations, then this local product structure also determines an orientation of each of the fibers $N_x A \cong \mathbb{R}^k$, which can be viewed as a family of orientations varying continuously with $x \in A$. The subset of NA consisting of all 0-vectors in the spaces $N_x A$ is a smooth submanifold that is canonically diffeomorphic to A , called the **zero-section**. We shall often regard A itself as a submanifold of NA and write

$$A \subset NA$$

by identifying A with the zero-section.

If $A \subset M$ and $B \subset M$ are submanifolds that intersect transversely, then for each $x \in A \cap B$, the diagonal map $T_x M \rightarrow T_x M \oplus T_x M$ descends to a canonical isomorphism

$$N_x^M(A \cap B) \xrightarrow{\cong} N_x^M A \oplus N_x^M B,$$

producing a vector bundle isomorphism of $N^M(A \cap B)$ with the direct sum of the restrictions of $N^M A$ and $N^M B$ along the submanifold $A \cap B$. Since $A \cap B$ is also a submanifold of A and B , we can similarly consider the normal bundles $N^A(A \cap B)$ and $N^B(A \cap B)$, and notice that for each $x \in A \cap B$, the inclusions $T_x A \hookrightarrow T_x M$ and $T_x B \hookrightarrow T_x M$ descend to canonical isomorphisms

$$N_x^A(A \cap B) \xrightarrow{\cong} N_x^M B, \quad N_x^B(A \cap B) \xrightarrow{\cong} N_x^M A,$$

giving vector bundle isomorphisms $N^A(A \cap B) \cong N^M B|_{A \cap B}$ and $N^B(A \cap B) \cong N^M A|_{A \cap B}$. In this way we can regard $N^B(A \cap B)$ as a subset of $N^M A$ and $N^A(A \cap B)$ as a subset of $N^M B$; in fact, both are smooth submanifolds whose codimensions match the codimensions of B and A respectively.

This is enough background to state the **tubular neighborhood theorem**. Its proof is a fairly straightforward matter of defining a smooth map $N^M A \rightarrow M$ whose derivative along the zero-section is the identity map, and then citing the inverse function theorem.

THEOREM 54.13. *There exists a smooth embedding $N^M A \hookrightarrow M$ that is a diffeomorphism onto a neighborhood of A and matches the inclusion $A \hookrightarrow M$ along the zero-section. Moreover, if $B \subset M$ is another smooth submanifold that intersects A transversely, then the embedding $N^M A \hookrightarrow M$ can be arranged so that it maps $N^B(A \cap B) \subset N^M A$ onto a neighborhood of $A \cap B$ in B . \square*

We shall use the tubular neighborhood theorem in the following to identify an open neighborhood of $A \subset M$ with the normal bundle NA . It will be useful also to shrink this to a smaller *compact* neighborhood that is a smooth manifold with boundary. To define this, choose a family of inner products $\langle \cdot, \cdot \rangle$ on the tangent spaces $T_x M$ that vary smoothly with x ; this is what is called a

Riemannian metric on M . The inner product on each $T_x M$ determines an isomorphism of $N_x^M A$ with the orthogonal complement of $T_x A \subset T_x M$ and thus (by restriction) also determines an inner product on $N_x^M A$. We can then define

$$\begin{aligned} \mathbb{D}N_x A &:= \{X \in N_x A \mid \langle X, X \rangle \leq 1\}, & \mathbb{D}N A &:= \bigcup_{x \in A} \mathbb{D}N_x A, \\ \mathbb{S}N_x A &:= \{X \in N_x A \mid \langle X, X \rangle = 1\}, & \mathbb{S}N A &:= \bigcup_{x \in A} \mathbb{S}N_x A. \end{aligned}$$

Each $\mathbb{D}N_x A$ is now a compact k -disk with boundary $\mathbb{S}N_x A \cong S^{k-1}$, and their union $\mathbb{D}N A$ is a compact n -manifold with boundary $\mathbb{S}N A$, called the **unit disk bundle** in NA . The tubular neighborhood theorem then identifies $\mathbb{D}N A$ with a compact neighborhood of A in M , while simultaneously identifying the subset $\mathbb{D}N^B(A \cap B) \subset \mathbb{D}N A$ with a compact neighborhood of $A \cap B$ in B . Since $\mathbb{D}N A$ is an n -dimensional submanifold of M and is also compact with boundary, it inherits an R -orientation from M and therefore has a relative fundamental class $[\mathbb{D}N A] \in H_n(\mathbb{D}N A, \mathbb{S}N A)$ and Poincaré duality isomorphism

$$\text{PD} : H^k(\mathbb{D}N A, \mathbb{S}N A) \xrightarrow{\cong} H_{n-k}(\mathbb{D}N A) : \varphi \mapsto \varphi \cap [\mathbb{D}N A].$$

In the following it will be important to specify unambiguously which homology classes belong to which spaces or pairs, so let us reserve the notation $[A]$ for the fundamental class in $H_{n-k}(A)$, and denote the class in $H_{n-k}(M)$ that appears in the statement of Theorem 54.7 by

$$(i_A^M)_*[A] \in H_{n-k}(M).$$

We shall use the notation $i_X^Y : X \hookrightarrow Y$ for the inclusion of any subspace $X \subset Y$, so for instance the homology class represented by the zero-section in NA is

$$(i_A^{\mathbb{D}N A})_*[A] \in H_{n-k}(\mathbb{D}N A).$$

Note that since orientations on NA and A determine orientations on the fibers, the individual disks $\mathbb{D}N_x A \cong \mathbb{D}^k$ for each $x \in A$ also have well-defined fundamental classes $[\mathbb{D}N_x A] \in H_k(\mathbb{D}N_x A, \mathbb{S}N_x A)$. For $x \in A$, let us abbreviate the inclusion of the fiber by

$$f_x^A : (\mathbb{D}N_x A, \mathbb{S}N_x A) \hookrightarrow (\mathbb{D}N A, \mathbb{S}N A),$$

so $(f_x^A)_*[\mathbb{D}N_x A] \in H_k(\mathbb{D}N A, \mathbb{S}N A)$.

DEFINITION 54.14. The **Thom class** of the normal bundle NA is

$$\tau(NA) := \text{PD}^{-1}((i_A^{\mathbb{D}N A})_*[A]) \in H^k(\mathbb{D}N A, \mathbb{S}N A).$$

The Thom class is determined by the vector bundle NA , i.e. it depends on its properties as a vector bundle, but not on the fact that it is the normal bundle of a submanifold $A \subset M$. More generally, for any vector bundle E with R -oriented k -dimensional fibers over a closed R -oriented manifold M , one can choose inner products on the fibers to define a unit disk bundle $\mathbb{D}E \subset E$ and define the Thom class $\tau(E) \in H^k(\mathbb{D}E, \mathbb{S}E)$ as the Poincaré dual of the homology class of the zero-section $M \subset E$. By excision, this can be identified with a class in $H^k(E, E \setminus M)$, thus it does not depend on the choice of inner products used to define $\mathbb{D}E$ and $\mathbb{S}E$. We have chosen to formulate the definition above in a less general way since it can then be understood without any concrete knowledge of the theory of vector bundles, i.e. you can simply think of $\mathbb{D}N A$ as a compact neighborhood of A in M and $\mathbb{S}N A$ as the boundary of that neighborhood.

The next theorem is similarly true for all vector bundles that have a well-defined Thom class, but we are stating it for the specific case that we need for applications.

THEOREM 54.15 (Thom isomorphism theorem). *The Thom class $\tau(NA) \in H^k(\mathbb{D}NA, \mathbb{S}NA)$ is uniquely determined by the condition*

$$\langle \tau(NA), (f_x^A)_*[\mathbb{D}N_xA] \rangle = 1 \quad \text{for every } x \in A.$$

Moreover, the map

$$H^m(A) \rightarrow H^{m+k}(\mathbb{D}NA, \mathbb{S}NA) : \varphi \mapsto \pi^* \varphi \cup \tau(NA)$$

is an isomorphism.

PROOF. One can contract all fibers $\mathbb{D}N_xA$ to their origins to produce a deformation retraction of $\mathbb{D}NA$ to its zero-section $A \subset \mathbb{D}NA$, thus the inclusion $i_A^{\mathbb{D}NA} : A \hookrightarrow \mathbb{D}NA$ and projection $\pi : \mathbb{D}NA \rightarrow A$ are homotopy inverses, implying that

$$\pi^* : H^*(A) \rightarrow H^*(\mathbb{D}NA) \quad \text{and} \quad (i_A^{\mathbb{D}NA})_* : H_*(A) \rightarrow H_*(\mathbb{D}NA)$$

are inverses. Given $\varphi \in H^m(A)$, we now set $\psi = \pi^* \varphi \in H^m(\mathbb{D}NA)$ and can feed $\psi \cup \tau(NA) \in H^{m+k}(\mathbb{D}NA, \mathbb{S}NA)$ into the duality isomorphism $\text{PD} : H^{m+k}(\mathbb{D}NA, \mathbb{S}NA) \rightarrow H_{n-(m+k)}(\mathbb{D}NA)$. Using the naturality of the cap product, this gives

$$\begin{aligned} \text{PD}(\pi^* \varphi \cup \tau(NA)) &= (\psi \cup \tau(NA)) \cap [\mathbb{D}NA] = \psi \cap (\tau(NA) \cap [\mathbb{D}NA]) = \psi \cap (i_A^{\mathbb{D}NA})_*[A] \\ &= (i_A^{\mathbb{D}NA})_*((i_A^{\mathbb{D}NA})^* \psi \cap [A]) = (i_A^{\mathbb{D}NA})^* \text{PD}(\varphi), \end{aligned}$$

which presents the map $\pi^*(\cdot) \cup \tau(NA) : H^m(A) \rightarrow H^{m+k}(\mathbb{D}NA, \mathbb{S}NA)$ as a composition

$$H^m(A) \xrightarrow{\text{PD}} H_{n-k-m}(A) \xrightarrow{(i_A^{\mathbb{D}NA})_*} H_{n-(k+m)}(\mathbb{D}NA) \xrightarrow{\text{PD}^{-1}} H^{k+m}(\mathbb{D}NA, \mathbb{S}NA).$$

$\searrow \pi^*(\cdot) \cup \tau(NA) \nearrow$

All three maps in this composition are isomorphisms, thus so is $\varphi \mapsto \pi^* \varphi \cup \tau(NA)$.

To check that $\langle \tau(NA), (f_x^A)_*[\mathbb{D}N_xA] \rangle = 1$ for all $x \in A$, it suffices to prove this for the case when $A \subset M$ is connected, as it is then an easy exercise to generalize to a finite disjoint union of connected submanifolds. The advantage of assuming A connected is that $H^0(A)$ is then canonically isomorphic to the coefficient ring R , with the unit $1 \in H^0(A)$ as a generator, so the isomorphism above then implies that $\tau(NA)$ generates $H^k(\mathbb{D}NA, \mathbb{S}NA) \cong R$. It is clear that at most one element $\tau \in H^k(\mathbb{D}NA, \mathbb{S}NA)$ can satisfy $\langle \tau, (f_x^A)_*[\mathbb{D}N_xA] \rangle = 1$ for any given $x \in A$; moreover, if this is satisfied for one $x \in A$ then it is also satisfied for every $y \in A$, as the path-connectedness of A produces a homotopy of inclusions $(\mathbb{D}^k, S^{k-1}) \hookrightarrow (\mathbb{D}NA, \mathbb{S}NA)$ relating f_x^A and some reparametrization of f_y^A , so that $(f_y^A)_*[\mathbb{D}N_yA] = (f_x^A)_*[\mathbb{D}N_xA]$. We shall now show that there exists a generator τ satisfying this relation for some $x \in A$. This will prove $\tau = \tau(NA)$ in the case $R = \mathbb{Z}_2$; if $R = \mathbb{Z}$, then there is a sign ambiguity $\tau = \pm \tau(NA)$ that can be resolved by paying more careful attention to orientation conventions, and the remaining cases follow from this via the universal coefficient theorem.

The idea is to realize the isomorphism $H^*(A) \rightarrow H^{*+k}(\mathbb{D}NA, \mathbb{S}NA)$ via a cell decomposition. Notice first that $(\mathbb{D}NA, \mathbb{S}NA)$ is a good pair, so its relative cohomology is naturally isomorphic to the reduced cohomology of the quotient space

$$\text{Th}(NA) := \mathbb{D}NA / \mathbb{S}NA.$$

This is known as the **Thom space** of the vector bundle NA . Since the interior of $\mathbb{D}NA$ is homeomorphic to NA , one can imagine $\text{Th}(NA)$ as the one-point compactification of NA , and we shall label the point represented by $\mathbb{S}NA = \partial(\mathbb{D}NA)$ accordingly as $\infty \in \text{Th}(NA)$. The isomorphism $\tilde{H}^*(\text{Th}(NA)) \cong H^*(\mathbb{D}NA, \mathbb{S}NA)$ is then equivalent to the isomorphism

$$H^*(\text{Th}(NA), \{\infty\}) \rightarrow H^*(\mathbb{D}NA, \mathbb{S}NA)$$

induced by the quotient map $(\mathbb{D}NA, \mathbb{S}NA) \rightarrow (\text{Th}(NA), \{\infty\})$; cf. Lecture 30.

Now since A is a smooth manifold, it has a triangulation, and after barycentric subdivision⁸² we can assume without loss of generality that every simplex in the triangulation is small enough to be contained in a region $\mathcal{U} \subset A$ where there exists a local trivialization identifying $\pi^{-1}(\mathcal{U}) \subset \mathbb{D}NA$ with $\mathcal{U} \times \mathbb{D}^k$. Regarding $\{\infty\}$ as the 0-skeleton of $\text{Th}(NA)$, we can then associate to each m -simplex σ in the triangulation of A a product $(m+k)$ -cell in $\text{Th}(NA)$, defined by identifying $\pi^{-1}(\sigma) \subset \mathbb{D}NA$ with $\sigma \times \mathbb{D}^k$ and attaching $\sigma \times \partial\mathbb{D}^k$ along the constant map to ∞ . This gives $(\text{Th}(NA), \{\infty\})$ the structure of a CW-pair, and while the characteristic maps of the cells in this decomposition depend on choices of local trivializations, their images do not, thus in a meaningful sense, our cell decomposition of $(\text{Th}(NA), \{\infty\})$ depends only on the chosen triangulation of A . The most important observation is that since $C_*^{\text{CW}}(\text{Th}(NA), \{\infty\})$ is generated by the cells in the interior $\mathbb{D}NA \setminus \mathbb{S}NA$ while ignoring the 0-cell at ∞ , the association of each simplex to its product with \mathbb{D}^k defines a chain isomorphism

$$C_*^\Delta(A) \rightarrow C_{*+k}^{\text{CW}}(\text{Th}(NA), \{\infty\}),$$

which can be dualized to define an isomorphism $H^{m+k}(\text{Th}(NA), \{\infty\}) \cong H^m(A)$ for every m . Take the generator $\tau \in H^k(\text{Th}(NA), \{\infty\})$ that corresponds to $1 \in H^0(A)$ under this isomorphism, then pull it back through the quotient map $q : (\mathbb{D}NA, \mathbb{S}NA) \rightarrow (\text{Th}(NA), \{\infty\})$ to define a generator $q^*\tau \in H^k(\mathbb{D}NA, \mathbb{S}NA)$. We now have an explicit cellular cochain representative of τ and can thus check that for any point $x \in A$ in the 0-skeleton of A , the associated product k -cell in $\text{Th}(NA)$ represents $q_*(f_x^A)_*[\mathbb{D}N_xA] \in H_k(\text{Th}(NA), \{\infty\})$ and satisfies

$$\langle q^*\tau, (f_x^A)_*[\mathbb{D}N_xA] \rangle = \langle \tau, q_*(f_x^A)_*[\mathbb{D}N_xA] \rangle = 1.$$

□

REMARK 54.16. The Thom space $\text{Th}(E)$ can also be defined for more general vector bundles $\pi : E \rightarrow M$ without any reference to a submanifold or normal bundle. Together with the cell decomposition constructed in the proof above, $\text{Th}(E)$ is known as the *Thom complex*.

Let us now associate to the closed R -oriented submanifold $A \subset M$ of codimension k its Poincaré dual class in M ,

$$\tau_A^M := \text{PD}^{-1}((i_A^M)_*[A]) \in H^k(M), \quad \text{i.e.} \quad \tau_A^M \cap [M] = (i_A^M)_*[A].$$

The Thom class provides a way of “localizing” τ_A^M , in the following sense. Consider the inclusions

$$(M, \emptyset) \xleftarrow{j_A^M} (M, M \setminus A) \xleftarrow{i_{\mathbb{D}NA}^M} (\mathbb{D}NA, \mathbb{S}NA),$$

where $i_{\mathbb{D}NA}^M$ is provided by the tubular neighborhood theorem. This is an excision map, so it induces an isomorphism on relative cohomology and thus identifies $\tau(NA)$ with a class

$$\hat{\tau}_A^M \in H^k(M | A) \quad \text{such that} \quad (i_{\mathbb{D}NA}^M)_*\hat{\tau}_A^M = \tau(NA).$$

LEMMA 54.17. $\tau_A^M = (j_A^M)_*\hat{\tau}_A^M$.

PROOF. The three fundamental classes $[M] \in H_n(M)$, $[M]_A \in H_n(M | A)$ and $[\mathbb{D}NA] \in H_n(\mathbb{D}NA, \mathbb{S}NA)$ in this picture are related by

$$(j_A^M)_*[M] = [M]_A = (i_{\mathbb{D}NA}^M)_*[\mathbb{D}NA].$$

⁸²The subdivision is not actually necessary at all if you know a little bit more about the properties of vector bundles: in particular, they are always trivialisable over a contractible space. For similar reasons, one could just as well work with an arbitrary cell decomposition of A instead of a triangulation.

Since $\text{PD}((j_A^M)^*\hat{\tau}_A^M) \in H_{n-k}(M)$ is an *absolute* homology class and $j_A^M : M \rightarrow M$ is the identity map, the naturality of the cap product with respect to $j_A^M : (M, \emptyset) \hookrightarrow (M, M \setminus A)$ and $i_{\mathbb{D}NA}^M : (\mathbb{D}NA, \mathbb{S}NA) \hookrightarrow (M, M \setminus A)$ gives

$$\begin{aligned} \text{PD}((j_A^M)^*\hat{\tau}_A^M) &= (j_A^M)_* \text{PD}((j_A^M)^*\hat{\tau}_A^M) = (j_A^M)_*((j_A^M)^*\hat{\tau}_A^M \cap [M]) = \hat{\tau}_A^M \cap (j_A^M)_*[M] \\ &= \hat{\tau}_A^M \cap (i_{\mathbb{D}NA}^M)_*[\mathbb{D}NA] = (i_{\mathbb{D}NA}^M)_*((i_{\mathbb{D}NA}^M)^*\hat{\tau}_A^M \cap [\mathbb{D}NA]) \\ &= (i_{\mathbb{D}NA}^M)_* \text{PD}(\tau(NA)) = (i_{\mathbb{D}NA}^M)_*(i_A^{\mathbb{D}NA})_*[A] = (i_A^M)_*[A], \end{aligned}$$

so $(j_A^M)^*\hat{\tau}_A^M$ satisfies the defining property of τ_A^M . \square

The message of this lemma is that the cohomology class Poincaré dual to A is determined by a class in $H^k(M|A)$, or equivalently $H^k(\mathbb{D}NA, \mathbb{S}NA)$; in either case it depends only on a neighborhood of A . It is now easy to see why Theorem 54.7 holds in the case $A \cap B = \emptyset$. It is equivalent in that case to $\tau_A^M \cup \tau_B^M = 0$, and by naturality of the cup product we have

$$\tau_A^M \cup \tau_B^M = (j_A^M)^*\hat{\tau}_A^M \cup (j_B^M)^*\hat{\tau}_B^M = (j_{A \cap B}^M)^*(\hat{\tau}_A^M \cup \hat{\tau}_B^M) = 0,$$

where the relative cup product

$$\hat{\tau}_A^M \cup \hat{\tau}_B^M \in H^{k+\ell}(M, (M \setminus A) \cup (M \setminus B)) = H^{k+\ell}(M, M \setminus (A \cap B)) = H^{k+\ell}(M, M) = 0$$

is well defined since $M \setminus A$ and $M \setminus B$ are both open in M and thus form an excisive couple. To handle the general case where $A \pitchfork B$ so that $A \cap B$ is a submanifold of A with normal bundle $N^A(A \cap B) = N^M B|_{A \cap B}$, we need the following application of the Thom isomorphism theorem:

LEMMA 54.18. $\tau_{A \cap B}^A = (i_A^M)^*\tau_B^M$.

PROOF. Let $f_{A \cap B}^B : (\mathbb{D}N^A A \cap B, \mathbb{S}N^A(A \cap B)) \hookrightarrow (\mathbb{D}N^M B, \mathbb{S}N^M B)$ denote the inclusion defined by identifying the fiber of $N^A(A \cap B)$ over each point in $A \cap B$ with the corresponding fiber of $N^M B$. For each $x \in A \cap B$ we then have

$$\begin{aligned} \langle (f_{A \cap B}^B)^*\tau(N^M B), (f_x^{A \cap B})_*[\mathbb{D}N_x^A(A \cap B)] \rangle &= \langle \tau(N^M B), (f_{A \cap B}^B \circ f_x^{A \cap B})_*[\mathbb{D}N_x^A(A \cap B)] \rangle \\ &= \langle \tau(N^M B), (f_x^B)_*[\mathbb{D}N_x^M B] \rangle = 1, \end{aligned}$$

so $(f_{A \cap B}^B)^*\tau(N^M B) = \tau(N^A(A \cap B))$ according to Theorem 54.15.⁸³ The result then follows from the commutative diagram

$$\begin{array}{ccccc} H^k(M) & \xleftarrow{(j_B^M)^*} & H^k(M|B) & \xrightarrow[\cong]{(i_{\mathbb{D}NB}^M)^*} & H^k(\mathbb{D}NB, \mathbb{S}NB) \\ \downarrow (i_A^M)^* & & \downarrow (i_A^M)^* & & \downarrow (f_{A \cap B}^B)^* \\ H^k(A) & \xleftarrow{(j_{A \cap B}^A)^*} & H^k(A|A \cap B) & \xrightarrow[\cong]{(i_{\mathbb{D}N^A(A \cap B)}^A)^*} & H^k(\mathbb{D}N^A(A \cap B), \mathbb{S}N^A(A \cap B)) \end{array}$$

\square

This is enough preparation to prove the main theorem.

⁸³What this argument actually shows is that for any vector bundle $\pi : E \rightarrow M$ and a map $f : N \rightarrow M$ for which the Thom classes of E and the *pullback bundle* $f^*E \rightarrow N$ are well defined, $\tau(f^*E) = \hat{f}^*\tau(f)$, where $\hat{f} : (\mathbb{D}(f^*E), \mathbb{S}(f^*E)) \rightarrow (\mathbb{D}(E), \mathbb{S}(E))$ is the natural map of disk bundles covering the map $f : N \rightarrow M$.

PROOF OF THEOREM 54.7. Combining Lemma 54.18 with the usual naturality and associativity properties of the cap and cup products, we compute:

$$\begin{aligned} (i_{A \cap B}^M)_*[A \cap B] &= (i_A^M)_*(i_{A \cap B}^A)_*[A \cap B] = (i_A^M)_* \text{PD}(\tau_{A \cap B}^A) \\ &= (i_A^M)_*(\tau_{A \cap B}^A \cap [A]) = (i_A^M)_*((i_A^M)^*\tau_B^M \cap [A]) = \tau_B^M \cap (i_A^M)_*[A] \\ &= \tau_B^M \cap \text{PD}(\tau_A^M) = \tau_B^M \cap (\tau_A^M \cap [M]) \\ &= (\tau_B^M \cup \tau_A^M) \cap [M] = \text{PD}(\tau_B^M \cup \tau_A^M) = (i_A^M)_*[A] \cdot (i_B^M)_*[B]. \end{aligned}$$

□

We conclude by mentioning a beautiful reinterpretation of the Lefschetz fixed point theorem in terms of intersection theory. Suppose M is a closed, connected and oriented smooth n -manifold and $f : M \rightarrow M$ is a smooth map. The fixed point set $\text{Fix}(f) = \{x \in M \mid f(x) = x\}$ is then in one-to-one correspondence with the intersection of the two smooth submanifolds

$$\begin{aligned} \Delta &:= \{(x, x) \mid x \in M\} \subset M \times M, \\ \Gamma_f &:= \{(x, f(x)) \mid x \in M\} \subset M \times M. \end{aligned}$$

Both are closed and inherit from M obvious orientations.

THEOREM 54.19. *The homological intersection number $[\Delta] \cdot [\Gamma_f] \in \mathbb{Z}$ is the Lefschetz number $L(f)$.*

The case of the Lefschetz fixed point theorem for smooth oriented closed manifolds follows from this immediately since $L(f) \neq 0$ now implies that Δ and Γ_f cannot be disjoint. The orientation condition can also be dropped by using coefficients in \mathbb{Z}_2 , in which case the analogous theorem identifies $[\Delta] \cdot [\Gamma_f] \in \mathbb{Z}_2$ with $L_{\mathbb{Z}_2}(f)$. In either case, the Lefschetz number is no longer just a criterion for the *existence* of a fixed point, but is actually a quantitative *count* of fixed points, so long as the concept of “counting” is understood in the proper generalized sense (e.g. counting with signs). For a proof of Theorem 54.19, see [Hut].

55. Higher homotopy groups

The last two lectures in this course will have more the character of a survey, as I want to mention several important things but will not have time to prove many of them.

The higher homotopy groups $\pi_n(X)$ were mentioned informally last semester in Lecture 20. 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

$$\pi_1(X, x_0) = [(S^1, \text{pt}), (X, x_0)] = [(I, \partial I), (X, x_0)],$$

where pt denotes an arbitrary choice of base point in S^1 , and I is the unit interval $[0, 1]$. Since the latter is homeomorphic to the 1-dimensional unit disk \mathbb{D}^1 , we could also equivalently write

$$\pi_1(X, x_0) = [(\mathbb{D}^1, \partial \mathbb{D}^1), (X, x_0)].$$

These definitions are equivalent to the definition in terms of S^1 because $S^1 \cong \mathbb{D}^1/\partial\mathbb{D}^1$. Note that there are also higher-dimensional analogues of this statement: S^n is homeomorphic to $\mathbb{D}^n/\partial\mathbb{D}^n$ and $I^n/\partial I^n$ for all $n \in \mathbb{N}$, where I^n here denotes the n -fold product of I , i.e. an n -dimensional unit cube.

DEFINITION 55.1. For each integer $n \geq 0$, we define the set

$$\pi_n(X, x_0) := [(S^n, \text{pt}), (X, x_0)].$$

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,

$$(I_1 \amalg I_2) / (I_1 \ni 1 \sim 0 \in I_2)$$

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 55.2. 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 55.3. For all $n \geq 2$, the operation in Definition 55.2 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 26. The shaded region in each picture represents a subset of I^n on which the map has constant value at the base point. The leftmost picture shows the map representing $[f] + [g]$ as specified in Definition 55.2,




FIGURE 26. The homotopy in the proof of Proposition 55.3.

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 $n = 1$. \square

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 $\text{pt} \in \partial\mathbb{D}^n$ and setting

$$\pi_n(X, A, x_0) := [(\mathbb{D}^n, \partial\mathbb{D}^n, \text{pt}), (X, A, x_0)].$$

This reduces to $\pi_n(X, x_0)$ if $A = \{x_0\}$, but in all other cases we need to be aware that it only makes sense for $n \geq 1$; 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 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 regarding the rest of ∂I^n as the base point: let

$$J_n := I^{n-1} \times \{0\} \subset \partial I^n$$

and redefine $\pi_n(X, A, x_0)$ as

$$\pi_n(X, A, x_0) := [I^n, \partial I^n, \overline{\partial I^n \setminus J_n}, (X, A, x_0)].$$

By this definition, the formula in Definition 55.2 still makes sense for $n \geq 2$ and defines a group structure on $\pi_n(X, A, x_0)$, though Proposition 55.3 no longer works in the $n = 2$ case. You can see why not if you look again at Figure 26 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 55.4. *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*




FIGURE 27. The isomorphism $\pi_n(X, y) \rightarrow \pi_n(X, x)$ determined by a path $x \rightsquigarrow y$.

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 27 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 $x \rightsquigarrow y$. The picture shows the $n = 2$ case, but if you draw the analogous picture for $n = 1$, you will find that it reproduces exactly the isomorphism $\Phi_\gamma : \pi_1(X, y) \rightarrow \pi_1(X, x)$ 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

$$f_* : \pi_n(X, x_0) \rightarrow \pi_n(Y, y_0) : [\varphi] \mapsto [f \circ \varphi].$$

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 55.5. *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 $x_0 \in X$. \square*

Since we've been talking about homology for the rest of this course, you may have forgotten why Theorem 55.5 is already a nontrivial statement in the $n = 1$ case, which took some effort to prove in *Topologie I*. The annoying detail is the base point: if $g : Y \rightarrow X$ 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 .

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 55.6. 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 55.7. 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 Lecture 40) 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 $dg(x) : T_x S^k \rightarrow T_y S^n$ 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 55.8. Viewing elements of $\pi_n(S^n)$ as represented by maps $f : S^n \rightarrow S^n$, there is an isomorphism

$$\pi_n(S^n) \xrightarrow{\cong} \mathbb{Z} : [f] \mapsto \deg(f)$$

for every $n \in \mathbb{N}$. This does not immediately follow from anything we've covered in this course, but here are two ways to see it: (1) Using differential topology as in [Mil97], the so-called *Pontryagin-Thom* construction elegantly defines a bijection for any closed, connected and oriented k -manifold M between the set of homotopy classes $[M, S^n]$ and the set of "framed bordism classes" in M , where the latter have a natural correspondence with the integers when $k = n$. In particular, when $\dim M = n$ this proves that the map $\deg : [M, S^n] \rightarrow \mathbb{Z}$ 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.) (2) In the next lecture we will state the *Hurewicz* theorem, which defines a natural homomorphism $\pi_n(X) \rightarrow H_n(X)$ and gives conditions for it to be an isomorphism, which hold in the case $X = S^n$ due to the computation of $H_*(S^n)$.

In Example 55.16 at the end of this lecture, we will discuss the interesting case of $\pi_3(S^2)$, which is fairly easy to compute, but the answer may contradict the intuition you've developed from homology, i.e. it is not trivial. Unlike $H_k(M)$, there is no reason in general why $\pi_k(M)$ should vanish when $k > \dim M$.

EXAMPLE 55.9. The following is way beyond the scope of this course, but just to give you a taste of what is studied in modern homotopy theory: it turns out that there are natural isomorphisms

$$\pi_{n+k}(S^n) \rightarrow \pi_{n+k+1}(S^{n+1})$$

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. They have been computed in many cases, but they are not known in general for $k > 64$. The computation of higher homotopy groups of spheres is considered one of the most important open problems in algebraic topology.

The following definition makes the notions of path-connectedness ($n = 0$) and simple connectedness ($n = 1$) into the first two items on an infinite hierarchy of conditions.

DEFINITION 55.10. 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 55.11. *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”.⁸⁴ 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

$$F : S^k \rightarrow Y : x \mapsto \begin{cases} f \circ \Phi_\alpha(x) & \text{if } x \in \mathbb{D}_+^k, \\ g \circ \Phi_\alpha(x) & \text{if } x \in \mathbb{D}_-^k. \end{cases}$$

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 55.11 has an obvious converse: if Y is not n -connected, then there clearly also exists a CW-complex X of dimension at most n (in 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 Problem

⁸⁴It 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 only heard the latter in mathematical contexts, e.g. one would not say that a politician with potentially damaging secrets has “skeleta in the closet”.

Set 12 #4). The standard procedure is to express the geometric structure of interest in terms of sections of some fiber bundle associated to the manifold, so that the important question to answer is whether a section of this bundle exists and under what conditions two such sections must be homotopic. By induction over the skeleta, these questions are typically equivalent to the vanishing of certain higher homotopy groups. For a detailed exposition of this subject, I recommend [Ste51].

We have not yet talked much about the relative homotopy groups, and we won't, but I should mention that they appear in a fairly obvious exact sequence. Given a pair of spaces (X, A) and a base point $x_0 \in A$, denote by

$$(A, x_0) \xrightarrow{i} (X, x_0) \quad \text{and} \quad (X, x_0, x_0) \xrightarrow{j} (X, A, x_0)$$

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 $f : (\mathbb{D}^n, S^{n-1}, \text{pt}) \rightarrow (X, A, x_0)$. You can easily check by translating this into the corresponding formula with $f : (I^n, \partial I^n, \partial I^n \setminus J_n) \rightarrow (X, A, x_0)$ that it really is a homomorphism.

THEOREM 55.12. 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 lecture we saw some examples of vector bundles, which one can imagine as families of vector spaces parametrized by an underlying space, carrying a topology determined by the notion of *local trivialization*. If one replaces vector spaces with arbitrary topological spaces in this picture, one arrives at the following notion.

DEFINITION 55.13. A **fiber bundle** consists of the following data: topological spaces E , B and F known as the **total space**, **base** and **standard fiber** respectively, and a continuous map $p : E \rightarrow B$, such that B can be covered by open sets \mathcal{U} that admit **local trivializations**, meaning homeomorphisms

$$\Phi : p^{-1}(U) \rightarrow U \times F$$

that send $p^{-1}(b)$ homeomorphically to $\{b\} \times F$ for each $b \in U$. The **fibers** of the bundle are the subspaces $p^{-1}(b) \cong F$ for $b \in B$.

Fiber bundles are often abbreviated with the notation

$$F \hookrightarrow E \xrightarrow{p} B,$$

where the inclusion $F \hookrightarrow E$ is not canonical but is defined by choosing any $b \in B$ and a local trivialization near b to identify $p^{-1}(b)$ with F . Note that while every fiber of a fiber bundle is homeomorphic to the standard fiber, there is typically no *canonical* homeomorphism since there may be many choices of local trivializations covering each $b \in B$. If we choose base points $b_0 \in B$ and $x_0 \in p^{-1}(b_0) \subset E$, then it is natural to identify F with $p^{-1}(b_0)$ so that we obtain base-point preserving maps

$$(F, x_0) \hookrightarrow (E, x_0) \xrightarrow{p} (B, b_0).$$

A **trivial fiber bundle** is one that admits a single trivialization covering all of B , so that E can be identified globally with $B \times F$ and the map $p : E \rightarrow B$ becomes the obvious projection map $B \times F \rightarrow B$. Here is a popular example of a fiber bundle that is 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 55.14. 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 $\mathbb{C}\mathbb{P}^1$ so that this becomes the map

$$p : S^3 \rightarrow \mathbb{C}\mathbb{P}^1 : (z_1, z_2) \mapsto [z_1 : z_2].$$

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 55.15. *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

$$\begin{aligned} \dots \rightarrow \pi_{n+1}(B, b_0) \xrightarrow{\partial} \pi_n(F, x_0) \xrightarrow{i_*} \pi_n(E, x_0) \xrightarrow{p_*} \pi_n(B, b_0) \xrightarrow{\partial} \pi_{n-1}(F, x_0) \rightarrow \dots \\ \dots \rightarrow \pi_1(E, x_0) \xrightarrow{p_*} \pi_1(B, b_0) \xrightarrow{\partial} \pi_0(F, x_0) \xrightarrow{i_*} \pi_0(E, x_0), \end{aligned}$$

where the maps $\partial : \pi_n(B, b_0) \rightarrow \pi_{n-1}(F, x_0)$ send each $[f]$ to $[f|_{S^{n-1}}]$ for $f : (\mathbb{D}^n, S^{n-1}) \rightarrow (B, b_0)$ and $\tilde{f} : (\mathbb{D}^n, S^{n-1}, \text{pt}) \rightarrow (E, F, x_0)$ solving the lifting problem

$$(55.1) \quad \begin{array}{ccc} & E & \\ & \nearrow \tilde{f} & \downarrow p \\ \mathbb{D}^n & \xrightarrow{f} & B \end{array}$$

□

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 (55.1). The proper formulation of this condition is something called the *homotopy lifting property*, and Theorem 55.15 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. A fiber bundle whose standard fiber is discrete is simply a covering map, and we saw in *Topologie I* that the lifting problem (55.1) is solvable for 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 55.16. Returning to the Hopf fibration of Example 55.14, the homotopy exact sequence has a segment of the form

$$0 = \pi_3(S^1) \rightarrow \pi_3(S^3) \xrightarrow{p_*} \pi_3(S^2) \rightarrow \pi_2(S^1) = 0,$$

proving that $p_* : \pi_3(S^3) \rightarrow \pi_3(S^2)$ 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 55.17. In obstruction theory, one often needs to know the homotopy groups of certain topological groups that arise as “structure groups” of fiber bundles. For example, the structure group of any oriented vector bundle with n -dimensional fibers is

$$\mathrm{GL}_+(n, \mathbb{R}) := \{ \mathbf{A} \in \mathrm{GL}(n, \mathbb{R}) \mid \det \mathbf{A} > 0 \}.$$

Here is a trick for computing $\pi_1(\mathrm{GL}_+(n, \mathbb{R}))$. Polar decomposition provides a deformation retraction of $\mathrm{GL}_+(n, \mathbb{R})$ to $\mathrm{SO}(n)$, the special orthogonal group, thus it suffices to compute $\pi_1(\mathrm{SO}(n))$. For $n = 1$ and $n = 2$ this is easy because $\mathrm{SO}(1) \cong \{\mathrm{pt}\}$ and $\mathrm{SO}(2) \cong S^1$. For $n = 3$, it is not hard to find a homeomorphism of $\mathrm{SO}(3)$ to \mathbb{RP}^3 : this arises from the fact that every element of $\mathrm{SO}(3)$ defines a rotation about some axis in \mathbb{R}^3 , so there is a natural map

$$\mathbb{D}^3 \rightarrow \mathrm{SO}(3)$$

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

$$\dots \rightarrow \pi_{k+1}(S^n) \rightarrow \pi_k(\mathrm{SO}(n)) \xrightarrow{i_*} \pi_k(\mathrm{SO}(n+1)) \rightarrow \pi_k(S^n) \rightarrow \dots,$$

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 $n \geq 3$.

56. The theorems of Hurewicz and Whitehead

I have more to say about higher homotopy groups, but I want to focus the discussion around a particular application:

THEOREM 56.1. *Every closed simply connected 3-manifold is homotopy equivalent to S^3 .*

You may have heard of the Poincaré conjecture, which was open for most of the 20th century and proved by Perelman early in the 21st: it strengthens the theorem above to the statement that every closed simply connected 3-manifold is *homeomorphic* to S^3 . Actually, Poincaré himself was originally more ambitious and suggested that every closed 3-manifold M with $H_*(M) \cong H_*(S^3)$ should be homeomorphic to S^3 , but he found a counterexample to this conjecture a few years later, now known as the *Poincaré homology sphere*. It was not simply connected and therefore, obviously, not homotopy equivalent to S^3 . Theorem 56.1 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. $\mathbb{C}P^2$ 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 $\mathbb{C}P^2$ from S^4 via its homology, of course. Part of the interest in Theorem 56.1, for our purposes, is the way that the condition $\pi_1(M) = 0$ in dimension three produces just enough constraints on $H_*(M)$ to make all the familiar obstructions to a homotopy equivalence between M and S^3 vanish, starting with the homology and cohomology groups, and then continuing with the higher homotopy groups. Several of the important theorems we've proved in this course have some role to play in the proof, thus it will serve both as a review of the course and as motivation to introduce two new and powerful theorems involving the higher homotopy groups.

Part 1: From simply connected to homology sphere. As mentioned above, this part will be a review of techniques developed in the course.

LEMMA 56.2. *If M is a closed and connected 3-manifold with $\pi_1(M) = 0$, then $H_*(M) \cong H_*(S^3)$.*

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 56.3. *$H_n(M)$ 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 50.8. The groups $H_k(M)$ for $k > n$ vanish by Corollary 51.15. 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 56.4. *$H_1(M) = 0$.*

PROOF. This is immediate from the isomorphism of $H_1(M)$ with the abelianization of $\pi_1(M)$. \square

LEMMA 56.5. 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 56.6. 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 51.15. \square

LEMMA 56.7. $H_2(M)$ is torsion free.

PROOF. This is true for $H_{n-1}(M)$ whenever M is a closed oriented n -manifold, and it was once assigned as a homework exercise (see Problem Set 12 #2(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

$$H_n(M; \mathbb{Z}_p) \cong (H_n(M) \otimes \mathbb{Z}_p) \oplus \text{Tor}(H_{n-1}(M), \mathbb{Z}_p).$$

Since $H_n(M; \mathbb{Z}_p) \cong \mathbb{Z}_p$ and $H_n(M) \cong \mathbb{Z}$ by Corollary 51.15, this isomorphism implies the vanishing of $\text{Tor}(H_{n-1}(M), \mathbb{Z}_p)$. Since $H_{n-1}(M)$ is finitely generated, we can then use the classification of finitely-generated abelian groups to write

$$H_{n-1}(M) \cong F \oplus \left(\bigoplus_{i=1}^N \mathbb{Z}_{k_i} \right)$$

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 $H_{n-1}(M)$ has torsion. Applying Theorem 41.2, we then have

$$0 = \text{Tor}(H_{n-1}(M), \mathbb{Z}_p) \cong \bigoplus_{i=1}^N \text{Tor}(\mathbb{Z}_{k_i}, \mathbb{Z}_p),$$

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 41.2 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

$$H^2(M) \cong H_1(M) = 0,$$

and the latter then implies

$$0 = H^2(M) \cong \text{Hom}(H_2(M), \mathbb{Z}) \oplus \text{Ext}(H_1(M), \mathbb{Z}),$$

hence $\text{Hom}(H_2(M), \mathbb{Z}) = 0$. Since $H_2(M)$ is torsion free, it follows that $H_2(M) = 0$. We already have isomorphisms $H_n(M) \cong H_n(S^3)$ for $n \geq 3$ by Lemmas 56.3 and 56.6, and $H_0(M) \cong H_0(S^3) \cong \mathbb{Z}$ is immediate since M is connected, so this completes the proof of Lemma 56.2.

Part 2: From homology sphere to homotopy sphere. The step from $\pi_1(M) = 0$ and $H_*(M) \cong H_*(S^3)$ to $M \underset{h.e.}{\simeq} S^3$ requires two theorems about homotopy groups that we will need to quote without proof, though the proofs (explained e.g. in [Hat02, Chapter 4]) do not require substantial machinery beyond what we have discussed in this course.

DEFINITION 56.8. 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 55.5 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 56.9 (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.* \square

While I do not intend to discuss the proof of this theorem, you will hopefully gain some intuition about it from Theorem 55.11 in the previous lecture; in particular, it should be clear why having cell decompositions of X and Y might be useful in the proof.

With Whitehead's theorem added to our toolbox, it would suffice to find a map $f : M \rightarrow S^3$ 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**,

$$h : \pi_n(X, x_0) \rightarrow H_n(X) : [f] \mapsto f_*[S^n].$$

We've seen that for $n = 1$, this map cannot generally be an isomorphism since $H_1(X)$ is always abelian while $\pi_1(X)$ is not, but the next best thing is true: when $\pi_0(X) = 0$, $h : \pi_1(X) \rightarrow H_1(X)$ 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 $h : \pi_n(X) \rightarrow H_n(X)$ actually being an isomorphism, though we've also seen cases where this is not true: e.g. $\pi_2(\mathbb{T}^2) = 0$ but $H_2(\mathbb{T}^2) \cong \mathbb{Z}$. 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 56.10 (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 $h : \pi_n(X, x_0) \rightarrow H_n(X)$ is an isomorphism.* \square

Here are a couple of applications before we get back to discussing 3-manifolds homotopy equivalent to S^3 .

COROLLARY 56.11. *If X is path-connected and has universal cover $\tilde{X} \rightarrow X$, then $\pi_2(\tilde{X}) \cong H_2(\tilde{X})$.*

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 $H_2(\tilde{X})$. \square

COROLLARY 56.12. *If X is a simply connected CW-complex with $\tilde{H}_*(X) = 0$, then X is contractible.*

PROOF. The Hurewicz theorem gives an isomorphism $\pi_2(X) \cong H_2(X) = 0$, proving X is 2-connected, so one can then apply the theorem again and conclude $\pi_3(X) \cong H_3(X) = 0$, and then again... by induction, we deduce $\pi_n(X) = 0$ for all $n \geq 0$. It follows that the unique map $\epsilon : X \rightarrow \{\text{pt}\}$ induces isomorphisms $\epsilon_* : \pi_n(X) \rightarrow \pi_n(\{\text{pt}\}) = 0$ 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 56.1: instead of the map $\epsilon : X \rightarrow \{\text{pt}\}$ in the proof of Corollary 56.12, 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 $H_*(M) \cong H_*(S^3)$ 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 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 $\tilde{H}_n(S^3) \neq 0$. 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

$$h : \pi_n(X, A, x_0) \rightarrow H_n(X, A) : [f] \mapsto f_*[\mathbb{D}^n],$$

where $[f] \in \pi_n(X, A, x_0)$ is represented by a map $f : (\mathbb{D}^n, \partial\mathbb{D}^n, \text{pt}) \rightarrow (X, A, x_0)$ and $[\mathbb{D}^n] \in H_n(\mathbb{D}^n, \partial\mathbb{D}^n)$ 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 $h : \pi_2(X, A, x_0) \rightarrow H_2(X, A)$ 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 56.10:

THEOREM 56.13. *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 $h : \pi_n(X, A, x_0) \rightarrow H_n(X, A)$ is an isomorphism. \square*

COROLLARY 56.14. *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 $f_* : H_n(X) \rightarrow H_n(Y)$ 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 $f_* : H_n(X) \xrightarrow{\cong} H_n(Y)$ into

$$H_n(Y, X) = 0 \quad \text{for all } n \geq 0.$$

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 $\pi_2(Y, X) \cong H_2(Y, X) = 0$ 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 has the following two properties:

- The map

$$g : Y \rightarrow M_f$$

defined by composing the inclusion $Y \hookrightarrow (X \times I) \amalg Y$ with the quotient map is a homotopy equivalence. Indeed, the map $M_f \rightarrow Y$ sending $[(x, t)] \mapsto f(x)$ for $(x, t) \in X \times I$ and $[y] \mapsto y$ for $y \in Y$ is a homotopy inverse.

- The inclusion

$$i : X \hookrightarrow M_f : x \mapsto [(x, 0)]$$

is homotopic to $g \circ f$.

If we can now prove that $i : X \hookrightarrow M_f$ is a homotopy equivalence, it will follow that $g \circ f : X \rightarrow M_f$ and therefore also $f : X \rightarrow Y$ are homotopy equivalences. Since Y is simply connected and homotopy equivalent to M_f , the latter is also simply connected. Moreover, $g_* : H_n(Y) \rightarrow H_n(M_f)$ is an isomorphism for every $n \geq 0$, so our assumption on f_* implies that $i_* = g_* \circ f_* : H_n(X) \rightarrow H_n(M_f)$ is also an isomorphism for every $n \geq 0$. This establishes that $i : X \hookrightarrow M_f$ satisfies the same hypotheses as $f : X \rightarrow Y$ except that it is also an inclusion, so the result follows. \square

CONCLUSION OF THE PROOF OF THEOREM 56.1. We have shown thus far that if M is a closed simply connected 3-manifold, then $H_*(M) \cong H_*(S^3)$. 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 33 implies $\deg(f) = 1$.

It is trivial that $f_* : H_0(M) \rightarrow H_0(S^3)$ is an isomorphism, and so is $f_* : H_3(M) \rightarrow H_3(S^3)$ due to the degree assumption. In all other dimensions, both homology groups vanish, so we conclude that $f_* : H_*(M) \rightarrow H_*(S^3)$ is an isomorphism. Since M and S^3 are both simply connected, Corollary 56.14 now implies that f is a homotopy equivalence. \square

Bibliography

- [Are46] R. Arens, *Topologies for homeomorphism groups*, Amer. J. Math. **68** (1946), 593–610. MR19916
- [Art91] M. Artin, *Algebra*, Prentice Hall, Inc., Englewood Cliffs, NJ, 1991.
- [AD14] M. Audin and M. Damian, *Morse theory and Floer homology*, Universitext, Springer, London; EDP Sciences, Les Ulis, 2014. Translated from the 2010 French original by Reinie Ern e.
- [Boo77] B. Booss, *Topologie und Analysis*, Springer-Verlag, Berlin-New York, 1977 (German). Einf uhrung 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.
- [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.
- [ES52] S. Eilenberg and N. Steenrod, *Foundations of algebraic topology*, Princeton University Press, Princeton, New Jersey, 1952.
- [FW99] G. K. Francis and J. R. Weeks, *Conway’s ZIP proof*, Amer. Math. Monthly **106** (1999), no. 5, 393–399.
- [Gal87] D. Gale, *The Teaching of Mathematics: The Classification of 1-Manifolds: A Take-Home Exam*, Amer. Math. Monthly **94** (1987), no. 2, 170–175.
- [Hat02] A. Hatcher, *Algebraic topology*, Cambridge University Press, Cambridge, 2002.
- [Hir94] M. W. Hirsch, *Differential topology*, Springer-Verlag, New York, 1994.
- [Hut] M. Hutchings, *Cup product and intersections*. lecture notes, available at <https://math.berkeley.edu/~hutching/teach/215b-2011/cup.pdf>.
- [J an05] K. J anich, *Topologie*, 8th ed., Springer-Verlag, Berlin, 2005 (German).
- [JW58] I. M. James and J. H. C. Whitehead, *Homology with zero coefficients*, Quart. J. Math. Oxford Ser. (2) **9** (1958), 317–320.
- [Kel50] J. L. Kelley, *The Tychonoff product theorem implies the axiom of choice*, Fund. Math. **37** (1950), 75–76.
- [Kel75] J. L. Kelley, *General topology*, Graduate Texts in Mathematics, vol. 27, Springer-Verlag, New York, 1975. Reprint of the 1955 edition [Van Nostrand, Toronto, Ont.]
- [KS69] R. C. Kirby and L. C. Siebenmann, *On the triangulation of manifolds and the Hauptvermutung*, Bull. Amer. Math. Soc. **75** (1969), 742–749.
- [Lan02] S. Lang, *Algebra*, 3rd ed., Graduate Texts in Mathematics, vol. 211, Springer-Verlag, New York, 2002.
- [Lee11] J. M. Lee, *Introduction to topological manifolds*, 2nd ed., Graduate Texts in Mathematics, vol. 202, Springer, New York, 2011.
- [LL01] E. H. Lieb and M. Loss, *Analysis*, 2nd ed., Graduate Studies in Mathematics, vol. 14, American Mathematical Society, Providence, RI, 2001.
- [Man14] C. Manolescu, *Triangulations of manifolds*, ICCM Not. **2** (2014), no. 2, 21–23.
- [Mas91] W. S. Massey, *A basic course in algebraic topology*, Graduate Texts in Mathematics, vol. 127, Springer-Verlag, New York, 1991.
- [MS12] D. McDuff and D. Salamon, *J-holomorphic curves and symplectic topology*, 2nd ed., American Mathematical Society Colloquium Publications, vol. 52, American Mathematical Society, Providence, RI, 2012.
- [Mil62] J. W. Milnor, *On axiomatic homology theory*, Pacific J. Math. **12** (1962), 337–341.
- [Mil97] ———, *Topology from the differentiable viewpoint*, Princeton Landmarks in Mathematics, Princeton University Press, Princeton, NJ, 1997. Based on notes by David W. Weaver; Revised reprint of the 1965 original.
- [MS74] J. W. Milnor and J. D. Stasheff, *Characteristic classes*, Princeton University Press, Princeton, N. J., 1974. Annals of Mathematics Studies, No. 76.

- [Moi77] E. E. Moise, *Geometric topology in dimensions 2 and 3*, Springer-Verlag, New York-Heidelberg, 1977. Graduate Texts in Mathematics, Vol. 47.
- [Pal66] R. S. Palais, *Homotopy theory of infinite dimensional manifolds*, *Topology* **5** (1966), 1–16.
- [Rad25] T. Radó, *Über den Begriff der Riemannschen Fläche*, *Acta Szeged* **2** (1925), 101–121.
- [Rud87] W. Rudin, *Real and complex analysis*, 3rd ed., McGraw-Hill Book Co., New York, 1987.
- [Sch93] M. Schwarz, *Morse homology*, Progress in Mathematics, vol. 111, Birkhäuser Verlag, Basel, 1993.
- [Spa95] E. H. Spanier, *Algebraic topology*, Springer-Verlag, New York, 1995. Corrected reprint of the 1966 original.
- [Spa48] ———, *Cohomology theory for general spaces*, *Ann. of Math. (2)* **49** (1948), 407–427.
- [Ste51] N. Steenrod, *The Topology of Fibre Bundles*, Princeton University Press, Princeton, N. J., 1951.
- [Ste67] N. E. Steenrod, *A convenient category of topological spaces*, *Michigan Math. J.* **14** (1967), 133–152.
- [Tho54] R. Thom, *Quelques propriétés globales des variétés différentiables*, *Comment. Math. Helv.* **28** (1954), 17–86 (French).
- [Tho92] C. Thomassen, *The Jordan-Schönflies theorem and the classification of surfaces*, *Amer. Math. Monthly* **99** (1992), no. 2, 116–130.
- [Vic94] J. W. Vick, *Homology theory*, 2nd ed., Graduate Texts in Mathematics, vol. 145, Springer-Verlag, New York, 1994. An introduction to algebraic topology.
- [Wen18] C. Wendl, *Holomorphic curves in low dimensions: from symplectic ruled surfaces to planar contact manifolds*, *Lecture Notes in Mathematics*, vol. 2216, Springer-Verlag, 2018.