Algebraic Topology IV, Michaelmas Term

James J. Walton

Preface

These notes cover the material of the first term of Algebraic Topology IV. The central goal is the construction of the *homology groups* of a space, and to establish their main properties along with a few applications and computations. You will go deeper into the theory in the second term where you shall also meet *cohomology*. Along with the notes here, I can suggest Armstrong [Arm] as a gentle introduction to some of the earlier areas we will see (and as a refresher on some point-set topology); see it for a nice account of simplicial homology and a hands-on approach to proving its homotopy invariance. Hatcher's book [Hat] must be recommended: it is nicely written, comprehensive and even freely available online. There are plenty of other good classical texts that you can head to, see Spanier [Spa], Fulton [Ful], Dold [Dol] or Bredon [Bre]. I suggest taking a look around.

In places these notes are more descriptive, but less concise than they could be. Because this entails a longer set of notes, I shall also provide a more typical distilled version which may be more to your taste in learning the subject, or perhaps used in conjunction with these notes which can then be read when you need the extra detail.

The introduction is there just for initial motivation to the subject—in particular, I won't expect knowledge of the homotopy groups in this term (except for a basic understanding of the fundamental group for the Hurewicz Theorem at the notes). The rest of the material can be read linearly, although sometimes results are stated before we have the tools to prove them, which we then return to later. Sometimes where technical details could get in the way of an otherwise simple idea, I have moved them to the appendix. These are there for your interest and completeness but, as for the introduction, these will not constitute details I expect for you to necessarily have internalised. The exception is the proof of the Snake Lemma, which I have hidden in the appendix to give you space to work out the proof on your own first.

A previously undefined term will appear in **bold** within the text when its definition appears nearby. Concepts which either are to be defined later on, or are being described but without necessarily implying that they are central to the course, will sometimes appear in *italics*.

Some but not all of the homework exercises can be found within the text. Usually the notes' exercises will be easier, they are intended to be thought about whilst reading through the notes. Often they will indicate an idea that will be helpful in thinking about a recently introduced topic, or sometimes even ask for a proof of a result to be used later.

All illustrations were produced by the author on Inkscape.

Finally: feedback, suggestions or criticisms on the notes is encouraged!

Round-up of important notation

- $\mathbb{N}$ the natural numbers (not including zero), $\mathbb{N}_0$ the natural numbers including zero, $\mathbb{Z}$ the integers, $\mathbb{Q}$ the rational numbers, $\mathbb{R}$ the real numbers, $\mathbb{C}$ the complex numbers. The cyclic group $\mathbb{Z}/n\mathbb{Z}$ of order n is denoted $\mathbb{Z}/n$.
- 0 sometimes means zero, and sometimes means the trivial group of one element, the context making which clear.
- $I = [0, 1]$ the unit interval.
- S^n the n-sphere.
- D^n the n-disc.
- $f : X \to Y$ a morphism f with domain X and codomain Y.
- $A \hookrightarrow B$, an inclusion morphism of A into B.
- $A \twoheadrightarrow B$, a surjection from A to B (rarely used. Will usually denote a quotient map).
- $x \mapsto y$ is read 'x maps to y' when discussing the definition of a function.
- $g \circ f$ the composition of f followed by g.
- f^{-1} the inverse of a morphism f.
- $\coprod_\alpha X_\alpha$ the disjoint union of spaces X_α, written $X \coprod Y$ for just two spaces.
- $X \times Y$ the product of spaces X and Y.
- $X/\sim$ the quotient space of X under the equivalence relation $\sim$.
- X/A the quotient space of X given by collapsing the subspace $A \subseteq X$ to a point.
- $X \cong Y$, the spaces X and Y are homeomorphic.
- $f \simeq g$, the maps f and g are homotopic.
- $X \simeq Y$, the spaces X and Y are homotopy equivalent.
- $G \cong H$, the groups G and H are isomorphic.
- $\bigoplus_\alpha G_\alpha$ the direct sum of groups G_α, written $G \oplus H$ for just two groups.
- C_*, and variants, indicate a chain complex.
- H_* homology, occasionally $H_*(C)$ to mean the homology of the chain complex C_*.
- $f_\sharp$ chain map, sometimes written (f_n).
- f_* indicates an induced map on homology of some chain map.
- $H_*(\mathcal{K})$ the simplicial homology of a simplicial complex $\mathcal{K}$, $H_*(X)$ the singular homology of a space X, $H_*(X^\bullet)$ the cellular homology of a CW complex $X^\bullet$, $H_*(X, A)$ the relative homology of a pair of spaces (X, A), $\widetilde{H}_*(X)$ the reduced homology of a space X (in short, what $H_*(-)$ is depends on what we put into the brackets).
- $\upharpoonright_j$ restriction to index j face, with identification of that face with standard simplex.

Contents

Chapter 0

Introduction

0.1 What is topology?

Topology is the study of abstract spaces which restricts attention to attributes which remain invariant under deformation. It cares not for angles or lengths, but for more 'rugged' features. For example, any loop inscribed on the surface of a ball can always be shrunk to a point within that surface, whereas on a torus there are plenty of non-trivial loops which cannot be pulled tight—the same is true for any torus, whether it looks more like the surface of a doughnut or the surface of a coffee cup. The fact that

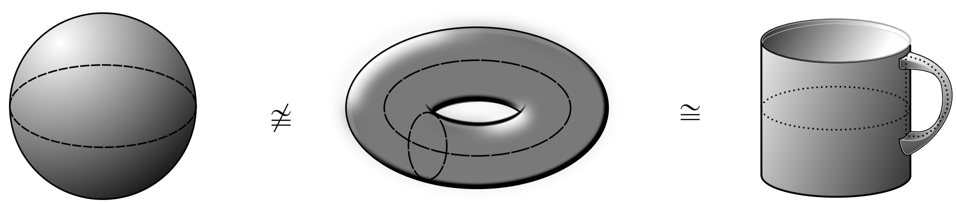

Figure 0.1.1: The 2-sphere, the 2-torus and a coffee mug, whose surface is homeomorphic to the 2-torus.

spaces are allowed to be freely deformed may at first seem like a weakness. If one is building a bridge, one cares about angles and lengths! But for some problems, it is natural to strip away some of the geometry to get to the heart of the matter. In finding a path crossing each of the bridges of Königsburg only once, the lengths of the bridges are irrelevant, all that matters is how many bridges there are and how they connect to each other. If you are investigating how a string is knotted, it doesn't matter what the length of the string is or exactly how you choose to embed it in space. If you want to know what kinds of vector fields (think of wind currents) you can have flowing on a surface, you can determine a lot without the exact proportions of the surface, but just its topological shape.

The utility of topology is not just in thinking about things which are obviously geometric already such as, say, asking how a string is knotted or what the shape of the universe might be. Often collections of mathematical objects naturally hold a geometry. Algebraic geometry, for example, studies the solutions to algebraic equations by arranging them into a moduli space of solutions. The solutions of differential equations can also sometimes be considered geometrically. One may wish to consider the state space of a physical system, perhaps shedding light on the types of energy functions it can support. To accurately plan the movement of a robot's arm or how goods are transported automatically around a warehouse, one considers 'configuration spaces'. An intriguing recent development is the study of large data sets topologically, so as to determine more qualitative rather than quantitative features, using tools such as persistent homology. Topology has for some time been a crucial component to many pure mathematicians' tool-kits, and increasingly now it finds uses within applied mathematics too.

0.2 What is *algebraic* topology?

By considering a geometric object merely *topologically*, one can simplify the setting of a problem at the cost of loosing some geometric information. In algebraic topology one goes further by applying tools which forget some topological information in return for algebraic information. These associated algebraic objects can sometimes be computed, or implemented in other theoretically useful ways. If one is lucky, enough information is retained to solve your problem.

As a rather crude example, one may associate to a space X the number of its path-connected components. This loses a lot of information! But we may usually quite easily compute this number, and if we find that two spaces have a different number of path components then we can be sure that they are different spaces.

A less trivial example (which in a sense naturally follows the above one as the next member of a family of invariants) is the fundamental group $\pi_1(X)$, which you likely met in the module Topology III. To slightly simplify, the fundamental group is determined by 'loops' in your space, maps $f\colon S^1 \to X$, where S^1 is the circle (considered *up to homotopy*). For example, for S^2 the 2-sphere we have that $\pi_1(S^2)$ is the trivial group consisting of one element, since all loops can be pulled tight to a 'constant' loop, but for the torus $\mathbb{T}^2$ we have that $\pi_1(\mathbb{T}^2) \cong \mathbb{Z}^2$, two generators may be given as a loop around the meridian of the torus, and around its longitude. By passing to the fundamental group we still lose information, two different spaces can quite easily have isomorphic fundamental groups. The tools of algebraic topology can often be used to distinguish spaces, but less often can they be used to show that two spaces are the same.

One is often interested not only in spaces, but also in continuous maps between spaces. The tools of algebraic topology are usually *functorial* (see Section 1.1). As a result, one of these gadgets not only assigns spaces X and Y algebraic objects $F(X)$ and $F(Y)$, it also assigns to a continuous map $f\colon X \to Y$ a homomorphism $f_*\colon F(X) \to F(Y)$. This can be useful in ruling out the existence of particular kinds of maps between spaces, as we shall see shortly in the proof of Brouwer's Fixed Point Theorem.

0.3 Homotopy and homology groups

Many spaces that one encounters can be built by glueing together 'cells': points, line-segments, discs, 3-dimensional balls, and their higher dimensional analogues. These spaces are called CW-complexes (see Section 1.2.1). It turns out that the fundamental group $\pi_1(X)$ of such a space does not 'see' the effect of adding cells above dimension two. So whilst $\pi_1(X)$ is an important invariant, it isn't too much good in isolation especially in analysing higher dimensional spaces.

There is a higher dimensional notion of a 'loop' in a space. Let S^n be the n-sphere, the set of points in $\mathbb{R}^{n+1}$ distance 1 from the origin. We could take an n-dimensional loop to mean a mapping $f\colon S^n \to X$. For example, for $n = 1$ you're just mapping the circle into your space as before in the fundamental group, and for $n = 2$ you would consider how one may map the 2-sphere into X, in particular whether you can do it in a way so as to 'catch' some hole in the space from which you cannot pull the sphere away.

This idea can be carried through and be used to define important invariants $\pi_n(X)$ called the *homotopy groups* of X. The homotopy groups are an incredibly powerful tool for distinguishing spaces. Unfortunately, it is often the case that they are *very* difficult to compute. It is a major problem in algebraic topology, for example, to calculate the homotopy groups of the n-spheres, which are centrepiece in the study of homotopy theory.

Much of our attention in this term will be devoted to defining and studying the *homology groups* of a space. For each topological space X and $n \in \mathbb{Z}_{\geqslant 0}$ one has the *degree n homology group*, denoted $H_n(X)$, which is an Abelian group. The nth homotopy group $\pi_n(X)$ is defined via probing the space X with n-dimensional spheres. In a loose sense, the nth homology group $H_n(X)$ may still be thought of as considering ways in which certain n-dimensional objects sit inside X, but unfortunately the definition is more complicated, certainly more so than for the definition of the homotopy groups. One refund for this effort, though, is that the homology groups are often readily computable.

0.4 The Brouwer Fixed Point Theorem, using homology as a 'black box'

To see the power of these techniques, we close this introduction by giving a proof of the famous Brouwer Fixed Point Theorem using homology. Of course, we have not defined homology yet, and have barely even attempted to describe it! However, one rarely uses the *definition* of homology directly, rather than its useful *properties*. Below are the properties that we shall use, where 'homology' should be interpreted as 'singular homology' (in fact, *reduced* singular homology if one is being picky, so that points 2 and 3 apply for $n = 0$):

1. 'functorality' (see Definition 1.1.2, Lemma 3.3.4):
   - for a space X and $n \in \mathbb{Z}_{\geqslant 0}$ we have an Abelian group $H_n(X)$;
   - for a continuous map $f\colon X \to Y$ we have induced homomorphisms $f_*\colon H_n(X) \to H_n(Y)$;
   - the identity map $\mathrm{id}\colon X \to X$ induces identity homomorphisms $\mathrm{id}_*\colon H_n(X) \to H_n(X)$;
   - for two continuous maps $f\colon X \to Y$ and $g\colon Y \to Z$, we have that $(g \circ f)_* = g_* \circ f_*$;
2. the homology groups $H_n(X)$ of a contractible space X are trivial in each degree (see e.g., Theorem 3.3.1 and Exercise 1.2.5);
3. the degree n homology group $H_n(S^n)$ of the n-sphere is isomorphic to $\mathbb{Z}$ (see Theorem 4.4.2).

Theorem 0.4.1 (Brouwer Fixed Point Theorem)**.** *Let $n \geqslant 1$ and X be homeomorphic to the closed n-disc D^n. Any continuous map $f\colon X \to X$ has a fixed point i.e., there exists $x \in X$ with $f(x) = x$.*

In the above theorem the *closed n-disc* D^n is the space of points $\{x \in \mathbb{R}^n \mid |x| \leqslant 1\}$ of $\mathbb{R}^n$ of distance at most 1 from the origin. As a rather visual example, consider a mug filled with coffee. The space X that the coffee occupies, cylindrical in shape, is homeomorphic to D^3 (unless you have a more jazzy kind of mug). Stirring the coffee and then setting it to rest, logging where each point moves from and to after stirring, one gets a map $f\colon X \to X$. The above theorem says that there is always a point x in the coffee which returns to its original location after stirring! If you tried to defy the theorem by stirring x to a different location then this guarantees that at least one other point y will now be in its original location. In reality, of course, the coffee consists of particles and is not a continuous medium. But it consists of so many particles that it is quite accurately modelled by a continuous medium, and two nearby particles will be stirred to relatively nearby locations, unless one causes splashes, so the stirring is still accurately modelled by a continuous map.

Proof of Brouwer Fixed Point Theorem. Firstly, we may as well take X as the n-disc D^n itself; if the result holds for D^n it holds for any homeomorphic space. Indeed, suppose that $h\colon X \to D^n$ is a homeomorphism and $f\colon X \to X$. Then $h \circ f \circ h^{-1}$ has some fixed point $x \in D^n$, so f has a fixed point $h(x)$.

The proof will follow from the properties of homology listed above, and one trick: supposing that there exists a map on D^n without fixed point, we claim that there must exist a retraction of D^n to its boundary sphere $S^{n-1} \subseteq D^n$, that is, a continuous map $r\colon D^n \to S^{n-1}$ for which $r(x) = x$ for all $x \in S^{n-1}$. To construct this map, consider the half-infinite ray $r_f(x)$ travelling from (but not including) $f(x)$ through x, which is well-defined since we are assuming that $f(x) \neq x$. Define $r(x)$ to be the point of intersection of $r_f(x)$ with S^{n-1}. Since f is continuous, the rays $r_f(x)$ and their intersections with S^{n-1} vary continuously too, and for $x \in S^{n-1}$ we clearly have that $r(x) = x$. So now we need only show that such a map r cannot exist. Consider

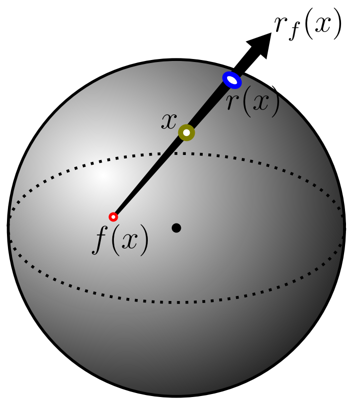

Figure 0.4.1: Defining retraction $r\colon D^n \to S^{n-1}$.

the identity map $\mathrm{id}\colon S^{n-1} \to S^{n-1}$. We may factorise it as $\mathrm{id} = r \circ i$, where i is the inclusion map $i\colon S^{n-1} \to D^n$ and $r\colon D^n \to S^{n-1}$ is the retraction we constructed above. Applying homology in degree $n-1$ we have that $\mathrm{id}_* = (r \circ i)_* = r_* \circ i_*$ by item 1 on the functorality of homology. Since D^n is contractible its homology is trivial in degree $(n-1)$ by item 2 above, and by item 3 we have that $H_{n-1}(S^{n-1}) \cong \mathbb{Z}$. Applying homology in degree $(n-1)$ to the commutative diagram of spaces and maps $\mathrm{id} = r \circ i$ thus leads to the following commutative diagram of groups and homomorphisms:

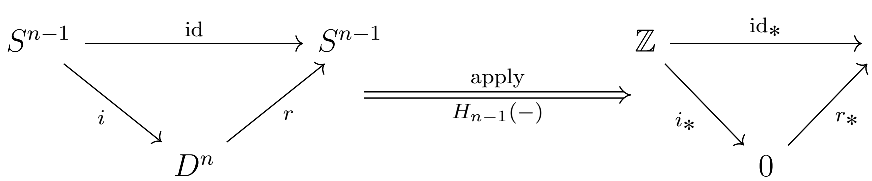

(see the next section on commutative diagrams; here it just means that following the top arrows from left to right is the same as the composition going down-right, then up-right). However, again by functorality, the map id_* should be the identity map $x \mapsto x$ on $\mathbb{Z}$. The only map which can factor through the trivial group as in the above diagram is the trivial map $x \mapsto 0$. This is a contradiction, and so the retraction r, and hence the map f without fixed points, cannot exist. $\square$

Chapter 1

Some foundations

1.1 Basic category theory

1.1.1 Categories and examples

Definition 1.1.1. A **category** $\mathcal{C}$ consists of a class $\mathrm{Ob}_{\mathcal{C}}$ of **objects** and a class $\mathrm{Hom}_{\mathcal{C}}$ of **morphisms**. A morphism $f \in \mathrm{Hom}_{\mathcal{C}}$ has some **domain** $X \in \mathrm{Ob}_{\mathcal{C}}$ and **codomain** $Y \in \mathrm{Ob}_{\mathcal{C}}$, denoted $f\colon X \to Y$. For $f\colon X \to Y$ and $g\colon Y \to Z$ there must be defined a **composition** $g \circ f\colon X \to Z$ in $\mathrm{Hom}_{\mathcal{C}}$. We require the following two properties, of *identities* and *associativity*, respectively:

- for each object $X \in \mathrm{Ob}_{\mathcal{C}}$ there is an **identity morphism** $\mathrm{id}_X\colon X \to X$, a morphism satisfying $f \circ \mathrm{id}_X = f$ and $\mathrm{id}_X \circ g = g$ for any $f\colon X \to Y$ and any $g\colon Y \to X$;
- for $f\colon W \to X$, $g\colon X \to Y$ and $h\colon Y \to Z$ we have that $(h \circ g) \circ f = h \circ (g \circ f)$.

Exercise 1.1.1. Show that for any category $\mathcal{C}$ identities are unique, that is, there is precisely one identity morphism for each object.

We'll list below some examples of categories. *Many of them we will not need again* (so don't go out of your way trying to commit them to memory!), but they should help to clarify the concept:

Example 1.1.1. One of the most important categories is **Set**, the category of sets. Its objects $\mathrm{Ob}_{\mathbf{Set}}$ are all sets, and the morphisms $\mathrm{Hom}_{\mathbf{Set}}$ are simply functions between sets. Identity morphisms $\mathrm{id}_S\colon S \to S$ are the functions $\mathrm{id}_S(s) = s$ for all $s \in S$. Composition is defined by standard function composition: if $f\colon A \to B$ and $g\colon B \to C$, then $g \circ f\colon A \to C$ is defined as the function with $g \circ f(a) \coloneqq g(f(a))$ for $a \in A$.

Example 1.1.2. There is a category **Grp** of groups, with objects all groups and morphisms group homomorphisms between them. There is another **Ab** whose objects are just the Abelian groups, with morphisms group homomorphisms. Similarly there is a category of vector spaces and linear maps, of rings and ring homomorphisms, of R-modules and R-module homomorphisms, of semigroups and semigroup homorphisms, of . . . .

Example 1.1.3. For any particular set S we have the (rather boring) category with $\mathrm{Ob}_{\mathcal{C}} = S$ and only identity morphisms $\mathrm{id}_x : x \to x$ for each $x \in S$.

Example 1.1.4. One may quite easily write down finite categories. One simple example is the category of two objects A and B, with single non-identity morphism $f : A \to B$. There are no choices as to the composition rule here: any valid composition in this category involves an identity morphism (e.g., $f \circ \mathrm{id}_A = f$).

The composition rule is not always uniquely defined from the objects and morphisms though, *the composition rule is part of the data defining a category.* For instance, consider again a category with two objects A and B, but three non-identity morphisms (so five including identities): $f_1 : A \to B$, $f_2 : B \to A$ and $g : B \to B$. The category is not uniquely defined by this data, one also needs the composition rule:

Exercise 1.1.2. Find two different categories for the example above of two objects with three non-identity morphisms. Also find a non-example with a composition rule which satisfies identities but not associativity.

Example 1.1.5. Let $(S, \leqslant)$ be a partially ordered set, so that some elements $x, y \in S$ are related, written $x \leqslant y$, in a way which is reflexive ($x \leqslant x$ for all $x \in S$), antisymmetric ($x \leqslant y$ and $y \leqslant x$ implies that $x = y$) and transitive ($x \leqslant y$ and $y \leqslant z$ implies that $x \leqslant z$). There is a category naturally associated to this, with objects $\mathrm{Ob}_{\mathcal{C}} = S$ and precisely one morphism $f : x \to y$ whenever $x \leqslant y$.

Exercise 1.1.3. There is only one possible composition rule for the above example. Check that this indeed defines a category.

Example 1.1.6. An important category for us will be the category **Top** with objects topological spaces and morphisms given by continuous maps. Composition of maps is defined in the usual way; note that a composition of continuous maps is continuous and identity maps are always continuous. There is also a category $\mathbf{Top}_*$ of pointed topological spaces (given by a pair of a topological space X and point $x \in X$), with morphisms continuous maps between spaces preserving base points. One can also define a category $\mathbf{Top}^2$ of *pairs* of topological spaces (X, A) with A a subspace of X. A morphism in this category between pairs (X, A) and (Y, B) is given by a continuous map $f : X \to Y$ with $f(A) \subseteq B$. Analogously there is a category $\mathbf{Top}^n$ for any $n \in \mathbb{N}$.

Also important is the *homotopy category* **hTop** whose objects are still topological spaces but with morphisms *homotopy classes* of continuous maps (see Section 1.2). Lemma 1.2.1 shows that composition is well-defined in this category.

Example 1.1.7. We may think of a single particular group G as its very own category, consisting of a single object $*$ and one morphism $g\colon * \to *$ for each group element $g \in G$. Composition is defined by the multiplication rule in G: we define $g \circ h \coloneqq g \cdot h$.

In this category every morphism is **invertible**, otherwise known as being an **isomorphism**, that is, for any morphism $f\colon A \to B$ there exists some morphism $f^{-1}\colon B \to A$ (called the **inverse** of f) with $f^{-1} \circ f = \mathrm{id}_A$ and $f \circ f^{-1} = \mathrm{id}_B$. A category consisting of a set of only invertible morphisms and a single object defines a group, by taking as elements the set of morphisms and defining multiplication according to the composition rule (**Exercise:** show that this does in fact result in a group). Generalising this, a neat definition of a 'groupoid', if you've heard of one, is a category with a set of morphisms which are all invertible.

Example 1.1.8. When we meet them (in Chapter 2), the chain complexes with chain maps as morphisms form a category **Ch**.

Example 1.1.9. For any category $\mathcal{C}$ we can form its *opposite category* $\mathcal{C}^{\mathrm{op}}$. We define $\mathrm{Ob}_{\mathcal{C}^{\mathrm{op}}} \coloneqq \mathrm{Ob}_{\mathcal{C}}$ but 'flip' the morphisms: we take one morphism $f^{\mathrm{op}}\colon Y \to X$ for each $f\colon X \to Y$. Composition of functions is defined from composition in $\mathcal{C}$ in the obvious way.

A **diagram** in a category $\mathcal{C}$ is a graph with nodes labelled by objects of $\mathcal{C}$ and (oriented) edges labelled by morphisms of $\mathcal{C}$, in a way so if an edge is labelled by a morphism $f\colon X \to Y$ then the tail of that edge is labelled by X and the head is labelled by Y. Such a diagram is called a **commutative diagram** if any two paths around the graph with the same origin and terminal nodes are such that their corresponding compositions of morphisms agree.

For example, the following diagram

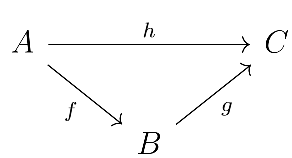

is commutative precisely when $h = g \circ f$: going straight from A to C should be the same as following f then g. You should be able to convince yourself that the following (quite randomly chosen) diagram

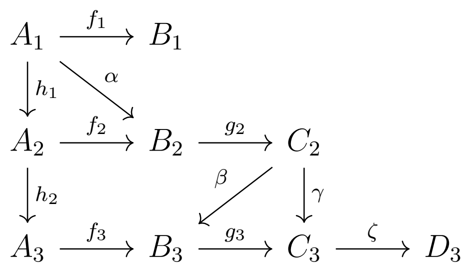

is commutative precisely when $f_2 \circ h_1 = \alpha$, $g_3 \circ \beta = \gamma$ and $\beta \circ g_2 \circ f_2 = f_3 \circ h_2$.

Exercise 1.1.4. Consider a diagram in some category of the following shape, where f is an isomorphism:

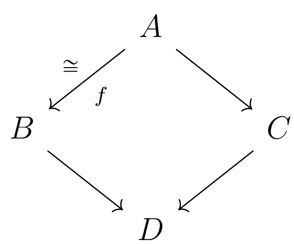

Show that this diagram is commutative if and only if the similar diagram with the arrow f replaced with its inverse is also commutative. See Example 1.1.7 for the definition of an isomorphism in a category and its inverse. The inverse f^{-1} is uniquely defined by a later Exercise 1.1.6.

Exercise 1.1.5. Don't assume that the behaviour above works for all shapes of diagrams though. Find a commutative diagram in some category (for example, in **Ab**) where one edge is an isomorphism but the diagram given by flipping that edge, labelling it instead with f^{-1}, is not commutative.

1.1.2 Maps between categories: functors

There is a natural notion of a morphism between two categories:

Definition 1.1.2. Let $\mathcal{C}$ and $\mathcal{D}$ be two categories. A (covariant) **functor** from $\mathcal{C}$ to $\mathcal{D}$ assigns to each object $X \in \mathrm{Ob}_{\mathcal{C}}$ an object $F(X) \in \mathrm{Ob}_{\mathcal{D}}$ and to each morphism $f\colon X \to Y$ in $\mathrm{Hom}_{\mathcal{C}}$ a morphism $F(f)\colon F(X) \to F(Y)$ in $\mathrm{Hom}_{\mathcal{D}}$. The following axioms of identities and composition, respectively, must be satisfied:

- an identity morphism id_X of $\mathcal{C}$ is sent to the identity morphism $\mathrm{id}_{F(X)} = F(\mathrm{id}_X)$ in $\mathcal{D}$;
- for $f\colon X \to Y$ and $g\colon Y \to Z$ in $\mathrm{Hom}_{\mathcal{C}}$, we have that $F(g \circ f) = F(g) \circ F(f)$.

Suppose that you have a commutative diagram over a category $\mathcal{C}$ and a functor F from $\mathcal{C}$ to $\mathcal{D}$. Then applying F to the diagram we also get a commutative diagram in $\mathcal{D}$ of the same shape, by replacing an object $A \in \mathrm{Ob}_{\mathcal{C}}$ labelling a node with $F(A) \in \mathrm{Ob}_{\mathcal{D}}$, and a morphism $f \in \mathrm{Hom}_{\mathcal{C}}$ labelling an edge by $F(f) \in \mathrm{Hom}_{\mathcal{D}}$ (quickly think about why this diagram is still commutative). Compare with the pair of diagrams in the proof of Brouwer's Fixed Point Theorem (Theorem 0.4.1) in the introduction.

Example 1.1.10. Interestingly, with the notion of a functor as a morphism, the categories themselves form a category! That is, we have a category[1] **Cat** with objects the categories and morphisms functors between categories. **Exercise:** check that this does define a category.

Exercise 1.1.6. Recall the definition of an isomorphism in a category from Example 1.1.7. Show that the inverse of an isomorphism in a category is uniquely defined. Moreover, show that if f^{-1} is the inverse of f and F is a functor, then $F(f)$ is an isomorphism with inverse $(F(f))^{-1}$ given by $F(f^{-1})$.

Example 1.1.11. We may associate to any group its underlying set, and to any homomorphism between groups the function between their underlying sets. This defines a functor $F\colon \mathbf{Grp} \to \mathbf{Set}$. For the category **Ring** of rings and ring homomorphisms, we could forget about the multiplicative structure keeping just the addition to get a group, defining a functor $F\colon \mathbf{Ring} \to \mathbf{Grp}$, or we could forget about the group structure too getting a functor $F\colon \mathbf{Ring} \to \mathbf{Set}$. We have a similar functor $F\colon \mathbf{Top} \to \mathbf{Set}$. Functors like this are sometimes called *forgetful functors*.

Example 1.1.12. I feel that it is quite instructive to visualise a category as a network of nodes (the objects) and arrows (the morphisms). There is a functor which realises this[2]. Let F be the functor from **Cat** to **Grph**, the category of directed graphs, which sends a category $\mathcal{C}$ to the graph with nodes the objects of $\mathcal{C}$ and one arrow for each morphism f of $\mathcal{C}$, with tail and head corresponding to the domain and codomain of f, respectively. This is also forgetful functor, in the sense that one may think of a category as a directed graph with the extra structure of a composition rule.

Example 1.1.13. Recall from Example 1.1.7 how we associated to a group G a category (let's call it $\mathcal{G}$). Such a category has just one object and morphisms in bijection with

[1]For those interested, there is also a notion of a map between functors, called a *natural transformation*. That is, one has a notion of maps between categories (functors) and maps between maps of categories (natural transformations). One may consider how to continue this process, leading to the subject of *higher category theory*.

[2]Technically here we should restrict to so-called *small categories*, so that the objects and morphisms form a set rather than a class (this avoids set-theoretic issues, such as Russell's Paradox which occurs when one tries to define the collection of all sets as itself a set). I would rather ignore these size issues here though.

the group elements. So a function $f\colon G \to H$ induces a function $F\colon \mathrm{Hom}_{\mathcal{G}} \to \mathrm{Hom}_{\mathcal{H}}$ in the obvious way. **Exercise:** by sending the single object of $\mathcal{G}$ to that of $\mathcal{H}$, show that this defines a functor $F\colon \mathcal{G} \to \mathcal{H}$ if and only if f is a group homomorphism.

Similarly, we associated to a partially ordered set a category. Show how order-preserving functions between partially ordered sets precisely correspond to functors between their associated categories.

Exercise 1.1.7. Show that a functor $F\colon \mathcal{G} \to \mathbf{Set}$, where $\mathcal{G}$ is the one object category associated to a group, corresponds precisely to a group action on a set. Recall that an action of G on a set S is an assignment $g \cdot s \in S$ for all $g \in G$ and $s \in S$, with $e \cdot s = s$ for all $s \in S$ (and e the identity of G) and $(gh) \cdot s = g \cdot (h \cdot s)$ for all $g, h \in G$ and $s \in S$.

Example 1.1.14. For each $n \in \mathbb{Z}_{\geqslant 0}$ taking the nth homotopy group defines a functor[3] from $\mathbf{Top}_*$ to $\mathbf{Grp}$. In particular, whenever one has a continuous map $f\colon X \to Y$ preserving base-points, one gets a group homomorphism $f_*\colon \pi_n(X, x_0) \to \pi_n(Y, y_0)$.

We shall see later how, for each $n \in \mathbb{Z}_{\geqslant 0}$, taking the degree n homology $H_n(X)$ defines a functor from $\mathbf{Top}$ to the category $\mathbf{Ab}$ of Abelian groups. As is the case for the homotopy groups, induced homomorphisms are invariant under homotopy, so taking homology even defines a functor from the homotopy category $\mathbf{hTop}$ to $\mathbf{Ab}$.

Remark 1.1.1. Recall that to a category $\mathcal{C}$ one may associate its opposite $\mathcal{C}^{\mathrm{op}}$, which is given by flipping the rôles of the domains and codomains of morphisms. A functor $F\colon \mathcal{C}^{\mathrm{op}} \to \mathcal{D}$ is called a *contravariant functor* from $\mathcal{C}$ to $\mathcal{D}$. Explicitly, such a thing still defines a map $F\colon \mathrm{Ob}_{\mathcal{C}} \to \mathrm{Ob}_{\mathcal{D}}$ and sends morphisms of $\mathrm{Hom}_{\mathcal{C}}$ to $\mathrm{Hom}_{\mathcal{D}}$. But it now reverses the directions of arrows: a contravariant functor sends a morphism $f\colon X \to Y$ to $F(f)\colon F(Y) \to F(X)$. It must still send compositions to (flipped) compositions and identities to identities.

Our functors will be covariant rather than contravariant in this term, so we shall drop that adjective in these notes. The definition of a contravariant functor may at first seem unnatural. But there are many important examples of them (loosely speaking, one often finds contravariant functors in situations where quantities "pull back", such as vector bundles under continuous maps, rather than "push forward"). One simple example is given by taking the dual of a vector space, which defines a *contravariant* functor $F\colon \mathbf{Vect} \to \mathbf{Vect}$ on the category of vector spaces and linear maps. Indeed a linear map $f\colon V \to W$ defines a dual linear map $\hat{f}\colon \hat{W} \to \hat{V}$ between the dual vector spaces (think of matrix transposition). Another important example of a contravariant functor is cohomology, which you will meet in the next term. For $n \in \mathbb{Z}_{\geqslant 0}$ it assigns to

[3] Again, a remark just for outside interest: one way of avoiding *pointed* topological spaces for $\pi_1(-)$ is to consider instead the so-called 'fundamental groupoid'. This removes the need for base-points, and defines the fundamental groupoid functor Π_1 from $\mathbf{Top}$ to the category $\mathbf{Grpd}$ of groupoids.

a topological space X an Abelian group $H^n(X)$ and to a continuous map $f\colon X \to Y$ a group homomorphism $f^*\colon H^n(Y) \to H^n(X)$.

1.2 Homotopy theory

Let $I \coloneqq [0,1]$ be the unit interval in $\mathbb{R}$. This is probably the most important space in homotopy theory! A close runner up might be the n**-sphere** (for $n \in \mathbb{N}_0$):

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

That is, the n-sphere is the set of points of $\mathbb{R}^{n+1}$ unit distance from the origin. For example, $S^0 = \{-1, +1\}$ is a disconnected two point space, S^1 is the circle and S^2 is the surface of a ball.

Whilst S^3 most naturally sits in $\mathbb{R}^4$—so one might expect it being hopeless to try to visualise it—in fact one can think of S^3 topologically as given by usual 3d space $\mathbb{R}^3$ with an extra point 'at infinity' (and in general S^n as $\mathbb{R}^n$ 'compactified with a point at infinity' c.f., Exercise 1.2.1 below).

The $(n-1)$-sphere is the boundary of the n**-disc**, defined as:

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

The disjoint union of a collection of spaces X_α is denoted by $\coprod_\alpha X_\alpha$, or just as $X \coprod Y$ for two spaces X and Y. For example, S^0 may be constructed as the disjoint union $* \coprod *$ of two copies of the one point space $*$. This operation is sometimes known as the *co-product* for category theoretic reasons that I won't get into.

I will expect you to be comfortable with the concept of a continuous map $f\colon X \to Y$ between topological spaces X and Y. Functions between spaces will always be continuous, so we just use the word 'map' rather than continuous map. I will also expect you to know about products $X \times Y$ of spaces and quotient spaces/maps. Given an equivalence relation $\sim$ on a space X, there is an associated quotient map $q\colon X \to X/\sim$ which glues together points in X identified via $\sim$. The standard example is that we can form the torus as a quotient of the unit square $[0,1] \times [0,1]$ by identifying opposite sides of the square, as in Figure 1.2.1.

One often defines the quotient space which collapses all of the points of a subspace $A \subseteq X$ to a single point, it is denoted X/A.

For any space X we have the identity map $\mathrm{id}_X\colon X \to X$ defined by $\mathrm{id}_X(x) = x$. Two spaces X and Y are **homeomorphic**, written $X \cong Y$, if there exists a continuous map $f\colon X \to Y$ which has a continuous inverse $f^{-1}\colon Y \to X$ ($f^{-1} \circ f = \mathrm{id}_X$ and $f \circ f^{-1} = \mathrm{id}_Y$). We should think of homeomorphic spaces as basically being 'the same', up to perhaps a different naming convention for their points.

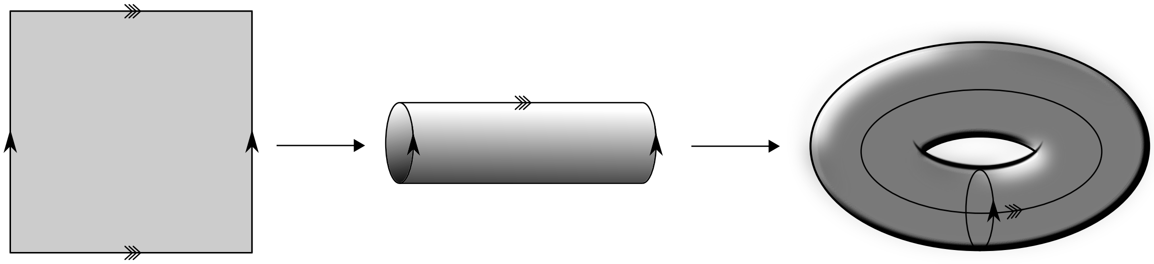

Figure 1.2.1: Glueing the 2-torus from I^2.

Exercise 1.2.1. Show that $S^n \cong D^n/S^{n-1}$. That is, the n-sphere is homeomorphic to an n-disc with its boundary collapsed to a point. Conclude that $S^n - \{x\} \cong \mathbb{R}^n$ for any point $x \in S^n$.

It is difficult make progress in trying to distinguish spaces up to homeomorphism. In fact, I think it's fair to say that it is usually *more interesting* to study them up to a weaker notion of equivalence. Enter homotopy theory:

Definition 1.2.1. Two maps $f, g \colon X \to Y$ are called **homotopic** if there exists a **homotopy** $F \colon X \times [0,1] \to Y$ between them, a map for which $F(x,0) = f(x)$ and $F(x,1) = g(x)$ for all $x \in X$. In this case we write $f \simeq g$.

A more dynamical way of thinking about a homotopy is that we have a continuously parametrised family of maps F_t, each given by $x \mapsto F(x,t)$. That F is continuous intuitively means that these F_t vary continuously between each other as t is varied continuously. This family of maps begins with $F_0 = f$ and ends at $F_1 = g$. See Figure 1.2.2.

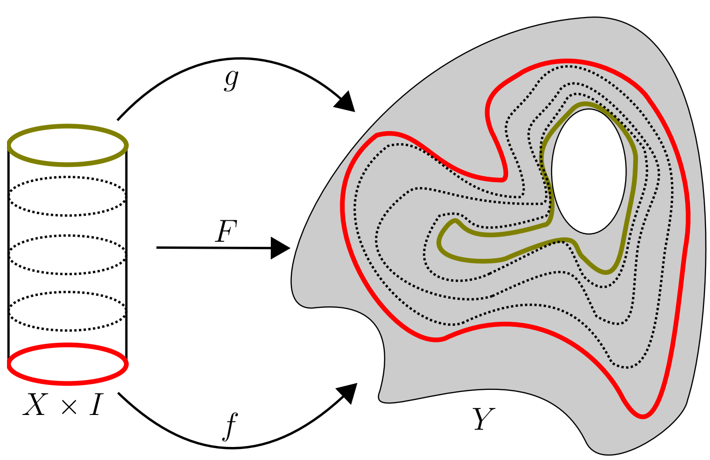

Figure 1.2.2: Homotopy between maps from $X = S^1$ to some space Y.

Definition 1.2.2. Two spaces X and Y are called **homotopy equivalent** if there exist maps $f\colon X \to Y$ and $g\colon Y \to X$ for which $g \circ f \simeq \mathrm{id}_X$ and $f \circ g \simeq \mathrm{id}_Y$. In this case we write $X \simeq Y$.

The notion of homotopy equivalence is far weaker than that of homeomorphism. For example, the n-discs are all homotopy equivalent to each other and to the one point space. Try the following basic exercises:

Exercise 1.2.2. Show that homotopy equivalence is indeed an equivalence relation, so $X \simeq X$ for any space X, if $X \simeq Y$ then $Y \simeq X$ and if $X \simeq Y$, $Y \simeq Z$ then $X \simeq Z$. Transitivity is where some small amount of work needs to be done.

Exercise 1.2.3. Show that S^n is homotopy equivalent to $\mathbb{R}^{n+1} - \{0\}$.

Exercise 1.2.4. Without being too rigorous about it, group the letters of the alphabet $\{A, B, C, \cdots, Z\}$ into their homeomorphism classes (this may depend on your choice of font), considering these as subspaces of $\mathbb{R}^2$. Then group these classes up to the weaker notion of homotopy equivalence. Does it make a difference whether some of your letters are infinitely thin (1-dimensional) instead of slightly thickened (2-dimensional)?

Exercise 1.2.5. A space is called **contractible** if $\mathrm{id}_X \simeq c_{x_0}$, where $c_{x_0}\colon X \to X$ is the *constant map*, defined by $x \mapsto x_0$ for all $x \in X$ and some fixed x_0. Show that X being contractible is equivalent to X being homotopy equivalent to the one point space.

There is a category called the **homotopy category**, denoted **hTop**, whose objects are topological spaces and whose morphisms are homotopy classes of continuous maps. Denote the homotopy class of a map by $[f]$ (i.e., the equivalence class of maps homotopy equivalent to f). By definition, the composition $[g] \circ [f]$ of continuous maps $f\colon X \to Y$ and $g\colon Y \to Z$ is given by $[g \circ f]$. We need to check that this is well-defined:

Lemma 1.2.1. *If $f \simeq f'\colon X \to Y$ and $g \simeq g'\colon Y \to Z$ then $g \circ f \simeq g' \circ f'$.*

Proof. Suppose that F is a homotopy realising $f \simeq f'$ and G realises $g \simeq g'$. Then the map $(x, t) \mapsto G(F(x, t), t)$ is a continuous map from $X \times I \to Z$. It is a homotopy from $g \circ f$ to $g' \circ f'$. Indeed for $t = 0$ we have that $G(F(x, 0), 0) = G(f(x), 0) = g(f(x)) = (g \circ f)(x)$ and for $t = 1$ we have that $G(F(x, 1), 1) = G(f'(x), 1) = g'(f'(x)) = (g' \circ f')(x)$. □

It follows that **hTop** is indeed a category: the identity morphism on X is the homotopy class $[\mathrm{id}_X]$ and associativity of usual function composition implies associativity in this category.

In terms of homotopy theory, perhaps the most boring (but surprisingly useful!) spaces are those homotopy equivalent to the one point space $*$. As mentioned in Exercise 1.2.5, these are the **contractible** spaces. Any space can be nicely embedded as a subspace of a contractible space by taking its **cone**:

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

That is, the cone is defined by taking the product with the interval I and collapsing the top of this 'cylinder' to a point, see Figure 1.2.3. A homeomorphic copy of X sits inside the base $X \times \{0\}$ of the cone, or indeed as $X \times \{t\}$ for any $t \in [0, 1)$.

Exercise 1.2.6. Show that for any space X its cone CX is contractible.

Although it will only make a small appearance towards the end of the notes, another very important construction in homotopy theory is the **suspension**. It is defined as:

$$\Sigma X \coloneqq (X \times I)/((x, 0) \sim (x', 0) \text{ and } (x, 1) \sim (x', 1) \text{ for all } x, x' \in X).$$

That is, you get the suspension by taking the product with the interval, pinching the top $X \times \{1\}$ to a point (as for the cone), and also pinching the base $X \times \{0\}$ to a point. You may like to think of the suspension as given by two copies of the cone of X, each identified at their bases, which are two copies of X.

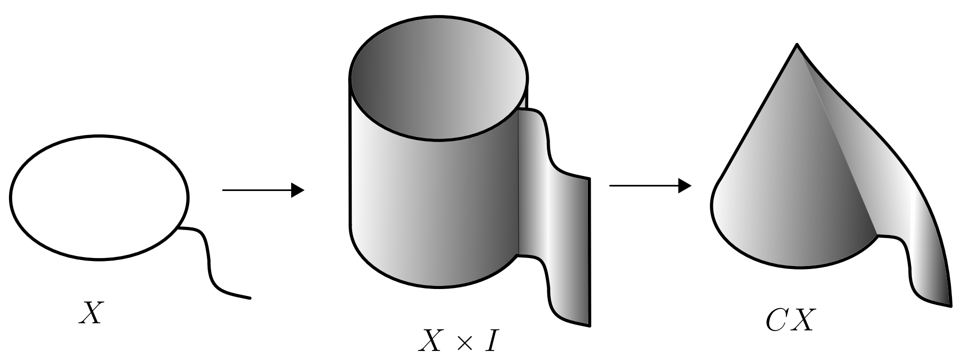

Figure 1.2.3: A space X, the 'cylinder' $X \times I$ and the cone CX.

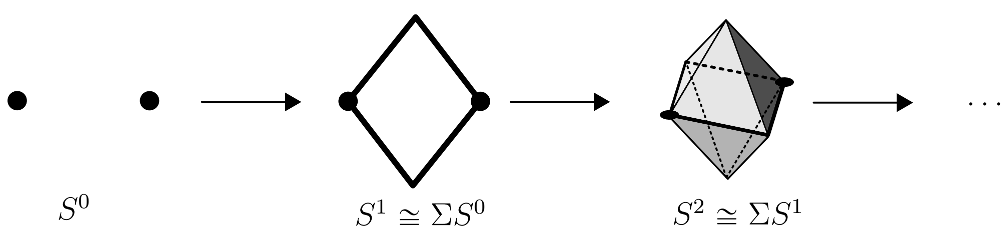

Figure 1.2.4: Suspending spheres to spheres.

Example 1.2.1. A natural easiest example to apply this to is the one point space $*$. It is easy to see that $\Sigma *$ is homeomorphic to $[0,1]$, equivalently D^1. In fact, in general, $\Sigma D^n \cong D^{n+1}$. So applying suspension just gives the sequence $*, *, *, \ldots$ up to homotopy equivalence.

Example 1.2.2. What's the next most natural example to apply suspension to? Probably the two point space S^0, surely that's not much more complicated... Stretching two points out along an interval and then glueing the top pair of points and bottom pair of points gives the circle $\Sigma S^0 \cong S^1$. And it isn't hard to see visually that $\Sigma S^1 \cong S^2$ (see Figure 1.2.4). In fact, in general we have that $\Sigma S^n \cong S^{n+1}$, so applying suspension to the two point space results in the sequence S^0, S^1, S^2, S^3,..., an interesting sequence of spaces indeed!

Exercise 1.2.7. Formally prove the above two examples.

1.2.1 CW complexes

A CW complex is a space that you can glue together from discs. It starts with a discrete set of points X^0 (which you can take as 0-dimensional discs). To construct the 1-skeleton X^1 you glue 1-discs (closed intervals) at their endpoints to points of X^0. The result is a 1-dimensional complex, sometimes known as a graph. Your complex may have 2-dimensional cells, in which case you attach 2-discs to X^1 by glueing along their boundary circles in some fashion. And so on, one continues adding n-discs for each n, defining X^n. Either this process stops at some finite n or it continues indefinitely so that the CW complex is 'infinite dimensional'.

Let's write up the above a little more precisely, and via a slightly different approach for some variety (see also the definition matching the above description a little more closely in Hatcher [Hat]). For a topological space Y, let a **characteristic map** be a continuous map $\sigma\colon D^n \to Y$ which is a homeomorphism when restricted to the **open n-disc** $\{x \in \mathbb{R}^n \mid |x| < 1\}$. An **open n-cell** of Y is a subset $e \subseteq Y$ which is homeomorphic to the open n-disc.

Then a **CW complex** $X^\bullet$ consists of the following data:

- a Hausdorff topological space X and
- a partition of X into open cells e_α for $\alpha \in \mathcal{I}$

for which:

1. for each $\alpha \in \mathcal{I}$ there is a characteristic map $\sigma_\alpha\colon D^n \to X$ mapping the open n-disc onto e_α and mapping the boundary S^{n-1} of the disc onto open cells of dimension $\leqslant n-1$;

2. the topology on X is such that a subset $C \subseteq X$ is closed if and only if the intersection $C \cap \text{im}(\sigma_\alpha)$ is closed for each $\alpha \in \mathcal{I}$ (equivalently the intersection of C and the closure of each e_α is closed).

The set $\mathcal{I}$ just indexes the open cells. The characteristic maps σ_α can be considered as the maps which attach the cells, as in the opening discussion. We define the n-**skeleton** X^n as the union of the n-cells e_α. One can show that X^n may be described as the quotient space of the disjoint union of X^{n-1} and copies of the disc D^n, one for each $\alpha \in \mathcal{I}$ corresponding to an n-cell, given by attaching the boundaries of these discs to the $(n-1)$-skeleton according to the characteristic maps.

Don't worry about condition 2 too much. If there are k-cells for dimensions k up to and including $n \in \mathbb{N}_0$, but not higher, we say that X is n-**dimensional**. In this case condition 2 turns out to be superfluous. It is there to make sure we have the correct topology in the infinite-dimensional case.

Exercise 1.2.8. Draw a CW decomposition of the 2-sphere with only two cells, and a CW decomposition with precisely one 0-cell, one 1-cell and two 2-cells.

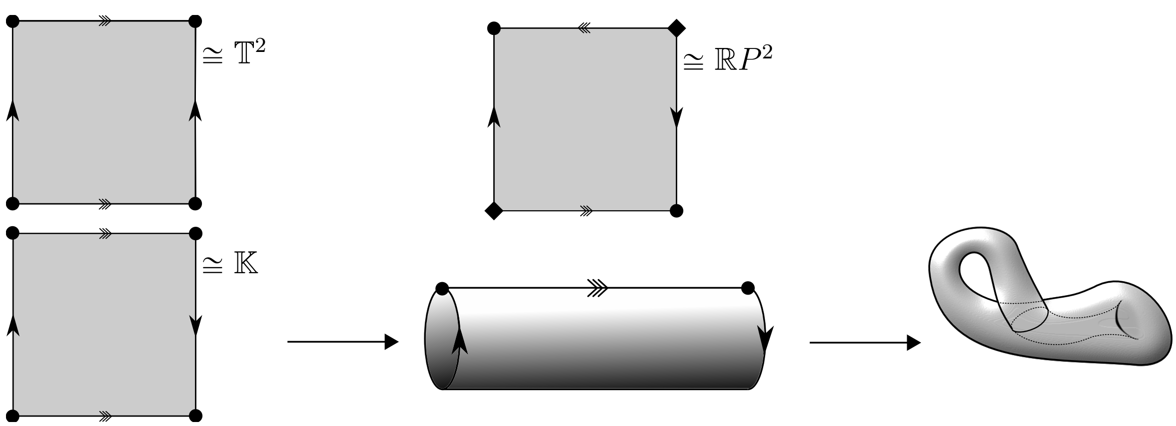

Figure 1.2.5: Torus (top left), real projective plane (top right) and Klein bottle (bottom).

Example 1.2.3. Recall the construction of the torus $\mathbb{T}^2 = S^1 \times S^1$ by glueing opposite sides of the square I^2, as in Figure 1.2.1. If at the second stage of that glueing one glues the ends of the cylinder by a reflection, one gets the famous non-orientable surface the *Klein bottle* $\mathbb{K}$, see Figure 1.2.5. So $\mathbb{K}$ is also given by glueing the edges of I^2 but flipping the direction of one of the edges. It should not be surprising that you cannot perform this glueing in $\mathbb{R}^3$—later we shall use homology to show that $\mathbb{K}$ does not embed smoothly into $\mathbb{R}^3$.

Flipping both edges, one gets $\mathbb{R}P^2$, the *real projective plane*. This is the space whose points are in bijection with lines passing through the origin in $\mathbb{R}^3$, which possesses a natural topology. A line in $\mathbb{R}^3$ is determined by a unit vector of S^2, but antipodal points determine the same line. As a result $\mathbb{R}P^2 \cong S^2/\sim$ where $\sim$ identifies antipodal points, i.e., via the antipodal map $x \mapsto -x$. More generally, one may define $\mathbb{R}P^n$ as the space of lines in $\mathbb{R}^{n+1}$, and it is homeomorphic to the quotient of S^n under the antipodal map.

Returning to $\mathbb{T}^2$, $\mathbb{K}$ and $\mathbb{R}P^2$, we see that the square models determine CW decompositions for these spaces. They each have two 1-cells and a single 2-cell; $\mathbb{T}^2$ and $\mathbb{K}$ have one 0-cell here, and $\mathbb{R}P^2$ has two.

Exercise 1.2.9. Explain how $\mathbb{R}P^n$ may be constructed from the n-disc D^n. Use this to illustrate a CW decomposition for $\mathbb{R}P^3$ (in analogy to the 2d pictures of Figure 1.2.5). Explain how we may give $\mathbb{R}P^n$ a CW decomposition with a single cell of each dimension.

Given a CW complex $X^\bullet$, a **subcomplex** $A^\bullet$ is a selection of cells from $X^\bullet$ whose union is closed in X. This turns out to be equivalent to asking that whenever we include a cell e_α of $X^\bullet$ into $A^\bullet$ then we have to also include in $A^\bullet$ all of the cells intersecting the boundary of e_α. In this case, $A^\bullet$ can be considered as a CW complex in its own right. Denote the union of cells in $A^\bullet$ by A. We call (X, A) a **CW pair**.

The following theorem should be mentioned in passing, although we won't use it later in these notes:

Theorem 1.2.1 (Cellular approximation theorem)**.** *Say that a map $f\colon X \to Y$ between CW complexes $X^\bullet$ and $Y^\bullet$ is* **cellular** *if f maps the k-skeleton of $X^\bullet$ into the k-skeleton of $Y^\bullet$ (that is, $f(X^k) \subseteq Y^k$) for each $k \in \mathbb{N}_0$.*

Consider two CW pairs (X, A) and (Y, B) and a map of pairs $f\colon (X, A) \to (Y, B)$ (which means that $f\colon X \to Y$ is continuous and $f(A) \subseteq B$); note that A and B are permitted to be empty. If the restriction of f to A is already cellular, then f is homotopic to a cellular map via a homotopy F which is stationary on A. That is, there exists $F\colon X \times I \to Y$ with $F(x, 0) = f(x)$, $F(-, 1)$ cellular and $F(a, t) = f(a)$ for all $a \in A$ and $t \in I$.

Proof. See [Hat]. Do try to think for a while, though, about why this result is completely believable. How might you approach constructing the homotopy? □

Here is a famous application:

Example 1.2.4. Let $k < n$ and consider a map $f\colon S^k \to S^n$. Choose base-points for both spheres preserved by f. There is a CW decomposition for each sphere with one 0-cell (the base point) and one cell of dimension that of the sphere, with boundary collapsed to the base point. By the cellular approximation theorem, f is homotopic to a cellular map through a homotopy which is stationary on the base point. But since $k < n$, the k-cell of S^k has nowhere to go in a cellular map other than the 0-cell of S^n, so all such maps are homotopic relative to their base-points to the constant map.

This example shows that the homotopy group $\pi_k(S^n)$ of an n-sphere is trivial for $k < n$. This is quite believable: one cannot tie a sphere around another of higher dimension in a way which cannot be unravelled (self-intersections being allowed) to the constant map.

1.2.2 Euler characteristic

Let $X^\bullet$ be a CW complex with only finitely many cells. We define its **Euler characteristic** to be

$$\chi(X) \coloneqq \sum_{n=0}^{\infty} (-1)^n c_n,$$

where c_n is the number of n-cells in $X^\bullet$. Whilst different CW decompositions for the same space X can have different numbers of cells, it turns out that choosing a different CW decomposition does not change the Euler characteristic. In fact, if X and Y are homotopy equivalent finite CW complexes then $\chi(X) = \chi(Y)$. This is one nice consequence that will follow from us proving that the homology groups of X and Y are homotopy invariants, see Section 4.6.

Example 1.2.5. There is a CW decomposition of S^n of a single 0-cell and a single n-cell (with boundary collapsed to the 0-cell). It follows that $\chi(S^n) = 1 + (-1)^n = 0$ for n odd and $\chi(S^n) = 2$ for n even.

Remark 1.2.1. That $\chi(S^n) = 0$ for n odd and $\chi(S^n) \neq 0$ for n even is related (for example, through the more general *Poincaré–Hopf Theorem*) to there being a non-vanishing vector field on S^n if and only if n is odd. For $n = 2$ this says that any vector field on S^2 has a zero. More poetically: "you cannot comb the hair of a hairy ball flat without a parting". This is the *Hairy Ball Theorem*. We won't go through the details of these things this term, but they would make for nice background reading.

1.2.3 Further topics

There are quite a few more important constructions in homotopy theory which we haven't covered but are worth mentioning. Included are: mapping cylinders mapping cones and wedge sums (which will be mentioned later) and mapping spaces. There are also *reduced* versions of these and previous constructions (e.g., the smash product in place of the product, or the reduced suspension in place of the suspension) which are useful when your spaces are pointed. The ideas of a cofibration and fibration are important if you want to explore homotopy theory more deeply.

Chapter 2

Homological algebra

2.1 Geometric motivation

The notions that we shall meet in this chapter are entirely algebraic: of chain complexes, chain maps, taking the homology of a chain complex, taking induced homomorphisms of chain maps and so on. Still, let's briefly try to motivate the core idea geometrically by how it will be applied in Chapter 3.

The approach will be to probe our space X with more fundamental geometric pieces called *simplicies* (think of points, line segments, triangles, tetrahedra and higher dimensional analogues). Doing this will define groups $C_n(X)$, whose elements are formal sums of these n-dimensional probes called *chains*. Given a chain of $C_n(X)$, there is a way of defining its *boundary*, which defines a group homomorphism $\partial_n \colon C_n(X) \to C_{n-1}(X)$. It turns out that the boundary of a boundary is always trivial, that is, $\partial_{n-1} \circ \partial_n = 0$; this is what makes a sequence of chain groups with boundary maps between them a *chain complex*. For example, the boundary of a chain based upon the 3-simplex is a certain oriented sum of its 4 bounding triangles. Applying the boundary map again, though, one gets zero, the boundary line segments of these bounding triangles all cancelling each other out, see Figure 2.1.1.

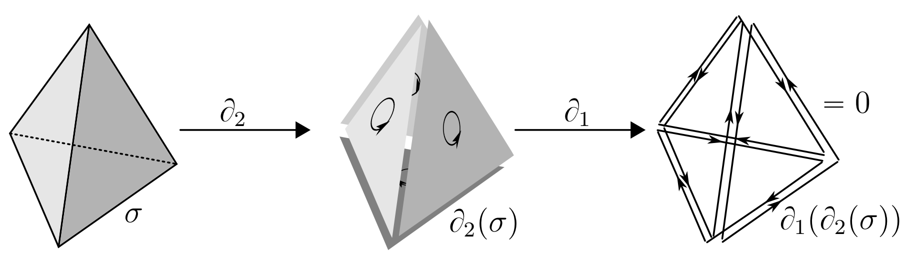

Figure 2.1.1: A 2-chain σ, its boundary $\partial\sigma$ and the boundary of its boundary $\partial^2\sigma = 0$.

The most geometrically interesting chains, those which one can use to "detect holes" in a space, have zero boundary i.e., they are elements $\sigma \in C_n(X)$ with $\partial_n(\sigma) = 0$. Such chains are called *cycles*. For example, in degree one you may think of cycles as sums of loops; exposed ends of line segments would represent non-trivial boundary. Since $\partial_{n-1} \circ \partial_n = 0$, any chain of the form $\partial_n(\sigma)$ is a cycle but it is not always the case that every cycle is the boundary of another chain.

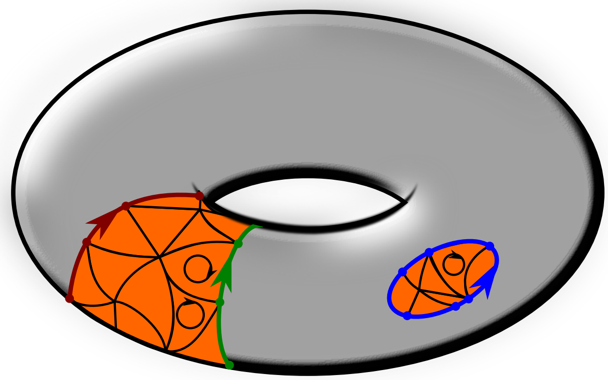

Figure 2.1.2: The green 1*-cycle plus the boundary of the left-hand part of the orange* 2*-chain gives the red* 1*-cycle, so these two elements are identified in homology. The blue* 1*-cycle plus the boundary of the right-hand part of the orange* 2*-chain results in zero, so this blue chain represents the trivial element in homology.*

We shall not care to distinguish many of the cycles of $\ker(\partial)$. For example, we would like to consider a meridian circle around a torus as unchanged if we slide it around, and we would like to think of a circle which bounds a disc as being shrinkable and hence trivial, see Figure 2.1.2. The clever idea is that we can phrase this relation in terms of higher dimensional chains: we shall identify any two n-cycles whenever their difference is a *boundary*, a chain of the form $\partial_{n+1}(\tau)$ for $\tau \in C_{n+1}(X)$. This is equivalent to saying that we quotient by $\mathrm{im}(\partial_{n+1})$. So our homology groups will turn out to be the group of cycles modulo boundaries: $H_n(X) := \ker(\partial_n)/\mathrm{im}(\partial_{n+1})$.

How one associates a chain complex to a space will be explored in Chapter 3. There is always a way of doing so (the so-called *singular chain complex*), but there are more computationally viable ways when your space has a simplicial or CW decomposition. Schematically, then, we have the following process:

$$\mathbf{Top} \xrightarrow[\text{a chain complex}]{\text{Assign space}} \mathbf{Ch} \xrightarrow[\text{of chain complex}]{\text{Take degree } n \text{ homology}} \mathbf{Ab}$$

In the case of singular homology, the first step above is easily made a functor. The second step is also functorial and takes place within the world of *homological algebra*. This is a completely algebraic setting and has application to other areas, not just topology. This section will be devoted to laying out the basics of homological algebra, and thus we will depart from spaces for a short while.

2.2 Chain complexes

2.2.1 Chain complexes and homology

Definition 2.2.1. A **chain complex** is a sequence of Abelian groups C_n (called the **chain groups**) and homomorphisms $\partial_n\colon C_n \to C_{n-1}$ (called the **boundary maps**), one for each $n \in \mathbb{Z}$. It is required that the composition of two consecutive boundary maps is trivial, i.e., $\partial_{n-1} \circ \partial_n = 0$ for all $n \in \mathbb{Z}$.

Diagrammatically, we write such a chain complex as

$$\cdots \xrightarrow{\partial_3} C_2 \xrightarrow{\partial_2} C_1 \xrightarrow{\partial_1} C_0 \xrightarrow{\partial_0} C_{-1} \xrightarrow{\partial_{-1}} C_{-2} \xrightarrow{\partial_{-2}} \cdots .$$

Elements $\sigma \in C_n$ are referred to as n**-chains**. We shall often refer to a chain complex symbolically by something like C_*, or (C_*, ∂_*) when we wish to emphasise the naming convention of the boundary map (which will be useful when we have more than one chain complex on the scene). It often reduces clutter to remove parentheses and subscripts from applications of the boundary map, so we will occasionally write, say, $\partial_n(\sigma)$ as $\partial\sigma$. The integer n of C_n is called the **degree**.

Note that it is possible to consider chain complexes in some other categories. For example, one often asks that the chain groups are actually modules and the boundary maps are homomorphisms of modules, or one has rings with ring homomorphisms as boundary maps, or vector spaces with linear boundary maps. Taking Abelian groups and homomorphisms as boundary maps will be general enough for our purposes here.

Example 2.2.1. Consider the following diagram

$$\cdots \to 0 \to 0 \to \mathbb{Z} \xrightarrow[\partial_2]{\binom{2}{2}} \mathbb{Z}^2 \xrightarrow[\partial_1]{\begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix}} \mathbb{Z}^2 \to 0 \to 0 \to \cdots$$

Here the chain groups C_n are assumed to be the trivial group for $n < 0$ or $n > 2$. The matrices induce homomorphisms ∂_1 and ∂_2 in the expected way. The composition $\partial_1 \circ \partial_2$ is easily seen to be the zero map from $\mathbb{Z}$ to $\mathbb{Z}^2$ and every other composition of two successive boundary maps involves a trivial group, so must also be a zero map. So this is a chain complex.

Example 2.2.2. It is easily checked that the following is a chain complex:

$$\cdots \to 0 \to 0 \to \mathbb{Z} \xrightarrow[\partial_3]{\times 2} \mathbb{Z} \xrightarrow[\partial_2]{\times 2} \mathbb{Z}/4 \xrightarrow[\partial_1]{\times 4} \mathbb{Z}/8 \to 0 \to 0 \to \cdots .$$

Here, the map $\times n$ sends the standard generator of the domain to n times the standard generator of the codomain.

Definition 2.2.2. Given a chain complex C_*, an element $\sigma \in \ker(\partial_n)$ is called an ***n*-cycle** and an element $\sigma \in \operatorname{im}(\partial_{n+1})$ is called an ***n*-boundary**. The **homology groups** of C_* are the quotient groups $H_n := \ker(\partial_n)/\operatorname{im}(\partial_{n+1})$ for $n \in \mathbb{Z}$.

Two n-chains $\sigma_1, \sigma_2 \in C_n$ are called **homologous** if there exists some $(n+1)$-chain $\tau \in C_{n+1}$ with $\sigma_1 - \sigma_2 = \partial_{n+1}(\tau)$. So the degree n homology group H_n is the group of n-cycles where homologous cycles are identified.

Remark 2.2.1. Again, some remarks on notation. The above is not completely notationally sound, since H_n does not refer to the chain complex whose homology has been taken. When this causes an issue we will write something like $H_n(C)$, to mean the degree n homology group of the chain complex C_*. Arguably the chain complex should then really be named C instead of C_* (as H_n is something we apply to our chain complex, like a function), but the lower $*$ is a helpful reminder that it is a chain complex, which will be particularly useful when you learn about cochain complexes in the next term (which are often denoted by something like C^*).

An element of $H_n(C)$ will sometimes be referred to as a **homology class**, and denoted, say, $[\sigma]$ where $\sigma \in C_n$ is a representative of the class. Remember that the homology class can be represented by different cycles, in particular $[\sigma]$ is represented precisely by chains of the form $\sigma + \partial_{n+1}(\tau)$ for $\tau \in C_{n+1}$, all of which are identified when we pass to homology.

Homology is well defined: the equation $\partial_{n-1} \circ \partial_n = 0$ is equivalent to saying that $\operatorname{im}(\partial_n) \subseteq \ker(\partial_{n-1})$, so the boundaries are a subgroup of the cycles. They are a normal subgroup, since everything here is Abelian, so the quotient H_n is a well defined Abelian group.

Example 2.2.3. Let's compute the homology of the simple chain complex from Example 2.2.1. In degrees $n \neq 0, 1, 2$ we have that C_n is the trivial group so it must be that $H_n \cong 0$ too. To compute H_0 note that everything in $C_0 \cong \mathbb{Z}^2$ is in the kernel of ∂_0 since it has codomain the trivial group. So $H_0 \cong \mathbb{Z}^2/\operatorname{im}(\partial_1)$. The image of ∂_1 is the subgroup

$$\operatorname{im}(\partial_1) = \{(n, -n) \mid n \in \mathbb{Z}\}.$$

It is easy to see then that the quotient group H_0 is isomorphic to $\mathbb{Z}$.

For H_1 we have that

$$\ker(\partial_1) = \{(n, n) \mid n \in \mathbb{Z}\},$$

a subgroup of $C_1 = \mathbb{Z}^2$ isomorphic to $\mathbb{Z}$; a generator can be taken as $(1, 1) \in \mathbb{Z}^2$. On the other hand

$$\operatorname{im}(\partial_2) = \{(2n, 2n) \mid n \in \mathbb{Z}\}.$$

This is the index 2 subgroup of $\ker(\partial_1)$ generated by $(2, 2)$ so $H_1 \cong \mathbb{Z}/2$. Finally, we have that $\operatorname{im}(\partial_3)$ is trivial, since C_3 is trivial, so $H_2 \cong \ker(\partial_2) \cong 0$.

Example 2.2.4. Check that for Example 2.2.2 we get homology $H_0 \cong \mathbb{Z}/4$ and $H_n \cong 0$ otherwise.

2.2.2 Exact sequences

Definition 2.2.3. A chain complex C_* is called a **long exact sequence (LES)**, or simply **exact**, if $\operatorname{im}(\partial_n) = \ker(\partial_{n-1})$ for all $n \in \mathbb{Z}$.

Note that a sequence is exact if and only if $H_n \cong 0$ for all $n \in \mathbb{Z}$. In this sense, you may think of the homology of a chain complex as measuring how far it is from being exact.

Many of our chain complexes will be trivial in negative degrees, or sufficiently large degrees. It is common to simply omit these trivial terms, for example a chain complex with all negative degrees trivial is usually written as

$$\cdots \xrightarrow{\partial_3} C_2 \xrightarrow{\partial_2} C_1 \xrightarrow{\partial_1} C_0 \to 0.$$

We say that smaller diagrams are exact if they are exact in each spot where this makes sense. For example,

$$A \xrightarrow{f} B \xrightarrow{g} C \xrightarrow{h} D$$

is exact precisely when $\operatorname{im}(f) = \ker(g)$ and $\operatorname{im}(g) = \ker(h)$.

Example 2.2.5. The diagram

$$0 \to A \xrightarrow{f} B$$

is exact precisely when $\ker f = \{0\}$, that is, when f is injective.

Example 2.2.6. The diagram

$$B \xrightarrow{g} C \to 0$$

is exact precisely when $\operatorname{im} g = C$, that is, when g is surjective.

An important example of an exact sequence is one of the form

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

called a **short exact sequence (SES)**. As above we see that f is injective and g is surjective. It is helpful to think of B as built from A and C: we have that $\operatorname{im}(f) \cong A$ and $C \cong B/\operatorname{im}(f)$, by the first isomorphism theorem of group theory. However, it is not necessarily true that $A \oplus C \cong B$. When A and C are specified, finding B is known as solving an *extension problem.* There are typically different possible groups B completing the SES given fixed A and C if nothing about the homomorphisms is specified:

Exercise 2.2.1. For a short exact sequence

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

show that there is a bijection between $A \times C$ and B *as sets*(!!), so their cardinalities agree. However, this bijection cannot always be chosen to be a homomorphism: B need not be isomorphic to $A \oplus C$ *as a group*(!!). Find such an example of a short exact sequence, with B not isomorphic to $A \oplus C$. Try to find an example where B has infinitely many elements (you may find Section 2.4.2 relevant in what you can hope to work) and one where B has finitely many elements (in which case the initial problem here will be relevant).

Whilst we don't always have that $A \oplus C \cong B$ for a SES, given just the fixed groups A and C there is of course always the short exact sequence

$$0 \to A \to A \oplus C \to C \to 0,$$

the first map being the inclusion $a \mapsto (a, 0)$ and the second the projection $(a, c) \mapsto c$. This is known as the **trivial extension**.

Example 2.2.7. Let B be an Abelian group and $A \leqslant B$ a subgroup. Then we get a SES

$$0 \to A \hookrightarrow B \to B/A \to 0$$

where the first map is the inclusion and the second is the quotient. By earlier comments, up to isomorphism one can essentially view any SES as one of the above form.

Another source of SESs is from group homomorphisms. For $f \colon B \to C$ we get an associated SES

$$0 \to \ker(f) \hookrightarrow B \to \mathrm{im}(f) \to 0$$

by the first isomorphism theorem of group theory. Again, up to isomorphism, you can essentially think of any SES as coming from this kind of situation.

2.2.3 Maps between chain complexes

Definition 2.2.4. A **chain map** between chain complexes (A_*, ∂_*^A) and (B_*, ∂_*^B) is determined by a commutative diagram of the form

$$\begin{array}{ccccccccccccc}
\cdots & \xrightarrow{\partial_3^A} & A_2 & \xrightarrow{\partial_2^A} & A_1 & \xrightarrow{\partial_1^A} & A_0 & \xrightarrow{\partial_0^A} & A_{-1} & \xrightarrow{\partial_{-1}^A} & A_{-2} & \xrightarrow{\partial_{-2}^A} & \cdots \\
 & & \downarrow f_2 & & \downarrow f_1 & & \downarrow f_0 & & \downarrow f_{-1} & & \downarrow f_{-2} & & \\
\cdots & \xrightarrow{\partial_3^B} & B_2 & \xrightarrow{\partial_2^B} & B_1 & \xrightarrow{\partial_1^B} & B_0 & \xrightarrow{\partial_0^B} & B_{-1} & \xrightarrow{\partial_{-1}^B} & B_{-2} & \xrightarrow{\partial_{-2}^B} & \cdots .
\end{array}$$

in the category **Ab** of Abelian groups. That is, a chain map is given by a sequence $(f_n)_{n \in \mathbb{Z}}$ of group homomorphisms $f_n \colon A_n \to B_n$ for which $f_{n-1} \circ \partial_n^A = \partial_n^B \circ f_n$.

We shall usually use a lower sharp to denote a chain map $(f_n)_{n\in\mathbb{Z}} = f_\sharp$. Given two chain maps $f_\sharp\colon A_* \to B_*$ and $g_\sharp\colon B_* \to C_*$, there is an obvious way to define their composition, namely we let $g_\sharp \circ f_\sharp \coloneqq (g_n \circ f_n)_{n\in\mathbb{Z}}$ so in degree n the chain map is the composition $g_n \circ f_n$.

Exercise 2.2.2. Check that chain complexes with chain maps as morphisms form a category.

Definition 2.2.5. For a chain map $f_\sharp\colon A_* \to B_*$ the **induced map** $f_*\colon H_n(A) \to H_n(B)$ is defined by setting $f_*([\sigma]) \coloneqq [f_n(\sigma)]$, where on both sides of this equation the square brackets refer to the homology class (on the left in A_* and on the right in B_*).

Exercise 2.2.3. Show that the induced map does not depend on the representatives of the homology classes we choose and so is well defined. That is, for $[\sigma] = [\sigma']$ in $H_n(A)$ show that $[f_n(\sigma)] = [f_n(\sigma')]$ in $H_n(B)$. Thus verify that the induced map is a well defined *homomorphism* of Abelian groups in each degree. This is a simple exercise but worth doing early to understand properly the definition of the induced map and why it works.

Lemma 2.2.1. *Taking the homology of a chain complex is functorial:*

- *the induced map* id_* *of the identity chain map* $\mathrm{id}_\sharp\colon C_* \to C_*$ *is the identity map* $\mathrm{id}\colon H_* \to H_*$*;*
- *given two chain maps* $f_\sharp\colon A_* \to B_*$ *and* $g_\sharp\colon B_* \to C_*$ *we have that* $(g_\sharp \circ f_\sharp)_* = g_* \circ f_*$*.*

Proof. By the definition of the induced map, we have that $\mathrm{id}_*([\sigma]) = [\mathrm{id}(\sigma)] = [\sigma]$ for a chain σ, so id_* is the identity homomorphism on H_* in each degree. Suppose then that $f_\sharp$ and $g_\sharp$ are chain maps as in the statement of the lemma. We have that

$$(g_\sharp \circ f_\sharp)_*([\sigma]) = [(g_\sharp \circ f_\sharp)_n(\sigma)] = [g_n(f_n(\sigma))] = g_*([f_n(\sigma)]) = g_*(f_*([\sigma])),$$

so $(g_\sharp \circ f_\sharp)_* = g_* \circ f_*$, as desired. □

2.2.4 Sub-, kernel, image and quotient chain complexes

Definition 2.2.6. Let (B_*, ∂_*) be a chain complex. A **sub-chain complex** $A_* \leqslant B_*$ is given by a sequence of subgroups $A_n \leqslant B_n$, for $n \in \mathbb{Z}$. We require that if $\sigma \in A_n$ then $\partial_n(\sigma) \in A_{n-1}$, so that (A_*, ∂_*) is itself a chain complex.

Just as a homomorphism $f\colon A \to B$ between groups has a kernel $\ker(f) \leqslant A$ and image $\mathrm{im}(f) \leqslant B$, so too a chain map $f_\sharp\colon A_* \to B_*$ has a kernel sub-chain complex

$\ker(f_\sharp) \leqslant A_*$ and image sub-chain complex $\operatorname{im}(f_\sharp) \leqslant B_*$. The definitions are obvious: $\ker(f_\sharp)$ has as degree n chain group $\ker(f_n) \leqslant A_n$ and $\operatorname{im}(f_\sharp)$ has as degree n chain group $\operatorname{im}(f_n) \leqslant B_n$.

Exercise 2.2.4. Check that the kernel and image of a chain map are indeed sub-chain complexes.

For an Abelian group B with subgroup $A \leqslant B$, we can form the quotient group B/A. Similarly, for chain complexes $A_* \leqslant B_*$ we can form the quotient complex B_*/A_*. It has as degree n chain group B_n/A_n, and degree n boundary map $\partial_n([\sigma]) \coloneqq [\partial_n(\sigma)]$ for $\sigma \in B_n$.

Exercise 2.2.5. Check that the quotient chain complex is a well-defined chain complex.

2.2.5 Chain homotopies

Definition 2.2.7. Let $f_\sharp$ and $g_\sharp$ be two chain maps from A_* to B_*. A **chain homotopy** between them is a sequence of group homomorphisms $h_n \colon A_n \to B_{n+1}$ satisfying:

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

In this case we call $f_\sharp$ and $g_\sharp$ **chain homotopic** and write $f_\sharp \simeq g_\sharp$.

We call A_* and B_* **chain homotopy equivalent** if there exist chain maps $f_\sharp \colon A_* \to B_*$ and $g_\sharp \colon B_* \to A_*$ for which $g_\sharp \circ f_\sharp \simeq \mathrm{id}_{A_*}$ and $f_\sharp \circ g_\sharp \simeq \mathrm{id}_{B_*}$.

Exercise 2.2.6. Show chain homotopy equivalence is an equivalence relation on chain maps, and that everything respects compositions of chain maps (so prove the counterparts of Exercise 1.2.2 and Lemma 1.2.1 from the geometric setting).

Remember that the *geometric* notion of homotopy relates continuous maps between topological spaces. The *algebraic* notion of chain homotopy relates chains maps between chain complexes. You may simply take the above as a definition for now, although to get some intuition it will help to prove the following lemma. For more geometric intuition, also see the proceeding exercises.

Lemma 2.2.2. *If $f_\sharp$ and $g_\sharp$ are chain homotopic then their induced maps on homology agree, that is, $f_* = g_*$.*

Proof. Another simple but instructive **Exercise.** □

Corollary 2.2.1. *If A_* and B_* are chain homotopy equivalent then they have isomorphic homology, that is, $H_n(A) \cong H_n(B)$ for all $n \in \mathbb{Z}$.*

Proof. Suppose that A_* and B_* are chain homotopy equivalent, so there exist chain maps $f_\sharp$ and $g_\sharp$ as described in Definition 2.2.7 for which $g_\sharp \circ f_\sharp \simeq \mathrm{id}_{A_*}$ and $f_\sharp \circ g_\sharp \simeq \mathrm{id}_{B_*}$. Applying induced maps and the above lemma, we see that $(g_\sharp \circ f_\sharp)_* = (\mathrm{id}_{A_*})_*$ and $(f_\sharp \circ g_\sharp)_* = (\mathrm{id}_{B_*})_*$. As taking homology and induced maps is a functor (see Lemma 2.2.1) it follows that f_* and g_* are isomorphisms: $f_* \circ g_* = \mathrm{id}_{H_*(A)}$ and $g_* \circ f_* = \mathrm{id}_{H_*(B)}$ so $H_n(A) \cong H_n(B)$ for all $n \in \mathbb{Z}$. □

Exercise 2.2.7. Let A_* be a chain complex. We have the *cylinder* chain complex A^c_*:

$$\cdots \xrightarrow{\partial^c_3} A_1 \oplus (A_2 \oplus A_2) \xrightarrow{\partial^c_2} A_0 \oplus (A_1 \oplus A_1) \xrightarrow{\partial^c_1} A_{-1} \oplus (A_0 \oplus A_0) \xrightarrow{\partial^c_0} \cdots$$

with degree n term $A_{n-1} \oplus (A_n \oplus A_n)$ and boundary maps

$$\partial^c_n(\sigma, (\tau_1, \tau_2)) := (\partial\sigma, ((-1)^{n+1}\sigma + \partial\tau_1, (-1)^n\sigma + \partial\tau_2)).$$

Show that this defines a chain complex.

Remark 2.2.2. On a first reading it's perhaps best to skip this remark and, again, to take the definition of a chain homotopy simply as a formal definition. There is some geometric intuition to the definition of the cylinder of a chain complex above. We want to think of it as the counterpart to the 'cylinder' $X \times I$ of a space X, where $I = [0, 1]$ is the interval. An element of the degree n chain group is given by an $(n-1)$-chain $\sigma \in A_{n-1}$, which we think of as n-dimensional after taking the product across the 1-dimensional interval I (write it as $\sigma \otimes e$), along with a specification of two n-chains $\tau_1, \tau_2 \in A_n$, chains which we think of as n-dimensional objects coming from the product with the bottom vertex of the interval (write $\tau_1 \otimes v_1$) and the top vertex (write $\tau_2 \otimes v_2$).

The boundary map then has a geometric picture. The various pieces which define the boundary map are

$$\begin{aligned}
\partial^c_n(\sigma \otimes e) &:= \partial\sigma \otimes e + (-1)^{n+1}\sigma \otimes v_1 + (-1)^n \sigma \otimes v_2;\\
\partial^c_n(\tau_1 \otimes v_1) &:= \partial\tau_1 \otimes v_1;\\
\partial^c_n(\tau_2 \otimes v_2) &:= \partial\tau_2 \otimes v_2.
\end{aligned}$$

The top equation says what is happening along the sides of the cylinder, the last two on the top and bottom[1].

[1] We may add algebraic detail here. There is a chain complex I_* of the form

$$\cdots \to 0 \to \mathbb{Z} \xrightarrow{\partial_1} \mathbb{Z} \oplus \mathbb{Z} \to 0 \to \cdots$$

where $\partial_1(x) = (x, -x)$; this chain complex may be thought of as corresponding to the interval (c.f., Example 3.2.4). There is a way of taking tensor products of chain complexes so that $A^c_* = A_* \otimes I_*$. This corresponds to our intuition that the cylinder should correspond to a product with the interval. However, we shall not cover the details of that here.

Exercise 2.2.8. In the spirit of the discussion above, show that a chain homotopy may be identified with a chain map $h_\sharp\colon A_*^c \to B_*$, where A_*^c is the cylinder chain complex above. Note the comparison to a geometric homotopy!

Exercise 2.2.9. For any map $F\colon X \times I \to Y$ we have a homotopy between uniquely specified maps, $F(-,0)$ to $F(-,1)$. In contrast, let $h_n\colon A_n \to B_{n+1}$ be any homomorphism for each $n \in \mathbb{Z}$, for chain complexes A_* and B_*. Show that (h_n) defines a chain homotopy from *any* given chain map $f_\sharp\colon A_* \to B_*$ to some other chain map $g_\sharp$, which is uniquely determined by $f_\sharp$ and (h_n).

Exercise 2.2.10. Show that the homology groups H_* of a chain complex C_* are trivial if there exists a chain homotopy between the zero map and identity map on C_*.

The other direction isn't quite true. But show that if C_* is such that each C_n is free Abelian (see Section 2.4) and $C_k \cong 0$ if $k < 0$, then $H_* \cong 0$ implies that there is a chain homotopy between the zero map and the identity on C_* (note: having all trivial homology groups is known as being *acyclic*, although we won't use that terminology elsewhere). *Hint: start by defining the chain homotopy in degree zero.*

2.3 The Snake Lemma

Definition 2.3.1. A diagram

$$0 \to A_* \xrightarrow{f_\sharp} B_* \xrightarrow{g_\sharp} C_* \to 0$$

of chain complexes and chain maps is called **exact** or a SES of chain complexes if the diagrams

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

are exact in each degree i.e., f_n is injective, g_n is surjective and $\mathrm{im}(f_n) = \ker(g_n)$ for all $n \in \mathbb{Z}$.

Example 2.3.1. Recall the notion of a sub-chain complex $A_* \leqslant B_*$ from Section 2.2.4 and the associated quotient complex B_*/A_*. Naturally associated to this there is a corresponding short exact sequence:

$$0 \to A_* \hookrightarrow B_* \to B_*/A_* \to 0. \tag{2.3.1}$$

The map from A_* to B_* is simply the inclusion chain map, and the map from B_* to B_*/A_* is the quotient map.

In the above example we see that there is a diagrammatic means of expressing the relationship between a chain complex, a sub-chain complex of it and the corresponding quotient complex. Sometimes one is interested in the homology of one of these, and in many cases it is possible to compute it (or at least say something about it) from the homology of the other two pieces. Unfortunately, simply applying homology to the SESs in each degree

$$0 \to H_n(A) \xrightarrow{f_*} H_n(B) \xrightarrow{g_*} H_n(C) \to 0 \quad ✗$$

is not typically a short exact sequence! The correct solution is not far from this though. From a short exact sequence of chain complexes one can naturally define a *long* exact sequence of the corresponding homology groups. So rather than the diagram of homologies above, which need not be a SES, the correct formulation requires one to consider all of the homology groups in each degree together in one large diagram. The way this diagram is usually drawn—with the new maps zig-zagging across it—is the reason for this construction being known as the *Snake Lemma*:

Lemma 2.3.1 (The Snake Lemma)**.** *Let*

$$0 \to A_* \xrightarrow{f_\sharp} B_* \xrightarrow{g_\sharp} C_* \to 0$$

be a SES of chain complexes. Then there is a corresponding LES of homology groups:

$$\begin{array}{ccccc} & & \cdots & \xrightarrow{g_*} & H_{n+1}(C) \\ & & \xrightarrow{\partial_*} & & \\ H_n(A) & \xrightarrow{f_*} & H_n(B) & \xrightarrow{g_*} & H_n(C) \\ & & \xrightarrow{\partial_*} & & \\ H_{n-1}(A) & \xrightarrow{f_*} & H_{n-1}(B) & \xrightarrow{g_*} & H_{n-1}(C) \\ & & \xrightarrow{\partial_*} & & \\ H_{n-2}(A) & \xrightarrow{f_*} & \cdots & & \end{array}$$

The maps $f_*\colon H_n(A) \to H_n(B)$ and $g_*\colon H_n(B) \to H_n(C)$ in each degree n are the usual induced maps in homology of the chain maps $f_\sharp$ and $g_\sharp$. The other maps

$$\partial_*\colon H_n(C) \to H_{n-1}(A)$$

are defined by $\partial_*([\gamma]) \coloneqq [\alpha]$ for an n-cycle $\gamma \in C_n$, where $\alpha \in A_{n-1}$ is such that $f_{n-1}(\alpha) = \partial_n^B(\beta)$ for $\beta \in B_n$ with $g_n(\beta) = \gamma$. Here, ∂_n^B the degree n boundary map of B_*.

Remark 2.3.1. One should check that the elements used to construct the map ∂_* above exist. That the resulting map is well defined requires more checking; all of this is verified in the proof of the Snake Lemma. It is known as the **connecting map**.

The connecting map is easier to understand in the setting of Example 2.3.1 where $f_\sharp$ is an inclusion of chain complexes $A_* \hookrightarrow B_*$ and $g_\sharp$ is the quotient $B_* \twoheadrightarrow B_*/A_*$. In this case, suppose that we have some n-cycle $[\gamma] \in B_*/A_*$. It is represented by some $\gamma \in B_*$ and being a cycle means that the boundary of $[\gamma]$ represents zero in the quotient B_*/A_*, that is, $\partial_n^B(\gamma) \in A_{n-1}$. Checking the definition of ∂_* in the Snake Lemma, we see that we may take $\partial_*([\gamma])$ as the homology class of $\partial_n^B(\gamma)$ in A_*.

There is quite a bit to check in proving the Snake Lemma. But one is essentially forced into the choices of elements to construct at each stage of the proof, leading one down a series of so-called 'diagram chases'. I *highly* recommend at this point that you attempt the proof yourself, which will be far more instructive than just reading the proof. You need to check that the elements defining the connecting map exist, that γ must be a cycle, that the connecting map does not depend on the particular choices we made for the representative of the homology class $[\gamma]$ or the chosen elements α and β and that the diagram is exact at every spot. In case you get stuck, the proof is included (in wordy detail) in Appendix A.

Example 2.3.2. Consider $A_* \leqslant B_*$, where these chain complexes have the form

$$A_* = 0 \to 0 \to \mathbb{Z}^3 \xrightarrow{\partial_1} \mathbb{Z}^2 \to 0$$

and

$$B_* = 0 \to \mathbb{Z} \xrightarrow{\partial_2} \mathbb{Z}^3 \xrightarrow{\partial_1} \mathbb{Z}^2 \to 0,$$

So the chain complexes A_* and B_* only differ in that B_* has an extra term in degree two and possibly non-trivial boundary map ∂_2 (the ∂_1 of the first complex is equal to that of second). Suppose that we have somehow determined $H_0(A) \cong \mathbb{Z}$. Can we work out the homology of B_*?

To get things going, first note that we can determine $H_*(A)$: by assumption $H_0(A) \cong \mathbb{Z}$ and $H_n(A) \cong 0$ for $n > 1$ (since the corresponding chain groups are trivial). In degree one we have that $H_1(A) \cong \ker(\partial_2)$ (no need to take a quotient, because $A_2 \cong 0$). So $H_1(A)$ fits into the following SES:

$$0 \to H_1(A) \to \mathbb{Z}^3 \xrightarrow{\partial_1} \mathbb{Z} \to 0.$$

It isn't hard to show from this that $H_1(A) \cong \mathbb{Z}^2$ (it must be a free Abelian group of rank 2, see the next section).

Consider now the quotient complex B_*/A_*. Since $A_n = B_n$ in each term except for in degree 2, clearly

$$B_*/A_* = 0 \to \mathbb{Z} \to 0 \to 0 \to 0$$

and $H_2(B_*/A_*) \cong \mathbb{Z}$, with $H_n(B_*/A_*) \cong 0$ in all other degrees.

Apply the snake lemma to the SES of chain complexes $0 \to A_* \to B_* \to B_*/A_* \to 0$:

$$\begin{array}{l} 0 \longrightarrow H_2(B) \xrightarrow{g_*} \mathbb{Z} \\ \xrightarrow{\partial_*} \mathbb{Z}^2 \xrightarrow{f_*} H_1(B) \xrightarrow{g_*} 0 \\ \xrightarrow{\partial_*} \mathbb{Z} \xrightarrow{f_*} H_0(B) \xrightarrow{g_*} 0 \end{array}$$

The final four terms

$$0 \to \mathbb{Z} \xrightarrow{f_*} H_0(B) \to 0$$

imply that f_* is an isomorphism there, so $H_0(B) \cong \mathbb{Z}$. The remaining relevant portion to consider is:

$$0 \to H_2(B) \xrightarrow{g_*} \mathbb{Z} \xrightarrow{\delta_*} \mathbb{Z}^2 \xrightarrow{f_*} H_1(B) \xrightarrow{g_*} 0.$$

Note that, since the degree three boundary map is trivial, we have that $H_2(B)$ may be identified with $\ker(\partial_2)$, where $\partial_2 \colon \mathbb{Z} \to \mathbb{Z}^3$. There are two cases: either $\ker(\partial_2) \cong \mathbb{Z}$ or $\ker(\partial_2) \cong 0$ (every subgroup of $\mathbb{Z}$ is of this form).

Take the former case. Since $\mathbb{Z}^3$ is free Abelian, the only way that $\ker(\partial_2) \cong \mathbb{Z}$ is for ∂_2 to be the zero map. Looking at the chain complex B_*, we thus see that we may essentially identify $H_n(B)$ and $H_n(A)$ for all $n \neq 2$; in particular, for $H_1(B)$ we have that $H_1(B) \cong \ker(\partial_1)/\operatorname{im}(\partial_2) \cong \ker(\partial_1) \cong H_1(A) \cong \mathbb{Z}^2$.

So suppose instead that $H_2(B) \cong \ker(\partial_2)$ is trivial. Then we get a short exact sequence from the snake diagram

$$0 \to \mathbb{Z} \xrightarrow{\partial_*} \mathbb{Z}^2 \xrightarrow{f_*} H_1(B) \to 0.$$

This means that $H_1(B) \cong \mathbb{Z}^2/\operatorname{im}(\partial_*)$. Note that ∂_* is injective, so $\operatorname{im}(\partial_*) \cong \mathbb{Z}$. However, we unfortunately do not have enough information to determine $H_1(B)$. For example, if $\operatorname{im}(\partial_*) = \mathbb{Z}\langle(1,0)\rangle$ then $H_1(B) \cong \mathbb{Z}$, but if $\operatorname{im}(\partial_*) = \mathbb{Z}\langle(n,0)\rangle$ then $H_1(B) \cong \mathbb{Z} \oplus \mathbb{Z}/n$. Here, $\mathbb{Z}\langle x\rangle \coloneqq \{n \cdot x \mid n \in \mathbb{Z}\} \leqslant \mathbb{Z}^2$. In fact, we may indeed choose the boundary maps ∂_1 and ∂_2 so that $H_2(B) \cong 0$ and $H_1(B) \cong \mathbb{Z} \oplus \mathbb{Z}/n$ for any desired $n \in \mathbb{N}_0$ (setting $\mathbb{Z}/0$ as the trivial group).

Summarising, either

$$H_k(B) \cong \begin{cases} \mathbb{Z} & \text{for } k = 0; \\ \mathbb{Z}^2 & \text{for } k = 1; \\ \mathbb{Z} & \text{for } k = 2; \\ 0 & \text{otherwise.} \end{cases}$$

or

$$H_k(B) \cong \begin{cases} \mathbb{Z} & \text{for } k = 0; \\ \mathbb{Z} \oplus \mathbb{Z}/n & \text{for } k = 1; \\ 0 & \text{otherwise.} \end{cases}$$

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

The Snake Lemma wasn't really necessary here (at some point we could have just thought directly about the chain complex B_*), but in more complicated situations it can conveniently organise what may be known on the relationships between the homology groups of the complexes.

Exercise 2.3.1. Suppose in the above example that $\partial_1(a, b, x) \coloneqq (x, -x)$. Complete the example by showing that $H_1(B) \cong \mathbb{Z} \oplus \mathbb{Z}/n$ can be realised for any $n \in \mathbb{N}$ by choosing ∂_2 appropriately.

2.4 Free Abelian groups, rank, split exact sequences and Euler characteristic

2.4.1 The rank of an Abelian group

Let A be an Abelian group. We write $0a \coloneqq 0$ (where on the right this denotes the identity of A), $na \coloneqq a + a + \cdots + a$ for $n \in \mathbb{N}$ (where the number of terms of the sum is n) and $(-n)a \coloneqq -(na) = n(-a)$ (the inverse of na, where we use additive notation because A is Abelian). So for any $n \in \mathbb{Z}$ and $a \in A$ the term na is defined and satisfies some obvious properties. We note that this is what makes an Abelian group the same thing as a $\mathbb{Z}$-module.

Take a subset $S \subseteq A$. A **$\mathbb{Z}$-linear sum** of elements of S is a finite sum $\sum n_i a_i$ where each $n_i \in \mathbb{Z}$ and $a_i \in S$. We say that S is **linearly independent** if such a sum equalling zero implies that each $n_i = 0$. If we can find a linearly independent subset S so that every element of A can be expressed (necessarily uniquely) as such a sum, then we call A **free Abelian**, and call S a **basis**. It is not hard to see that in this case $A \cong \bigoplus_S \mathbb{Z}$. The cardinality $|S|$ of S is called the **rank** of the free Abelian group A, denoted $\mathrm{rk}(A)$. Free Abelian groups are determined up to isomorphism by their rank. For example, $\mathbb{Z}^n$ is free Abelian of rank n and any free Abelian group of rank n is isomorphic to it.

Let $S \subseteq A$ be linearly independent. We call S **maximal** if we cannot add another element to S and still have a linearly independent set (in other words, for all $a \in A$ we have that na is a $\mathbb{Z}$-linear combination in S for some non-zero $n \in \mathbb{Z}$). We can extend the definition of rank to an arbitrary Abelian group: we let $\mathrm{rk}(A)$ be the cardinality of a maximal set of linearly independent elements (note that this agrees with the definition above in the free Abelian case, presuming both are well defined). For example:

- $\mathrm{rk}(A) = 0$ if and only if A is torsion, that is for all $a \in A$ there exists some non-zero $n \in \mathbb{Z}$ with $na = 0$.

- if A is a *finitely generated* Abelian group, then a fundamental theorem says that $A \cong \mathbb{Z}^n \oplus T$ for some $n \in \mathbb{N}_0$ and torsion group $T \cong \mathbb{Z}/k_1 \oplus \mathbb{Z}/k_2 \oplus \cdots \mathbb{Z}/k_\ell$, a finite direct sum of finite cyclic groups. Then $\mathrm{rk}(A) = n$.
- (**Exercise:**) show that $\mathrm{rk}(\mathbb{Q}) = 1$ and $\mathrm{rk}(\mathbb{R})$ is uncountable.

The rank behaves additively with respect to short exact sequences:

Lemma 2.4.1. *Let*

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

be a SES of Abelian groups. Then $\mathrm{rk}(A) + \mathrm{rk}(C) = \mathrm{rk}(B)$.

We leave the proof of this, and the proof that rank is well-defined, to Appendix B.

Note that since $\ker(g) = \mathrm{im}(f) \cong A$ and $\mathrm{im}(g) = C$, thinking of the rank of an Abelian group as an analogue of the dimension of a vector space, this gives a generalisation of the rank–nullity theorem of linear algebra. In fact, if you know about tensor products, a fancier approach is to tensor such a SES by $\mathbb{Q}$:

$$0 \to A \otimes_{\mathbb{Z}} \mathbb{Q} \to B \otimes_{\mathbb{Z}} \mathbb{Q} \to C \otimes_{\mathbb{Z}} \mathbb{Q} \to 0.$$

This turns the original SES into a SES of vector spaces (over $\mathbb{Q}$). The dimensions of these vector spaces are exactly the ranks of the groups, and then the above lemma just follows from the standard rank–nullity theorem. However, we shall avoid defining tensor products in these notes.

2.4.2 Split exact sequences

Any subgroup of a free Abelian group is free Abelian. Another important property of free Abelian groups is the following: if

$$0 \to A \to B \to F \to 0$$

is a SES with F free Abelian, then $B \cong A \oplus F$. This will follow from the following *splitting lemma*:

Lemma 2.4.2. *Let* $0 \to A \xrightarrow{f} B \xrightarrow{g} C \to 0$ *be a short exact sequence. The following properties are equivalent:*

1. *(Left Split): there is a map* $l\colon B \to A$ *for which* $l \circ f = \mathrm{id}_A$*;*
2. *(Right Split): there is a map* $r\colon C \to B$ *for which* $g \circ r = \mathrm{id}_C$*;*

3. *(Triviality): the short exact sequence is a trivial extension. That is, there is a commutative diagram*

$$\begin{array}{ccccccccc}
0 & \longrightarrow & A & \xrightarrow{f} & B & \xrightarrow{g} & C & \xrightarrow{f} & 0 \\
 & & \Big\| \mathrm{id}_A & & \Big\downarrow \cong & & \Big\| \mathrm{id}_C & & \\
0 & \longrightarrow & A & \xrightarrow{i} & A \oplus C & \xrightarrow{p} & C & \longrightarrow & 0
\end{array}$$

where the outside vertical maps are the identity maps[2]*, the central one is an isomorphism and the bottom row is the trivial extension (so i is the canonical inclusion and p the projection to the second factor).*

A SES satisfying one of the above equivalent properties is called **split**. Note that for a split diagram as above we have that $B \cong A \oplus C$, but the existence of *some* such isomorphism may not be enough to imply that a diagram is split (one needs the commutative diagram as given).

Proof. Try the proof of the splitting lemma as an **Exercise**. I would suggest proving $1 \Rightarrow 3$, $2 \Rightarrow 3$ and $3 \Rightarrow 1$, $3 \Rightarrow 2$. Given 3, there's a reasonably obvious choice for the maps l and r for 1 and 2, making the final two implications are a bit easier to check than the first two.

Try also showing the consequence below by constructing a right splitting: □

Corollary 2.4.1. *A SES $0 \to A \to B \to F \to 0$ for F free Abelian is split, so in particular $B \cong A \oplus F$.*

2.4.3 Euler characteristic of chain complexes

Let C_* be a chain complex whose chain groups C_n are all finite rank and trivial for all but finitely many $n \in \mathbb{Z}$. Define the **Euler characteristic** of C_* as the alternating sum

$$\chi(C_*) \coloneqq \sum_{n \in \mathbb{Z}} (-1)^n \mathrm{rk}(C_n).$$

Then we can also compute this number as the alternating sum of ranks of the homology groups:

Theorem 2.4.1. *For a chain complex C_* as above, we have that*

$$\chi(C_*) = \sum_{n \in \mathbb{Z}} (-1)^n \mathrm{rk}(H_n).$$

[2]It is actually enough for the outside vertical maps to be just isomorphisms rather than identity maps. In fact, it follows from the five lemma (Homework 2) that in either case the middle vertical map being a homomorphism implies that it must be an isomorphism already.

Proof. Consider the following two SESs:

$$0 \to \ker(\partial_n) \to C_n \to \operatorname{im}(\partial_n) \to 0$$

and

$$0 \to \operatorname{im}(\partial_{n+1}) \to \ker(\partial_n) \to H_n \to 0.$$

So $\operatorname{rk}(C_n) = \operatorname{rk}(\ker(\partial_n)) + \operatorname{rk}(\operatorname{im}(\partial_n))$ and $\operatorname{rk}(H_n) = \operatorname{rk}(\ker(\partial_n)) - \operatorname{rk}(\operatorname{im}(\partial_{n+1})$ by Lemma 2.4.1.

Now just plug everything in to the definition of χ:

$$\chi(C_*) \coloneqq \sum_{n\in\mathbb{Z}}(-1)^n \operatorname{rk}(C_n) = \sum_{n\in\mathbb{Z}}(-1)^n\big(\operatorname{rk}(\ker(\partial_n)) + \operatorname{rk}(\operatorname{im}(\partial_n))\big) =$$
$$\sum_{n\in\mathbb{Z}}(-1)^n\big(\operatorname{rk}(\ker(\partial_n)) - \operatorname{rk}(\operatorname{im}(\partial_{n+1}))\big) = \sum_{n\in\mathbb{Z}}(-1)^n \operatorname{rk}(H_n).$$

□

Chapter 3

Homology of spaces

3.1 Overview

We shall look at three approaches to assigning homology groups to spaces, in the following order:

1. simplicial homology;
2. singular homology;
3. cellular homology.

Simplicial and cellular homology apply only to spaces with particular kinds of nice decompositions into cells (they apply to *simplicial* complexes and *CW* complexes, respectively). A CW cellular decomposition is a far more general structure than a simplicial one, and where a space *can* be triangulated by a simplicial complex it is typically the case that the number of cells required is far greater than for a CW decomposition, and more cells translates into more time-consuming calculations. However, the definition of simplicial homology is simpler and will be helpful in getting us going and getting a feel for the idea of how to define and find the homology of a space.

Singular homology is rather different. Technically, simplicial and cellular homology do not apply to spaces rather than to simplicial complexes and CW complexes, respectively, whereas singular homology can be applied to any topological space, regardless of whether it is (or can be) equipped with a cellular decomposition. Unfortunately it is completely unsuitable for computational purposes, but it is useful as a bridge between other homology theories, and for showing that they are genuine homotopy invariants. One may show that the simplicial and cellular homology groups are isomorphic to the singular homology groups, and so all three approaches lead to the same answers on the spaces to which they apply and are homotopy invariants.

3.2 Simplicial homology

3.2.1 Abstract simplicial complexes

Definition 3.2.1. An **(abstract) simplicial complex** $\mathcal{K}$ is a set of non-empty finite subsets of some set V such that:

- the singleton set $\{v\} \in \mathcal{K}$ for all $v \in V$;
- if $\beta \in \mathcal{K}$ and $\varnothing \neq \alpha \subseteq \beta$ then $\alpha \in \mathcal{K}$ too.

An abstract simplicial complex is a combinatorial (often even finite) object. However, we think of it geometrically: the set V lists the vertices of the complex. For each one-element set $\{v\} \in \mathcal{K}$ we imagine a vertex. For two vertices $\{v_0\}, \{v_1\} \in \mathcal{K}$, if $\{v_0, v_1\} \in \mathcal{K}$ then we think of there being a line-segment, a 1-cell, connecting v_0 to v_1. Suppose that we have a triangular cycle of such 1-cells, $\{v_0, v_1\}$, $\{v_1, v_2\}$, $\{v_2, v_0\}$. Then it is permitted that $\{v_0, v_1, v_2\} \in \mathcal{K}$. If this is the case, then we imagine our complex as having a 2-dimensional triangle filled in, with vertices $\{v_0\}$, $\{v_1\}$ and $\{v_2\}$ bounded by the edges between them. A four-element set $\{v_0, v_1, v_2, v_3\} \in \mathcal{K}$ is thought of as a 3-dimensional tetrahedron whose four faces (i.e., the three element subsets with some v_i removed) should all also be in the complex. And so on, we continue filling in n-dimensional objects for elements of size $n + 1$ in $\mathcal{K}$.

3.2.2 Geometric simplices

The process described above, converting an abstract simplicial complex into an associated space, is the *geometric realisation*. To construct it, we firstly define our fundamental kinds of cells, the vertices, edges, triangles, tetrahedra and higher dimensional analogues, so-called *simplices*.

Definition 3.2.2. The **convex hull** of a finite set of points $\alpha = \{v_0, \ldots, v_n\} \subseteq \mathbb{R}^N$ is the subspace

$$\{\sum_{i=0}^{n} \lambda_i v_i \mid \lambda_i \in [0,1], \sum_{i=0}^{n} \lambda_i = 1\} \subseteq \mathbb{R}^N.$$

If the set $\{v_1 - v_0, v_2 - v_0, \ldots, v_n - v_0\}$ is linearly independent then we say that the $n + 1$ points of α are in **general position** and call the convex hull Δ_α an **n-simplex**. The **standard n-simplex** Δ^n is the convex hull of the standard $(n + 1)$ basis vectors of $\mathbb{R}^{n+1}$.

Note that a subset $X \subseteq \mathbb{R}^N$ is called *convex* if the straight line segment $[x, y]$ between any two points x, $y \in X$ is wholly contained in X. One may show that the convex hull

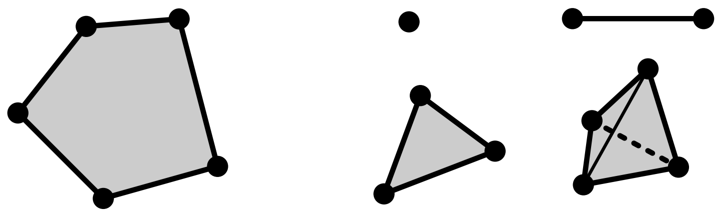

Figure 3.2.1: The convex hull of 5 points not in general position (left), and 0, 1, 2 *and* 3*-simplex (right)*

of a set of points is the smallest convex set containing those points.

An n-simplex Δ_α is homeomorphic to D^n, the closed unit disc of $\mathbb{R}^n$. We denote its **boundary** by $\partial\Delta_\alpha$, which is the union of **faces** Δ_β for $\beta \in \mathcal{K}$ with $\varnothing \neq \beta \subset \alpha$. The boundary of an n-simplex is homeomorphic to the $(n-1)$-sphere S^{n-1}.

In lots of what is to follow, it will help to take the computer science convention of indexing from zero. For example, we will list the n standard basis vectors of $\mathbb{R}^n$ as

$$\{e_0, e_1, \ldots, e_{n-1}\}.$$

When referring to, say, the element e_i of the above list, we shall call it the "index i element" so as to not confuse it with the ith element of the list (which one would probably agree is e_{i-1}).

3.2.3 Geometric realisation

The **geometric realisation** $|\mathcal{K}|$ of a simplicial complex $\mathcal{K}$ is a CW complex with the set of k-cells in bijection with elements $\alpha \in \mathcal{K}$ with $(k+1)$ elements. The idea of the construction is simple enough, so we'll leave a couple of details to the appendix; the example below should make the process clear. Let us assume for notational simplicity that the vertices of $\mathcal{K}$ are totally ordered, so for every v_0, $v_1 \in \mathcal{K}$ either $v_0 \leqslant v_1$ or $v_1 \leqslant v_0$; if both then $v_0 = v_1$ and if $v_0 \leqslant v_1$ and $v_1 \leqslant v_2$ then $v_0 \leqslant v_2$. So whenever $\{v_0, v_1, \ldots, v_n\} \in \mathcal{K}$ we can and will assume that $v_0 \leqslant v_1 \leqslant \cdots \leqslant v_n$ are listed in order.

One starts with the 0-skeleton $|\mathcal{K}|^0$, which is the discrete space with points in bijection with singletons $\{v\} \in \mathcal{K}$. For each point of this space we have the characteristic map $\sigma_{\{v\}} \colon \Delta^0 \to |\mathcal{K}|^0$, where Δ^0 is the standard 0-simplex (the one point space) and the map $\sigma_{\{v\}}$ simply determines which point corresponds to which singleton $\{v\} \in \mathcal{K}$.

We may build the k-skeleton from the $(k-1)$-skeleton from similar information. For each $\alpha = \{v_0, \ldots, v_k\} \in \mathcal{K}$ with $(k+1)$ elements we wish to attach a k-disc. We may as well attach the standard k-simplex Δ^k, which is homeomorphic. The boundary of Δ^k is a union $(k-1)$-simplex faces. Each such may be specified by removing an index

j element from $\{e_0, \ldots, e_k\}$. We attach that face (after canonically identifying it with the standard $(k-1)$-simplex) to $|\mathcal{K}|^{k-1}$ according to how we attach the $(k-1)$-simplex $\hat{\alpha}_j = \{v_0, \ldots, \hat{v}_j, \ldots, v_k\}$, where $\hat{v}_j$ indicates that this term is to be omitted from the list. That cell is attached via $\sigma_{\hat{\alpha}_j}$, defined one step down the construction.. Doing the same for each face determines a map attaching the boundary of Δ^k to the $(k-1)$-skeleton. Repeat for every other element of $\mathcal{K}$ of $(k+1)$-elements. By glueing the k-simplices to the $(k-1)$-skeleton along their boundaries in this way, this builds the k-skeleton and defines the characteristic maps attaching the cells, which determines the topological space $|\mathcal{K}|$ as a CW complex.

Example 3.2.1. Figure 3.2.2 illustrates the geometric realisation of the simplicial complex $\mathcal{K}$, given as follows:

$$\mathcal{K} = \{\{v_0\}, \{v_1\}, \{v_2\}, \{v_3\}, \{v_4\}, \{v_5\}, \{v_6\},$$
$$\{v_0, v_1\}, \{v_1, v_2\}, \{v_0, v_2\}, \{v_2, v_4\}, \{v_3, v_4\}, \{v_4, v_5\}, \{v_3, v_5\}, \{v_3, v_6\}, \{v_4, v_6\}, \{v_5, v_6\}, \{v_0, v_6\},$$
$$\{v_0, v_1, v_2\}, \{v_3, v_4, v_5\}, \{v_3, v_4, v_6\}, \{v_4, v_5, v_6\}, \{v_3, v_5, v_6\},$$
$$\{v_3, v_4, v_5, v_6\}\}.$$

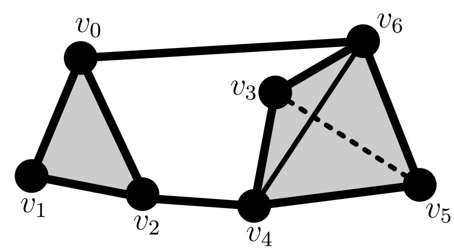

Figure 3.2.2: The geometric realisation $|\mathcal{K}|$ of the simplicial complex $\mathcal{K}$.

Definition 3.2.3. A topological space X is called **triangulable** if it permits a **triangulation**, which is a simplicial complex $\mathcal{K}$ along with a homeomorphism

$$h\colon |\mathcal{K}| \to X.$$

It is helpful to think of a triangulated space as a space which has a subdivision into distorted simplices, see Figure 3.2.3.

Rather than abstractly glueing together simplices according to the combinatorial data of $\mathcal{K}$, usually one can build $|\mathcal{K}|$ by filling in simplicies within an ambient Euclidean space $\mathbb{R}^N$ using appropriately placed vertices, which is really what Figure 3.2.2 is doing; see Appendix C for details.

Remark 3.2.1. Simplicial decompositions are more awkward than you may first expect. Figure 3.2.4 suggests two simplicial-ish decompositions of the torus. The one

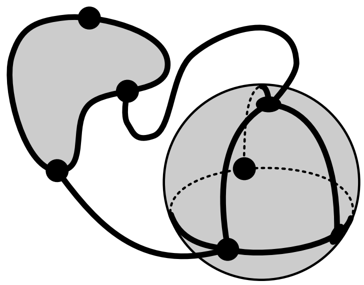

Figure 3.2.3: Space triangulated by simplicial complex $\mathcal{K}$ from above example.

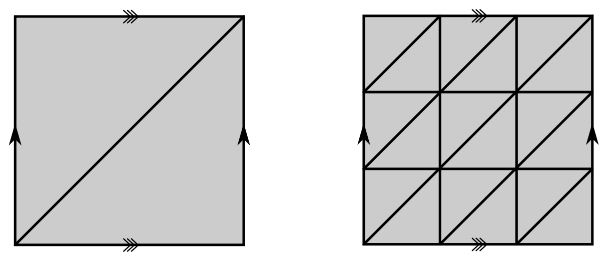

Figure 3.2.4: A non-simplicial decomposition of $\mathbb{T}^2$ (left) and a valid simplicial one (right).

on the left is *not* a simplicial decomposition as defined here. Note that in simplicial complexes every simplex $\alpha \in \mathcal{K}$ is uniquely defined by its set of vertices. The characteristic maps σ_α of the n-simplices into the geometric realisation always turn out to be homeomorphisms; in particular, no identifications are made on the boundaries of the simplices. In the left-hand decomposition there is only one vertex, so there are lots of identifications along the boundaries of cells; for example, the 1-cells give embedded circles rather than intervals.

One can make mild alterations so as to allow the flexibility offered by decompositions like that on the left. In [Hat] these are called Δ-complexes. The definition of the simplicial chain complex, as we are about to see, works essentially identically for them. However, we shall be satisfied for now with the more classical simplicial complexes as defined. The flexibility offered by Δ-complexes will be achieved later in the yet more general setting of cellular homology for CW complexes.

3.2.4 The simplicial chain complex

In this section, for a simplicial complex $\mathcal{K}$, let us refer to an element $\alpha \in \mathcal{K}$ with $n+1$ elements as an n**-simplex**. We shall assume for ease of notation, as before, that our simplicial complexes $\mathcal{K}$ come equipped with a total order on their set of vertices and that when we write down an n-simplex $\alpha = \{v_0, \ldots, v_n\}$ we list the elements $v_0, \ldots, v_n$ in order. Thus, given also some $j \in \{0, \ldots, n\}$, there is an associated $(n-1)$ simplex (considered as a *face* of α) given by

$$\hat{\alpha}_j \coloneqq \{v_0, \dots, \hat{v}_j, \dots, v_n\},$$

where the fancy hatted term $\hat{v}_j$ indicates that it is to be omitted from the list.

Chains will be given by $\mathbb{Z}$-linear sums of the simplices. Boundaries will be defined by sums over faces $\hat{\alpha}_j$ of simplices, consistently signed to get the orientations right:

Definition 3.2.4. Let $\mathcal{K}$ be a simplicial complex. Define the **degree n simplicial chain group** $C_n(\mathcal{K})$ to be the free Abelian group generated by the n-simplices of $\mathcal{K}$, so

$$C_n(\mathcal{K}) \cong \bigoplus_{n\text{-simplicies of } \mathcal{K}} \mathbb{Z}.$$

Its elements (called **n-chains**) can thus be considered as finite formal sums

$$\sigma = \sum \ell_\alpha \alpha$$

where the sum is taken over the set of n-simplicies α of $\mathcal{K}$, the coefficients $\ell_\alpha \in \mathbb{Z}$ and are zero for all but finitely many n-simplicies $\alpha \in \mathcal{K}$. The sum of such a chain with another $\sigma' = \sum \ell'_\alpha \alpha$ is given by

$$\sigma + \sigma' = \sum (\ell_\alpha + \ell'_\alpha)\alpha,$$

(i.e., the coefficient of an n-simplex α in the chain $\sigma + \sigma'$ is just the sum of the coefficients of α in σ and σ'). The chain $1 \cdot \alpha \in C_n(\mathcal{K})$, the sum of just one n-simplex $\alpha \in \mathcal{K}$, is called an **elementary chain** and may sometimes, by a slight abuse of notation, also be denoted simply by α.

For $n \geqslant 1$, the **degree n (simplicial) boundary map** $\partial_n \colon C_n(\mathcal{K}) \to C_{n-1}(\mathcal{K})$ is defined on an elementary chain α by

$$\partial_n(\alpha) \coloneqq \sum_{j=0}^{n} (-1)^j \hat{\alpha}_j.$$

Since every chain of $C_n(\mathcal{K})$ is a linear sum of elementary chains, we may thus define the boundary map ∂_n by extending linearly to all chains, i.e., by setting

$$\partial_n\left(\sum \ell_\alpha \alpha\right) \coloneqq \sum \ell_\alpha (\partial_n(\alpha)).$$

This defines the **simplicial chain complex**

$$C_*(\mathcal{K}) \coloneqq \cdots \xrightarrow{\partial_3} C_2(\mathcal{K}) \xrightarrow{\partial_2} C_1(\mathcal{K}) \xrightarrow{\partial_1} C_0(\mathcal{K}) \xrightarrow{\partial_0} 0.$$

The homology of this chain complex is denoted by $H_*(\mathcal{K})$ and is called the **simplicial homology of $\mathcal{K}$**.

Example 3.2.2. Consider again the simplicial complex $\mathcal{K}$ from Example 3.2.1. Let $\sigma \in C_2(\mathcal{K})$ be given by assigning coefficient -2 to $\{v_0, v_1, v_2\}$ (in red in Figure 3.2.5), coefficient $+1$ to $\{v_3, v_4, v_6\}$ (in green), $+1$ to $\{v_4, v_5, v_6\}$ (in blue) and zero to the other 2-simplices. Applying the definition of the boundary map we get the sum

$$\partial_2(\sigma) = \Big(-2\{v_1, v_2\} + 2\{v_0, v_2\} - 2\{v_0, v_1\}\Big) + \Big(\{v_3, v_4\} - \{v_3, v_6\}\Big) + \Big(\{v_4, v_5\} + \{v_5, v_6\}\Big).$$

Note that the coefficients of the $\{v_4, v_6\}$ terms cancel. The arrows in the figure indicate the orientations on the edges of the complex of non-zero coefficient in the boundary, pointing from v_i to v_j for $i < j$.

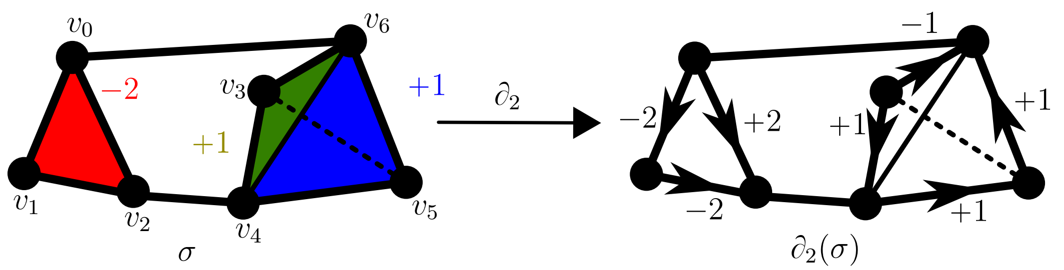

Figure 3.2.5: Illustration of $\sigma \in C_2(\mathcal{K})$ and its boundary chain.

Notice that flipping arrows on edges with negative coefficient results in a sum of cyclic loops. This is a visual representation of the fact that the boundary is a *cycle*, that is, remember, a chain with zero boundary. To make sure that the simplicial chain complex is well defined, we need to check more generally that two consecutive applications of the boundary map $\partial_{n-1} \circ \partial_n$ (∂^2 for short) results in the zero homomorphism:

Lemma 3.2.1. *We have that $\partial^2 = 0$ in the simplicial chain complex $C_*(\mathcal{K})$.*

Proof. Obviously if $n \leqslant 1$ then $\partial_{n-1} \circ \partial_n = 0$ automatically, since in that case ∂_{n-1} maps into the trivial group. Otherwise, by linearity, we just need to check that $\partial^2 = 0$ holds when applied to elementary chains α. The simplices involved in the sum defining $\partial^2(\alpha)$ are given by removing two vertices, some v_i and v_j, from α, where $i, j \in \{0, \ldots, n\}$ denote the respective indices in α. Moreover, such a simplex β occurs precisely twice in the sum. One occurrence corresponds to when v_i is removed first, by ∂_n, and then v_j is removed by ∂_{n-1}; the other case is when v_j is removed first and then v_i is removed second. Interchanging the rôles of i and j if necessary, we may as well assume that $i < j$. Then v_j is the index $(j-1)$ element of $\hat{\alpha}_i$ (since removing v_i drops v_j down one spot), and v_i is the index i element of $\hat{\alpha}_j$. So the occurrence of β in the sum with v_i removed first has coefficient $(-1)^i(-1)^{j-1}$, and the only other occurrence, with v_j removed first, has coefficient $(-1)^j(-1)^i$. These coefficients have opposite signs, so it follows that β has coefficient zero in $\partial^2(\alpha)$. The same logic applies to every other $(n-2)$-simplex of the sum defining $\partial^2(\alpha)$, so $\partial^2(\alpha) = 0$. □

At the moment, simplicial homology as we have defined it is not defined on spaces but rather on abstract simplicial complexes. So we make the following definition:

Definition 3.2.5. Let X be a topological space which is triangulated by $\mathcal{K}$. We define the **simplicial homology** of X as the simplicial homology of $\mathcal{K}$:

$$H_*(X) \coloneqq H_*(\mathcal{K}).$$

There is a glaring issue with the above definition: how can we be sure that the simplicial homology of a space does not depend upon which particular triangulation we use? Thankfully, it turns out not to:

Theorem. *The simplicial homology of a space is isomorphic to its* singular homology.

Corollary 3.2.1. *The simplicial homology of a space is well defined up to isomorphism. Moreover, homotopy equivalent triangulable spaces have isomorphic simplicial homology.*

Singular homology will be defined in the next section. We shall see why the above theorem holds in Section 4.5.2, although will skip a couple of details since a later theorem equating *cellular* and singular homology is in spirit a generalisation of it (Theorem 4.5.1).

The above corollary follows from the theorem and Corollary 3.3.1, that homotopy equivalent spaces have isomorphic singular homology.

Equipped with at least some faith that simplicial homology is a well-defined homotopy invariant for triangulable spaces (I really do promise), let's look at some examples:

Example 3.2.3. Consider the simplicial complex $\mathcal{K}$ consisting of just a single 0-simplex, which triangulates the one point space $*$. Then the simplicial chain complex is trivial outside of degree zero, and looks like:

$$\cdots \to 0 \to 0 \xrightarrow{\partial_1} \mathbb{Z} \to 0.$$

So $H_0(*) \cong \mathbb{Z}$ and $H_n(*) = 0$ for $n \neq 0$.

Example 3.2.4. Let $\mathcal{K}$ be the obvious simplicial complex with which one triangulates the interval, consisting of two 0-simplicies and a 1-simplex: $\mathcal{K} = \{\{v_0\}, \{v_1\}, \{v_0, v_1\}\}$. The degree 0 chain group is isomorphic to $\mathbb{Z}^2$ and the degree 1 chain group is isomorphic to $\mathbb{Z}$, the chain complex looks like:

$$\cdots \to 0 \to \mathbb{Z} \xrightarrow{\partial_1} \mathbb{Z}^2 \to 0.$$

The chain group $C_1(\mathcal{K})$ is generated by the elementary chain of the single 1-simplex $e \coloneqq \{v_0, v_1\}$. The boundary map on it is defined by:

$$\partial_1(e) = \sum_{j=0}^{1}(-1)^j \hat{e}_j = \hat{e}_0 - \hat{e}_1 = \{v_1\} - \{v_0\}.$$

This final term is a difference of the generators of $C_1(\mathcal{K}) \cong \mathbb{Z}^2$. So the ∂_1 boundary map can be thought of as

$$\partial_1 \colon \mathbb{Z} \to \mathbb{Z}^2, \quad \partial_1(n) = (-n, n).$$

Hence $\ker(\partial_1) = 0$ and $H_1(\mathcal{K}) \cong 0$. We have that $\operatorname{im}(\partial_1)$ is the subgroup of elements of the form $(-n, n)$ for $n \in \mathbb{Z}$. Since ∂_0 is the zero map, $\ker(\partial_0) \cong \mathbb{Z}^2$. So it is easy to see that

$$H_0(\mathcal{K}) \coloneqq \ker(\partial_0)/\operatorname{im}(\partial_1) \cong \mathbb{Z},$$

since taking the quotient of $\mathbb{Z}^2$ by the subgroup of elements of the form $(-n, n)$ is to identify precisely the elements $(m, 0) \simeq (m - n, n)$.

This complex has the same homology as the one point space. We already expected that: the interval is contractible i.e., homotopy equivalent to the one point space.

Example 3.2.5. Let $n \in \mathbb{N}$ and consider a set $V = \{v_0, \ldots, v_n\}$ of $n + 1$ elements. Construct the 1-dimensional simplicial complex $\mathcal{K}$ over V which has 1-simplicies $\{v_0, v_1\}$, $\{v_1, v_2\}$, ..., $\{v_{n-1}, v_n\}$, $\{v_n, v_0\}$ and no higher dimensional simplicies.

This simplicial complex gives a triangulation of the circle S^1 by cutting it into $n + 1$ intervals, so whatever n we pick we should get the same answer: "the" homology of the circle.

The degree zero and one chain groups are isomorphic to $\mathbb{Z}^{n+1}$ (since there are $n + 1$ 0-simplicies and 1-simplicies), and the chain complex looks like

$$\cdots \to 0 \to \mathbb{Z}^{n+1} \xrightarrow{\partial_1} \mathbb{Z}^{n+1} \to 0.$$

On an elementary 1-chain $e_k \coloneqq \{v_k, v_{k+1}\}$ (where we take $v_{n+1} \coloneqq v_0$) we have[1] that $\partial_1(e_k) = \{v_{k+1}\} - \{v_k\}$. So for a 1-chain

$$\sigma = c_0e_0 + c_1e_1 + \cdots + c_ne_n$$

we have $\sigma \in \ker(\partial_1)$ if and only if all of the coefficients $c_i \in \mathbb{Z}$ are equal. Indeed, if $c_j \neq c_{j+1}$ then the coefficient of v_{j+1} (which is a face of only e_j and e_{j+1}) will be non-zero in the boundary. It follows that

$$\ker(\partial_1) = \{(m, m, m, \ldots, m) \in \bigoplus_{\{e_0,\ldots,e_n\}} \mathbb{Z} \mid m \in \mathbb{Z}\} \cong \mathbb{Z}.$$

[1] Here we should order $\{v_n, v_0\}$ with v_n before v_0. With everything else having the obvious ordering, this doesn't come from a total order on the vertices V, as requested in our definition of the simplicial boundary, but one can make a mild adjustment here to allow for just partial orders on the vertex set which gives total orders on the simplices. Alternatively we just have a flip in sign for this one term.

Think of such a cycle as a bit like an unbased loop around the circle, which carries some weight $m \in \mathbb{Z}$.

Since away from degrees zero and one $C_n(\mathcal{K}) \cong 0$, we have that $H_n(\mathcal{K}) \cong 0$ for $n \neq 0, 1$. For degree zero it's easy enough to see that $H_0(\mathcal{K}) \cong \mathbb{Z}$; modding out by $\mathrm{im}(\partial_1)$ here is to identify the elementary chains of $\{v_k\}$ and $\{v_{k+1}\}$. Generalising this argument, you may like to try your hand at the following:

Exercise 3.2.1. Say that the simplicial complex $\mathcal{K}$ is *connected* if, for any two 0-simplices $\{v\}$, $\{w\} \in \mathcal{K}$, there exists a sequence of 0-simplices $\{v_0\}, \{v_1\}, \{v_2\}, \ldots, \{v_k\}$ with $v_0 = v$, $v_k = w$ and each $\{v_i, v_{i+1}\} \in \mathcal{K}$ a 1-simplex of $\mathcal{K}$. Show that if $\mathcal{K}$ is connected then $H_0(\mathcal{K}) \cong \mathbb{Z}$.

Try extending this by proving a statement on the relationship between $H_0(\mathcal{K})$ and the number of path-connected components of $|\mathcal{K}|$.

Example 3.2.6. Let's compute the homology of a more interesting example, the 2-sphere S^2. As a simplicial complex we may take it as $\partial\Delta^3$, the set of proper faces of a 3-simplex. So it has four 0-simplicies:

$$V = \{\{v_0\}, \{v_1\}, \{v_2\}, \{v_3\}\},$$

six 1-simplicies:

$$E = \{\{v_0, v_1\}, \{v_0, v_2\}, \{v_0, v_3\}, \{v_1, v_2\}, \{v_1, v_3\}, \{v_2, v_3\}\}$$

and four 2-simplicies

$$F = \{\{v_1, v_2, v_3\}, \{v_0, v_2, v_3\}, \{v_0, v_1, v_3\}, \{v_0, v_1, v_2\}\}.$$

So the simplicial chain complex looks like:

$$\cdots \to 0 \to \mathbb{Z}^4 \xrightarrow{\partial_2} \mathbb{Z}^6 \xrightarrow{\partial_1} \mathbb{Z}^4 \to 0.$$

Whilst perfectly manageable (and you may like to try it), it would take a bit of time to write down all of the boundary maps here, and to work out the corresponding homology groups. Let's try a more clever way.

Enlarge the above simplicial complex by also adding in the 3-simplex $\alpha := \{v_0, v_1, v_2, v_3\}$, giving the simplicial complex corresponding to Δ^3. Let us denote the simplicial chain complexes of $\partial\Delta^3$ and Δ^3 by $C_*(S^2)$ and $C_*(D^3)$, respectively. We thus have an inclusion of simplicial chain complexes $C_*(S^2) \hookrightarrow C_*(D^3)$. The corresponding quotient complex only has one generator in degree 3, corresponding to α, since all of the other cells of Δ^3 are in $\partial\Delta^3$. So the quotient complex looks like

$$\cdots \to 0 \to \mathbb{Z} \xrightarrow{\partial_3} 0 \to 0 \to 0 \to 0.$$

Clearly the homology of that is $H_3 \cong \mathbb{Z}$ and $H_k \cong 0$ for $k \neq 3$. Applying the Snake Lemma (Lemma 2.3.1) to the short exact sequence

$$0 \to C_*(S^2) \to C_*(D^3) \to C_*(D^3)/C_*(S^2) \to 0$$

we get the diagram

$$\begin{array}{ccccc}
\cdots & \longrightarrow & H_3(D^3) & \longrightarrow & \mathbb{Z} \\
 & & \xrightarrow{\partial_*} & & \\
H_2(S^2) & \longrightarrow & H_2(D^3) & \longrightarrow & 0 \\
 & & \xrightarrow{\partial_*} & & \\
H_1(S^2) & \longrightarrow & H_1(D^3) & \longrightarrow & 0 \\
 & & \xrightarrow{\partial_*} & & \\
H_0(S^2) & \longrightarrow & H_0(D^3) & \longrightarrow & 0
\end{array}$$

Since D^3 is contractible and simplicial homology should return identical answers for homotopy equivalent triangulable spaces, we have that $H_0(D^3) \cong \mathbb{Z}$ and $H_k(D^3) \cong 0$ for $k \neq 0$ (the homology groups of the one point space). Filling that information in to the above diagram, we see that the entries for $H_k(S^2)$ occur between zeros for $k \neq 0, 2$, and for $k = 0, 2$ we get the exact sequences

$$0 \to \mathbb{Z} \to H_2(S^2) \to 0$$

and

$$0 \to H_0(S^2) \to \mathbb{Z} \to 0.$$

It follows that $H_k(S^2) \cong \mathbb{Z}$ for $k = 0$, 2 and $H_k(S^2) \cong 0$ otherwise.

Exercise 3.2.2. Extend the above example by computing the simplicial homology of S^k for any $k \in \mathbb{N}_0$. As above, you may assume that the homology of a simplicial complex triangulating a contractible space is that of the one point space.

3.3 Singular homology

3.3.1 What singular homology is and is not good for

We now introduce singular homology so as to have a theory which will apply directly to *any* topological space. It is immediate from its definition that homeomorphic spaces have isomorphic singular homology, which is not immediately apparent for simplicial homology (in fact, with some extra work we will show that *homotopy equivalent* spaces have isomorphic singular homology, see Corollary 3.3.1). Moreover, it will be clear how to make singular homology *functorial* over all continuous maps, so whenever we have a continuous map $f\colon X \to Y$ we have homomorphisms $f_*\colon H_n(X) \to H_n(Y)$. Sounds great. The downside is that singular homology is essentially useless for direct computations. To actually compute, one typically uses tools such as long exact sequences arising from decomposing a space into simpler pieces that are already understood (which we will see in Chapter 4), or finds a simplicial/CW decomposition of the space of interest, applying simplicial/cellular homology instead. Rather than a direct *computational* tool, singular homology is a useful *theoretical* tool for showing that we at least have some homotopy invariant functor, which one can use as a bedrock with which to compare the other theories and prove that they are genuine homotopy invariants.

3.3.2 Restricting functions on simplices

Let $f\colon \Delta^n \to X$ be a function. Given $j \in \{0, \ldots, n\}$ we have the $(n-1)$-dimensional face $\hat{\Delta}^n_j \subset \Delta^n$, with vertices $\{e_0, \ldots, \hat{e}_j, \ldots, e_n\}$, where the term $\hat{e}_j$ is omitted. We can restrict f to this face, denoted

$$f \mid \hat{\Delta}^n_j\colon \hat{\Delta}^n_j \to X.$$

Sometimes it will be useful, though, to first canonically identify $\hat{\Delta}^n_j$ with the *standard* $(n-1)$-simplex Δ^{n-1}, so as to compare two such maps. So we denote by

$$f \restriction^n_j\colon \Delta^{n-1} \to X$$

the restriction of f to the index j face of Δ^n, but first canonically identifying this face with the standard $(n-1)$-simplex by the obvious affine map which takes the vertices of Δ^{n-1} bijectively and in order to those of $\hat{\Delta}^n_j$. Explicitly, we can read $\restriction^n_j$ as simply the map $\restriction\colon \Delta^{n-1} \to \Delta^n$ defined by

$$\restriction^n_j(t_0, \ldots, t_{n-1}) \coloneqq (t_0, \ldots, t_{j-1}, 0, t_j, \ldots, t_{n-1}),$$

where the '0' is inserted at index j. We shall sometimes drop the superscript of $\restriction^n_j$.

For $i < j$ note the special formula

$$\restriction_j \restriction_i = \restriction_i \restriction_{j-1}.$$

The composition on the left first inserts a zero at index i, then at index j. The second one inserts a zero at index $j-1$ and then at index i, which bumps up the zero inserted before by one spot to index j, since $i < j$. In particular, the index i face of the index j face is the same as the index $(j-1)$ face of the index i face.

3.3.3 Definition of singular homology and basic examples

Definition 3.3.1. A continuous map $\alpha\colon \Delta^n \to X$ is called a **singular n-simplex of X**. We define the singular chain group $C_n(X)$ as the free Abelian group of singular n-simplicies of X, so

$$C_n(X) \cong \bigoplus_{\text{singular } n\text{-simplicies of } X} \mathbb{Z}.$$

An element of $C_n(X)$ is thus a finite sum

$$\sigma = \sum \ell_\alpha \alpha$$

where the sum is taken over the set of singular n-simplices α of X, the coefficients $\ell_\alpha \in \mathbb{Z}$ and are zero for all but finitely many singular n-simplices α. The sum of such a chain with another $\sigma' = \sum \ell'_\alpha \alpha$ is given by

$$\sigma + \sigma' = \sum (\ell_\alpha + \ell'_\alpha)\alpha$$

(i.e., the coefficient of a singular n-simplex α in the chain $\sigma + \sigma'$ is just the sum of coefficients of α in σ and σ'). The chain $1 \cdot \alpha \in C_n(X)$, the sum of just one singular n-simplex, is called an **elementary chain** and may sometimes, by a slight abuse of notation, also be denoted simply by α.

For $n \geqslant 1$ the **degree n (singular) boundary map ∂_n**: $C_n(X) \to C_{n-1}(X)$ is defined on an elementary chain α by

$$\partial_n(\alpha) := \sum_{j=0}^{n} (-1)^j (\alpha\restriction_j),$$

(remember, $\restriction_j$ means: restrict $\alpha\colon \Delta^n \to X$ to the index j face of Δ^n, but first implicitly identifying this with the standard $(n-1)$-simplex Δ^{n-1}). Since every chain of $C_n(X)$ is a linear sum of elementary chains, we may thus define the boundary map ∂_n by extending linearly to all chains, i.e., by setting

$$\partial_n\Big(\sum \ell_\alpha \alpha\Big) := \sum \ell_\alpha (\partial_n(\alpha)).$$

This defines the **singular chain complex**

$$C_*(X) := \cdots \xrightarrow{\partial_3} C_2(X) \xrightarrow{\partial_2} C_1(X) \xrightarrow{\partial_1} C_0(X) \xrightarrow{\partial_0} 0.$$

The homology of this chain complex is denoted by $H_*(X)$ and is called the **singular homology of X**.

Remark 3.3.1. *Any* continuous map $\alpha\colon \Delta^n \to X$ counts as a singular n-simplex. The map could be far from injective for example (hence *singular* simplex). For most spaces the singular chain complex is *ginormous*. For example, the generators of $C_0(I)$ can be identified with the set of points of I (a map $\Delta^0 \to X$ is a map from the one point space to X, so it's just a choice of point). So $C_0(I) \cong \mathbb{Z}^{\mathfrak{c}}$, where $\mathfrak{c}$ is the uncountable cardinality of the continuum (the cardinality of $\mathbb{R}$ or I). In fact, the number of singular n-simplicies $\Delta^n \to I$ for other values of n is also $\mathfrak{c}$ (unimportant **Exercise:** why?).

As before, we should check that this really does define a chain complex:

Lemma 3.3.1. *We have that $\partial^2 = 0$ in the singular chain complex $C_*(X)$.*

Proof. The trick of the proof is essentially the same as that of Lemma 3.2.1. Like there, when considering $\partial^2(\alpha)$ for an elementary n-chain α (with $n \geqslant 2$), the sum involves singular $(n-2)$ simplicies β given by restrictions to faces by deleting two vertices. Such a singular simplex β appears twice in the sum, depending on which order the vertices are deleted, appearing with opposite signs as in the proof of Lemma 3.2.1. These are the terms (for $i < j$):

$$(\alpha\restriction_j)\restriction_i = (\alpha\restriction_i)\restriction_{j-1}.$$

This equality has already been mentioned, see the discussion at the bottom of Section 3.3.2. Thinking the other way around, you can see this as coming from the fact that the face corresponding to removing the index i term of $\{e_0, \ldots, \hat{e}_j, \ldots, e_n\}$ is the same as that of removing the index $(j-1)$ term from $\{e_0, \ldots, \hat{e}_i, \ldots, e_n\}$, since e_j moves down one spot after we've removed e_i. □

There aren't too many examples that we can compute by direct means. Consider the following though:

Exercise 3.3.1. Let $*$ be the one point space. Find the singular chain complex of $*$ and compute its homology.

Exercise 3.3.2. Show that the degree zero singular homology $H_0(X)$ is isomorphic to $\mathbb{Z}^n$, where n is the cardinality of the set of path-components of X.

Exercise 3.3.3. Show that $H_*(\coprod_\alpha X_\alpha) \cong \bigoplus_\alpha H_*(X_\alpha)$. That is, the homology of a disjoint union of spaces is the direct sum of the homologies of those spaces.

Exercise 3.3.4. This one is more tricky: try to compute by direct means the degree one singular homology of the circle as $H_1(S^1) \cong \mathbb{Z}$. It may help to think in terms of winding numbers: a map $\alpha\colon \Delta^1 \to S^1$ determines a 'starting point' in S^1, where the left endpoint of the interval is mapped to by α, and some number of radians through which the map winds around the circle as it traverses the interval to reach the 'terminal point', given by where α maps the right end point to.

Lemma 3.3.2. *Let X be a convex subspace of $\mathbb{R}^N$. Then*

$$H_k(X) \cong \begin{cases} \mathbb{Z} & \text{for } k = 0; \\ 0 & \text{otherwise.} \end{cases}$$

(In the language of later Section 4.2, the 'reduced homology' of X is trivial).

Proof.

Exercise 3.3.5. There aren't many spaces that one may compute the singular homology of directly, so I think proving the above is an instructive exercise. Let me give the main hints though, and we can discuss further details in class.

Just for notational ease, replace the singular chain complex of X with the 'reduced' version, given by 'augmenting' it in degree -1:

$$\widetilde{C}_* = \cdots \xrightarrow{\partial_3} C_2(X) \xrightarrow{\partial_2} C_1(X) \xrightarrow{\partial_1} C_0(X) \xrightarrow{\varepsilon} \mathbb{Z} \to 0.$$

The 'augmentation' map ε is the map which sends a 0-chain $\sum \ell_\alpha \alpha$ to $\sum \ell_\alpha$. This has the effect of just killing a $\mathbb{Z}$ term of $H_0(X)$; for more discussion see Section 4.2. If at present this seems confusing, you may prefer to just think about degree zero in the following discussion separately instead.

To prove the lemma you can then find a chain homotopy from the identity chain map on the above complex to the zero map (c.f., Exercise 2.2.10). That is (Definition 2.2.7), you should show that there exist homomorphisms $c_n \colon C_n(X) \to C_{n+1}(X)$ for which:

$$\sigma = c_{n-1}(\partial_n \sigma) + \partial_{n+1} c_n(\sigma). \tag{3.3.1}$$

Write in short:

$$\partial(c\sigma) = \sigma - c(\partial\sigma).$$

We have used the notation c_n instead of h_n because you can construct the map c_n via a 'cone' like construction. The final equation above represents the intuitive idea that the boundary of the cone is given by its 'base' σ and (appropriately signed) cone of its boundary $c(\partial\sigma)$, which is its 'sides'.

To give more details so you can start the proof: given a singular n-simplex $\alpha \colon \Delta^n \to X$, there is a simple way of defining an associated singular $(n+1)$-simplex $c\alpha$. The idea is to just fix any old point $x_0 \in X$ throughout and define $c\alpha$ as a singular $(n+1)$-simplex which is just α when restricted to the 'base' face of Δ^{n+1}, and is defined on the rest of the simplex by stretching out α in a linear fashion so that the tip of Δ^{n+1} (the vertex opposite to the base face) is sent to x_0. This can be done by convexity, see Figure 3.3.1.

□

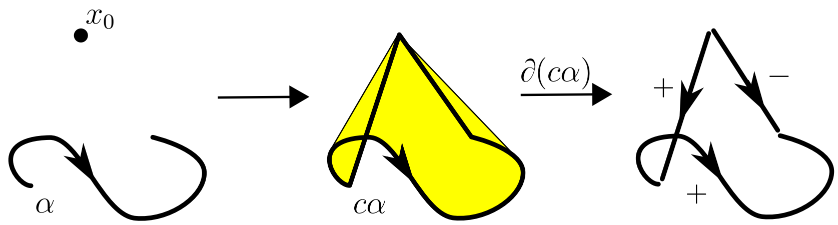

Figure 3.3.1: A 1-simplex α in some convex space (left), the 2-simplex $c\alpha$ (centre) and its boundary $\partial(c\alpha) = \alpha - c(\partial\alpha)$ (right).

3.3.4 Functorality: induced maps

Definition 3.3.2. Let $f\colon X \to Y$ be a continuous map. We define the **degree n chain map** $f_n\colon C_n(X) \to C_n(Y)$ on an elementary n-chain α by

$$f_n(\alpha) \coloneqq f \circ \alpha \colon \Delta^n \to Y.$$

This defines a homomorphism $f_n\colon C_n(X) \to C_n(Y)$ by extending linearly i.e., given an arbitrary chain $\sigma = \sum \ell_\alpha \alpha$ we define

$$f_n(\sum \ell_\alpha \alpha) \coloneqq \sum \ell_\alpha f_n(\alpha).$$

This defines a chain map $f_\sharp\colon C_*(X) \to C_*(Y)$. The **induced map** of this chain map is denoted $f_*\colon H_*(X) \to H_*(Y)$.

Lemma 3.3.3. *The above definition is well defined: $f_\sharp$ is a chain map, that is, $\partial_n \circ f_n = f_{n-1} \circ \partial_n$.*

Proof. As ever, it suffices to check the equality on elementary chains. Given an elementary n-chain α we have that

$$\partial_n(f_n(\alpha)) = \partial_n(f \circ \alpha) = \sum_{j=0}^{n} (-1)^j (f \circ \alpha)\restriction_j,$$

using the definition of ∂_n on the elementary n-chain $f \circ \alpha$. Composing the other way we get

$$f_{n-1}(\partial_n(\alpha)) = f_{n-1}(\sum_{j=0}^{n} (-1)^j (\alpha\restriction_j)) = \sum_{j=0}^{n} (-1)^j f \circ (\alpha\restriction_j),$$

as f_{n-1} is defined on the sum from its definition on the elementary chains $\alpha\restriction_j$ and extending linearly. By associativity $(f \circ \alpha)\restriction_j = f \circ (\alpha\restriction_j)$ and the result follows. □

Lemma 3.3.4. *Singular homology defines a functor $H_n(-)$ from the category* **Top** *of topological spaces and continuous maps to the category* **Ab** *of Abelian groups and group homomorphisms. That is:*

- *for a topological space X we have an Abelian group $H_n(X)$;*
- *for a continuous map $f\colon X \to Y$ we have an induced map, a homomorphism $f_*\colon H_n(X) \to H_n(Y)$;*
- *the induced map $\mathrm{id}_*\colon H_n(X) \to H_n(Y)$ of the identity map is the identity homomorphism;*
- *for two continuous maps $f\colon X \to Y$ and $g\colon Y \to Z$ we have that $(g \circ f)_* = g_* \circ f_*$.*

Proof. The first two items were included just as a reminder of what data one needs for a functor, they hold by our definitions. Consider the identity map $\mathrm{id}_X\colon X \to X$. By the definition of the corresponding chain map $\mathrm{id}_\sharp$ we have that $\mathrm{id}_\sharp(\alpha) = \mathrm{id}_X \circ \alpha = \alpha$ on an elementary n-chain α, so $\mathrm{id}_\sharp$ is the identity chain map and id_* is the identity homomorphism on $H_n(X)$ (taking homology is a functor, Lemma 2.2.1). Suppose then that f and g are two continuous maps as in the statement of the lemma. On an elementary chain $\alpha \in C_n(X)$ we have that

$$(g \circ f)_n(\alpha) = (g \circ f) \circ \alpha = g \circ (f \circ \alpha) = g_n(f \circ \alpha) = g_n(f_n(\alpha)).$$

So by extending linearly to all of $C_n(X)$ we have that the chain maps $(g \circ f)_\sharp$ and $g_\sharp \circ f_\sharp$ agree. It follows (homology is a functor, Lemma 2.2.1) that the compositions of induced maps also agree. □

3.3.5 Homotopy invariance

It turns out that we can do a lot better than what is stated in the above lemma: homology actually defines a functor from the *homotopy category* (Section 1.2) **hTop** to **Ab**:

Theorem 3.3.1. *Suppose that $f, g\colon X \to Y$ are homotopic maps. Then the induced maps f_*, $g_*\colon H_*(X) \to H_*(Y)$ are equal.*

Before proving this theorem, let's see what good this can do for us. Suppose that X and Y are homotopy equivalent spaces, so there exist $f\colon X \to Y$ and $g\colon Y \to X$ with $g \circ f \simeq \mathrm{id}_X$ and $f \circ g \simeq \mathrm{id}_Y$. Then

$$g_* \circ f_* = (g \circ f)_* = (\mathrm{id}_X)_* = \mathrm{id}_{H_*(X)}$$

and

$$f_* \circ g_* = (f \circ g)_* = (\mathrm{id}_Y)_* = \mathrm{id}_{H_*(Y)}\,.$$

The first equalities are just functorality (of composition), the second homotopy invariance, and the final ones functorality again (of identities). So $f_*\colon H_*(X) \to H_*(Y)$ is an isomorphism with inverse g_* and we have shown the following:

Corollary 3.3.1. *If X and Y are homotopy equivalent spaces then $H_*(X) \cong H_*(Y)$.*

This corollary implies, for example, that the homology $H_*(X)$ of a contractible space X is the same as that of a point, namely isomorphic to $\mathbb{Z}$ in degree zero and trivial in other degrees (see Example 3.2.3 or Exercise 3.3.1). Following Exercise 3.2.2, we may compute the simplicial homology of the n-sphere ($n > 0$):

$$H_k(S^n) \cong \begin{cases} \mathbb{Z} & \text{for } k = 0, n; \\ 0 & \text{otherwise.} \end{cases}$$

Allowing ourselves to assume that simplicial and singular homology agree (which they do, but we have not proved that yet), we have thus developed enough tools to cover all of the ingredients of the proof in the introduction of the Brouwer Fixed Point Theorem. Homology is no longer a black box.

Proof of homotopy invariance

Rather than showing that all homotopic maps induce the same maps on homology, it will save notation to restrict attention to those of the following sort:

Lemma 3.3.5. *For any space X, the maps ι_0, $\iota_1\colon X \hookrightarrow X \times I$ given by $\iota_0(x) \coloneqq (x, 0)$ and $\iota_1(x) = (x, 1)$ give chain homotopic maps $C_*(X) \to C_*(X \times I)$.*

If we can prove this lemma then the theorem follows. Indeed, suppose that f and g are homotopic, with $F\colon X \times I \to Y$ the homotopy with $F(x, 0) = f(x)$ and $F(x, 1) = g(x)$. We can write $f = F \circ \iota_0$ and $g = F \circ \iota_1$. Hence, assuming that ι_0 and ι_1 induce the same maps on homology and simply using functorality:

$$f_* = (F \circ \iota_0)_* = F_* \circ (\iota_0)_* = F_* \circ (\iota_1)_* = (F \circ \iota_1)_* = g_*.$$

Proof of lemma. For $n \in \mathbb{N}_0$, let l_n (l for 'lower') denote the singular n-simplex $\Delta^n \to \Delta^n \times I$ given by $x \mapsto (x, 0)$. Similarly, let u_n (u for 'upper') be given by $x \mapsto (x, 1)$. The proof will largely come down to the following: we shall inductively define singular $(n+1)$-chains $\tau_{n+1} \in C_{n+1}(\Delta^n \times I)$ so that

$$\partial_{n+1}(\tau_{n+1}) = u_n - l_n - s_n,$$

where s_n is the singular n-chain given by placing (with appropriate signs) the chains τ_n on the sides of the prism $\Delta^n \times I$. More precisely, let $\upharpoonleft_j\colon \Delta^{n-1} \to \Delta^n$ denote the canonical inclusion of Δ^{n-1} as the index j face of Δ^n (notation as in Section 3.3.2). Then consider $(\upharpoonleft_j, \mathrm{id})\colon \Delta^{n-1} \times I \to \Delta^n \times I$ as the inclusion of the index j side of the prism. We wish to define our chains τ_{n+1} so that

$$\partial_{n+1}(\tau_{n+1}) = u_n - l_n - \sum_{j=0}^{n} (-1)^j (\upharpoonleft_j, \mathrm{id})_\sharp(\tau_n). \tag{3.3.2}$$

Denote the summation on the right by s_n.

To start the induction, define $\tau_1\colon \Delta^1 \to \Delta^0 \times I$ to be the obvious singular 1-simplex whose boundary is $u_0 - l_0$ (it's a path starting at $(*, 0)$ and ending at $(*, 1)$). So Equation 3.3.2 is satisfied (taking by definition the sum s_0 to be empty in this bottom degree). Suppose then that τ_k have been constructed for $k \leqslant n$ satisfying the equation. The right side of Equation 3.3.2, which we want to show is the boundary of some τ_{n+1}, is a cycle. Indeed

$$\partial_n(u_n - l_n - s_n) = \partial_n(u_n) - \partial_n(l_n) - \partial_n(s_n) = \partial u_n - \partial l_n - \sum_{j=0}^{n} (\restriction_j, \mathrm{id})_\sharp(\partial \tau_n) =$$

$$(\partial u_n - \partial l_n) - \left(\sum_{j=0}^{n} (-1)^j (\restriction_j, \mathrm{id})_\sharp (u_{n-1} - l_{n-1}) \right) + \left(\sum_{j=0}^{n} (-1)^j (\restriction_j, \mathrm{id})_\sharp (s_{n-1}) \right).$$

The second equality just applies Equation 3.3.2 one step down. The first two bracketed terms above cancel as is easily checked by applying the boundary maps to u_n and l_n. On the other hand, the right bracketed term is zero, for precisely the same reason that $\partial^2 = 0$ (as in Lemma 3.3.1). That's easy to see after expanding it out using the definition of s_{n-1}:

$$\sum_{j=0}^{n} (-1)^j (\restriction_j, \mathrm{id})_\sharp \sum_{i=0}^{n-1} (-1)^i (\restriction_i, \mathrm{id})_\sharp \tau_{n-1} = \sum_{j=0}^{n} \sum_{i=0}^{n-1} (-1)^j (-1)^i (\restriction_j \restriction_i, \mathrm{id})_\sharp \tau_{n-1} = 0.$$

Since the right side of Equation 3.3.2 (which is already defined by induction) is a cycle, it must be a boundary: recall from Lemma 3.3.2 that convex regions of $\mathbb{R}^N$, such as $\Delta^n \times I$, have trivial homology in higher degrees, so if the right hand side of the equation was not a boundary then it would represent a non-zero element of $H_n(\Delta^n \times I)$. So we may construct the elements τ_n satisfying the equation by induction.

We now use the chains τ_n to construct a chain homotopy between the chain maps induced by ι_0 and ι_1. For a singular simplex $\alpha\colon \Delta^n \to X$ we simply define

$$h_n(\alpha) \coloneqq (\alpha, \mathrm{id})_\sharp(\tau_{n+1}) \tag{3.3.3}$$

and extend linearly. Checking that this defines a chain homotopy is easy:

$$\partial_{n+1}(h_n(\alpha)) = \partial_{n+1}((\alpha, \mathrm{id})_\sharp(\tau_{n+1})) = (\alpha, \mathrm{id})_\sharp \partial_n(\tau_{n+1}) = (\alpha, \mathrm{id})_\sharp(u_n - l_n - s_n).$$

The final term above is $(\iota_1)_\sharp(\alpha) - (\iota_0)_\sharp(\alpha) - (\alpha, \mathrm{id})_\sharp(s_n)$ (just write down the definitions!). On the other hand

$$h_{n-1}\partial_n(\alpha) = h_{n-1}\left(\sum_{j=0}^{n} (-1)^j \alpha \restriction_j \right) = \sum_{j=0}^{n} (-1)^j (\alpha \restriction_j, \mathrm{id})_\sharp \tau_n = \sum_{j=0}^{n} (-1)^j (\alpha, \mathrm{id})_\sharp \circ (\restriction_j, \mathrm{id})_\sharp \tau_n.$$

This final term is $(\alpha, \mathrm{id})_\sharp(s_n)$, which verifies that

$$\partial_{n+1} h_n + h_{n-1}\partial_n = (\iota_1)_\sharp - (\iota_0)_\sharp,$$

so $(\iota_0)_\sharp$ and $(\iota_1)_\sharp$ are chain homotopic and hence induce the same maps on homology by Lemma 2.2.2. □

Remark 3.3.2. The proof above looks a little different to the typical proof of homotopy invariance. What is usually done is to decompose the prism $\Delta^n \times I$ combinatorially into simplexes in a methodical way. I recommend also looking at this more standard style of proof (see for instance [Hat, Theorem 2.10]). I like the above proof because it avoids the particularities of a clever way of decomposing the prism associated to working with simplices, although this is partially hidden in the proof of Lemma 3.3.2 showing that convex subspaces have trivial homology. Equation 3.3.2 is essentially 'decomposing' the prism $\Delta^n \times I$ as a singular chain, a $\mathbb{Z}$-linear sum τ_{n+1} of singular $(n+1)$-simplices whose boundary is the appropriately signed top, bottom and sides of the prism (the sides being determined one step down the construction), so the two approaches are ultimately the same idea.

3.4 Cellular homology

Just as simplicial homology could be applied to any simplicial complex $\mathcal{K}$, and the degree n chain group was the free Abelian group with generators corresponding to the n-simplices, cellular homology is something which we can apply to any CW complex $X^\bullet$ and will have degree n chain group the free Abelian group with generators corresponding to the n-cells. This makes cellular homology an efficient tool, because one can often find nice cellular decompositions of spaces without too many cells. For example, an n-sphere can be given a CW decomposition with just two cells, of dimensions 0 and n. The 2-torus can be given a CW decomposition into a 0-cell, two 1-cells and a 2-cell; the usual simplicial decomposition (Figure 3.2.4) consists of nine 0-cells, twenty seven 1-cells and eighteen 2-cells! (one can do a little better, but not significantly).

For low dimensional pictures, the action of the boundary maps are 'what one would expect'. Rather than giving a formal definition of the boundary maps right away, we shall first look at some illustrating examples. To give the formal definition we will need to develop some more tools (mostly *relative* homology), which shall be done in the next chapter. These tools firstly allow one to define the boundary maps formally, and then their properties that we establish also serve as the main ingredients in the proof that cellular and singular homology agree.

Definition 3.4.1. Let $X^\bullet$ be a CW complex. The **degree n cellular chain group**

$C_n(X^\bullet)$ is the free Abelian group generated by the n-cells of $X^\bullet$, so

$$C_n(X^\bullet) \cong \bigoplus_{n\text{-cells of } X^\bullet} \mathbb{Z}.$$

The **cellular chain complex** of X is given by

$$\cdots \xrightarrow{\partial_3} C_2(X^\bullet) \xrightarrow{\partial_2} C_1(X^\bullet) \xrightarrow{\partial_1} C_0(X^\bullet) \to 0.$$

The boundary maps ∂_n will be informally explained later in this section, and formally defined in Definition 4.5.1. The homology of this chain complex is the **cellular homology of $X^\bullet$**, denoted $H_*(X^\bullet)$.

Analogously to simplicial homology, the above definition does not apply to spaces but to CW complexes, so we make the following definition:

Definition 3.4.2. Let X be a space with CW decomposition $X^\bullet$. We define the **cellular homology** of the space X as the cellular homology of the CW complex $X^\bullet$:

$$H_*(X) \coloneqq H_*(X^\bullet).$$

We shall see later that the cellular homology does not depend upon which particular CW decomposition we chose for X:

Theorem 4.5.1. *For a CW complex $X^\bullet$ its singular homology is isomorphic to its cellular homology.*

Remark 3.4.1. As for simplicial homology, the theorem above means that our clash of notation between singular, simplicial and cellular homology will not cause us issues.

Example 3.4.1. Let $n \geqslant 2$. The n-sphere S^n can be given a CW decomposition of a single 0-cell and a single n-cell, attached to the 0-cell by collapsing its whole boundary to the point. So the cellular chain complex looks like:

$$\cdots \xrightarrow{\partial_{n+2}} 0 \xrightarrow{\partial_{n+1}} \mathbb{Z} \xrightarrow{\partial_n} 0 \xrightarrow{\partial n-1} \cdots \xrightarrow{\partial_2} 0 \xrightarrow{\partial_1} \mathbb{Z} \to 0.$$

There's no choice in the boundary maps here, they have to be the zero homomorphisms and we can read off the homology as $H_k(S^n) \cong \mathbb{Z}$ for $k = 0$ or n, and $H_k(S^n)$ is trivial otherwise. That agrees with our earlier calculations.

Example 3.4.2. Recall from Example 1.2.3 that real projective space $\mathbb{R}P^k$ is the space of lines through the origin of $\mathbb{R}^{k+1}$. One can also define **complex projective space** as the space $\mathbb{C}P^k$ of 'complex lines' in $\mathbb{C}^{k+1}$. By a 'complex line' we mean a set of points of the form λz for some fixed non-zero $z \in \mathbb{C}^{k+1}$ and $\lambda \in \mathbb{C}$ (so a 'complex line' is topologically actually a plane). Then $\mathbb{C}P^k$ is given by identifying $w \in \mathbb{C}^{k+1} - \{0\}$ with $z \in \mathbb{C}^{k+1} - \{0\}$ if there exists some λ with $\lambda w = z$; we give $\mathbb{C}P^k$ the corresponding quotient topology. By normalising, we can also define $\mathbb{C}P^k = S^{2k+1}/\sim$, where $\sim$ identifies points $w, z \in S^{2k+1}$ (the sphere of unit vectors of $\mathbb{C}^{k+1} \cong \mathbb{R}^{2k+2}$) if there is a (necessarily unit length) $\lambda \in \mathbb{C}$ with $\lambda w = z$.

Just as how $\mathbb{R}P^k$ can be given a CW decomposition with one cell in each dimension up to k, complex projective space $\mathbb{C}P^k$ can be given a CW decomposition with one cell in each *even* dimension up to $2k$. One nice consequence[2] of this is that $\mathbb{C}P^1 = (S^3/\sim) \cong S^2$. Here's a proof that such a CW decomposition exists.

Proof. We may consider $\mathbb{C}P^0$ as the one point space (consisting of $\mathbb{C}$ as the only complex line in $\mathbb{C}$), so this satisfies the claim. Suppose for induction then that $\mathbb{C}P^{k-1}$ has such a CW decomposition. Consider points $(z, \sqrt{1-|z|^2}) \in \mathbb{C}^{k+1}$ with $|z| \leqslant 1$, where $z \in \mathbb{C}^k$ and $|z| \in \mathbb{R} \subset \mathbb{C}$ is just the Euclidean norm of z, thought of as an element of $\mathbb{R}^{2k}$. Every point of $\mathbb{C}P^{k+1} - \{0\}$ is represented by an element of that form. Indeed, given $(z_1, \ldots, z_{k+1})$, firstly multiply through by $\overline{z_{k+1}}$ (if non-zero) to make the final entry real and non-negative. Then rescale the vector by dividing through by its norm (that is, so that it belongs to S^{2k+1}). For the resulting $(z, t) \in \mathbb{C}^{2k} \times \mathbb{C}$ we have that $1 = |(z,t)|^2 = |z|^2 + t^2$ so $t = \sqrt{1-|z|^2}$. The space of points of this form, just a choice of $|z| \in \mathbb{C}^{2k}$ with $|z| \leqslant 1$, is homeomorphic to a $2k$-disc. The boundary of this disc consists of those points with final entry zero, so points $(z, 0) \in \mathbb{C}^{2k} \times \mathbb{C}$ with $z \in S^{2k-1}$. In the quotient defining $\mathbb{C}P^k$, this subspace of points corresponds to $\mathbb{C}P^{k-1}$. On the other hand, there are no identifications on the interior of the $2k$-disc, since $\lambda(z, \sqrt{1-|z|^2})$ can only have final entry real if $\lambda \in \mathbb{R}$, and to preserve the modulus we would need $\lambda = 1$ or $\lambda = -1$; in the latter case we get a negative entry. So we have constructed $\mathbb{C}P^k$ by attaching a $2k$-disc to $\mathbb{C}P^{k-1}$ and the result holds by induction. □

It follows that the cellular chain complex for $\mathbb{C}P^k$ has $\mathbb{Z}$ entries in all even degrees $n \leqslant 2k$, and are zero elsewhere. Thus the boundary maps either have domain or codomain the trivial group and so are all trivial. Again, for this rather special example, we don't need to know the definition of the boundary map! It follows that

$$H_n(\mathbb{C}P^k) \cong \begin{cases} \mathbb{Z} & \text{for } n \text{ even and } n \leqslant 2k; \\ 0 & \text{otherwise.} \end{cases}$$

3.4.1 The cellular boundary map, informally

Now we shall explain how the boundary map works, and then look at a couple of low dimensional examples. Since ∂_n is a homomorphism, to work out the boundary map on

[2] Interesting aside: the quotient $S^3 \to \mathbb{C}P^1 \cong S^2$ is known as the *Hopf fibration*, and realises the 3-sphere as a 'twisted product' (more precisely, a *fibre bundle*) of S^2 with S^1 'fibres'.

an arbitrary chain one just needs to know the coefficients $d_{\alpha\beta} \in \mathbb{Z}$ in the expression

$$\partial_n(e_\alpha) = \sum_{(n-1)-\text{cells } e_\beta} d_{\alpha\beta} e_\beta,$$

where we identify an n-cell e_α of the CW complex with its associated generator of $C_n(X^\bullet)$. We shall have more to say on what these coefficients are in Section 4.5.1. Loosely speaking one proceeds as follows: firstly we have an orientation on e_α determined by the map σ_α attaching the n-cell e_α, and setting an orientation on D^n. The orientation on D^n induces one on S^{n-1}. The characteristic map σ_α sweeps the boundary S^{n-1} over $(n-1)$-cells e_β, perhaps with some multiplicity, which either preserves or reverses orientations. Counted according to whether orientations are preserved or reversed, this determines the coefficient $d_{\alpha\beta}$. Again, more details in Section 4.5.1, see also Figure 3.4.1.

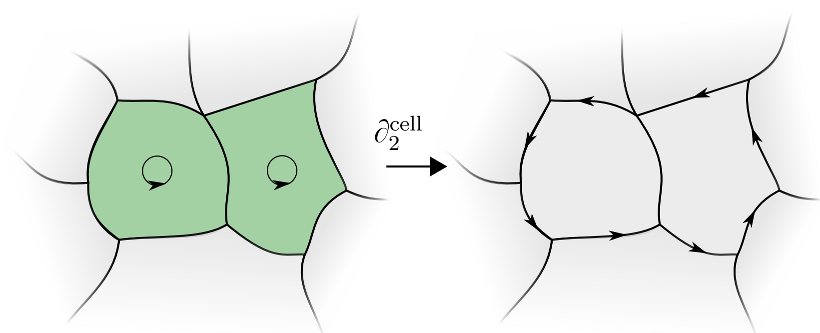

Figure 3.4.1: A cellular 2-chain, assigning coefficients 1 to the two green oriented 2-cells (and zero to the others) and its boundary, assigning coefficient 1 to the edges with orientations assigned according to the arrows.

For example, take a 1-cell e_α. It is mapped into X via $\sigma_\alpha \colon D^1 = [-1, 1] \to X$. With the standard orientation on $[-1, 1]$, this positively orients the right endpoint and negatively the left one. The right endpoint is mapped to some vertex e_r, and the left endpoint is mapped to some vertex e_l. Then

$$\partial_1(e_\alpha) = e_r - e_l. \tag{3.4.1}$$

In degree two we attach a 2-cell e_α via a map $\sigma_\alpha \colon D^2 \to X$. We imagine the standard orientation of D^2 orienting its boundary circle going anticlockwise. Suppose that σ_α maps the boundary of D^2 onto X^1 nicely, in the sense that S^1 can be subdivided into finitely many open intervals, separated by points, with σ_α mapping those points into the 0-skeleton and homeomorphically mapping each open interval onto some 1-cell e_β. Each of these intervals is either mapped onto some e_β in a direction agreeing with the orientation on e_β, contributing $+1$ to $d_{\alpha\beta}$, or disagreeing, contributing coefficient -1 to $d_{\alpha\beta}$.

Example 3.4.3. Take S^1 with CW decomposition of a single 0-cell v and single 1-cell e. By Equation 3.4.1, $\partial_1(e) = v - v = 0$. So the cellular chain complex is

$$\cdots \to 0 \to 0 \to \mathbb{Z} \xrightarrow[\partial_1]{\times 0} \mathbb{Z} \to 0,$$

hence $H_0(S^1) \cong \mathbb{Z}$, $H_1(S^1) \cong \mathbb{Z}$ and $H_n(S^1) \cong 0$ otherwise.

Example 3.4.4. The square model for the torus, as in Figure 1.2.1, gives a CW decomposition of the 2-torus $\mathbb{T}^2$ with one 0-cell, two 1-cells and one 2-cell. Each 1-cell is attached so its head and tail is the single 0-cell (of course), so as above $\partial_1 = 0$. Orientations for the 1-cells can be assigned, say, according to the arrows along the edges of Figure 1.2.1. Attaching the 2-cell, we see that its boundary is glued in a way which traverses each 1-cell twice: once in a direction which agrees with the orientation of the cell, and once with the opposite orientation. The two contributions for cancel on each 1-cell, so $\partial_2 = 0$. So we get chain complex

$$\cdots \to 0 \to 0 \to \mathbb{Z} \xrightarrow[\partial_2]{\times 0} \mathbb{Z}^2 \xrightarrow[\partial_1]{\times 0} \mathbb{Z} \to 0,$$

hence $H_0(\mathbb{T}^2) \cong \mathbb{Z}$, $H_1(\mathbb{T}^2) \cong \mathbb{Z}^2$, $H_2(\mathbb{T}^2) \cong \mathbb{Z}$ and $H_k(\mathbb{T}^2) \cong 0$ otherwise.

Exercise 3.4.1. Work out $H_*(\mathbb{K})$ and $H_*(\mathbb{R}P^2)$.

Exercise 3.4.2. Try to think about what the boundary map ∂_3 should be for CW decompositions where the 3-cells are attached 'nicely'. One can construct the 3-torus $\mathbb{T}^3 \cong S^1 \times S^1 \times S^1$ from a cube whose opposite square faces are identified in a similar way to the 2-torus (e.g., the 'left-hand' face is translated right 1 unit to identify it with the 'right-hand' face). What should the cellular chain complex and homology be? Try experimenting with this construction by glueing the faces in different ways (for example, glueing opposite faces with a 90 degrees twist) and working out the homology.

Chapter 4

Computational machinery for homology

Homology would be an awkward tool if we had to start from scratch in computing it for each new space encountered. What is needed are tools relating the homologies of spaces which are associated via some form of construction. For example, if a space can be decomposed into a union of subspaces A and B, it would be useful to know the relationship between the homologies of the spaces A, B, $A \cup B$ and $A \cap B$. Later in this chapter we shall see that there is such a gadget, the Mayer–Vietoris sequence.

More important still is the case that we have a pair of spaces (X, A), of a space X with subspace $A \subseteq X$. It would be nice to know the relationship between the homologies of X, A and the quotient X/A. In the next section we shall introduce the *relative homology* $H_*(X, A)$ of the pair (X, A), which morally plays the rôle of the homology of the quotient space X/A. In fact, $H_k(X/A) \cong H_k(X, A)$ when $k > 0$ so long as the subspace A is not embedded too pathologically into X, and in degree zero $H_0(X, A) \oplus \mathbb{Z} \cong H_0(X/A)$ (so the relative homology corresponds to the so-called *reduced homology* of X/A, Section 4.2). Just as for a pair of a chain complex and sub-chain complex we get a long exact sequence from the Snake Lemma, for the pair (X, A) we will get a long exact sequence relating $H_*(X)$, $H_*(A)$ and $H_*(X, A)$.

One particular use of this comes from applying it to spaces with cellular decompositions. For a CW complex $X^\bullet$, consider the pair (X^k, X^{k-1}). Since X^{k-1} is nicely embedded into X^k, it will turn out that the relative homology $H_*(X^k, X^{k-1})$ can be thought of as the (reduced) homology of the quotient X^k/X^{k-1}, which is a wedge of k-spheres, in correspondence with the k-cells. This has homology concentrated in degree k as the free Abelian group with generators in bijection with the k-cells, exactly the description of the degree k chain group of the cellular chain complex. It is via this route that one defines the boundary maps of the cellular chain complex and then proves that cellular homology is isomorphic to singular homology.

4.1 Relative homology

4.1.1 Definition of relative homology and the LES of a pair

Throughout this section we let (X, A) be a topological pair, that is a space X with subspace $A \subseteq X$. The singular chain complex $C_*(A)$ of A can be naturally considered as a sub-chain complex of $C_*(X)$, namely as the sub-chain complex whose degree n chain groups are freely generated by singular n-simplices $\alpha\colon \Delta^n \to X$ with images contained wholly in A. This really is a sub-chain complex, since if the image of α is contained in A then so are the images of the terms of the boundary $\partial\alpha$.

We may consider the quotient complex $C_*(X)/C_*(A)$. By definition, remember, the boundary map of the quotient complex will essentially correspond to the usual one, but we identify any two n-chains σ_1, $\sigma_2 \in C_n(X)$ if $\sigma_1 - \sigma_2 \in C_n(A)$ (in particular, σ represents the zero element if $\sigma \in C_*(A)$). Let's improve our notation by writing $C_*(X, A)$ instead for the quotient complex.

The homology $H_*(X, A)$ of this complex is called the **relative homology** of (X, A). Elements of $H_n(X, A)$ are represented by **relative cycles**, n-chains $\sigma \in C_n(X)$ for which $\partial\sigma \in C_n(A)$. Two relative cycles σ_1, σ_2 are identified in $H_n(X, A)$ if there exists some $(n + 1)$-chain $\tau \in C_{n+1}(X)$ with $(\sigma_1 - \sigma_2) - \partial\tau \in C_n(A)$, that is, if up to a boundary their difference is a chain in A. See Figure 4.1.1. It's sometimes useful to let $A = \varnothing$, in which case the relative homology $H_n(X, \varnothing)$ is just the usual singular homology $H_n(X)$.

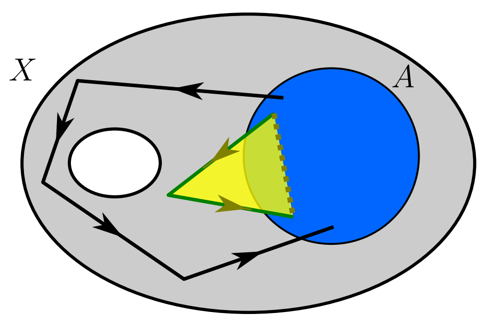

Figure 4.1.1: The black 1*-chain is a relative cycle, its boundary is in* A*. The green* 1*-chain is also a relative cycle. In fact, unlike for the black* 1*-chain, it is a relative boundary so represents zero in relative homology: its difference with the boundary of the yellow* 2*-chain is a* 1*-chain, indicated by the dotted line, lying inside* A*.*

The complexes $C_*(A)$, $C_*(X)$ and $C_*(X, A)$ are of course related by a short exact sequence of chain complexes (inclusion followed by quotient):

$$0 \to C_*(A) \hookrightarrow C_*(X) \to C_*(X, A) \to 0.$$

We may apply the Snake Lemma (Lemma 2.3.1) to this to obtain the **long exact sequence of the pair** (X, A):

$$\cdots \xrightarrow{q_*} H_{k+1}(X, A) \xrightarrow{\partial_*}$$

$$H_n(A) \xrightarrow{\iota_*} H_n(X) \xrightarrow{q_*} H_n(X, A) \xrightarrow{\partial_*}$$

$$H_{n-1}(A) \xrightarrow{\iota_*} H_{n-1}(X) \xrightarrow{q_*} H_{n-1}(X, A) \xrightarrow{\partial_*}$$

$$H_{n-2}(A) \xrightarrow{\iota_*} \cdots$$

It's often important to know not only the groups involved in the above LES, but also the homomorphisms between them. The map ι_* is the induced map of the inclusion chain map $i\colon C_*(A) \hookrightarrow C_*(X)$. This is the same thing as the induced map of the inclusion map $\iota\colon A \hookrightarrow X$. The map q_* is the induced map of the quotient chain map $q\colon C_*(X) \to C_*(X, A)$. It will turn out that in most cases of interest this corresponds to the induced map of the quotient map $f\colon X \to X/A$ (Theorem 4.4.1). Finally, following Remark 2.3.1, the boundary map ∂_* also has a nice description: any element of $H_n(X, A)$ is represented by a relative cycle $\sigma \in C_n(X)$. Then the connecting map $\partial_*([\sigma])$ applied to the homology class of σ is the homology class in $H_{n-1}(A)$ of $\partial(\sigma)$. Remember that since σ is a relative cycle, we do indeed have that $\partial(\sigma) \in C_{n-1}(A)$.

Exercise 4.1.1. Describe $H_0(X, A)$ for any pair (X, A).

Exercise 4.1.2. Let $x \in X$. Describe $H_*(X, \{x\})$ in terms of $H_*(X)$ (try a LES).

Exercise 4.1.3. Consider a triple of spaces $B \subseteq A \subseteq X$. Show that there is a **long exact sequence of the triple** relating the relative homologies of the pairs (X, A), (X, B) and (A, B).

4.1.2 Induced maps

It's easy to make relative homology functorial. Given a map of pairs $f\colon (X, A) \to (Y, B)$, so $f\colon X \to Y$ is continuous and $f(A) \subseteq B$, we get induced maps between the relative homology groups $f_*\colon H_*(X, A) \to H_*(Y, B)$. This is induced by the chain map from the non-relative case: in degree n we just set $f_n([\sigma]) := [f_n(\sigma)]$ for $\sigma \in C_n(X)$.

It is then easily verified that relative singular homology $H_n(-, -)$ defines a functor from the category $\mathbf{Top}^2$ of pairs of topological spaces to the category **Ab** of Abelian groups.

Here is a simple application of the LES of a pair and the LES of a triple:

Lemma 4.1.1. *Consider a triple of spaces (X, U, A), so $A \subseteq U \subseteq X$. Suppose that the inclusion $\iota\colon A \hookrightarrow U$ is a homotopy equivalence. Then the inclusion $(X, A) \hookrightarrow (X, U)$ induces an isomorphism in relative homology.*

Proof. First consider the LES of the pair (U, A):

$$\cdots \xrightarrow{q_*} H_{n+1}(U, A) \xrightarrow{\partial_*} H_n(A) \xrightarrow{\iota_*} H_n(U) \xrightarrow{q_*} H_n(U, A) \xrightarrow{\partial_*} H_{n-1}(A) \xrightarrow{\iota_*} \cdots .$$

Since ι is a homotopy equivalence, ι_* is an isomorphism on homology. It follows from the above diagram that $H_n(U, A)$ is trivial for all $n \in \mathbb{Z}$.

Now consider the LES of the triple (X, U, A) from Exercise 4.1.3. If you write that down, you will see that the trivial terms $H_n(U, A)$ imply that the induced map of the inclusion $(X, A) \hookrightarrow (X, U)$ is an isomorphism on relative homology. □

Homotopy invariance

Remember from Theorem 3.3.1 that for two homotopic maps $f, g\colon X \to Y$ the induced maps f_*, $g_*\colon H_*(X) \to H_*(Y)$ agreed. There is a relative version of this too:

Theorem 4.1.1. *Suppose that f and g are maps of pairs from (X, A) to (Y, B) i.e., they are continuous maps from X to Y mapping A into B. Suppose that f and g are homotopic, so that there exists $F\colon X \times I \to Y$ with $F(x, 0) = f(x)$ and $F(x, 1) = g(x)$. Moreover, suppose that each intermediate map $x \mapsto F(x, t)$ for $t \in (0, 1)$ in the homotopy is also a map between the pairs (X, A) and (Y, B). Then $f_* = g_*$ for the induced maps f_*, $g_*\colon H_*(X, A) \to H_*(Y, B)$ between relative homology groups.*

Proof. We have the following maps of pairs:

$$\begin{aligned}
&f\colon (X, A) \to (Y, B),\ \ g\colon (X, A) \to (Y, B);\\
&\iota_0\colon (X, A) \to (X \times I, A \times I),\ \ \iota_1\colon (X, A) \to (X \times I, A \times I);\\
&F\colon (X \times I, A \times I) \to (Y, B).
\end{aligned}$$

The final map is a map of pairs by the assumption of the theorem. The maps ι_0 and ι_1 are just the inclusions $x \mapsto (x, 0)$ and $x \mapsto (x, 1)$, respectively. Then as maps of pairs $f = F \circ \iota_0$ and $g = F \circ \iota_1$.

Remember that to prove homotopy invariance in the non-relative case, we showed that $(\iota_0)_\sharp$ and $(\iota_1)_\sharp$ were chain homotopy equivalent. The chains τ_n constructed in that proof can be picked as before. We can use the same Equation 3.3.2 to define a chain homotopy, since for a singular simplex $\sigma\colon \Delta^n \to A$ we have that $(\sigma, \mathrm{id})_\sharp(\tau_{n+1})$ is clearly a singular chain of $C_*(A \times I)$. It follows that (h_n) still gives a well defined chain homotopy, now between the relative chain maps

$$(\iota_0)_\sharp \simeq (\iota_1)_\sharp \colon C_*(X, A) \to C_*(X \times I, A \times I).$$

The proof is then complete via the same argument following the statement of Lemma 3.3.5.

□

4.2 Reduced homology

Under certain conditions we want to think of the relative homology of (X, A) as corresponding to the homology of the quotient X/A. If you followed through the above exercises, you will have noticed that this can't be so in degree zero. For example, the relative homology of (X, X) is going to be trivial in each degree (as all the chain groups are trivial), but the quotient space is the one point space, which we computed before as having $H_0 \cong \mathbb{Z}$.

This issue in degree zero is fixed by using instead *reduced homology*. We modify the singular chain complex by 'augmenting' in degree -1:

$$\cdots \xrightarrow{\partial_3} C_2(X) \xrightarrow{\partial_2} C_1(X) \xrightarrow{\partial_1} C_0(X) \xrightarrow{\varepsilon} \mathbb{Z} \to 0.$$

Here, the map ε is defined by $\varepsilon(\sum \ell_\alpha \alpha) \coloneqq \sum \ell_\alpha$. Remember that in degree zero an elementary chain α is essentially just a choice of point in X, so a chain is a finite collection of $\mathbb{Z}$-weighted points and ε just takes the sum of these weights. The homology of this chain complex is called the **reduced homology** of X, denoted $\widetilde{H}_*(X)$.

Exercise 4.2.1. Consider the above construction but without any geometry involved: take a chain complex C_* which has C_n trivial for $n < 0$ and suppose that there exists some surjective map $\varepsilon \colon C_0 \to \mathbb{Z}$ making

$$\cdots \xrightarrow{\partial_3} C_2 \xrightarrow{\partial_2} C_1 \xrightarrow{\partial_1} C_0 \xrightarrow{\varepsilon} \mathbb{Z} \to 0$$

a chain complex, with homology $\widetilde{H}_*$. Implementing the Snake Lemma, or otherwise, show that there are canonical isomorphisms $\widetilde{H}_n \cong H_n$ for $n > 0$, and in degree 0 we have a short exact sequence

$$0 \to \widetilde{H}_0 \to H_0 \to \mathbb{Z} \to 0.$$

This is a SES with final term free Abelian so, by Corollary 2.4.1, $H_0 \cong \widetilde{H}_0 \oplus \mathbb{Z}$.

Exercise 4.2.2. Another way to think about the reduced homology: for any space X, consider the unique map $p \colon X \to *$, where $*$ is the one point space. Show that we may identify $\widetilde{H}_n(X)$ as the subgroup of $H_n(X)$ given as the kernel of the induced map $p_* \colon H_n(X) \to H_n(*)$.

Summing up the above exercises, we have isomorphisms

$$H_n(X) \cong \begin{cases} \widetilde{H}_n(X) & \text{for } n \neq 0; \\ \widetilde{H}_0(X) \oplus \mathbb{Z} & \text{for } n = 0. \end{cases}$$

So taking the reduced homology simply reduced the rank of $H_0(X) \cong \mathbb{Z}^n$ by one (namely, to the number of path-components of X less one, see Exercise 3.3.2).

Remark 4.2.1. This may seem an unnatural homology group to assign to degree zero. But in some ways it isn't: one should think of $H_n(X)$ as detecting n-dimensional kinds of 'holes' in X. In degree zero this corresponds to path-connectedness, so it's perhaps persuasive to think of a path-connected space as having no interesting homology in degree zero. For the 0-sphere S^0, in contrast, we have that $\widetilde{H}_0(S^0) \cong \mathbb{Z}$ and is trivial in other degrees. So the reduced homology is 'generated' by the principal feature of disconnectedness of the two point space S^0. As an aside, if you now take a look at the proof of Brouwer's Fixed Point Theorem in the introduction, you can see how this disconnectedness is precisely the obstruction used to show the non-existence of the retract of an interval to its endpoints. You've likely seen before the proof that a continuous map $f\colon [0,1] \to [0,1]$ has a fixed point. In this proof one considers $g(x) \coloneqq f(x) - x$. Assuming f doesn't have a fixed point is to assume that $g(x) \neq 0$ for all x. But $g(0) > 0$ and $g(1) < 0$, so one gets a contradiction by the intermediate value theorem. The IVT is proved using the notion of connectedness. So really, the proof of Brouwer's Fixed Point Theorem as in the introduction is based upon the same topological argument, but using higher dimensional obstructions than that of connectedness for the higher dimensional cases.

Exercise 4.2.3. Let X be a space and $x \in X$. Show that the reduced homology $\widetilde{H}_*(X)$ is isomorphic to the relative homology $H_*(X, \{x\})$ (use Exercise 4.1.2).

Remark 4.2.2. Whilst we have an isomorphism $\widetilde{H}_*(X) \cong H_*(X, \{x\})$, it's more natural to consider $\widetilde{H}_0(X)$ as a *subgroup* of $H_0(X)$ (the subgroup of 0-cycles whose signed sums are zero) versus $H_0(X, \{x\})$ as a *quotient* of $H_0(X) \cong \mathbb{Z}^n$ (given by collapsing the generator corresponding to the path-component of x).

Following the above exercises, we see that there's an explicit isomorphism between the two given by the composition

$$\widetilde{H}_*(X) \hookrightarrow H_*(X) \twoheadrightarrow H_*(X, \{x\}),$$

where the first map is induced by the inclusion from Exercise 4.2.2 and the second is induced by the inclusion $(X, \varnothing) \hookrightarrow (X, \{x\})$. It's easy to see how to make reduced homology functorial using induced maps as in the non-reduced case and one can essentially replace appearances of reduced homology in diagrams with relative homologies of a space relative to a point.

For example, consider the triple $\{x\} \subseteq A \subseteq X$. Using the above exercise and Exercise 4.1.3, we have a reduced version of the LES of a pair:

$$\cdots \xrightarrow{q_*} H_{k+1}(X, A) \xrightarrow{\partial_*} \widetilde{H}_k(A) \xrightarrow{\iota_*} \widetilde{H}_k(X) \xrightarrow{q_*} H_k(X, A) \xrightarrow{\partial_*} \widetilde{H}_{k-1}(A) \xrightarrow{\iota_*} \widetilde{H}_{k-1}(X) \xrightarrow{q_*} H_{k-1}(X, A) \xrightarrow{\partial_*} \widetilde{H}_{k-2}(A) \xrightarrow{\iota_*} \cdots$$

Of course, this only makes a difference in degree zero. This essentially just tidies away some boring initial terms. The same LES results from applying the Snake Lemma to the SES of chain complexes:

$$0 \to \widetilde{C}_*(A) \xrightarrow{\iota_\sharp} \widetilde{C}_*(X) \xrightarrow{q_\sharp} C_*(X, A) \to 0.$$

The chain map $\iota_\sharp$ is induced by the inclusion map $A \hookrightarrow X$ (which we equate with $(A, \varnothing) \hookrightarrow (X, \varnothing)$), and $q_\sharp$ is induced by the inclusion of pairs $(X, \varnothing) \hookrightarrow (X, A)$. The relative complex $C_*(X, A)$ is 'already reduced' for $A \neq \varnothing$.

4.3 Excision

4.3.1 Warm-up: simplicial excision

Let $(\mathcal{X}, \mathcal{A})$ be a simplicial pair, that is, a simplicial complex $\mathcal{X}$ with sub-complex $\mathcal{A}$, meaning a subset of simplices of $\mathcal{X}$ forming a simplicial complex in its own right. Just as we defined relative singular homology, it is easy to define relative simplicial homology for such pairs (and similarly, once we've defined the boundary maps, it is easy to define relative cellular homology). The relative simplicial chain complex $C_*(\mathcal{X}, \mathcal{A})$ is the quotient of the simplicial chain complexes $C_*(\mathcal{X})/C_*(\mathcal{A})$. In this chain complex we identify chains lying in $\mathcal{A}$ with zero.

Consider a triple $(\mathcal{X}, \mathcal{A}, \mathcal{B})$ of simplicial complexes, so that $\mathcal{X}$ is a simplicial complex with a sub-complex $\mathcal{A}$ which in turn has sub-complex $\mathcal{B}$. Let $\mathcal{X} - \mathcal{B}$ be the sub-complex of $\mathcal{X}$ consisting of simplices which are neither in $\mathcal{B}$ nor are a face of an element of $\mathcal{B}$ (check that this is a subcomplex). One can think of $\mathcal{X} - \mathcal{B}$ as given by removing 'open cells' of $\mathcal{B}$ from $\mathcal{X}$. Define $\mathcal{A} - \mathcal{B}$ analogously.

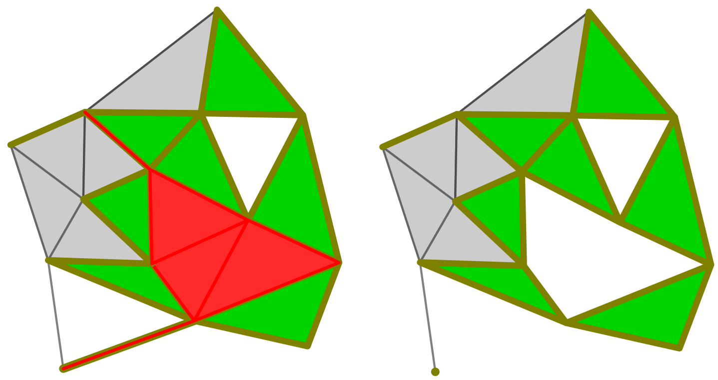

Figure 4.3.1: On the left, a simplicial complex $\mathcal{X}$ (in grey, green and red) with sub-complex $\mathcal{A}$ (in green and red) which in turn has a sub-complex $\mathcal{B}$ (in red). The relative homology of $(\mathcal{X}, \mathcal{A})$ does not change if we remove from $\mathcal{X}$ and $\mathcal{A}$ open cells from $\mathcal{B}$, as in the right-hand picture.

Consider the two relative simplicial homology groups $H_*(\mathcal{X}, \mathcal{A})$, $H_*(\mathcal{X} - \mathcal{B}, \mathcal{A} - \mathcal{B})$. In the second one we have 'excised' $\mathcal{B}$. Doing so does not effect the homology calculations, since the excised chains, those falling properly within $\mathcal{B}$, already represent zero in $C_*(\mathcal{X}, \mathcal{A})$. That is, $C_*(\mathcal{X}, \mathcal{A})$ and $C_*(\mathcal{X} - \mathcal{B}, \mathcal{A} - \mathcal{B})$ are isomorphic chain complexes via the chain map induced by the inclusion of pairs $(\mathcal{X} - \mathcal{B}, \mathcal{A} - \mathcal{B}) \hookrightarrow (\mathcal{X}, \mathcal{A})$.

4.3.2 Excision for singular homology

What makes the above argument work is that if a simplex belongs to $\mathcal{B}$ then it definitely belongs to $\mathcal{A}$. We want a similar theorem to hold for singular homology: for a triple (X, A, B) we would like that $H_*(X, A)$ is isomorphic to $H_*(X - B, A - B)$—we can *excise* B—realised through the inclusion of pairs $(X - B, A - B) \hookrightarrow (X, A)$. Unfortunately the above argument won't work since a singular simplex $\alpha \colon \Delta^n \to X$ can easily have image intersecting B but also not lying entirely within A.

We at least need to assume that B has some 'wiggle room' within A; to make this precise: for the closure of B to be contained in the interior of A. In this case, any singular simplex intersecting B intuitively has some way to travel to escape A. Sufficiently 'small' singular simplices which intersect B would then have to be within A. To prove excision for such pairs, then, one must firstly show that no harm is done by restricting to 'small' singular simplices (this is Lemma 4.3.1 below).

We don't have a metric, but we can make precise what we mean by 'small' here using covers. Let X be a topological space and $\mathscr{U}$ be a collection of subsets whose interiors cover X. For a singular simplex $\alpha \colon \Delta^n \to X$ we say that α is **$\mathscr{U}$-small** if the image of α is contained in some $U \in \mathscr{U}$; similarly, a chain given as a $\mathbb{Z}$-linear sum of $\mathscr{U}$-small singular simplices is also called $\mathscr{U}$-small. For example, if X *did* happen to be a metric space, we could take $\mathscr{U}$ to be the set of balls of some radius $\epsilon > 0$, then a chain is $\mathscr{U}$-small if each singular simplex in its defining sum has image ('support') contained in some ϵ-ball.

It is easy to see that the boundary of a $\mathscr{U}$-small chain is still $\mathscr{U}$-small, so these chains form a chain complex, denoted $C_*^{\mathscr{U}}(X)$. For $A \subseteq X$ we also have the relative complex $C_*^{\mathscr{U}}(X, A)$ given as you'd expect by the quotient of $C_*^{\mathscr{U}}(X)$ by the sub-complex of $\mathscr{U}$-small chains in $C_*(A)$.

The main technical hurdle in proving excision is that one may freely pass to $\mathscr{U}$-small chains without changing the homology:

Lemma 4.3.1. *The inclusion $\iota\colon C_*^{\mathscr{U}}(X) \hookrightarrow C_*(X)$ induces isomorphisms on homology.*

The above is proved using a process called 'barycentric subdivision'. We will need to devote some space to the proof, so before doing that we give the following corollary and then the proof of excision from the above lemma and the below relative version:

Corollary 4.3.1. *For a pair of spaces the obvious inclusion $\iota\colon C_*^{\mathscr{U}}(X, A) \hookrightarrow C_*(X, A)$ induces isomorphisms on homology.*

Proof. The 'obvious inclusion' simply takes $[\sigma] \in C_n^{\mathscr{U}}(X, A)$ to $[\sigma] \in C_n(X, A)$; note that this is well-defined since $[\sigma] = [\sigma']$ in the $\mathscr{U}$-small complex precisely if $\sigma - \sigma' \in C_n(A)$, which is the identical equivalence relation defining $C_n(X, A)$.

We have a commutative diagram of short exact sequences of chain complexes:

$$\begin{array}{ccccccccc} 0 & \longrightarrow & C_*^{\mathscr{U}}(A) & \longrightarrow & C_*^{\mathscr{U}}(X) & \longrightarrow & C_*^{\mathscr{U}}(X, A) & \longrightarrow & 0 \\ & & \downarrow & & \downarrow & & \downarrow & & \\ 0 & \longrightarrow & C_*(A) & \longrightarrow & C_*(X) & \longrightarrow & C_*(X, A) & \longrightarrow & 0 \end{array}$$

The maps here are the obvious inclusions and quotients, everything commutes. One may apply the Snake Lemma, implementing also its naturality (Appendix A.1):

$$\begin{array}{ccccccccccc} \cdots & \longrightarrow & H_{n+1}^{\mathscr{U}}(X, A) & \longrightarrow & H_n^{\mathscr{U}}(A) & \longrightarrow & H_n^{\mathscr{U}}(X) & \longrightarrow & H_n^{\mathscr{U}}(X, A) & \longrightarrow & \cdots \\ & & \downarrow & & \downarrow & & \downarrow & & \downarrow & & \\ \cdots & \longrightarrow & H_{n+1}(X, A) & \longrightarrow & H_*(A) & \longrightarrow & H_n(X) & \longrightarrow & H_n(X, A) & \longrightarrow & \cdots \end{array}$$

From Lemma 4.3.1 the vertical non-relative maps are isomorphisms. In this situation the remaining vertical maps must be isomorphisms too (this follows from the Four Lemma, Homework 2). $\square$

Theorem 4.3.1 (Excision)**.** *Let $X \supseteq A \supseteq B$ be a triple of spaces for which the closure of B is contained in the interior of A. Then the inclusion of pairs $i\colon (X-B, A-B) \hookrightarrow (X, A)$ induces is an isomorphism in relative homology:*

$$H_*(X-B, A-B) \xrightarrow[i_*]{\cong} H_*(X, A).$$

Proof. Take the couple of subspaces $\mathscr{U} \coloneqq \{X-B, A\}$. Their interiors cover X, since $\mathrm{int}(X-B) = X - \mathrm{cl}(B)$ and by assumption $\mathrm{cl}(B) \subseteq \mathrm{int}(A)$.

Consider the following commutative diagram:

$$\begin{array}{ccc} H_*^{\mathscr{U}}(X-B, A-B) & \xrightarrow{i_*^{\mathscr{U}}} & H_*^{\mathscr{U}}(X, A) \\ \downarrow{\scriptstyle \iota_*} & & \downarrow{\scriptstyle \iota_*} \\ H_*(X-B, A-B) & \xrightarrow{i_*} & H_*(X, A) \end{array}$$

Each map here is just one induced by a natural inclusion of chain complexes, so the diagram certainly commutes. The vertical maps are isomorphisms by Corollary 4.3.1, and we wish to show that the bottom horizontal map is an isomorphism. This follows from the fact that $\iota_*^{\mathscr{U}}$ is an isomorphism, which is the case for the analogous reasons of the toy example in the introduction for simplicial complexes. Indeed the inverse chain map to $\iota_\sharp^{\mathscr{U}}$ is given by sending a generator chain $[\alpha] \in C_*^{\mathscr{U}}(X, A)$, for α a $\mathscr{U}$-small singular simplex, to $[\alpha] \in C_*^{\mathscr{U}}(X-B, A-B)$ if the support of α does not intersect B, and to 0 if it does—we can safely do this since α is $\mathscr{U}$-small, so in the latter case has support wholly in A and hence $[\alpha] = 0$ in $C_*^{\mathscr{U}}(X, A)$ already anyway. □

4.3.3 Proof of Subdivision Lemma 4.3.1

The main work in proving excision is the content of Lemma 4.3.1. With full details the proof is relatively long-winded and I don't expect you to be able to reproduce it from memory. I haven't moved the proof to the appendix though, I think it is a worthwhile exercise to go through it carefully and try to understand how and why it works. Hopefully the outline below gives enough sign posts to indicate the main ideas. Figures 4.3.2 and 4.3.3 should be persuasive in demonstrating that we may 'cut up' a singular chain to a smaller one.

Outline

1. For all spaces X, we inductively construct *subdivision operators*, chain maps $\mathrm{Sd}_*^X\colon C_*(X) \to C_*(X)$. These are defined in terms of barycentric subdivision of the standard n-simplex (Figure 4.3.2) which is then extended in a natural way to all chains (Figure 4.3.3).

2. We show that these are indeed chain maps (Equation 4.3.4) and are natural with respect to maps between spaces (Equation 4.3.3).
3. We identify how to construct a chain homotopy $h_n^X\colon C_n(X) \to C_{n+1}(X)$ from the subdivision operator to the identity, imposing again a naturality condition with respect to continuous maps (Equation 4.3.5), and the chain homotopy formula (Equation 4.3.6). As for the construction of the subdivision operator, everything is determined by what happens to the standard simplex.
4. The existence of the chain homotopy is proved by exploiting the vanishing homology of Δ^n (c.f., the proof of homotopy invariance, Theorem 3.3.1).
5. Assuming that repeated application of subdivision always eventually results in $\mathscr{U}$-small chains, we prove the statement of the lemma using subdivision and the chain homotopy.
6. Finally, we prove that repeated application of subdivision does indeed always eventually result in $\mathscr{U}$-small chains.

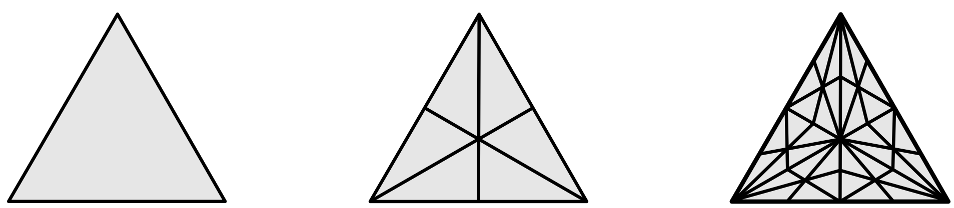

Figure 4.3.2: The standard 2-simplex Δ^2, re-projected to $\mathbb{R}^2$ (left), the images of the singular 2-simplices defining the first barycentric subdivision of λ^2 (centre) and a second application of barycentric subdivision (right).

Step 1: constructing the barycentric subdivision operator

Let $K \subseteq \mathbb{R}^N$ be a convex subspace (later in this proof $K = \Delta^n$). We will briefly reuse a cone-type construction from the proof of Lemma 3.3.2. Given $x_0 \in K$ and a singular n-simplex $\alpha\colon \Delta^n \to K$, we define $c_n^{x_0}(\alpha)$ as the singular $(n+1)$-simplex given by 'stretching α linearly from base α to tip x_0', see Figure 3.3.1. I'll give you the explicit formula now: for $x = (t_0, t_1, \ldots, t_{n+1}) \in \Delta^{n+1}$, define

$$(c_n^{x_0}(\alpha))(x) \coloneqq \begin{cases} t_0 x_0 + (1-t_0)\alpha(\frac{(t_1,\ldots,t_{n+1})}{1-t_0}) & \text{for } t_0 \neq 1; \\ x_0 & \text{for } t_0 = 1. \end{cases}$$

This is extended to $\mathbb{Z}$-linear sums of singular simplices, as usual, $\mathbb{Z}$-linearly. Hopefully you wrote something similar before and proved the following:

$$\partial_{n+1}(c_n^{x_0}\sigma) = \sigma - c_{n-1}^{x_0}(\partial_n \sigma). \tag{4.3.1}$$

Although not relevant for the proof to follow, technically this formula does not hold for $n = 0$, but it does in the reduced complex; for $n = 0$ the occurrence of ∂_0 above should be instead the augmentation map and $c_{-1}^{x_0}$ sends $n \in \mathbb{Z}$ to n times the singular 0-simplex mapping the point of Δ^0 to x_0.

Now we define our barycentric subdivision operator in an inductive fashion. Let X be a space. In degree zero barycentric subdivision does nothing, $\mathrm{Sd}_0^X(\sigma) = \sigma$ for $\sigma \in C_0(X)$.

Suppose then that Sd_k^X has been constructed for $k < n$ (for all spaces X). The singular n-simplex $\mathrm{id}_{\Delta^n}\colon \Delta^n \to \Delta^n$, which we denote by λ^n, acts as a seed for all others, in the sense that for any singular n-simplex $\alpha\colon \Delta^n \to X$ we have that as an elementary chain $\alpha = \alpha_\sharp(\lambda^n)$. So we start with explicitly defining the subdivision of λ^n:

$$\mathrm{Sd}_n^{\Delta^n}(\lambda^n) \coloneqq c_{n-1}^{\odot^n}(\mathrm{Sd}_{n-1}^{\Delta^n}(\partial_n \lambda^n)). \tag{4.3.2}$$

Here, $\odot^n$ stands for the 'barycentre' $(\frac{1}{n+1}, \ldots, \frac{1}{n+1})$ of Δ^n. In English: *the subdivision of the n-simplex is defined by first subdividing its boundary $(n-1)$-simplices, and then extending these via cones to the barycentre.* Figure 4.3.2 shows a couple of applications of barycentric subdivision in dimension 2.

We then define Sd_n^X for any singular n-simplex $\alpha\colon \Delta^n \to X$ by

$$\mathrm{Sd}_n^X(\alpha) \coloneqq \alpha_\sharp(\mathrm{Sd}_n^{\Delta^n}(\lambda^n)).$$

That is, we first subdivide the standard simplex and *then* map from it with α. See Figure 4.3.3. This defines Sd_n^X on any n-chain by extending linearly. Note that there is no clash in the definitions to Equation 4.3.2 if we apply the above formula to $\mathrm{Sd}_n^{\Delta^n}(\lambda^n)$ since in that case $\lambda_\sharp^n$ is the identity chain map on $C_*(\Delta^n)$.

Exercise 4.3.1. Draw the barycentric subdivision $\mathrm{Sd}_2^{\Delta^n}(\lambda^2)$ as seen in Figure 4.3.2 yourself, but by carefully following the definitions (for example, you will need to also draw the barycentric subdivision of λ^1). I *really* recommend doing this exercise!

Step 2: Naturality and chain map property of subdivision

Subdivision is natural with respect to continuous maps $f\colon X \to Y$, which means:

$$f_\sharp \circ \mathrm{Sd}_n^X = \mathrm{Sd}_n^Y \circ f_\sharp. \tag{4.3.3}$$

Indeed, for a singular simplex $\alpha\colon \Delta^n \to X$ we have:

$$\begin{aligned} f_\sharp(\mathrm{Sd}_n^X(\alpha)) &\coloneqq \\ f_\sharp \alpha_\sharp \mathrm{Sd}_n^{\Delta^n}(\lambda^n) = (f \circ \alpha)_\sharp \mathrm{Sd}_n^{\Delta^n}(\lambda^n) \eqqcolon \mathrm{Sd}_n^Y(f \circ \alpha) &= \\ \mathrm{Sd}_n^Y(f_\sharp \alpha). \end{aligned}$$

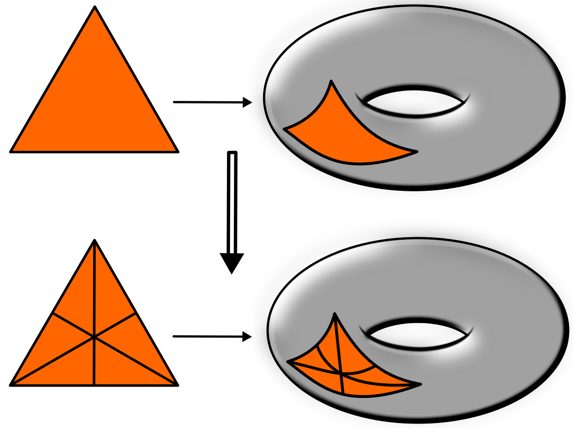

Figure 4.3.3: A singular simplex α on $X = \mathbb{T}^2$, and what it looks like after applying barycentric subdivision.

Secondly, subdivision defines chain maps:

$$\partial_n \circ \mathrm{Sd}_n^X = \mathrm{Sd}_{n-1}^X \circ \partial_n. \tag{4.3.4}$$

We prove this by induction on n.

This is trivially true for $n = 0$ (both sides are zero). Suppose then that Equation 4.3.4 holds for $k < n$. For a singular n-simplex $\alpha\colon \Delta^n \to X$ we have:

$$\begin{aligned}
& & \partial_n[\mathrm{Sd}_n^X(\alpha)] &\coloneqq \\
&(\textit{definition}): & \partial_n(\alpha_\sharp[\mathrm{Sd}_n^{\Delta^n}(\lambda^n)] &\coloneqq \\
&(\textit{definition}): & [\partial_n\alpha_\sharp](c_{n-1}^{\odot^n}(\mathrm{Sd}_{n-1}^{\Delta^n}(\partial_n\lambda^n))) &= \\
&(\alpha_\sharp\ \textit{chain map}): & \alpha_\sharp[\partial_n c_{n-1}^{\odot^n}](\mathrm{Sd}_{n-1}^{\Delta^n}(\partial_n\lambda^n)) &= \\
&(\textit{Cone formula 4.3.1}): & \alpha_\sharp\big(\mathrm{Sd}_{n-1}^{\Delta^n}(\partial_n\lambda^n) - c_{n-2}^{\odot^n}([\partial_{n-1}\mathrm{Sd}_{n-1}^{\Delta^n}]\partial_n\lambda^n)\big) &= \\
&(\textit{Induction for ch. map Eq. 4.3.4}) & \alpha_\sharp\big(\mathrm{Sd}_{n-1}^{\Delta^n}(\partial_n\lambda^n) - c_{n-2}^{\odot^n}\mathrm{Sd}_{n-2}^{\Delta^n}[\partial_{n-1}\partial_n]\lambda^n)\big) &= \\
&(\partial^2 = 0): & [\alpha_\sharp](\mathrm{Sd}_{n-1}^{\Delta^n}\partial_n\lambda^n) &= \\
&(\textit{naturality Eq. 4.3.3}): & \mathrm{Sd}_{n-1}^X([\alpha_\sharp\partial_n\lambda^n]) &= \\
&(\alpha_\sharp\ \textit{chain map}): & \mathrm{Sd}_{n-1}^X(\partial_n\alpha) &
\end{aligned}$$

The elements in square brackets indicate what is changed to proceed to the line below, with a description on that next line of what property is used.

Step 3: How to find a chain homotopy between subdivision and identity

We want a chain homotopy $h_n^X\colon C_n(X) \to C_{n+1}(X)$ between Sd_*^X and the identity chain map. We want these homomorphisms to be natural, in the sense that for a continuous map $f\colon X \to Y$ we have that:

$$f_\sharp \circ h_n^X = h_n^Y \circ f_\sharp. \tag{4.3.5}$$

The homomorphisms defining a chain homotopy from Sd to id means that:

$$\mathrm{Sd}_n^X - \mathrm{id} = h_{n-1}^X \circ \partial_n + \partial_{n+1} \circ h_n^X. \tag{4.3.6}$$

Again, the singular simplex λ^n has a lot to say. Naturality implies that, for any singular n-simplex $\alpha\colon \Delta^n \to X$, we have:

$$h_n^X(\alpha) = h_n^X(\alpha_\sharp \lambda^n) = \alpha_\sharp h_n^{\Delta^n}(\lambda^n),$$

so everything is determined by choosing $h_n^{\Delta^n}(\lambda^n)$ appropriately; denote this choice by $\tau_{n+1} \coloneqq h_n^{\Delta^n}(\lambda^n)$. After choosing τ_{n+1} we simply define $h_n^X(\alpha) \coloneqq \alpha_\sharp \tau_{n+1}$ for any n-simplex α, and extend linearly. Equation 4.3.5 then holds since:

$$f_\sharp h_n^X(\alpha) \coloneqq f_\sharp \alpha_\sharp \tau_{n+1} = (f \circ \alpha)_\sharp \tau_{n+1} = h_n^Y(f \circ \alpha) = h_n^Y(f_\sharp \alpha).$$

Now, to satisfy the chain homotopy formula (Equation 4.3.6), we need the following of our choices of τ_{n+1}:

$$\partial_{n+1}(\tau_{n+1}) \eqqcolon \partial_{n+1} h_n^{\Delta^n}(\lambda^n) = \mathrm{Sd}_n^{\Delta^n}(\lambda^n) - \lambda^n - h_{n-1}^{\Delta^n}(\partial_n \lambda^n). \tag{4.3.7}$$

Assuming that we have picked our τ_{n+1} so that the above holds then the chain homotopy formula Equation 4.3.6 holds for everything else too in the corresponding degree, since then for a singular n-simplex $\alpha\colon \Delta^n \to X$:

$$\begin{aligned}
& & \partial_{n+1}(h_n^X(\alpha)) \coloneqq \\
(\textit{Definition}): & & \partial_{n+1}\alpha_\sharp \tau_{n+1} = \\
(\alpha_\sharp \textit{ chain map}): & & \alpha_\sharp \partial_{n+1} \tau_{n+1} = \\
(\textit{Equation 4.3.7}): & & \alpha_\sharp(\mathrm{Sd}_n^{\Delta^n}(\lambda^n) - \lambda^n - h_{n-1}^{\Delta^n}(\partial_n \lambda^n)) = \\
(\textit{Naturality of } \mathrm{Sd}_n \text{ and } h_n): & & \mathrm{Sd}_n^X(\alpha_\sharp \lambda^n) - \alpha_\sharp \lambda^n - h_{n-1}^X(\alpha_\sharp \partial_n \lambda^n)) = \\
(\alpha_\sharp \textit{ chain map}): & & \mathrm{Sd}_n^X(\alpha) - \alpha - h_{n-1}^X(\partial_n \alpha).
\end{aligned}$$

Step 4: Proving the existence of the chain homotopy

As a result, we have constructed our chain homotopy so long as we can find chains $\tau_{n+1} \in C_{n+1}(\Delta^n)$ satisfying Equation 4.3.7. For $n = 0$ the equation simply says:

$$\partial_1(\tau_1) = \mathrm{Sd}_0^{\Delta^0}(\lambda^0) - \lambda^0 - h_{-1}^{\Delta^0}(\partial_0 \lambda^0)) = 0,$$

since Sd_0^X is the identity for degree zero. We can take τ_1 as anything here (e.g., $\tau_1 \coloneqq 0$), the boundary of a singular 1-simplex in the one point space is always trivial (c.f., Exercise 3.3.1). So suppose the chains τ_{k+1} have been constructed satisfying Equation 4.3.7 for $k < n$. The right hand side of Equation 4.3.7 is a cycle:

$$\partial_n(\mathrm{Sd}_n^{\Delta^n}(\lambda^n) - \lambda^n - h_{n-1}^{\Delta^n}(\partial_n\lambda^n)) = \mathrm{Sd}_{n-1}^{\Delta^n}(\partial_n\lambda^n) - \partial_n(\lambda^n) - \partial_n h_{n-1}^{\Delta^n}(\partial_n\lambda^n)) =$$
$$\mathrm{Sd}_{n-1}^{\Delta^n}(\partial_n\lambda^n) - \partial_n(\lambda^n) - [\mathrm{Sd}_{n-1}^{\Delta^n}(\partial_n\lambda^n) - \partial_n\lambda^n - h_{n-2}^{\Delta^n}(\partial_{n-1}\partial_n\lambda^n)] = 0.$$

In the first equality we use that $\mathrm{Sd}_n^{\Delta^n}$ is a chain map and in the second we use the chain homotopy formula Equation 4.3.6 one step down where it already holds by the induction hypothesis and the reasoning following Equation 4.3.7. Since Δ^n is contractible, we have that $H_{n+1}(\Delta^n) \cong 0$, so since we have established that the right hand side of Equation 4.3.7 is a cycle it must also be a boundary and our desired τ_{n+1} exists. This completes the induction step and constructs our chain homotopy.

Step 5: Proof of Lemma assuming subdivisions eventually small

Now, suppose the following holds:

Claim: for any chain $\sigma \in C_n(X)$ there exists some $k_\sigma \in \mathbb{N}_0$ with $(\mathrm{Sd}_n^X)^{k_\sigma}(\sigma) \in C_n^{\mathscr{U}}(X)$.

That is, for any chain σ, sufficiently many applications of the subdivision operator produces a $\mathscr{U}$-small chain. We will show that the lemma follows from this. To see that the induced map on homology

$$\iota_*\colon H_n^{\mathscr{U}}(X) \to H_n(X)$$

is onto, let σ be a cycle. Then $\sigma' \coloneqq (\mathrm{Sd}_n^X)^{k_\sigma}(\sigma)$ is a cycle in $C_n^{\mathscr{U}}(X)$ and $\iota_*([\sigma']) = [\sigma]$ in $H_n(X)$. Indeed, σ and $\mathrm{Sd}_n^X(\sigma)$ are homologous by Equation 4.3.6 (their difference is the boundary of $h_n^X(\sigma)$). The subdivision of σ is homologous to its double barycentric subdivision by the same argument. And so on; by inductive application of subdivision, σ and σ' are homologous.

We also need to show that ι_* is injective. So suppose that $\iota_*([\sigma]) = 0$. This means that there exists some $\tau \in C_{n+1}(X)$ with $\partial_{n+1}(\tau) = \sigma$; the issue is that τ need not be $\mathscr{U}$-small. Consider a sufficiently large subdivision of τ though: let $\tau' = (\mathrm{Sd}_{n+1}^X(\tau))^{k_\tau}$. Then

$$(\mathrm{Sd}_n^X)^{k_\tau}\sigma = (\mathrm{Sd}_n)^{k_\tau}\partial_{n+1}\tau = \partial_{n+1}(\mathrm{Sd}_{n+1})^{k_\tau}\tau = \partial_{n+1}(\tau').$$

Now τ' is $\mathscr{U}$-small, so we at least have that some iterated subdivision of σ is the boundary of a $\mathscr{U}$-small chain. But $\sigma' \coloneqq (\mathrm{Sd}_n^X)^{k_\tau}\sigma$ and σ are homologous in $C_*^{\mathscr{U}}(X)$. This follows from the argument above, the difference of the subdivision of σ and itself is the boundary of $h_n^X(\sigma)$. This latter chain is easily seen to be $\mathscr{U}$-small. Indeed, remember that we defined $h_n^X(\alpha) \coloneqq \alpha_\sharp\tau_{n+1}$ for a singular simplex α. This is evidently $\mathscr{U}$-small if α is (this also shows that the subdivision of a $\mathscr{U}$-small chain must still be $\mathscr{U}$-small). By linearity and iterating, σ is homologous via the boundary of a $\mathscr{U}$-small chain to σ', which in turn is homologous via a $\mathscr{U}$-small chain to zero, so σ represents zero in $H_*^{\mathscr{U}}(X)$.

Step 6: Subdivisions eventually small

To conclude the proof we just need to prove the above claim on the existence of $k_\sigma \in \mathbb{N}_0$. The cover $\mathscr{U}$ of X gives a cover $\mathscr{U}'$ of Δ^n by taking pre-images, which by continuity still has interiors covering Δ^n. Since $\mathrm{Sd}_n^X(\alpha) := \alpha_\sharp \mathrm{Sd}_n^{\Delta^n}(\lambda^n)$ for a singular simplex α, we see that it is sufficient to show that for any such covering $\mathscr{U}'$ of Δ^n, we have that $(\mathrm{Sd}_n^{\Delta^n})^k(\lambda^n)$ is $\mathscr{U}'$-small for sufficiently large k. Call $\mathscr{U}'$ instead $\mathscr{U}$ from here on.

It is sufficient to show that the supports of the singular simplices of the sum defining $(\mathrm{Sd}_n^{\Delta^n})^k(\lambda^n)$ become arbitrarily small in radius, because of the following:

Lemma 4.3.2 (Lebesgue Covering Lemma)**.** *Let (Y, d) be a compact metric space and $\mathscr{U}$ be a collection of subsets with interiors covering Y. Then there exists some number $L(\mathscr{U})$ satisfying the following: whenever $S \subseteq Y$ is a subset of at most $L(\mathscr{U})$ in radius (i.e., $d(x, y) \leqslant L(\mathscr{U})$ for all x, $y \in S$) then $S \subseteq U$ for some $U \in \mathscr{U}$.*

Exercise 4.3.2. If you don't remember or haven't seen it, give a proof of this lemma (it's fun).

Clearly there's nothing to do in degree zero (0-chains are always $\mathscr{U}$-small). As a result of the above lemma, we now need to show that, for any $r > 0$, there always exists some k_r with the singular simplices defining $(\mathrm{Sd}_n^{\Delta^n})^{k_r}(\lambda^n)$ of radius of support at most r.

So suppose by induction that this has been proved for $k < n$. Going right back to the definition of the subdivision of λ^n, defined through a cone construction, we see that it is a $\mathbb{Z}$-sum of affine maps $\Delta^n \to \Delta^n$. To see that iterates of these affine maps eventually have arbitrarily small support, it is enough to show that they are contractions; that is, for each such $\alpha\colon \Delta^n \to \Delta^n$, we want to show that $|\alpha(y) - \alpha(x)| < R|y - x|$ for some $R < 1$. It is not hard to see that this will be the case when each α maps vertices of Δ^n closer together.

The map $\restriction_j$ is an isometry, so by induction distinct vertices of a term of $\mathrm{Sd}_n^{\Delta^n}(\lambda^n)$ which lie in a boundary face of Δ^n are closer to each other than in the standard simplex. The only other vertex, by the definition of this subdivision, is sent to the barycentre of Δ^n. It is easily checked that the barycentre is closer to points of any given face of Δ^n than any two distinct vertices of Δ^n are. So the vertices of the terms of the subdivision are mapped closer together, and being affine that means that each is an affine contraction. So for sufficiently large iterates the singular simplices have diameter smaller than the Lebesgue covering number $L(\mathscr{U})$ and the proof is complete.

Remark 4.3.1. We proved that the inclusion map induces isomorphisms on homology, but in fact more is true: $\iota\colon C_*^{\mathscr{U}}(X) \hookrightarrow C_*(X)$ is a chain homotopy equivalence, that is, there is a chain map ρ in the reverse direction with $\iota\rho$ and $\rho\iota$ chain homotopic to the identities. We will not need this fact, but you may like to prove it, see Homework 4.

4.4 Relative homology and quotients

Excision allows us to link together relative homology and homology of quotient spaces. Consider a pair (X, A) of a topological space X and a subspace $A \subseteq X$ with the following property: we assume that A is non-empty, closed and a deformation retract of some neighbourhood $U \supseteq A$ in X. Being a **deformation retract** means that there exists a *deformation retraction* of U to A, which is a map $F\colon U \times I \to U$ with $F(u, 0) = u$, $F(u, 1) \in A$ and $F(a, t) = a$ for all $u \in U$, $a \in A$ and $t \in I$. We then call (X, A) a **good pair**[1]. Most pairs of spaces that you come across will be good pairs, the following gives an important example:

Exercise 4.4.1. Show that for a CW complex $X^\bullet$ the pair (X^k, X^{k-1}) is a good pair.

Example 4.4.1. As an example of something which is *not* a good pair, consider $(\mathbb{R}, A)$ where

$$A = \{0\} \cup \{1/n \mid n \in \mathbb{N}\}.$$

Exercise 4.4.2. Prove that this is not a good pair: consider what would happen for the deformation retraction near 0.

Another example of a pair which is not good is $(\mathbb{R}^n, S)$ where S is the famous 'Topologist's Sine Curve'. The case for no deformation retraction here is to do with S not being 'locally connected'; try to work out the details.

Lemma 4.4.1. *Let (X, A) be a good pair and $q\colon (X, A) \to (X/A, A/A)$ be the quotient map as a map of pairs. Then the induced map of q induces an isomorphism on relative homology:*

$$q_*\colon H_*(X, A) \xrightarrow[q_*]{\cong} H_*(X/A, A/A).$$

Proof. Consider the triple (X, U, A) where $U \supseteq A$ is a neighbourhood of A which deformation retracts to it. The quotient map above fits into the following commutative diagram:

$$\begin{array}{ccccc}
H_*(X, A) & \longrightarrow & H_*(X, U) & \longleftarrow & H_*(X - A, U - A) \\
\downarrow & & \downarrow & & \downarrow \\
H_*(X/A, A/A) & \longrightarrow & H_*(X/A, U/A) & \longleftarrow & H_*(X/A - A/A, U/A - A/A)
\end{array}$$

[1] If you are browsing other sources the more high-tech jargon to look out for here is of a *cofibration*.

The horizontal maps are just given from the corresponding inclusions of pairs of spaces, and the vertical ones are induced by the quotient map. Our goal is to prove that the left-hand vertical arrow is an isomorphism.

By assumption, U deformation retracts onto the subspace A, which induces a deformation retraction of U/A to A/A in the quotient space X/A. So by Lemma 4.1.1 the left-hand horizontal maps are isomorphisms. The right-hand horizontal maps are isomorphisms by excision, Theorem 4.3.1.

Since the horizontal arrows are isomorphisms, the theorem follows so long as one of the vertical maps is also an isomorphism. And indeed the right-hand arrow must be, since it is induced by the restriction of the quotient map as the following map of pairs:

$$q\colon (X - A, U - A) \to (X/A - A/A, U/A - A/A).$$

This map is a homeomorphism of pairs as taking the quotient by A does not identify points; obviously a point $[x] \in X/A$ with $x \notin A$ is represented only by x. So the obvious inverse to the map is $[x] \mapsto x$ for $x \in (X/A - A/A)$. Continuity follows from A being closed, which is easily checked. □

The subspace A/A of X/A is just a single point, so following Remark 4.2.2 the relative homology $H_*(X/A, A/A)$ is just the reduced homology $\widetilde{H}_*(X/A)$. This observation, combined with the above lemma and the LES of the pair gives the following **long exact sequence of a quotient**:

Theorem 4.4.1. *Let (X, A) be a good pair. Then we have a LES*

$$\cdots \xrightarrow{q_*} \widetilde{H}_{k+1}(X/A) \xrightarrow{\partial_*}$$
$$\widetilde{H}_k(A) \xrightarrow{i_*} \widetilde{H}_k(X) \xrightarrow{q_*} \widetilde{H}_k(X/A) \xrightarrow{\partial_*}$$
$$\widetilde{H}_{k-1}(A) \xrightarrow{i_*} \widetilde{H}_{k-1}(X) \xrightarrow{q_*} \widetilde{H}_{k-1}(X/A) \xrightarrow{\partial_*}$$
$$\widetilde{H}_{k-2}(A) \xrightarrow{i_*} \cdots$$

Here, $i\colon A \hookrightarrow X$ is the inclusion and $q\colon X \to X/A$ is the quotient map.

Proof. Just apply the (reduced) LES of a pair to (X, A) and replace the relative terms of $H_*(X, A)$ by $\widetilde{H}_*(X/A)$ using the Lemma 4.4.1.

Although it may seem pedantic, technically we're not quite done because we want to equate q_* with the map coming from the (reduced) LES of the pair. We have the following commutative square:

$$\begin{array}{ccc} \widetilde{H}_*(X) & \xrightarrow{q'_*} & H_*(X,A) \\ \downarrow{\scriptstyle q_*} & & \downarrow{\scriptstyle l} \\ \widetilde{H}_*(X/A) & \xrightarrow{r} & H_*(X/A, A/A) \end{array}$$

The map q'_* is induced by the inclusion $(X, \varnothing) \hookrightarrow (X, A)$ coming from the (reduced) LES of the pair. The map l is the isomorphism of Lemma 4.4.1. The map r is the isomorphism of reduced homology with homology relative to a point discussed in Remark 4.2.2. Thus when we replace $H_*(X, A)$ in the LES of the pair by $\widetilde{H}_*(X)$ we can swap the map q'_* with q_* via the isomorphism $r^{-1} \circ l$ identifying the two homologies. □

4.4.1 Some applications of the LES of a quotient

Homology of spheres

Theorem 4.4.2. *Let $n \in \mathbb{N}_0$. The reduced singular homology of the n-sphere S^n is:*

$$\widetilde{H}_k(S^n) \cong \begin{cases} \mathbb{Z} & \text{for } k = n; \\ 0 & \text{otherwise.} \end{cases}$$

Proof. We may compute the singular homology of the two point space S^0 directly (see Exercise 3.3.1). So the theorem holds for $n = 0$. We now proceed by induction. We have the good pair (D^n, S^{n-1}) of the n-disc and its boundary $(n-1)$-sphere. The quotient is $D^n/S^{n-1} \cong S^n$. Apply the LES from Theorem 4.4.1. Since D^n is contractible its reduced homology is trivial, so the terms involving homologies of spheres occur between successive trivial groups:

$$0 \xrightarrow{q_*} \widetilde{H}_k(S^n) \xrightarrow{\partial_*} \widetilde{H}_{k-1}(S^{n-1}) \xrightarrow{i_*} 0.$$

Hence $\widetilde{H}_k(S^n) \cong \widetilde{H}_{k-1}(S^{n-1})$ for all $k \in \mathbb{Z}$ and the result follows by induction. □

This agrees with our earlier simplicial and cellular calculations (see Exercise 3.2.2 and Example 3.4.1).

Homology of suspensions

Recall (Section 1.2) the definition of the suspension ΣX of a space X (take the product of X with I, and pinch the top and bottom copies of X in $X \times [0, 1]$ to points).

Theorem 4.4.3. *For any space X we have that $\widetilde{H}_{k+1}(\Sigma X) \cong \widetilde{H}_k(X)$.*

Proof. Take the closed subspace $X \times \{0\} \subseteq CX$, i.e., the base of the cone CX, which is a homeomorphic copy of X. It isn't hard to see that this is a good pair with quotient homeomorphic to ΣX. The cone CX of X is contractible (Exercise 1.2.6) and so has reduced homology that of the one point space, which is trivial in all degrees. The result follows from the LES of the quotient. □

Note that since $\Sigma S^{n-1} \cong S^n$ (Example 1.2.2) this recovers the homology of the spheres as above (really the proof here is just an extension of that argument, where D^n played the rôle of CS^{n-1}).

Homology of wedge spaces

Take a collection of pointed spaces (X_α, x_α) (so $x_\alpha \in X_\alpha$ for each α). We define their **wedge sum** $\bigvee_\alpha X_\alpha$ as the quotient of the disjoint union $\coprod_\alpha X_\alpha$ by the disjoint union of their base points $\coprod_\alpha \{x_\alpha\}$. That is, we join the spaces X_α together at a single point. Of course we have inclusions $i_\alpha \colon X_\alpha \hookrightarrow \bigvee_\alpha X_\alpha$ for each of the spaces defined by the quotient map.

Theorem 4.4.4. *Suppose that each $(X_\alpha, \{x_\alpha\})$ above is a good pair. The inclusion maps induce an isomorphism*

$$\bigoplus_\alpha (i_\alpha)_* \colon \bigoplus_\alpha \widetilde{H}_*(X_\alpha) \xrightarrow{\cong} \widetilde{H}_*(\bigvee_\alpha X_\alpha).$$

Proof. It is easy to see that since each $(X_\alpha, \{x_\alpha\})$ is good, so is $(\coprod_\alpha X_\alpha, \coprod_\alpha \{x_\alpha\})$. So we apply the LES of the quotient of Theorem 4.4.1.

$$\begin{array}{l}
\cdots \xrightarrow{q_*} \widetilde{H}_{k+1}(\bigvee_\alpha X_\alpha) \xrightarrow{\partial_*} \\
\widetilde{H}_k(\coprod_\alpha \{x_\alpha\}) \xrightarrow{i_*} \widetilde{H}_k(\coprod_\alpha X_\alpha) \xrightarrow{q_*} \widetilde{H}_k(\bigvee_\alpha X_\alpha) \xrightarrow{\partial_*} \\
\widetilde{H}_{k-1}(\coprod_\alpha \{x_\alpha\}) \xrightarrow{i_*} \widetilde{H}_{k-1}(\coprod_\alpha X_\alpha) \xrightarrow{q_*} \widetilde{H}_{k-1}(\bigvee_\alpha X_\alpha) \xrightarrow{\partial_*} \\
\widetilde{H}_{k-2}(\coprod_\alpha \{x_\alpha\}) \xrightarrow{i_*} \cdots
\end{array}$$

For $k \neq 0$ we may replace the terms $\widetilde{H}_k(\coprod_\alpha X_\alpha)$ with $\bigoplus_\alpha \widetilde{H}_k(X_\alpha)$, see Exercise 3.3.3; it is easy to check under this correspondence that q_* translates to the homomorphism in the statement of the theorem. We can work out the reduced homology of the disjoint union of points directly (Exercise 3.3.1). This has trivial homology in positive degrees. Moreover, we can show directly that i_* is injective in degree zero. Putting this information into the diagram, with a small extra adjustment in degree zero, the result follows. $\square$

Application: mapping cones

Let $f\colon X \to Y$. The **mapping cylinder of** f is the space

$$M_f \coloneqq \big((X \times I) \coprod Y\big)/\sim$$

where the equivalence relation defining the quotient identifies points $(x, 0) \sim f(x)$ for all $x \in X$. That is, we glue the bottom of the 'cylinder' $X \times I$ to a disjoint copy of Y via the map f.

Remark 4.4.1. The inclusion i of X into the top of the cylinder is a 'cofibration' (in the language here, the corresponding subspace and the mapping cylinder are a good pair; **Exercise:** why?). There is a map f' which pushes the cylinder down onto Y (so it leaves Y fixed and sends (x, t) to $(x, 0) \sim f(x)$). This map is a homotopy equivalence (**Exercise:** why?). So this construction shows that every map f can be factored as $f' \circ i$ where i is a so-called cofibration (similar to an inclusion for a good pair) and f' is a homotopy equivalence.

The **mapping cone** C_f is the result of pinching the top of the mapping cylinder to a point, so it is the quotient space

$$C_f \coloneqq M_f/\sim$$

where $(x, 1) \sim (x', 1)$ for all x, $x' \in X$. It is also sometimes called the 'homotopy cofibre'.

Example 4.4.2. Let $X = Y$ and $f = \mathrm{id}_X$. Then $C_f = CX$. So the mapping cone is a generalisation of the usual cone of a space.

Example 4.4.3. Suppose that $X = S^n$ and $f\colon S^n \to Y$. Then C_f is the space formed by attaching an n-cell to Y, as is done in constructing a CW complex.

Exercise 4.4.3. Show that if $f\colon A \to X$ is an inclusion then its mapping cone C_f satisfies $\widetilde{H}_n(C_f) \cong H_n(X, A)$. *Hint: consider the special subspace of the cone of A sitting inside of C_f.*

This example indicates that taking the mapping cone of an inclusion is the 'homotopy theoretic friendly way of taking the quotient'. In fact, it turns out that if (X, A) is a good pair then $C_f \simeq X/A$.

Theorem 4.4.5. *Let $f\colon X \to Y$. We have a LES*

$$
\begin{array}{l}
\cdots \xrightarrow{q_*} \widetilde{H}_{k+1}(C_f) \xrightarrow{\partial_*} \\
\widetilde{H}_k(X) \xrightarrow{f_*} \widetilde{H}_k(Y) \xrightarrow{q_*} \widetilde{H}_k(C_f) \xrightarrow{\partial_*} \\
\widetilde{H}_{k-1}(X) \xrightarrow{f_*} \widetilde{H}_{k-1}(Y) \xrightarrow{q_*} \widetilde{H}_{k-1}(C_f) \xrightarrow{\partial_*} \\
\widetilde{H}_{k-2}(X) \xrightarrow{f_*} \cdots
\end{array}
$$

Hence the induced map f_* is an isomorphism on homology if and only if the reduced homology of C_f is trivial.

Proof. Just apply everything that was established above to the pair $(M_f, X \times \{1\})$. □

Application: invariance of dimension

Theorem 4.4.6. *Let $U \subseteq \mathbb{R}^m$ and $V \subseteq \mathbb{R}^n$ be non-empty open sets. If U and V are homeomorphic, then $m = n$. In particular, $\mathbb{R}^m \cong \mathbb{R}^n$ if and only if $m = n$.*

Proof. Let $x \in U$ be any point. We claim that

$$H_k(U, U - \{x\}) \cong \begin{cases} \mathbb{Z} & \text{for } k = n; \\ 0 & \text{otherwise.} \end{cases}$$

The analogous thing holds for V. If this is the case then clearly U and V can only be homeomorphic if $m = n$, since a homeomorphism h induces a homeomorphism of pairs $(U, U - \{x\}) \to (V, V - \{h(x)\})$ and so an isomorphism of relative homologies.

To prove the above claim, first note that $\widetilde{H}_*(U, U - \{x\}) \cong \widetilde{H}_*(\mathbb{R}^m, \mathbb{R}^m - \{x\})$ by excision, so we may as well work with the latter. Since $\mathbb{R}^m$ is contractible, the LES of the pair implies that $H_k(\mathbb{R}^m, \mathbb{R}^m - \{x\}) \cong \widetilde{H}_{k-1}(\mathbb{R}^m - \{x\})$. The space $\mathbb{R}^m - \{x\}$ is homotopy equivalent to the sphere S^{m-1}, so the result now follows from our previous calculation (Theorem 4.4.2) of the homology of the sphere. □

Whilst on the topic, an alternative way of proving this theorem is as a consequence of *invariance of domain*[2]:

[2] It may be better to call this theorem "invariance of openness". Domain here comes from an older term for an open subset of $\mathbb{R}^n$. It says, in particular, that two homeomorphic subsets of $\mathbb{R}^n$ are either both open or both not open. The statement of the theorem could be read as 'an embedding of an open set of $\mathbb{R}^n$ into $\mathbb{R}^n$ is an open map'; this need not be true replacing $\mathbb{R}^n$ with other spaces.

Theorem 4.4.7. *Let $U \subseteq \mathbb{R}^n$ be open and $f\colon U \to \mathbb{R}^n$ be continuous and injective. Then the image $f(U)$ is also open in $\mathbb{R}^n$.*

See [Hat, Theorem 2B.3] for a proof of this from the Jordan–Brouwer Separation Theorem, which we we will see later on.

4.5 Equating cellular and singular homology

Throughout this section, let $X^\bullet$ be some fixed CW complex. The following three lemmas establish some basic properties on the homology groups associated to the skeleta of a CW complex, which we then combine them for a definition of the cellular chain complex and proof that its homology agrees with the singular homology of X.

Lemma 4.5.1. *The relative homology groups of successive pairs of the skeleta are as follows:*

$$H_k(X^n, X^{n-1}) \cong \begin{cases} \bigoplus_{n\text{-cells}} \mathbb{Z} & \text{for } k = n; \\ 0 & \text{otherwise.} \end{cases}$$

Proof. We have that (X^n, X^{n-1}) is a good pair (see Exercise 4.4.1) so its relative homology is that of the reduced homology of the quotient by Lemma 4.4.1. The quotient space is a wedge of spheres, one for each n-cell, so the result follows from our calculation of the homology of spheres (Theorem 4.4.2) and the homology of a wedge of spaces (Theorem 4.4.4). □

Lemma 4.5.2. *$H_k(X^n)$ is trivial for $k > n$.*

Proof. We certainly know this is true for X^0, which is just a disjoint union of points. Assume then, for induction, that the lemma holds up to the skeleton X^{n-1}. Apply the LES of the pair (X^n, X^{n-1}):

$$\cdots \xrightarrow{\partial_*} H_k(X^{n-1}) \xrightarrow{\iota_*} H_k(X^n) \xrightarrow{q_*} H_k(X^n, X^{n-1}) \xrightarrow{\partial_*} \cdots.$$

If $k > n$ we know the right-hand term is trivial by the previous lemma, and the left-hand term is trivial by the induction assumption. So the result holds for X^n too. □

Lemma 4.5.3. *Consider the inclusion $i\colon X^n \hookrightarrow X$ and its induced map*

$$i_*\colon H_k(X^n) \to H_k(X).$$

Then i_* here is an isomorphism if $k < n$ and is surjective if $k = n$.

Proof. Consider a LES like that of the previous proof:

$$\cdots \xrightarrow{q_*} H_{k+1}(X^{n+1}, X^n) \xrightarrow{\partial_*} H_k(X^n) \xrightarrow{\iota_*} H_k(X^{n+1}) \xrightarrow{q_*} H_k(X^{n+1}, X^n) \xrightarrow{\partial_*} \cdots$$

If $k \leqslant n$ then the right-hand term is trivial so ι_* is surjective. If $k < n$ then the left-hand term is trivial so ι_* is injective. Now, for large $N \in \mathbb{N}$ we can express the induced map $i_* \colon H_k(X^n) \to H_k(X^N)$ as the composition of maps ι_* above, as the composition:

$$H_k(X^n) \to H_k(X^{n+1}) \to H_k(X^{n+2}) \to \cdots \to H_k(X^N).$$

So from the above this composition is an isomorphism for $k < n$ and surjective for $k = n$ (not necessarily injective, because of the first map). If $X^\bullet$ is finite dimensional then $X = X^N$ for some N and we are done.

For the infinite dimensional case one can prove directly that any chain $\sigma \in C_k(X)$ must be supported on some finite dimensional skeleta. This is because it is a finite sum of singular k-simplices $\Delta^k \to X$ and so have compact image, and it turns out that any compact subset of a CW complex meets only finitely many cells, so is contained in some X^N. As an **Exercise** you may like to fill in the details here: to prove that a compact subset of a CW complex intersects only finitely many (open) cells, and then use this to tie up the proof of the lemma. □

Consider the LES of the pair (X^n, X^{n-1}), and similarly for (X^{n+1}, X^n). The term $H_n(X^n)$ appears in both; line up the two sequences as follows:

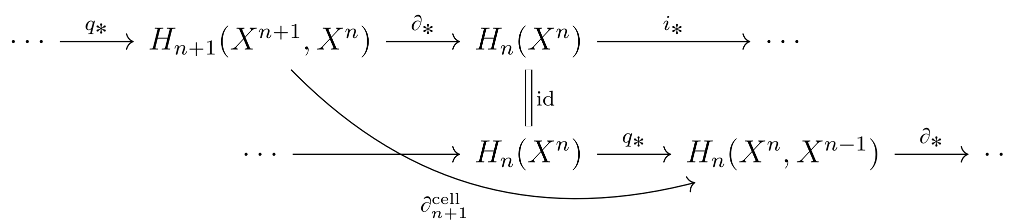

The new arrow $\partial^{\text{cell}}_{n+1} \colon H_{n+1}(X^{n+1}, X^n) \to H_n(X^n, X^{n-1})$ is simply defined as the composition $q_* \circ \partial_*$. This defines the cellular chain complex (and it really does defines it this time!):

Definition 4.5.1. The **cellular chain complex** is the chain complex

$$\cdots \xrightarrow{\partial^{\text{cell}}_3} H_2(X^2, X^1) \xrightarrow{\partial^{\text{cell}}_2} H_1(X^1, X^0) \xrightarrow{\partial^{\text{cell}}_1} H_0(X^0) \to 0$$

where the boundary maps $\partial^{\text{cell}}_n \colon H_n(X^n, X^{n-1}) \to H_{n-1}(X^{n-1}, X^{n-2})$ are defined as above. The homology of this complex is called the **cellular homology of $X^\bullet$.**

Note that the chain groups above in degree k are just the free Abelian groups with generators in correspondence to the k-cells by Lemma 4.5.1, just as was the case for our earlier description in Definition 3.4.1. It is easily seen that the above is a chain complex: if you consider $\partial_{n+1}^{\text{cell}} \circ \partial_n^{\text{cell}}$ then you get the composition $(q_* \circ \partial_*) \circ (q_* \circ \partial_*)$. This is slightly sloppy notation, because the two q_* maps and two ∂_* maps occur in different snake diagrams. Ignoring that notational issue, rewriting as $q_* \circ (\partial_* \circ q_*) \circ \partial_*$, note that the middle composition *does* appear as two consecutive maps of a snake diagram, so must be zero.

We now have the tools to show that the cellular and singular homology agree:

Theorem 4.5.1. *For a CW complex $X^\bullet$ its singular homology is isomorphic to its cellular homology.*

Proof. Let's add some missing terms of the LES to the diagram above in blue (we shall, just here, also add indices for the connecting maps of the snake diagrams):

$$\begin{array}{ccccccc} H_{n+1}(X^{n+1}, X^n) & \xrightarrow{\partial_{n+1}} & H_n(X^n) & \xrightarrow{i_*} & H_n(X^{n+1}) & \xrightarrow{q_*} & 0 \\ & & \Vert\,\text{id} & & & & \\ & & H_n(X^n) & \longrightarrow & H_n(X^n, X^{n-1}) & \longrightarrow & \cdots \end{array}$$

The '0' term is there by the fact that $H_n(X^{n+1}, X^n) \cong 0$ (Lemma 4.5.1). Moreover, $H_n(X^{n+1}) \cong H_n(X)$ by Lemma 4.5.3, the homology group we want to find! It follows from exactness that:

$$H_n(X) \cong H_n(X^n)/\operatorname{im}(\partial_{n+1}). \tag{4.5.1}$$

Now let's investigate the next row down, adding extra blue terms:

$$\begin{array}{ccccccc} H_{n+1}(X^{n+1}, X^n) & \xrightarrow{\partial_{n+1}} & H_n(X^n) & \xrightarrow{i_*} & H_n(X^{n+1}) & \longrightarrow & 0 \\ & & \Vert\,\text{id} & & & & \\ 0 & \xrightarrow{i_*} & H_n(X^n) & \xrightarrow{q_*} & H_n(X^n, X^{n-1}) & \xrightarrow{\partial_n} & H_{n-1}(X^{n-1}) \end{array}$$

The '0' is there because $H_n(X^{n-1}) \cong 0$ (Lemma 4.5.2). Hence q_* is injective here, so it maps $H_n(X^n)$ and $\operatorname{im}(\partial_{n+1})$ isomorphically onto their images. By Equation 4.5.1: $H_n(X) \cong q_*(H_n(X^n))/q_*(\operatorname{im}(\partial_{n+1}))$. Now, $q_*(\operatorname{im}(\partial_{n+1})) = \operatorname{im}(q_* \circ \partial_{n+1}) = \operatorname{im}(\partial_{n+1}^{\text{cell}})$ just by the definition of the cellular boundary map. On the other hand, $q_*(H_n(X^n)) = \operatorname{im}(q_*) = \ker(\partial_n)$ by exactness, so:

$$H_n(X) \cong \ker(\partial_n)/\operatorname{im}(\partial_{n+1}^{\text{cell}}) \tag{4.5.2}$$

Finally, looking at the next row down with new blue terms:

$$
\begin{array}{ccccccccc}
H_{n+1}(X^{n+1}, X^n) & \xrightarrow{\partial_{n+1}} & H_n(X^n) & \xrightarrow{i_*} & H_n(X^{n+1}) & \longrightarrow & 0 & & \\
 & & \Big\| \text{id} & & & & & & \\
0 & \xrightarrow{i_*} & H_n(X^n) & \xrightarrow{q_*} & H_n(X^n, X^{n-1}) & \xrightarrow{\partial_n} & H_{n-1}(X^{n-1}) & & \\
 & & & & & & \Big\| \text{id} & & \\
 & & & & 0 & \xrightarrow{i_*} & H_{n-1}(X^{n-1}) & \xrightarrow{q_*} &
\end{array}
$$

The '0' comes from the fact that $H_{n-1}(X^{n-2}) \cong 0$ (Lemma 4.5.2). So the q_* at the bottom of the picture is injective and $\ker(\partial_n) = \ker(q_* \circ \partial_n) = \ker(\partial_n^{\text{cell}})$. Combining with Equation 4.5.2 completes the proof. □

4.5.1 What is the cellular boundary map, really?

The above description of the cellular boundary map is a little abstract. It has a more digestible form as the 'cellular boundary formula', which shall be stated below, given in terms of the degree of a map on an n-sphere. We shall not give full proofs in this section, see [Hat, Section 2.2] for details.

Let $f\colon S^n \to S^n$. Then f induces a homomorphism f_* from $\widetilde{H}_*(S^n)$ to itself. We determined (Theorem 4.4.2) that the homology is concentrated in degree n where it is isomorphic to $\mathbb{Z}$. So in this degree $f_*(\ell) = d(f) \cdot \ell$ for some unique number $d(f) \in \mathbb{Z}$ called the **degree of** f.

For a map $f\colon S^1 \to S^1$ the degree $d(f)$ is just the *winding number* of f, which you've likely met before. For a map $f\colon S^2 \to S^2$, I like to imagine the map as given by taking a bin bag (the torn domain S^2), putting the target sphere into it, wrapping the bag around it a bit (allowing the bag to pass through itself), and then pulling the bag tight and then fusing the hole of the bag shut. For a map which is not too weird, locally the sheet of plastic passes flat in layers over a typical point x on the target sphere, either keeping or reversing orientations. The number of times we lay the plastic down on such a point, signed according to orientation, is the degree of the map. For more rigorous details, see [Hat, Section 2.2].

One can use the degree, and some basic properties of it, to prove the famous result that there is a non-zero vector-field on S^n if and only if n is odd. It can also be used to describe the cellular boundary map. Let e_α be a k-cell of $X^\bullet$, attached via the characteristic map $\sigma_\alpha\colon D^k \to X$. Take a $(k-1)$-cell e_β of $X^\bullet$. Consider the following composition:

$$S^{k-1} \xrightarrow{\sigma_\alpha} X^{k-1} \xrightarrow{q} X^{k-1}/(X^{k-1} - e_\beta) \cong S^{k-1}. \tag{4.5.3}$$

The first map is the restriction of the characteristic map to the boundary sphere, which maps into the $(k-1)$-skeleton, and the second is the quotient of X^{k-1} collapsing $(X^{k-1} -$

e_β) to a point. This space is homeomorphic to the $(k-1)$-sphere (a homeomorphism is determined by σ_β). The degree of this map is precisely the coefficient that the cell e_β receives when taking the cellular boundary of e_α. That is, if we refer to the map above as $f_{\alpha\beta}$ then

$$\partial_k^{\text{cell}}(e_\alpha) = \sum_{k\text{-cells } e_\beta} d(f_{\alpha\beta})e_\beta,$$

which is extended linearly to define the boundary map on an arbitrary cellular chain. This is called the **cellular boundary formula**.

It is important to note that the signs of the coefficients of the cellular boundary formula depend on the orientations of the cells, which are set by the choices of characteristic maps attaching discs.

We've already looked at some basic examples, but just for a refresher let's do another couple:

Example 4.5.1. Take the closed surface Σ_g of genus g (with $g \geqslant 1$, the 'torus with g holes'). There is a nice CW decomposition of this surface given by identifying the sides of a $4g$-gon. This results in a CW decomposition of one 0-cell, $2g$ 1-cells and one 2-cell, so the chain complex looks like:

$$\cdots \to 0 \to \mathbb{Z} \xrightarrow{\partial_2} \mathbb{Z}^{2g} \xrightarrow{\partial_1} \mathbb{Z} \to 0$$

We know that $H_0(\Sigma_g) \cong \mathbb{Z}$ (it is a path-connected space), so ∂_1 is the trivial homomorphism. Moreover, if you draw the CW decomposition you have that every edge contributes once in one direction to the boundary of the 2-cell, and then occurs once more pointing in the opposite direction. This means that the contributions cancel and $\partial_2 = 0$ as well. Hence:

$$H_k(\Sigma_g) \cong \begin{cases} \mathbb{Z} & \text{for } k = 0; \\ \mathbb{Z}^{2g} & \text{for } k = 1; \\ \mathbb{Z} & \text{for } k = 2; \\ 0 & \text{otherwise.} \end{cases}$$

Example 4.5.2. The space $\mathbb{R}P^2$ is obtained from the disc D^2 by identifying the points of the boundary circle under the antipodal map $x \mapsto -x$. You can give S^1 a CW decomposition into one 0-cell and one 1-cell, giving the following chain complex:

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

As in the above example we know that $H_0(\mathbb{R}P^n) \cong \mathbb{Z}$ so $\partial_1 = 0$. To work out ∂_2, note that D^2 is attached to the 1-skeleton of $\mathbb{R}P^2$ by winding twice around the 1-cell, which with the 0-cell sits as a copy of $\mathbb{R}P^1 \cong S^1$ in $\mathbb{R}P^2$. Hence $\partial_2(x) = \pm 2x$, depending on choices of orientations, so

$$H_k(\mathbb{R}P^2) \cong \begin{cases} \mathbb{Z} & \text{for } k = 0; \\ \mathbb{Z}/2 & \text{for } k = 1; \\ 0 & \text{otherwise.} \end{cases}$$

More generally, $\mathbb{R}P^n$ has the following homology groups:

$$H_k(\mathbb{R}P^n) \cong \begin{cases} \mathbb{Z} & \text{for } k = 0 \\ \mathbb{Z}/2 & \text{for } 0 < k < n \text{ and } k \text{ odd;} \\ \mathbb{Z} & \text{for } k = n \text{ if } k \text{ is odd;} \\ 0 & \text{otherwise.} \end{cases}$$

Homework Sheet 4 will lead you through the computation of this.

4.5.2 Comparing simplicial and cellular homology

Remember that a triangulation of a space X by $\mathcal{K}$ also defines a CW decomposition of X. The simplicial chain complex that you get from $\mathcal{K}$ is isomorphic to the cellular chain complex.

It is clear that the chain groups of the two are isomorphic (each are freely generated by the simplicial cells). To prove that the cellular boundary maps correspond to the simplicial boundary maps one needs to explicitly work out generators for the singular homology groups $H_k(\Delta^k, \partial\Delta^k) \cong \mathbb{Z}$. It shouldn't surprise you (and it is proved in Homework 4) that this group is generated by the singular simplex $\mathrm{id}\colon \Delta^n \to \Delta^n$. So for a simplex $\alpha \in \mathcal{K}$, with corresponding simplex Δ^k_α of the geometric realisation, one has that $H_k(\Delta^k_\alpha, \partial\Delta^k_\alpha) \cong \mathbb{Z}$ is generated by the map $\sigma_\alpha\colon \Delta^k \to \Delta^k_\alpha$ used to attached the k-simplex (which comes up in the explicit construction of $|\mathcal{K}|$). Given a general simplicial chain $\sum n_\alpha \alpha$ this defines the cellular chain $\sum n_\alpha e_\alpha$, and one may check that this defines the required chain isomorphism. Following these ideas through leads to:

Theorem 4.5.2. *For a space X triangulated by $\mathcal{K}$, the simplicial chain complex of $\mathcal{K}$ is isomorphic to the corresponding cellular chain complex, with CW decomposition naturally induced by the triangulation. Hence, for a triangulated space, its simplicial, singular and cellular homology groups all agree.*

4.6 Homotopy invariance of Euler characteristic

Recall from Section 1.2.2 that the Euler characteristic of a CW complex $X^\bullet$ of finitely many cells is given by the sum

$$\chi(X) \coloneqq \sum_{n=0}^{\infty} (-1)^n c_n,$$

where c_n is the number of n-cells in $X^\bullet$. The homotopy invariance of singular homology proves the following:

Theorem 4.6.1. *Let $X^\bullet$ and $Y^\bullet$ be two finite CW complexes. Suppose that $X \simeq Y$. Then $\chi(X) \cong \chi(Y)$.*

Proof. The Euler characteristic of X as defined above is equal to the Euler characteristic of the cellular chain complex of $X^\bullet$, as defined in Section 2.4, and similarly for Y. By Theorem 2.4.1 the alternating sums of ranks of their cellular homology groups must agree with the Euler characteristic. Since the cellular homology groups of $X^\bullet$ and $Y^\bullet$ agree with the singular homology groups (Theorem 4.5.1) and the singular homology groups of homotopy equivalent spaces agree (Corollary 3.3.1) the result follows. $\square$

Example 4.6.1. We already ran through some examples of Euler characteristic calculations in Section 1.2.2. Here's an extra one: we have a CW decomposition of Σ_g of one 0-cell, $2g$ 1-cells and one 2-cell, so $\chi(\Sigma_g) = 2 - 2g$. You may remember that from Geometric Topology II or Topology III, and now you have a proof that this number really is a genuine invariant.

4.7 The Mayer–Vietoris sequence

4.7.1 Cellular version

Let $X^\bullet$ be a CW complex with subcomplexes $A^\bullet$, $B^\bullet$ whose union is all of $X^\bullet$. Their intersection $(A \cap B)^\bullet$ is also sub-complex of A. We can relate the homologies of these spaces:

Theorem 4.7.1 (Cellular Mayer–Vietoris Theorem)**.** *The cellular homologies of X, A, B and $A \cap B$ are related via the following long exact sequence:*

$$\cdots \xrightarrow{\partial_{n+1}} H_n(A \cap B) \xrightarrow{(\eta_*^A, \eta_*^B)} H_n(A) \oplus H_n(B) \xrightarrow{i_*^A - I_*^B} H_n(X) \xrightarrow{\partial_*} H_{n-1}(A \cap B) \xrightarrow{(\eta_*^A, \eta_*^B)} \cdots.$$

Here, the map η_*^A is induced by the inclusion $\eta^A \colon A \cap B \hookrightarrow A$ and i_*^A is induced by the inclusion $i^A \colon A \hookrightarrow X$; the maps η_*^B and i_*^B are defined analogously. The map

∂_* is defined on a homology class $[\sigma] \in H_n(X)$ by first expressing $\sigma = \sigma_A - \sigma_B$, for $\sigma_A \in C_n(A)$ and $\sigma_B \in C_n(B)$, and then setting $\partial_*([\sigma])$ to be the homology class of $\partial_n^{\text{cell}}(\sigma_A) = \partial_n^{\text{cell}}(\sigma_B)$ in $H_{n-1}(A \cap B)$.

Proof. We have the following SES of chain complexes:

$$0 \to C_*(A \cap B) \xrightarrow{(\eta_\sharp^A, \eta_\sharp^B)} C_*(A) \oplus C_*(B) \xrightarrow{i_\sharp^A - i_\sharp^B} C_*(X) \to 0.$$

It is simple to verify that this is indeed a SES. The middle group is the direct sum of two chain complexes, with degree n chain group the direct sum of the two degree n chain groups and boundary defined by the usual one on each summand. It is easily shown in that situation the homology is isomorphic in a natural way to the homology of the individual chain complexes. Applying the Snake Lemma (Lemma 2.3.1), along with Remark 2.3.1 on the connecting map completes the proof. □

4.7.2 Singular version

Theorem 4.7.2 (Singular Mayer–Vietoris Theorem)**.** *Let X be a topological space with subspaces A, $B \subseteq X$ whose interiors cover X. Then the singular homologies of X, A, B and $A \cap B$ are related via the following long exact sequence:*

$$\cdots \xrightarrow{\partial_{n+1}} H_n(A \cap B) \xrightarrow{(\eta_*^A, \eta_*^B)} H_n(A) \oplus H_n(B) \xrightarrow{i_*^A - I_*^B} H_n(X) \xrightarrow{\partial_*} H_{n-1}(A \cap B) \xrightarrow{(\eta_*^A, \eta_*^B)} \cdots .$$

The homomorphisms in this LES have analogous descriptions to those in the cellular version given earlier.

Proof. Let $\mathscr{U} = \{A, B\}$. Consider the following SES:

$$0 \to C_*(A \cap B) \xrightarrow{(\eta_\sharp^A, \eta_\sharp^B)} C_*(A) \oplus C_*(B) \xrightarrow{i_\sharp^A - i_\sharp^B} C_*^{\mathscr{U}}(X) \to 0.$$

Note that we use $\mathscr{U}$-small chains on the right here, which was necessary for surjectivity of the right-hand map (note that the elements of $C_*^{\mathscr{U}}$ are precisely those chains which are sums of chains lying in $C_*(A)$ and chains lying in $C_*(B)$). Again, checking this is a SES is easy. Apply the Snake Lemma and replace the terms $H_n^{\mathscr{U}}(X)$ with $H_n(X)$ by Lemma 4.3.1 to complete the proof. □

Note that there is also a reduced version of the Mayer–Vietoris sequence, given by just replacing each occurrence of homology group with a reduced one.

4.7.3 Some applications

Basic examples

I'll leave it as an exercise to apply the Mayer–Vietoris sequence in these examples:

- S^n from $D^k_+ \cup D^k_-$: the union discs (upper and lower hemispheres) intersecting along an equator homotopy equivalent to S^{n-1}.
- ΣX from $CX_+ \cup CX_-$: the suspension as a union of two copies of the cone of X, intersecting on a subspace homotopy equivalent to X.
- $X \bigvee Y$ from $X \cup Y$: the wedge sum of X and Y, intersecting in a subspace homotopy equivalent to a point (you need to have good pairs for the inclusions of the point into X and Y here);
- $\mathbb{T}^2$ from $(S^1 \times I)_1 \cup (S^1 \times I)_2$: the torus as a union of two annuli (or cylinders), intersecting in the union of two homotopy equivalent S^1 subspaces.
- $\mathbb{K} = \ddot{\mathrm{M}}_1 \cup \ddot{\mathrm{M}}_2$: the Klein bottle as a union of two Möbius bands, intersecting along a cylinder.

Try playing around with some decompositions of spaces into subspaces, perhaps try finding one, say, for $\mathbb{R}P^2$.

No embedding of Klein bottle in $\mathbb{R}^3$

We shall give a proof that the Klein bottle $\mathbb{K}$ does not embed 'nicely' into $\mathbb{R}^3$. Explicitly, we shall prove that there is no injective continuous map $i\colon \mathbb{K}\times I \hookrightarrow \mathbb{R}^3$ from a 'thickened' version[3] of $\mathbb{K}$ into $\mathbb{R}^3$.

?? picture

Denote $\mathbb{K}_\epsilon \coloneqq \mathrm{im}(i)$. This is the thickened version of $\mathbb{K}$ sitting inside $\mathbb{R}^3$. Consider a very slight thickening C of the complement $\mathbb{R}^3 - \mathbb{K}_\epsilon$. We can work out the topology of the intersection $\mathbb{K}_\epsilon \cap C$. The figure gives a model for $\mathbb{K}\times I$ as a 'slab' with identifications along its faces. The intersection of the copy of this in $\mathbb{R}^3$ with C is a small thickening of the front and back square faces with identifications. We can easily show that these identifications give a copy of $\mathbb{T}^2$.

Apply the Mayer–Vietoris sequence to the decomposition of $\mathbb{R}^3$ as the union of $\mathbb{K}_\epsilon$ and C. We know the homology of $\mathbb{R}^3$ (that of a point), of K_ϵ (the same as that of a Klein bottle,

[3]Don't worry too much about this 'thickening' part. This map would exist, say, if there was a *smooth* embedding of the Klein bottle into $\mathbb{R}^3$. Most proofs I've seen on this end up making a similar such simplifying assumption.

in particular $H_1(\mathbb{K}_\epsilon) \cong \mathbb{Z} \oplus \mathbb{Z}/2$) and of the intersection $\mathbb{K}_\epsilon \cap C$ (homotopy equivalent to a torus, in particular $H_1(\mathbb{T}^2) \cong \mathbb{Z}^2$). Here is one portion of that diagram:

$$\cdots \to H_2(\mathbb{R}^3) \to H_1(\mathbb{K}_\epsilon \cap C) \to H_1(\mathbb{K}_\epsilon) \oplus H_1(C) \to H_1(\mathbb{R}^3) \to \cdots .$$

Filling in the information we know:

$$\cdots \to 0 \to \mathbb{Z}^2 \to (\mathbb{Z} \oplus \mathbb{Z}/2) \oplus H_1(C) \to 0 \to \cdots .$$

This implies that $\mathbb{Z}^2 \cong (\mathbb{Z} \oplus \mathbb{Z}/2) \oplus H_1(C)$. But there is no way this can be, since the group on the right has 2-torsion (a non-trivial element x with $2x = 0$) but $\mathbb{Z}^2$ does not. This is a contradiction, so the embedding of $\mathbb{K}$ into $\mathbb{R}^3$ cannot exist.

Exercise 4.7.1. Try a similar proof to show that $\mathbb{R}P^2$ does not nicely embed into $\mathbb{R}^3$.

Application: Jordan Curve Theorem

Let $f\colon S^1 \to \mathbb{R}^2$ be continuous and injective. So $f(S^1)$ is an embedded circle in the plane. *Obviously* it splits the plane into two connected components, right? Yes and no: the claim is true but it is not obvious—the first decent stab at a proof was in 1887 by Camille Jordan—there has been debate on whether or not his proof was sufficiently complete. Either way, this result is still usually referred to as the *Jordan Curve Theorem.* Note that for f a smooth map the claim is not difficult to prove, the issue is that continuous maps can still be fairly wild (c.f., space-filling curves). The Alexander horned sphere demonstrates what kinds of strange behaviours one can see in the higher dimensional generalisation of this theorem, stated below:

Theorem 4.7.3 (Jordan–Brouwer Separation Theorem)**.** *Let $f\colon S^{n-1} \to \mathbb{R}^n$ be continuous and injective, and define $S \coloneqq f(S^{n-1})$. Then $\mathbb{R}^n - S$ consists of precisely two path-connected components, one of which is bounded and the other is unbounded.*

We essentially follow the proof from [Hat] here. To start with we re-frame the question slightly: since we may view S^n as $\mathbb{R}^n$ with an extra point 'added at infinity' (as $S^n \cong D^n/S^{n-1}$), we may consider S as an embedded $(n-1)$-sphere in S^n rather than $\mathbb{R}^n$. The complement $S^n - S$ is open so its connected components are also open in S^n and hence are the same as its path components (it is easy to check that open connected subspaces of S^n are also path-connected). Removing the single point (corresponding to 'infinity') from an open component doesn't change whether or not it is path-connected, so it will suffice to prove that $S^n - S$ has precisely two path-components. This is the same as proving that[4]:

$$\widetilde{H}_0(S^n - S) \cong \mathbb{Z}.$$

[4] This isomorphism actually follows from a more general theorem known as *Alexander duality* which you will meet next term. Alexander duality allows one to relate the homology of a subspace of the sphere to the *cohomology* of its complement, which, by a *universal coefficient theorem*, can be related to the homology of the complement.

This will follow from the below lemma:

Lemma 4.7.1. *Let $f\colon D^d \to S^n$ be injective and continuous, and denote the image by $Q \coloneqq f(D^d)$. Then the reduced homology of its complement is trivial: $\widetilde{H}_k(S^n - Q) \cong 0$ for all k.*

Corollary 4.7.1. *Let $f\colon S^d \to S^k$ be injective and continuous, and $d < n$, and denote the image by $S \coloneqq f(S^d)$. Then the reduced homology of the quotient is:*

$$\widetilde{H}_k(S^n - S) \cong \begin{cases} \mathbb{Z} & \text{for } k = n - d - 1; \\ 0 & \text{otherwise.} \end{cases}$$

Proof of Lemma 4.7.1. The lemma is proved by induction on the dimension d of the embedded disc. For $d = 0$ the result is obvious: Q is a single point of S^n so $S^n - Q \cong \mathbb{R}^n$, which is contractible and thus has trivial reduced homology.

Suppose then that the result is true for discs up to dimension $d - 1$ and let Q be an embedded d-disc in S^n. It is the same to think of it as an embedded cube, as $D^d \cong I^d$, so replace f with a map $f\colon I^d \to S^n$ for $Q = \operatorname{im}(f)$. To prove that $S^n - Q$ has trivial homology, we must show that any cycle $\sigma \in \widetilde{C}_k(S^n - Q)$ is the boundary of some $\tau \in \widetilde{C}_{k+1}(S^n - Q)$.

We know the result is true for the $(d-1)$-dimensional slices $Q[t]$, given by the images of $f \mid (I^{d-1} \times \{t\}) \to S^n$. The chain σ is still a chain of $S^n - Q[t]$, so there exist $(k+1)$-chains τ_t lying in the complement of $Q[t]$ with $\partial(\tau_t) = \sigma$. Each τ_t is a finite sum of singular simplices, so are supported on a closed subset of S^n not intersecting $Q[t]$. By continuity and compactness, $f(I^{d-1} \times (t - \epsilon_t, t + \epsilon_t))$ still lies in the complement of $Q[t]$ for some sufficiently small ϵ_t. Again by compactness, finitely many of the $I^{d-1} \times (t - \epsilon_t, t + \epsilon_t)$ cover the cube I^d, so we can pick $0 = t_0 < t_1 < t_2 < \cdots < t_N = 1$ with each $[t_i, t_{i+1}]$ contained in some $(t - \epsilon_t, t + \epsilon_t)$.

Denote the image of f restricted to $I^{d-1} \times [u, v]$ by $Q[u, v]$. Since σ avoids Q it defines a chain $\sigma \in C_k(S^n - Q[u, v])$ for any $u \leqslant v$. Similarly, by the above construction, for each i there is some $t \in [t_i, t_{i+1}]$ with τ_t avoiding $Q[t_i, t_{i+1}]$. Since $\sigma = \partial(\tau_t)$ we have that σ represents zero in $H_k(S^n - Q[t_i, t_{i+1}])$.

We claim that σ is a boundary in $\widetilde{C}_k(S^n - Q[t_0, t_i])$ for all i; recall that $t_0 = 0$ and $t_N = 1$ so $Q[t_0, t_N] = Q$ so deducing this claim proves the lemma. For $i = 0$ the result is already known to be true by induction, $Q[t_0, t_0] = Q[0]$ (indeed, we defined the boundary already as τ_0). So suppose this holds up to $i - 1$. Consider $A \coloneqq S^n - Q[t_0, t_{i-1}]$ and $B \coloneqq S^n - Q[t_{i-1}, t_i]$. Note that $Q[t_0, t_{i-1}] \cap Q[t_{i-1}, t_i] = Q[t_{i-1}]$ and $Q[t_0, t_{i-1}] \cup Q[t_{i-1}, t_i] = Q[t_0, t_i]$ so:

$$\begin{aligned} A \cup B &= S^n - Q[t_{i-1}] \\ A \cap B &= S^n - Q[t_0, t_i]. \end{aligned}$$

Moreover, it is easy to see that A and B are open subsets of $A \cup B$ (the $Q[u,v]$ are compact in a Hausdorff space, hence closed). Applying the Mayer–Vietoris sequence:

$$\cdots \longrightarrow \widetilde{H}_{k+1}(S^n - Q[t_{i-1}]) \longrightarrow$$
$$\longrightarrow \widetilde{H}_k(S^n - Q[t_0, t_i]) \xrightarrow{(\eta_*^A, \eta_*^B)} \widetilde{H}_k(S^n - Q[t_0, t_{i-1}]) \oplus \widetilde{H}_k(S^n - Q[t_{i-1}, t_i]) \longrightarrow \cdots$$

By our original induction hypothesis, the first term is trivial, so the map (η_*^A, η_*^B) is injective. Apply this to the homology class of σ. We know that its image in each of the two summands is trivial; in the first by induction, and in the second from our construction of the τ_t chains above. Hence σ must represent zero in $\widetilde{H}_k(S^n - Q[t_0, t_i])$ too, which completes the proof. □

Proof of Corollary 4.7.1 from Lemma 4.7.1. As in the above lemma, we prove the result by induction. For $d = 0$ we have that $S^n - S$ is a sphere with two points removed, which is homeomorphic to $\mathbb{R}^n$ with one point removed, which is homotopy equivalent to S^{n-1}. By our calculation of the homologies of spheres, the result holds for $d = 0$.

Suppose then that the result holds up to dimension $d - 1$. The d-sphere is the union of two hemispheres, D_+^d and D_-^d (with final coordinates in $\mathbb{R}^{d+1}$ being $\geqslant 0$ and $\leqslant 0$, respectively), each homeomorphic to discs and intersecting along an equator $S^{d-1} \subset S^d$. Denote $D_+ \coloneqq f(D_+^d)$, $D_- \coloneqq f(D_-^d)$ and $S' \coloneqq f(S^{d-1})$. The subspaces $A \coloneqq S^n - D_+$ and $B \coloneqq S^n - D_-$ give an open cover of $A \cup B = S^n - S'$, with $A \cap B = S^n - S$. Applying the Mayer–Vietoris sequence:

$$\cdots \longrightarrow \widetilde{H}_{k+1}(S^n - D_+) \oplus \widetilde{H}_{k+1}(S^n - D_-) \longrightarrow \widetilde{H}_{k+1}(S^n - S') \longrightarrow$$
$$\longrightarrow \widetilde{H}_k(S^n - S) \longrightarrow \widetilde{H}_k(S^n - D_+) \oplus \widetilde{H}_k(S^n - D_-) \longrightarrow \cdots$$

By the lemma, the middle summand terms here are trivial, hence $\widetilde{H}_k(S^n - S) \cong \widetilde{H}_{k+1}(S^n - S')$. Remember that S' was an embedded sphere one dimension less than that of S, so the result follows by induction. □

A nice corollary of all of this is Brouwer's Invariance of Domain, stated previously and restated here:

Theorem 4.4.7. *Let $U \subseteq \mathbb{R}^n$ be open and $f\colon U \to \mathbb{R}^n$ be continuous and injective. Then the image $f(U)$ is also open in $\mathbb{R}^n$.*

Proof. As for the proof of the Jordan–Brouwer separation theorem, it is the same to compactify $\mathbb{R}^n$ to the n-sphere, and show instead that $f(U)$ is open in S^n. Pick a point of $h(U)$, say $h(x)$ for $x \in U$. We need to show that there is an open subset of S^n contained in $h(U)$ and containing $h(x)$.

Since U is open, we may pick a small enough closed disc D centred at x contained in U. The boundary S of D is an $(n-1)$-sphere. We claim that $f(D-S)$ is open in $f(U)$, from which the result follows.

By the Jordan–Brouwer Separation Theorem, $S^n - f(S)$ has exactly two path-connected components. These must be the obvious candidates: $f(D-S)$ and $S^n - f(D)$. Indeed they are disjoint, cover $S^n - f(S)$ and are path-connected: $f(D-S)$ is path-connected because $D-S$ is, and $S^n - f(D)$ is path-connected because of Lemma 4.7.1. By compactness of S, we have that $S^n - f(S)$ is open in S^n from which it easily follows that the two *path*-connected components of $S^n - f(S)$ are also the connected components. Remember that connected components are closed so, being as there are only two of them, $f(D-S)$ must also be open in $S^n - f(S)$. Since $S^n - f(S)$ is itself open in S^n, we have that $f(D-S)$ is open in S^n, as desired. □

4.8 Further topics

4.8.1 Eilenberg–Steenrod axioms for homology

Usually when using homology we work from its fundamental properties rather than its basic definition. Eilenberg and Steenrod outlined the following important properties of homology:

A homology theory consists of a sequence of functors $H_n(-,-)$ ($n \in \mathbb{Z}$) from the category $\mathbf{Top}^2$ *of pairs of topological spaces to the category* $\mathbf{Ab}$ *of Abelian groups. We write $H_n(X, \varnothing)$ for the non-relative homology $H_n(X)$, and for a map of pairs $f\colon (X, A) \to (Y, B)$ we denote by f_* the application of the functor to f (omitting here the degree). These functors should satisfy:*

—**(Homotopy):** *If f, $g\colon (X, A) \to (Y, B)$ are homotopic as maps of pairs (in the sense of Theorem 4.1.1) then $f_* = g_*$.*

—**(LES of pair):** *There exist natural homomorphisms $\partial_*\colon H_n(X, A) \to H_{n-1}(A)$ (meaning that $f_*\partial_* = \partial_* f_*$ for $f\colon (X, A) \to (Y, B)$) for which we have the LES of* *the pair:*

$$\cdots \xrightarrow{q_*} H_{k+1}(X,A) \xrightarrow{\partial_*}$$

$$H_n(A) \xrightarrow{\iota_*} H_n(X) \xrightarrow{q_*} H_n(X,A) \xrightarrow{\partial_*}$$

$$H_{n-1}(A) \xrightarrow{\iota_*} H_{n-1}(X) \xrightarrow{q_*} H_{n-1}(X,A) \xrightarrow{\partial_*}$$

$$H_{n-2}(A) \xrightarrow{\iota_*} \cdots$$

The maps ι_* and q_* are induced by the obvious inclusions of pairs.

—**(Excision):** *If X has subspace A, which in turn has subspace B with the closure of B contained in the interior of A, then the inclusion $(X-B, A-B) \hookrightarrow (X,A)$ has isomorphisms as induced maps between homology groups.*

—**(Dimension)** *The homology $H_n(*)$ of the one-point space $*$ is trivial for $n \neq 0$.*

For the dimension axiom, the group $H_0(*)$ is sometimes referred to as the *coefficient* of the homology theory. Singular homology satisfies all of the above, and has coefficient $\mathbb{Z}$.

Milnor noticed that by adding the following axiom, also satisfied by singular homology, the homology theory is completely determined on the category of pairs of spaces which are homotopy equivalent to CW complexes:

—**(Additivity):** For a family of spaces X_α the inclusions $i_\alpha \colon X_\alpha \hookrightarrow \coprod_\alpha X_\alpha$ induce isomorphisms

$$\bigoplus_\alpha H_n(X_\alpha) \xrightarrow[\oplus(i_\alpha)_*]{\cong} H_n(\coprod_\alpha X_\alpha).$$

Next term you will meet *cohomology*. There are Eilenberg–Steenrod axioms for cohomology theories too. They are almost identical: one essentially just needs to flip arrows of induced maps, and replace the appearance of the direct sum $\bigoplus$ of groups in the above additivity axiom with the product Π (in the product of groups, one allows sequences where infinitely many of the terms are non-identity elements).

Dropping the dimension axiom allows for different examples of (co)homology theories called *generalised* or *extraordinary (co)homology theories.* One of particular importance is K-theory, which is based upon complex vector bundles over a space. It assigns the following cohomology groups to the one point space: $K_n(*) \cong \mathbb{Z}$ for n even and $K_n(*) \cong 0$ for n odd.

4.8.2 Homology with coefficients

Keeping the dimension axiom, the question arises of what homology theory corresponds to the case where $H_0(*) \cong G$ for a given Abelian group G. By the result of Milnor, such a homology theory satisfying all of the above axioms is determined on the class of spaces homotopy equivalent to CW complexes (which are the spaces one typically deals with in practice in most settings).

One may construct such a theory directly with only a mild alteration to how we constructed usual homology. In the definition of the singular chain complex $C_*(X)$, instead of considering $\mathbb{Z}$-linear sums of singular simplices, consider instead G-linear sums. This defines a chain complex $C_*(X;G)$. That is, an element $\sigma \in C_n(X;G)$ can be thought of as a formal sum

$$\sigma = \sum_\alpha g_\alpha \alpha$$

where each $g_\alpha \in G$, each α is a singular n-simplex and g_α is the zero element (the additive identity) for all but finitely many α.

The boundary map has the obvious definition from the usual one: for α a singular n-simplex and $g_\alpha \in G$ we let

$$\partial_n(g_\alpha \alpha) \coloneqq \sum_{j=0}^{n} (-1)^j g_\alpha \alpha\restriction_j,$$

(here, we read $-g$ as the additive inverse of g). For a more general sum $\sigma = \sum g_\alpha \alpha$ we let

$$\partial_n(\sigma) = \sum_\alpha \partial_n(g_\alpha \alpha).$$

This defines the **singular chain complex with G coefficients** $C_*(X;G)$. Note the semi-colon rather than comma, as used for the relative homology. More abstractly, one may define $C_*(X;G)$ via a tensor product $C_*(X) \otimes G$. The homology of this chain complex is denoted $H_*(X;G)$, the **homology of X with G coefficients**. Using the tensor product description of the chain complex one may relate the usual homology (i.e., with $\mathbb{Z}$ coefficients) to the homology with G coefficients via the *universal coefficient theorem*, which you will see next term.

Similar constructions apply to give relative singular homology $H_*(X, A; G)$, or reduced homology $\widetilde{H}_*(X;G)$ with G coefficients. One can also easily modify simplicial and cellular homology to work with G coefficients. In this case, for cellular homology say, rather than n-chains being given by assigning finitely many cells $\mathbb{Z}$ coefficients, one assigns finitely many cells elements in G.

Example 4.8.1. Let's calculate the cellular homology with $\mathbb{Z}/2$ coefficients for $\mathbb{R}P^2$. We take the CW decomposition with one cell in each dimension. So the $\mathbb{Z}/2$ coefficient chain complex looks like:

$$\cdots \xrightarrow{\partial_4} 0 \xrightarrow{\partial_3} \mathbb{Z}/2 \xrightarrow{\partial_2} \mathbb{Z}/2 \xrightarrow{\partial_1} \mathbb{Z}/2 \to 0.$$

Let the 0-cell be denoted by v, the 1-cell be denoted by e and the 2-cell be denoted by f. Identify such a cell with the non-trivial element of the corresponding chain group above. As for $\mathbb{Z}$ coefficient homology, we still get $\partial_1(e) = v - v = 0$. But now in degree two $\partial_2(f) = 2e = 0$, since $2g = 0$ for $g \in \mathbb{Z}/2$. So we get the following homology groups:

$$H_n(\mathbb{R}P^2; \mathbb{Z}/2) \cong \begin{cases} \mathbb{Z}/2 & \text{for } n = 0, 1, 2; \\ 0 & \text{otherwise.} \end{cases}$$

Exercise 4.8.1. Calculate the homologies of $\mathbb{K}$ and $\mathbb{T}^2$ with coefficients in $\mathbb{Z}/2$, $\mathbb{Z}/3$ and $\mathbb{R}$.

Example 4.8.2. Cellular homology with coefficients in $\mathbb{Z}/2$ is quite easy to visualise: an n-chain is determined by choosing a finite number of n-cells to assign the non-trivial element of $\mathbb{Z}/2$. I like to think of a light-bulb on each cell, which you can either turn on or off. When the cells are attached nicely, the degree n boundary map works by sequentially turning off the light of each n-cell e, flipping the switches of the $(n-1)$-cells in the boundary of e when those cells are covered an odd number of times by the boundary of e.

For example, take a cellular decomposition of a compact, connected surface S. Each 1-cell has precisely two 2-cells sitting either side of it. So there are precisely two 2-cycles: the one where all the lights of 2-cells are off, and the one where all the lights are on. This shows that $H_2(S; \mathbb{Z}/2) \cong \mathbb{Z}/2$. More generally, for a compact and connected d-manifold M (loosely, a space which 'locally looks like $\mathbb{R}^d$'), one may show that $H_d(M; \mathbb{Z}/2) \cong \mathbb{Z}/2$. In contrast, for the usual $\mathbb{Z}$ coefficient homology, we have that $H_d(M) \cong \mathbb{Z}$ when M is orientable, and $H_d(M) \cong \mathbb{Z}$ when M is not orientable. Remember from the $d = 2$ case that being 'orientable' corresponded to M having 'two-sides' and the orientable surfaces are the surfaces of genus g (the 'tori with g holes'). The two famous non-orientable ones are $\mathbb{R}P^2$ or $\mathbb{K}$ (in fact, you can get all of the non-orientable ones like you get all of the orientable ones, starting instead with $\mathbb{R}P^2$ and 'attaching handles').

Example 4.8.3. There is a computational advantage to computing homology with coefficients in $\mathbb{F} = \mathbb{Z}/p$, $\mathbb{Q}$ or $\mathbb{R}$. In that case the cellular chain complex

$$\cdots \xrightarrow{\partial_4} \mathbb{F}^{k_3} \xrightarrow{\partial_3} \mathbb{F}^{k_2} \xrightarrow{\partial_2} \mathbb{F}^{k_1} \xrightarrow{\partial_1} \mathbb{F}^{k_0} \to 0$$

is actually a chain complex of vector spaces over $\mathbb{F}$ considered as a field, with boundary maps linear maps between vector spaces. The homology groups are $H_n \cong \mathbb{F}^{z_n - b_n}$ where

z_n is the dimension of the kernel of $\partial_n \colon \mathbb{F}^{k_n} \to \mathbb{F}^{k_{n-1}}$ and b_n is dimension of the image of ∂_{n+1} (this is just the rank–nullity theorem).

For example, over $\mathbb{R}$ coefficients we cannot get torsion like we can for $\mathbb{Z}$ coefficients. This can sometimes be a real loss of information: convince yourself that over $\mathbb{R}$ coefficients the homology of $\mathbb{R}P^2$ coincides with that of the one-point space, which is certainly not the case over $\mathbb{Z}$ coefficients. For certain problems, though, the $\mathbb{R}$ coefficient homology may be sufficient and far easier to compute.

4.8.3 Relation between degree one homology and fundamental group

Let X be a topological space and $x_0 \in X$ some fixed 'base-point'. The *fundamental group* $\pi_1(X, x_0)$ is the group of homotopy classes of 'based loops', continuous maps $\gamma \colon S^1 \to X$ sending $1 \in S^1$ to x_0. Given two such loops γ_1, γ_2, one can define their product as $\gamma_2 \cdot \gamma_1$, given by 'following γ_1 then γ_2'.

Let's make this more precise: firstly we can think of a based loop as equivalently a mapping $\gamma \colon I \to X$ sending the endpoints 0 and 1 to x_0. Then the product $\gamma_2 \cdot \gamma_1$ is defined by:

$$\gamma_2 \cdot \gamma_1(t) \coloneqq \begin{cases} \gamma_1(2t) & \text{for } 0 \leqslant t < 1/2; \\ \gamma_2(2t-1) & \text{for } 1/2 \leqslant t \leqslant 1. \end{cases}$$

In other words, you follow γ_1 and then γ_2 at twice the speed. We identify two such loops if there is a continuous map $F \colon I \times I \to X$ with $F(0, s) = F(1, s) = x_0$, $F(t, 0) = \gamma_1(t)$ and $F(t, 1) = \gamma_2(t)$ for all s, $t \in I$. Write $[\gamma]$ for the class of loops equivalent to γ and $\pi_1(X, x_0)$ for the set of equivalence classes. One may show that $[\gamma_2] \cdot [\gamma_1] \coloneqq [\gamma_2 \cdot \gamma_1]$ gives a well defined binary product and makes $\pi_1(X, x_0)$ a group.

It is easily shown that $\pi_1(X, x_0)$ only depends up to isomorphism on the path-component of X from which x_0 is chosen. So from now on we restrict to path-connected spaces and drop the mention of x_0.

Some examples indicate that there is a close relationship between $\pi_1(X)$ and $H_1(X)$. For example, $H_1(\mathbb{T}^2) \cong \mathbb{Z}^2$, $H_1(\mathbb{R}P^2) \cong \mathbb{Z}/2$, $H_1(S^1) \cong \mathbb{Z}$, $H_1(S^n) \cong 0$ for $n > 1$ are all isomorphic to the corresponding fundamental groups. Are they always the same? They cannot always be: the fundamental group can be non-Abelian, for example $\pi_1(S^1 \vee S^1) \cong \mathbb{Z} * \mathbb{Z}$, the free group (not free *Abelian* group!) on two generators; also, $\pi_1(\mathbb{K})$ of the Klein bottle is non-Abelian. The homology of a space is always Abelian. But given this restriction, they agree as much as they possible could do: $H_1(X)$ is the *Abelianisation* of $\pi_1(X)$, which we will see below.

Since $\Delta^1 \cong I$, we may think of a based loop γ as above as instead a singular 1-simplex σ_γ. Since $\gamma(0) = \gamma(1)$, it even defines a cycle. So define a map $h_1 \colon \pi_1(X) \to H_1(X)$ by

$h([\gamma]) \coloneqq [\sigma_\gamma]$. One may show that this is a well defined homomorphism of groups. We have the following theorem of Hurewicz:

Theorem 4.8.1. *Let X be a path-connected space. The homomorphism $h_1 \colon \pi_1(X) \to H_1(X)$ described above is surjective and has kernel the commutator subgroup*

$$[\pi_1(X), \pi_1(X)] \leqslant \pi_1(X).$$

Hence h_1 induces an isomorphism from the Abelianisation of $\pi_1(X)$ to $H_1(X)$.

In the above, the commutator subgroup $[G, G]$ is the normal subgroup of a group G generated by *commutators* $[g, h] \coloneqq ghg^{-1}h^{-1}$ for g, $h \in G$ (so $[G, G]$ consists of arbitrary products of commutators and their inverses in G). The quotient sending G to $G/[G, G]$ is called the Abelianisation; this quotient group is Abelian and the subgroup $[G, G]$ is the smallest normal subgroup by which one can quotient G to get an Abelian group. Note that if G is Abelian then each $[g, h]$ is the identity element so $[G, G]$ is the trivial subgroup of just the identity. Taking the Abelianisation is essentially adding the relations that $ghg^{-1}h^{-1}$ is the trivial element for any g, $h \in G$, which is necessary to be Abelian.

Exercise 4.8.2. Try to prove the above theorem. There are a few slightly tricky technical details to sort in the proof; if you get too bogged down check the literature.

Example 4.8.4. The Abelianisation of $\mathbb{Z} * \mathbb{Z}$ adds the relation that $aba^{-1}b^{-1}$ is the identity for a and b generators of the first and second copy of $\mathbb{Z}$; one may show that this results in the free Abelian group $\mathbb{Z}^2$. This agrees with the above Hurewicz theorem, since $\pi_1(S^1 \vee S^1) \cong \mathbb{Z}^2$, whereas $H_1(S^1 \vee S^1) \cong \mathbb{Z}^2$ by Theorem 4.4.4.

There also exist higher Hurewicz homomorphisms $h_n \colon \pi_n(X) \to H_n(X)$ from the higher homotopy groups of X (based on 'higher dimensional loops', maps $S^n \to X$, as alluded to in the introduction). These higher homotopy groups, for $n > 1$, are always Abelian. Hurewicz showed that if $\pi_k(X) \cong 0$ for $k < n$ then h_n is an isomorphism and h_{n+1} is surjective.

Example 4.8.5. Since $\pi_1(S^n) \cong 0$ for $n > 1$, it follows that $\pi_k(S^n) \cong H_k(S^n) \cong 0$ for $k < n$ and $\pi_n(S^n) \cong H_n(S^n) \cong \mathbb{Z}$.

The higher homotopy groups of a space are typically very difficult to compute, so Hurewicz's theorem can give interesting information in terms the easier to compute homology groups in some special cases. For example, it can be used to compute a few higher homotopy groups of the n-spheres.

Appendix A

Proof of the Snake Lemma

There are two sorts of thing to check: (1) we must check that all of the maps in the LES are well defined and (2) we must show that the LES is indeed exact in each position. Being chain maps, we already know that the induced maps f_* and g_* are well defined maps on homology. So let us firstly check that the connecting map ∂_* is well defined.

1(a) Well defined (α and β exist):

Let $\gamma \in C_n$ be an n-cycle; firstly we show that the elements $\alpha \in A_{n-1}$ and $\beta \in B_n$ as described in the statement of the lemma exist. The element $\beta \in B_n$ is chosen so that $g_n(\beta) = \gamma$. Such a β exists since g_n is surjective, due to the exactness of

$$B_n \xrightarrow{g_n} C_n \to 0.$$

We must thus show that there exists $\alpha \in A_{n-1}$ for which $f_{n-1}(\alpha) = \partial_n^B(\beta)$. Since $\mathrm{im}(f_{n-1}) = \ker(g_{n-1})$ by exactness of

$$A_{n-1} \xrightarrow{f_{n-1}} B_{n-1} \xrightarrow{g_{n-1}} C_{n-1},$$

we can find such an α precisely if $\partial_n^B(\beta) \in \ker(g_{n-1})$. Since (g_i) is a chain map

$$g_{n-1}(\partial_n^B(\beta)) = \partial_n^C(g_n(\beta)) = \partial_n^C(\gamma) = 0,$$

as γ was taken to be a cycle. So the elements α and β exist, as required.

1(b) Well defined (α is a cycle):

The map ∂_* is supposed to induce maps between homology groups, so we should check that α is a cycle, i.e., that $\partial_{n-1}^A(\alpha) = 0$. Since f_{n-2} is injective by the exactness of

$$0 \to A_{n-2} \xrightarrow{f_{n-2}} B_{n-2},$$

this is equivalent to showing that $f_{n-2}(\partial^A_{n-1}(\alpha)) = 0$. This is the case, since (f_i) is a chain map and so

$$f_{n-2}(\partial^A_{n-1}(\alpha)) = \partial^B_{n-1}(f_{n-1}(\alpha)) = \partial^B_{n-1}(\partial^B_n(\beta)) = 0,$$

as $\partial \circ \partial = 0$ in a chain complex.

1(c) Well defined (homology class of α does not depend on choices made):

For ∂_* to be well defined, there is a little more to do: we need to check that taking a different representative of γ' of $[\gamma] \in H_n(C)$ or different choices of $\alpha \in A_{n-1}$ and $\beta \in B_n$ (α' and β', say) does not make a difference, so that $[\alpha] = [\alpha'] \in H_{n-1}(A)$.

That $[\gamma] = [\gamma']$ in $H_n(C)$ is equivalent to saying that there exists some $\tau \in C_{n+1}$ with $\gamma - \gamma' = \partial^C_{n+1}(\tau)$ (i.e., γ and γ' differ by a boundary). It follows that

$$g_n(\beta - \beta') = g_n(\beta) - g_n(\beta') = \gamma - \gamma' = \partial^C_{n+1}(\tau). \tag{A.0.1}$$

This says that $\beta - \beta'$ is in the kernel of g_n, up to a boundary in C_*. We'd rather have, up to a boundary in B_*, that $\beta - \beta'$ is an element of the kernel of g_n, so as to utilise exactness; this prompts us to find some $\zeta \in B_{n+1}$ for which

$$g_n(\partial^B_{n+1}(\zeta)) = \partial^C_{n+1}(\tau).$$

We can find such a ζ, since $g_n \circ \partial^B_{n+1} = \partial^C_{n+1} \circ g_{n+1}$ as (g_i) is a chain map, so we pick $\zeta \in B_{n+1}$ with $g_{n+1}(\zeta) = \tau$ (which we can do, since g_{n+1} is surjective by exactness). By Equation A.0.1,

$$g_n(\beta - \beta' - \partial^B_{n+1}(\zeta)) = 0.$$

Exactness implies that $\operatorname{im}(f_n) = \ker(g_n)$, so there must exist some $\nu \in A_n$ with

$$f_n(\nu) = \beta - \beta' - \partial^B_{n+1}(\zeta). \tag{A.0.2}$$

Remember that we wish to show that $\alpha - \alpha'$ is a boundary, and it seems that $\partial^A_n(\nu)$ would be a good candidate. From Equation A.0.2

$$f_{n-1}(\partial^A_n(\nu)) = \partial^B_n(f_n(\nu)) = \partial^B_n(\beta - \beta' - \partial^B_{n+1}(\zeta)) = \partial^B_n(\beta) - \partial^B_n(\beta') = f_{n-1}(\alpha) - f_{n-1}(\alpha').$$

So we have that $f_{n-1}(\alpha - \alpha' - \partial^A_n(\nu)) = 0$. By exactness f_{n-1} is injective and we conclude that $\alpha - \alpha' - \partial^A_n(\nu) = 0$, so $\alpha - \alpha'$ is a boundary, as desired.

1(d) Well defined (connecting map is a homomorphism):

This one is easy. Take γ, $\gamma' \in C_n$. We must check that $\partial_*([\gamma] + [\gamma']) = \partial_*([\gamma]) + \partial_*([\gamma'])$ in A_{n-1}. So let α, $\alpha' \in A_{n-1}$ and β, $\beta' \in B_n$ be such that $f_{n-1}(\alpha) = \partial^B_n(\beta)$ and $g_n(\beta) = \gamma$, and similarly for the primed versions, as in the definition of the connecting map. Then

$$f_{n-1}(\alpha + \alpha') = f_{n-1}(\alpha) + f_{n-1}(\alpha') = \partial^B_n(\beta) + \partial^B_n(\beta') = \partial^B_n(\beta + \beta')$$

and

$$g_n(\beta + \beta') = g_n(\beta) + g_n(\beta') = \gamma + \gamma',$$

as everything involved is a homomorphism, so we may take $\partial_*([\gamma]+[\gamma']) = \partial_*([\gamma+\gamma']) \coloneqq \alpha + \alpha' = \partial_*([\gamma]) + \partial_*([\gamma'])$.

2(a) Exactness ($\mathrm{im}(f_*) = \ker(g_*)$**):**

There are three types of location in the LES for which we must check exactness. We shall cover two of them and leave the third as an exercise. Firstly we wish to show that the diagram

$$H_n(A) \xrightarrow{f_*} H_n(B) \xrightarrow{g_*} H_n(C)$$

is exact. So suppose that $\sigma \in B_n$ is a cycle. To say that $[\sigma] \in \ker(g_*)$ is to say that $g_n(\sigma)$ is a boundary i.e., there exists some $\tau \in C_{n+1}$ with $\partial^C_{n+1}(\tau) = g_n(\sigma)$. Since g_{n+1} is surjective there exists $\rho \in B_{n+1}$ with $g_{n+1}(\rho) = \tau$. It follows that

$$g_n(\sigma - \partial^B_{n+1}(\rho)) = g_n(\sigma) - g_n(\partial^B_{n+1}(\rho)) = \partial^C_{n+1}(\tau) - \partial^C_{n+1}(g_{n+1}(\rho)) = 0,$$

in other words σ, up to a boundary in B_*, is in $\ker(g_n)$. By exactness, there exists $\nu \in A_n$ with $f_n(\nu) = \sigma - \partial^B_{n+1}(\rho)$, which means that $f_*([\nu]) \coloneqq [f_n(\nu)] = [\sigma]$, so $[\sigma] \in \mathrm{im}(f_*)$.

For the other direction, i.e., to show that $\mathrm{im}(f_*) \subseteq \ker(g_*)$, it is equivalent to show that $g_* \circ f_* = 0$, the zero map from $H_*(A)$ to $H_*(C)$. However, this is pretty obvious by functorality: $g_* \circ f_* = ((g_n) \circ (f_n))_* = (0)_*$, where the 0 here represents the trivial map from A_* to C_* sending every element to zero (which we have from exactness of the chain maps). Clearly the induced map of the zero chain map is the zero map between the corresponding homology groups.

2(b) Exactness ($\mathrm{im}(\partial_*) = \ker(f_*)$**):**

Next we show that

$$H_n(C) \xrightarrow{\partial_*} H_{n-1}(A) \xrightarrow{f_*} H_{n-1}(B)$$

is exact. So let $\alpha \in A_{n-1}$ be an $(n-1)$-cycle. If $[\alpha] \in \ker(f_*)$ that means that $f_n(\alpha)$ is a boundary, i.e., there exists some $\beta \in B_n$ with $f_{n-1}(\sigma) = \partial^B_n(\beta)$. Setting $\gamma \coloneqq \partial^B_n(\beta)$, we see that $[\alpha] = \partial_*([\gamma])$, by the definition of the connecting map, and so $[\alpha] \in \mathrm{im}(\partial_*)$, as desired.

Conversely, suppose that $[\alpha] \in \mathrm{im}(\partial_*)$. Hence, there exists some $\alpha' \in A_{n-1}$, $\beta \in B_n$ and $\gamma \in C_n$ so that: (1) $[\alpha] = [\alpha']$, (2) $f_{n-1}(\alpha') = \partial^B_n(\beta)$ and (3) $g_n(\beta) = \gamma$. We wish to show that the *homology class* of α is in $\ker(f_*)$, so we may as well just show that $f_{n-1}(\alpha')$ represents a trivial element in $H_{n-1}(B)$, which is not different problem as α' represents the same homology class by (1). But it follows directly from (2) that $f_{n-1}(\alpha')$ is a boundary, so $f_*([\alpha']) = 0$ and hence $[\alpha] \in \ker(f_*)$.

2(a) Exactness ($\mathrm{im}(g_*) = \ker(\partial_*)$**):**

All that's left is to show that

$$H_n(B) \xrightarrow{g_*} H_n(C) \xrightarrow{\partial_*} H_{n-1}(A)$$

is exact. If you heeded our advice above you will have already been making your own way through the proof. If not, at least try to wrap up this final small part as an **Exercise**.

A.1 Naturality

There is a natural notion of a morphism between two short exact sequences of chain complexes, a commutative diagram like the following:

$$\begin{array}{ccccccccc}
0 & \longrightarrow & A_* & \overset{f}{\hookrightarrow} & B_* & \xrightarrow{g} & C_* & \longrightarrow & 0 \\
 & & \downarrow{\scriptstyle \alpha} & & \downarrow{\scriptstyle \beta} & & \downarrow{\scriptstyle \gamma} & & \\
0 & \longrightarrow & A'_* & \overset{f'}{\hookrightarrow} & B'_* & \xrightarrow{g'} & C'_* & \longrightarrow & 0
\end{array}$$

It turns out that the Snake Lemma is functorial: when applied to such a diagram, you not only get two long exact sequences from the Snake Lemma, you also get a map between those long exact sequences:

$$\begin{array}{ccccccccccc}
\cdots & \xrightarrow{f_*} & H_{k+1}(B) & \xrightarrow{g_*} & H_k(C) & \xrightarrow{\partial_*} & H_k(A) & \xrightarrow{f_*} & H_k(B) & \xrightarrow{g_*} & \cdots \\
 & & \downarrow{\scriptstyle \beta_*} & & \downarrow{\scriptstyle \gamma_*} & & \downarrow{\scriptstyle \alpha_*} & & \downarrow{\scriptstyle \beta_*} & & \\
\cdots & \xrightarrow{f'_*} & H_{k+1}(B') & \xrightarrow{g'_*} & H_{k+1}(C') & \xrightarrow{\partial_*} & H_k(A') & \xrightarrow{f'_*} & H_k(B') & \xrightarrow{g'_*} & \cdots
\end{array} \tag{A.1.1}$$

This property is more usually referred to as 'naturality', which can be given a rigorous category-theoretic meaning.

Exercise A.1.1. Prove the assertion, namely on the commutativity of the diagram A.1.1. You will need to look back at the definition of the connecting map ∂_* defining the LES from the Snake Lemma 2.3.1.

Appendix B

Rank and short exact sequences

Here we shall briefly explain why the rank of an Abelian group as given in Section 2.4.1 is well defined and behaves additively with respect to short exact sequences.

Lemma B.0.1. *The rank of an Abelian group is well defined.*

Proof. One can show that there is at least one maximal linearly independent set using Zorn's Lemma. Given two, we just need to show then that they are the same size. The proof is just like that for proving that the dimension of a vector space is well defined. Take two maximal linearly independent sets S_1, $S_2 \subseteq A$ and suppose, without loss of generality, that $|S_1| < |S_2|$.

For any $a \in S_1$ there exists some non-zero $k_a \in \mathbb{Z}$ so that

$$k_a a = \sum_{S_a} n_b b \tag{B.0.1}$$

for $n_b \in \mathbb{Z}$ and $b \in S_a$ for some finite subset $S_a \subseteq S_2$ (otherwise we can add a to S_2 which is thus not minimal).

Suppose first that S_2 is infinite. As the sums as in Equation B.0.1 are finite sums and $|S_1| < |S_2|$, some $b \in S_2$ does not appear in S_a for any $a \in A$. As above we can still write ℓb as a finite $\mathbb{Z}$-linear combination of terms in S_1 for some non-zero $\ell \in \mathbb{Z}$, and each of these terms have multiples which can be expressed as finite sums in S_2 not involving b. This means that we can express a multiple of b as a finite $\mathbb{Z}$-linear combination of others in S_2, which is thus not linearly independent, a contradiction.

If S_2 is finite then so is S_1. Write $S_1 = \{a_1, \ldots, a_m\}$ and $S_2 = \{b_1, \ldots, b_n\}$ with $m < n$. We may express a non-zero multiple of a_1 as a $\mathbb{Z}$-linear combination of the b_is and show that $\{a_1, b_1, \ldots, \hat{b}_j \ldots, b_n\}$ is linearly independent and maximal for some omitted $\hat{b}_j$. One can repeat this process, replacing b_i terms with a_i terms in S_2 (c.f., the proof for vector spaces). As a result one constructs a linearly independent set which contains S_1 and with strictly more elements, contradicting that S_1 is maximal. □

Proof of Lemma 2.4.1. Let $0 \to A \xrightarrow{f} B \xrightarrow{g} C \to 0$ be a short exact sequence. Take a maximal linearly independent set S_A of elements of A, and S_C of C. For each $c_i \in S_C$ pick $c_i' \in B$ with $g(c_i') = c_i$. Then we claim that the following set is linearly independent and maximal in B:

$$S_B = \{f(a_i) \mid a_i \in S_A\} \cup \{c_i' \mid c_i \in S_C\}.$$

To see that S_B is linearly independent, suppose that $\sum_{S_A} m_i f(a_i) + \sum_{S_C} n_i c_i' = 0$. Since $g \circ f = 0$, applying g to the sum we have that $\sum_{S_C} n_i c_i = 0$, which by linear independence of S_C implies that each $n_i = 0$, so $\sum_{S_A} m_i f(a_i) = 0$. But f is injective, so $\sum_{S_A} m_i a_i = 0$. By linear independence of S_A we have that each $m_i = 0$ too and S_C is linearly independent.

To see that S_B is maximal, suppose that $b \in B$. Consider $g(b) \in C$. By maximality of S_C we have that $g(\ell b) = \ell g(b) = \sum n_i c_i$ for some non-zero $\ell \in \mathbb{Z}$. Hence $b' = \sum_{S_C} n_i c_i'$ is such that $g(\ell b) = g(\ell b')$, so $\ell(b - b') \in \ker(g) = \operatorname{im}(f)$ by exactness. Let $a \in A$ be such that $f(a) = \ell(b - b')$. By maximality of S_A there is some non-zero $k \in \mathbb{Z}$ with $ka = \sum_{S_C} m_i a_i$. Hence $k\ell(b - b') = k(f(a)) = f(k(a)) = \sum_{S_C} m_i f(a_i)$. It follows that $(k\ell)b$ is a linear combination of terms from S_B, which is thus maximal. Since $|S_B| = |S_A \cup S_C| = |S_A| + |S_C|$, we have shown that $\operatorname{rk}(A) + \operatorname{rk}(C) = \operatorname{rk}(B)$. □

Appendix C

Geometric realisation

Let α and β be finite subsets of $\mathbb{R}^n$ in general position and let $f\colon \alpha \to \beta$. We define the corresponding **affine extension** $L_f\colon \Delta_\alpha \to \Delta_\beta$ by the formula

$$L_f(\sum_{v\in\alpha} \lambda_i v_i) \coloneqq \sum_{v\in\alpha} \lambda_i f(v).$$

Remark C.0.1. This is functorial in the following sense:

- for id_α the identity map on α we have that $L_{\mathrm{id}_\alpha} = \mathrm{id}_{\Delta_\alpha}$, the identity map on Δ_α;
- for $f\colon \alpha \to \beta$ and $g\colon \beta \to \gamma$, we have that $L_{g\circ f} = L_g \circ L_f$.

For $\alpha \subseteq \beta$ we have that $\Delta_\alpha \subseteq \Delta_\beta$ (Δ_α is a face of Δ_β), and the inclusion map $\Delta_\alpha \hookrightarrow \Delta_\beta$ is given by L_ι, where $\iota\colon \alpha \hookrightarrow \beta$ is the inclusion of α into β.

One use of these affine maps is that they allow us to canonically identify simplices. In what follows, as usual, the underlying set V of a simplicial complex $\mathcal{K}$ is equipped with a total order, and simplices of $\mathcal{K}$ are always written with elements in order.

Suppose that $f\colon \Delta_\alpha \to X$ is a function, for an n-simplex Δ_α whose vertex set α is totally ordered. Let $J \subseteq \alpha$. Generalising the notation of Section 3.3.2, let

$$f\restriction^n_J\colon \Delta^{n-|J|} \to X$$

be the function given by restricting Δ_α to the face spanned by vertices $\alpha - J$ but first implicitly identifying this simplex with the standard $(n-|J|)$-simplex using the affine extension $L(f)$ of the unique order-preserving bijection f mapping $\{e_0, \ldots, e_{n-|J|}\}$ to $\alpha - J$. In other words $\restriction^n_J = L(\iota) \circ L(f)$ where ι is the injection $\iota\colon \alpha - J \hookrightarrow \alpha$. Again, we will occasionally drop the superscript of $\restriction^n_J$. Note that it is only well defined in context (we need to know the domain of f).

Lemma C.0.1. *Let $J_1 \subseteq J_2 \subseteq \alpha$. We have the equality*

$$f\restriction_{J_1} = (f\restriction_{J_2})\restriction_J$$

where $J \subseteq \{e_0, \ldots, e_{n-|J_2|}\}$ lists the indices of elements of J_2 that are not in J_1.

Proof. We are really just showing here that $L(\iota_1) \circ L(f_1) = (L(\iota_2) \circ L(f_2)) \circ (L(\iota) \circ L(f))$ associated to various maps between finite sets. By functorality one just checks that these compositions of functions agree. The one on the left of this equation just injects a finite set of $(n - |J_1|)$ elements into that of n elements by skipping the index J_1 elements. The one on the right first skips over an element of index $j \in J$ if the index j element of J_2 is not in J_1, and then skips over the index J_2 elements. By construction these compositions agree. $\square$

Note that for $i < j$ this recovers the simple equation $\restriction_j \restriction_i = \restriction_i \restriction_{j-1}$.

C.1 Geometric realisation

The 0-skeleton $|\mathcal{K}|^0$ is taken simply as the discrete space with points in bijection with V (really, in bijection with the singleton subsets of V). We have maps $\sigma_{\{v\}} \colon \Delta^0 \to |\mathcal{K}|^0$, where Δ^0 is the standard 0-simplex (i.e., a point) which record which points are associated to which elements $\{v\} \in \mathcal{K}$.

So suppose that we have constructed the k-skeleton $|\mathcal{K}|^k$, which we shall denote throughout by X and, moreover, for each n-simplex $\alpha \in \mathcal{K}$ for $n \leqslant k$ we have an associated continuous map (even, it turns out, a homeomorphism onto its image) $\sigma_\alpha \colon \Delta^n \to X$, where Δ^n is the standard n-simplex. For each $\beta = \{v_0, \ldots, v_{k+1}\} \in \mathcal{K}$ we wish to attach a $(k+1)$-cell to the k-skeleton. The simplex Δ^{k+1} is homeomorphic to the disc D^{k+1} and its boundary $\partial\Delta^{k+1}$ is the union of faces of Δ^{k+1}. To define the attaching map

$$\partial_\beta \colon \partial\Delta^{k+1} \to X$$

it is enough to know how to map each of the faces of Δ^{k+1} into the k-skeleton.

There is a canonical way to do this: a face of $\partial\Delta^{k+1}$ is given by deleting some non-empty proper subset J of indices from $\{e_0, \ldots, e_{k+1}\}$ and taking the convex hull. We define ∂_β so that

$$\partial_\beta \restriction_J^{k+1} = \sigma_\alpha,$$

where $\alpha \in \mathcal{K}$ is the subset of β given by deleting index J elements (so the re-parametrised restriction to the face spanned by vertices of β not in J).

It needs to be checked that ∂_β has the same definition when defined as above on the overlap of two different faces of $\partial\Delta^{k+1}$. Since any two such faces, Δ_{α_1} and Δ_{α_2}, corresponding to $\alpha_1, \alpha_2 \subset \beta \in \mathcal{K}$, intersect in another (i.e., $\Delta_{\alpha_1 \cap \alpha_2}$) it is enough to check this for $\Delta_{\alpha_1} \subset \Delta_{\alpha_2}$, i.e., with $\alpha_1 \subset \alpha_2 \subset \beta$. Applying the definition straight to Δ_{α_1}, we have that $\partial_{\alpha_1}\restriction_{J_1} = \sigma_{\alpha_1}$, where α_1 is given by removing the index J_1 elements of β. Similarly, α_2 is given by deleting elements of β with indices in some J_2, and α_1 is given by deleting elements of α_2 with indices in some J. So the restriction of $\partial_\beta\restriction_{J_2}$ to Δ_{α_1} (with a canonical identification with a standard simplex) is really $(\partial_\beta\restriction_{J_2})\restriction_J = (\sigma_{\alpha_2})\restriction_J = (\partial_{\alpha_2})\restriction_J = \sigma_{\alpha_1}$. The final two equalities come from the definition of ∂_{α_2}, which we assume to already be well defined by induction further down the skeleta.

So ∂_β is well defined. Repeating for each $\beta \in \mathcal{K}$ of size $(k+2)$, this defines the attaching maps of the $(k+1)$-cells, and so this defines a CW complex as described in Section 1.2.1.

C.2 Geometric geometric realisation

Usually the geometric realisation can be done in a more geometric way i.e., it can all be done inside Euclidean space using actual simplices. This is the case, in particular, if the set of vertices is finite (or even countable, although then one uses $\mathbb{R}^\infty$). We'll quickly cover how that goes here.

Let $\mathcal{K}$ be a simplicial complex over the set V (so $\mathcal{K}$ has vertices $\{v\}$ for $v \in V$). Suppose, moreover, that we have some Euclidean space $\mathbb{R}^N$ and for each $v \in V$ some chosen point $x_v \in \mathbb{R}^N$. For $\alpha \in \mathcal{K}$ we let Δ_α be the convex hull of the vertices x_v, for $v \in \alpha$, and let e_α be the corresponding open cell, i.e.,

$$e_\alpha := \{\sum_{v\in\alpha} \lambda_v x_v \mid \lambda_v \in (0,1), \sum_v \lambda_v = 1\}$$

for α consisting of more than 1 point (and we set $e_\alpha = \{x_v\}$ if $|\alpha| = 1$). Suppose that the vertices x_v are in sufficiently general position with respect to $\mathcal{K}$, in the following sense: for any $\alpha \in \mathcal{K}$ we have that the vertices defining Δ_α are in general position, and that for $\alpha \neq \beta \in \mathcal{K}$ we have that $e_\alpha \cap e_\beta = \varnothing$. Then we may define the **geometric realisation** of $\mathcal{K}$ as the subspace

$$|\mathcal{K}| := \bigcup_{\alpha\in\mathcal{K}} \Delta_\alpha \subseteq \mathbb{R}^N.$$

In other words, we have defined the geometric realisation just as the union of the simplicies Δ_α, as we usually imagine it when drawing pictures. Of course, the k-skeleton is given by the union of k-simplicies Δ_α (i.e., with $|\alpha| = k+1$).

Note that if V is finite then we can always find an embedding into a Euclidean space like this. We simply choose each x_v as a different standard basis vector in $\mathbb{R}^{|V|}$ (c.f., the standard n-simplex). If you've met $\mathbb{R}^\infty$ before, one can also consider the analogue of the above process through embeddings into it. Some authors choose to define the geometric realisation in this way, at the small cost of disallowing simplicial complexes with uncountably many vertices. One advantage, beyond being a little more visual, is that this shows that simplicial complexes with countably many cells have geometric realisations which are metrisable.

Bibliography

[Arm] M. A. Armstrong, *Basic topology*, McGraw-Hill Book Co. (UK), Ltd., London, 1979.

[Bre] G. E. Bredon, *Topology and geometry*, Graduate Texts in Mathematics, 139, Springer-Verlag, New York, 1993.

[Dol] A. Dold, *Lectures on algebraic topology*, Springer-Verlag, New York, 1972.

[Ful] W. Fulton, *Algebraic topology*, Graduate Texts in Mathematics, 153, Springer-Verlag, New York, 1995.

[Hat] A. Hatcher, *Algebraic topology*, Cambridge University Press, Cambridge, 2002.

[Spa] E. H. Spanier, *Algebraic topology*, McGraw-Hill Book Co., New York, 1966.