Part II Algebraic Topology

Henry Wilton

November 8, 2019

0 Introduction

Topology is often loosely defined as 'rubber-band geometry'. Perhaps a more rigorous definition is that topology is the study of continuous maps.

Here is a question that the mathematical tools we've seen so far in the tripos aren't particularly good at answering.

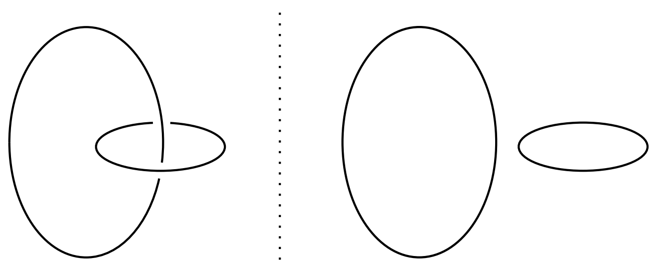

Figure 1: The Hopf link on the left and the 2-component unlink on the right.

Question 0.1. *Is the Hopf link really linked? More precisely, is there a homeomorphism of $\mathbb{R}^3$ that takes the Hopf link H to the two-component unlink U?*

How can we think about attacking a *qualitative* question like this? We can phrase it as an *extension* problem. The Hopf link is a particular embeddding $\eta : S^1 \sqcup S^1 \hookrightarrow \mathbb{R}^3$, and the unknot is another embedding $\upsilon : S^1 \sqcup S^1 \hookrightarrow \mathbb{R}^3$. It's easy to see that the unlink *extends* to a continuous map of two closed discs $D^2 \sqcup D^2$ into $\mathbb{R}^3$, where we identify S^1 with the boundary ∂D^2.

$$\begin{tikzcd}
S^1 \sqcup S^1 \arrow[rr, "v"] \arrow[dr] & & \mathbb{R}^3 \\
& D^2 \sqcup D^2 \arrow[ur] &
\end{tikzcd}$$

So we can ask a more specific question.

Question 0.2. *Does the Hopf link $\eta : S^1 \sqcup S^1 \to \mathbb{R}^3$ extend to a map of $D^2 \sqcup D^2$ into $\mathbb{R}^3$?*

This poses knottedness as an *extension problem.* Here's another extension problem that seems difficult, but is even easier to state. Let

$$S^{n-1} := \{x = (x_1, \ldots, x_n) \in \mathbb{R}^n \mid \sum_i x_i^2 = 1\}$$

be the $(n-1)$-sphere, which is naturally the boundary of the n-ball (or disc) D^n.

Question 0.3. *Does the identity map* $\mathrm{id} : S^{n-1} \to S^{n-1}$ *extend to a continuous map $D^n \to S^{n-1}$?*

$$\begin{tikzcd}
S^{n-1} \arrow[rr, "\mathrm{id}"] \arrow[dr] & & S^{n-1} \\
& D^n \arrow[ur] &
\end{tikzcd}$$

When $n = 1$, this is answered by the Intermediate Value Theorem! When $n = 2$, you may have seen questions like this answered using winding number. In general, this kind of problem seems difficult to answer, because the set of continuous maps $D^n \to S^{n-1}$ is very complicated.

A similar question in algebra, on the other hand, seems very easy.

Question 0.4. *Does the identity map* $\mathrm{id} : \mathbb{Z} \to \mathbb{Z}$ *factor through the trivial group?*

$$\begin{tikzcd}
\mathbb{Z} \arrow[rr, "\mathrm{id}"] \arrow[dr] & & \mathbb{Z} \\
& 0 \arrow[ur] &
\end{tikzcd}$$

The goal of this course is to develop tools to translate difficult problems like Questions 0.2 and 0.3 into easy problems like Question 0.4.

1 The fundamental group

As mentioned above, *winding number* is the first really interesting instance of a tool to solve this kind of problem. But what does winding number really mean? We'll start to develop some definitions to make sense of it. The most important definition in this section will be the fundamental group, which we will get to after some preliminaries.

1.1 Deforming maps and spaces

To study 'rubber geometry', we need to be able to *continuously deform* spaces, and also maps. How do we make sense of this? We're going to use the closed interval $[0,1]$ so often that we will write $I := [0,1]$.

Definition 1.1. Let $f_0, f_1 : X \to Y$ be (continuous) maps between topological spaces. A *homotopy* between f_0 and f_1 is a continuous map

$$F : X \times I \to Y$$

with $F(x,0) = f_0(x)$ an $F(x,1) = f_1(x)$ for all $x \in X$. We often write $f_t(x) := F(x,t)$. If such an F exists, we say that f_0 *is homotopic to* f_1 and write $f_0 \simeq_F f_1$, or just $f_0 \simeq f_1$.

Informally, we think of F as *deforming* f_0 into f_1.

Example 1.2. If $Y \subseteq \mathbb{R}^2$ is a convex region in the plane, then any pair of maps $f_0, f_1 : X \to Y$ are homotopic, via the *straight-line homotopy* defined by

$$F(x,t) = tf_0(x) + (1-t)f_1(x)$$

for all $x \in X$ and $t \in I$.

Sometimes it will be useful to have a technical strengthening.

Definition 1.3. Suppose that $f_0 \simeq_F f_1 : X \to Y$ as above. If $Z \subseteq X$ and $F(z,t) = f_0(z) = f_1(z)$ for all $z \in Z$ and $t \in I$ then we say that $f_0 \simeq f_1$ *relative to* Z.

The notation $\simeq$ strongly suggests that homotopy is an equivalence relation, and it's easy to prove that it is.

Lemma 1.4. *Let $Z \subseteq X, Y$ be topological spaces. The relation $\simeq$ (relative to Z) on the set of continuous maps $X \to Y$ is an equivalence relation.*

Proof. The map

$$(x,t) \mapsto f_0(x)$$

is a homotopy from f_0 to itself, so $\simeq$ is reflexive. For symmetry, note that if F is a homotopy from f_0 to f_1 then

$$(x,t) \mapsto F(x, 1-t)$$

is a homotopy from f_1 to f_0.

For transitivity, suppose that $f_0 \simeq_{F_0} f_1$ and $f_1 \simeq_{F_1} f_2$ (both relative to Z). We can get the idea of how to construct a homotopy from f_0 to f_1 by gluing together the domains of F_0 and F_1, then rescaling. Formally, this gives us the homotopy

$$F(x,t) = \begin{cases} F_0(x, 2t) & t \leq 1/2 \\ F_1(x, 2t-1) & t \geq 1/2 \end{cases}$$

gives a well-defined homotopy between f_0 and f_2 (relative to Z), since $F_0(x,1) = f_1(x) = F_1(x,0)$. □

Having seen how to deform maps, we next need to see how to deform spaces. Recall that a *homeomorphism* between two spaces X and Y is a continuous map $f : X \to Y$ with a continuous inverse $g : Y \to X$; that is, $f \circ g = \mathrm{id}_Y$ and $g \circ f = \mathrm{id}_X$. In order to deform spaces, we replace equality by homotopy in this definition.

Definition 1.5. Let X, Y be topological spaces. A *homotopy equivalence between X and Y* is a map $f : X \to Y$ with a *homotopy inverse* $g : Y \to X$; that is, $f \circ g \simeq \mathrm{id}_Y$ and $g \circ f \simeq \mathrm{id}_X$. In this case, we say that X *is homotopy equivalent to* Y and write $X \simeq Y$.

Example 1.6. Let $*$ be the space with one point. Let $f : \mathbb{R}^n \to *$ be the only map, and let $g : * \to \mathbb{R}^n$ send the point to 0. Then $f \circ g = \mathrm{id}_*$ and $g \circ f = 0$, the 0 map. The straight-line homotopy shows that $0 \simeq \mathbb{R}^n$, so $\mathbb{R}^n$ is homotopy equivalent to $*$.

The moral is that, unlike homeomorphism, homotopy equivalence is a violent operation on a space, which destroys a lot of structure. In fact, being homotopy equivalent to a point is a very common phenomenon in topology, and has a name.

Definition 1.7. If $X \simeq *$ then we say that X is *contractible.*

Here's another, slightly less trivial, example.

Example 1.8. Let $X = S^1$ and let $Y = \mathbb{R}^2 - 0$. Let $f : S^1 \to \mathbb{R}^2 - 0$ be the natural inclusion and define $g : \mathbb{R}^2 - 0 \to S^1$ by

$$g(x) = \frac{x}{\|x\|}$$

for all x. Then $g \circ f = \mathrm{id}_{S^1}$. If we look at the straight-line homotopy

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

between id_Y and $f \circ g$ in $\mathbb{R}^2$, we can see that $F(x,t) \neq 0$ for all $x \neq 0$ and $t \in I$. Therefore $f \circ g \simeq \mathrm{id}_Y$, and so $X \simeq Y$.

Remark 1.9. The same argument shows that $S^{n-1} \simeq \mathbb{R}^n - 0$ for all $n \geq 1$.

The homotopy equivalences that we constructed in the two examples had some special features. They are in fact examples of *deformation retractions.*

Definition 1.10. Let X, Y be topological spaces and $f : X \to Y$, $g : Y \to X$ continuous maps. If $g \circ f = \mathrm{id}_X$ then we say that X is a *retract* of Y, and g is a *retraction.* If, in addition, $f \circ g \simeq \mathrm{id}_Y$ relative to $f(X)$ then we say that X is a *deformation retract* of Y.

Again, our choice of notation suggests that homotopy equivalence should be an equivalence relation, and this is true.

Lemma 1.11. *Homotopy equivalence is an equivalence relation on topological spaces.*

Proof. Identity maps are homotopy equivalences, so $\simeq$ is reflexive, and symmetry is built into the definition. It remains to prove transitivity. Suppose therefore that $X \simeq Y$ and $Y \simeq Z$ via pairs of homotopy equivalences f, g and f', g' respectively. We need to prove that $f' \circ f$ and $g' \circ g$ define a pair of homotopy equivalences between X and Z. Indeed,

$$(g \circ g') \circ (f' \circ f) = g \circ (g' \circ f') \circ f \ .$$

But we know that $g' \circ f' \simeq \mathrm{id}_Y$ via some homotopy F', and so the composition $(x,t) \mapsto g \circ F'(f(x), t)$ defines a homotopy

$$g \circ (g' \circ f') \circ f \simeq g \circ \mathrm{id}_Y \circ f = g \circ f \ .$$

Since $g \circ f \simeq \mathrm{id}_X$, we have shown that $(g \circ g') \circ (f' \circ f) \simeq \mathrm{id}_X$, as required. Similarly, we see that $(f' \circ f) \circ (g \circ g') \simeq \mathrm{id}_Z$ which completes the proof. ☐

1.2 The definition of the fundamental group

Now that we have talked about deforming maps in general, we will focus specifically on continuous maps from the *interval* I into a space X. We think of these as *paths*.

Definition 1.12. Let X be a space and $x_0, x_1 \in X$. A *path in* X is a continuous map $\gamma : I \to X$. A *path from* x_0 *to* x_1 is a path γ in X so that $\gamma(0) = x_0$ and $\gamma(1) = x_1$. A *loop in* X *(based at* x_0*)* is a path in X from x_0 to itself.

Recall that a space X is *path-connected* if, for every $x_0, x_1 \in X$, there is a path from x_0 to x_1. This is a mild hypothesis on a space, since we can usually pass to path components. Since I is contractible, any two paths in a path-connected space are homotopic. To define an interesting theory of deformations of paths and loops, we consider homotopies *relative to endpoints*.

Definition 1.13. Let $\gamma_0, \gamma_1 : I \to X$ be paths. A *homotopy of paths* between γ_0 and γ_1 is a homotopy $F : I \times I \to X$ between γ_0 and γ_1 relative to $\{0, 1\}$. If this exists, we say that γ_0 and γ_1 are homotopic *as paths*, but we will often abuse notation and write $\gamma_0 \simeq \gamma_1$. We write $[\gamma]$ for the homotopy class of a path γ.

We are going to see that we can give a group structure to a set of loops in a space X. To do this, we define some operations and structures on paths which are reminiscent of the operations and structures enjoyed by groups.

Definition 1.14. Let X be a space, let x, y, z be points of X and let γ be a path from x to y and δ a path from y to z.

(i) The *concatenation* of γ and δ is the path from x to z defined by

$$(\gamma \cdot \delta)(t) := \begin{cases} \gamma(2t) & t \leq 1/2 \\ \delta(2t-1) & t \geq 1/2 \end{cases}$$

for all t.

(ii) The *constant path* at x is the path c_x in X such that $c_x(t) = x$ for all t.

(iii) The *inverse* of γ is the path from y to x defined by

$$\bar{\gamma}(t) = \gamma(1-t)$$

for all t.

We can now state the existence of the group that we have been referring to.

Theorem 1.15. *Let X be a space and $x_0 \in X$. Let $\pi_1(X, x_0)$ be the set of homotopy classes of loops in X based at x_0. Then $\pi_1(X, x_0)$ admits the structure of a group, with binary operation*

$$[\gamma][\delta] = [\gamma \cdot \delta] \ ,$$

identity given by the constant path $[c_{x_0}]$, and inverses given by $[\gamma]^{-1} = [\bar{\gamma}]$. This group is called the fundamental group of X (based at x_0).

To prove the theorem, we need to check that the group operation is well defined, and then check the group axioms. We start with the well-definedness.

Lemma 1.16. *If $\gamma_0 \simeq \gamma_1$ are paths to y and $\delta_0 \simeq \delta_1$ are paths from y, then $\gamma_0 \cdot \delta_0 \simeq \gamma_1 \cdot \delta_1$ and $\bar{\gamma}_0 \simeq \bar{\gamma}_1$.*

Proof. The homotopy $\gamma_0 \simeq \gamma_1$ is a map F from a square $P \cong I \times I$ to X, equal to c_y along the right-hand side. The homotopy $\delta_0 \simeq \delta_1$ is a map G from a square $Q \cong I \times I$ to X, equal to c_y along the left-hand side. We may therefore glue F and G along a side to obtain a map from the rectangle $R = P \cup_I Q \to X$. Rescaling R to a square defines the required homotopy H. Note that we could have written H down explicitly.

$$H(s,t) := \begin{cases} F(2s,t) & s \le 1/2 \\ G(2s-1,t) & s \ge 1/2 \end{cases}$$

For the second assertion, note that $(s,t) \mapsto F(s, 1-t)$ is a homotopy of paths from $\bar{\gamma}_0$ to $\bar{\gamma}_1$.

□

The group axioms now follow from the following observations.

Lemma 1.17. *Consider paths α, β, γ in X from w to x to y to z respectively.*

(i) $(\alpha \cdot \beta) \cdot \gamma \simeq \alpha \cdot (\beta \cdot \gamma)$

(ii) $\alpha \cdot c_x \simeq \alpha \simeq c_w \cdot \alpha$

(iii) $\alpha \cdot \bar{\alpha} \simeq c_w$

Proof. (i) Consider the path

$$\delta(t) = \begin{cases} \alpha(3t) & t \le 1/3 \\ \beta(3t-1) & 1/3 \le t \le 2/3 \\ \gamma(3t-2) & t \ge 2/3 \end{cases}$$

in X. Consider the functions $I \to I$ defined by

$$f_0(t) = \begin{cases} \frac{4}{3}t & t \le 1/2 \\ \frac{1}{3} + \frac{2}{3}t & t \ge 1/2 \end{cases}$$

and

$$f_1(t) = \begin{cases} \frac{1}{3}t & t \le 1/2 \\ -\frac{1}{3} + \frac{4}{3}t & t \ge 1/2 \end{cases}$$

The straight-line homotopy F in the interval shows that $f_0 \simeq f_1$ as paths. But

$$(\alpha \cdot \beta) \cdot \gamma = \delta \circ f_0$$

and

$$\alpha \cdot (\beta \cdot \gamma) = \delta \circ f_1$$

so $\delta \circ F$ is the required homotopy of paths.

(ii) Consider the functions $I \to I$ defined by

$$g_0(t) = \begin{cases} 2t & t \le 1/2 \\ 1 & t \ge 1/2 \end{cases}$$

and

$$g_1(t) = \begin{cases} 1 & t \le 1/2 \\ 2t-1 & t \ge 1/2 \end{cases}$$

for $t \in I$. Again, straight-line homotopy shows that $g_0 \simeq \mathrm{id}_I \simeq g_1$ as paths and, since $\alpha \cdot c_x = \alpha \circ g_0$ and $c_w \cdot \alpha = \alpha \circ g_1$, the result follows.

(iii) Consider the function $I \to I$ defined by

$$h(t) = \begin{cases} 2t & t \le 1/2 \\ 2-2t & t \ge 1/2 \end{cases}$$

for all $t \in I$. Straight-line homotopy shows that $h \simeq c_0$ as paths in I. Since $\alpha \cdot \bar{\alpha} = \alpha \circ h$ and $c_w = \alpha \circ c_0$, the result follows.

□

This completes the proof of Theorem 1.15, so we can talk about the fundamental group of a *based space*, i.e. a space equipped with a base point.

Example 1.18. Let $X = \mathbb{R}^n$ and $x_0 = 0$. Consider a based loop γ. The deformation retraction of Example 1.6 shows that $\gamma \simeq c_0$. Since γ was arbitrary, it follows that $\pi_1(\mathbb{R}^n, 0)$ is trivial.

Let's generalise this computation to some formal properties of fundamental groups.

Lemma 1.19. *Let $f : X \to Y$ be a map with $f(x_0) = y_0$. The induced map $f_* : \pi_1(X, x_0) \to \pi_1(Y, y_0)$ defined by*

$$f_*[\gamma] = [f \circ \gamma]$$

is a well-defined group homomorphism. Furthermore:

(i) if $f' \simeq f$ relative to x_0 then $f'_* = f_*$;

(ii) if $g : Y \to Z$ is a map with $g(y_0) = z_0$ then $g_* \circ f_* = (g \circ f)_*$;

(iii) $(\mathrm{id}_X)_* = \mathrm{id}_{\pi_1(X, x_0)}$.

Proof. The map f_* is well defined because, if F exhibits a homotopy between γ and δ as paths, then $f \circ F$ exhibits a homotopy of paths between $f \circ \gamma$ and $f \circ \delta$. It's a group homomorphism because $f \circ (\gamma \cdot \delta) = (f \circ \gamma) \cdot (f \circ \delta)$. We now check the other properties.

(i) If F exhibits a homotopy between f' and f relative to x_0 then, for any loop γ based at x_0,
$$(s, t) \mapsto F(\gamma(s), t)$$
is a homotopy of paths between $f \circ \gamma$ and $f' \circ \gamma$. Therefore $f_*[\gamma] = f'_*[\gamma]$ as required.

(ii) This is immediate from associativity of composition of functions.

(iii) This is also immediate.

$\square$

At this point, we know that the fundamental group is unchanged by homotopy equivalences *which fix the base point.* But the dependence on the base point is disappointing. We'd like to know that different choices of base points give isomorphic fundamental groups, at least when X is path connected. So we need to think about what happens when we change base points, and to this end, we define another natural map on fundamental groups.

Lemma 1.20. *Let X be a space. A path α from x_0 to x_1 defines a group isomorphism*

$$\alpha_\# : \pi_1(X, x_0) \to \pi_1(X, x_1)$$

via

$$\alpha_\#[\gamma] = [\bar{\alpha} \cdot \gamma \cdot \alpha] \ .$$

Furthermore:

(i) if $\alpha \simeq \alpha'$ then $\alpha_\# = \alpha'_\#$;

(ii) $(c_{x_0})_\# = \mathrm{id}_{\pi_1(X,x_0)}$;

(iii) if β is a path from x_1 to x_2 then $(\alpha \cdot \beta)_\# = \beta_\# \cdot \alpha_\#$;

(iv) if $f : X \to Y$ is a map with $y_i = f(x_i)$ then $(f \circ \alpha)_\# \circ f_* = f_* \circ \alpha_\#$.

Proof. To see that $\alpha_\#$ is a group homomorphism, we notice that

$$\begin{aligned} (\bar{\alpha} \cdot \gamma \cdot \alpha) \cdot (\bar{\alpha} \cdot \delta \cdot \alpha) &= \bar{\alpha} \cdot \gamma \cdot (\alpha \cdot \bar{\alpha}) \cdot \delta \cdot \alpha \\ &\simeq \bar{\alpha} \cdot \gamma \cdot c_{x_0} \cdot \delta \cdot \alpha \\ &\simeq \bar{\alpha} \cdot (\gamma \cdot \delta) \cdot \alpha \end{aligned}$$

for any loops γ and δ based at x_0. By Lemma 1.17 (iii), we see that $\alpha_\#^{-1} = \bar{\alpha}_\#$ and so $\alpha_\#$ is an isomorphism.

(i) Let γ be any loop based at x_0. If F is a homotopy of paths between α and α', we can glue together two copies of F with a copy of $\gamma \times \mathrm{id}_I$ to obtain a homotopy of paths $\bar{\alpha} \cdot \gamma \cdot \alpha \simeq \bar{\alpha}' \cdot \gamma \cdot \alpha'$, which proves (i).

(ii) This follows from Lemma 1.17(i-ii).

(iii) This follows immediately from the definition and Lemma 1.17(i).

(iv) This also follows immediately from the definition.

∎

This helps to reduce dependence on the base point. In particular, it makes sense to talk about the isomorphism type of the fundamental group of a path-connected space: if X is path-connected then $\pi_1(X, x_0) \cong \pi_1(X, x_1)$ for any $x_0, x_1 \in X$. But care is needed! If you want to understand more precisely what happens when we change base points, you should also meditate on the following simple observation.

Remark 1.21. If α is a loop in X based at x_0 then $\alpha_\#$ is conjugation by $[\alpha]$ in $\pi_1(X, x_0)$.

Since the isomorphism type of the fundamental group of a path-connected space is well defined, the following definition now makes sense.

Definition 1.22. If X is path-connected and $\pi_1(X, x_0)$ is trivial for some (any) $x_0 \in X$, we say that X is *simply connected.*

We now need to understand what happens to fundamental groups under homotopies that do not fix base points. Let X be a topological space, $x_0 \in X$, and let $f, g : X \to Y$ be continuous maps, homotopic via a homotopy F. Then

$$\alpha(t) = F(x_0, t)$$

is a path in Y from $f(x_0)$ to $g(x_0)$. The next lemma explains the different maps induced on fundamental group.

Lemma 1.23. *The triangle*

$$\begin{tikzcd} & \pi_1(Y, f(x_0)) \arrow[d, "\alpha_\#"] \\ \pi_1(X, x_0) \arrow[ur, "f_*"] \arrow[r, "g_*"] & \pi_1(Y, g(x_0)) \end{tikzcd}$$

commutes, which is to say that

$$\alpha_\# \circ f_* = g_* \ .$$

Proof. We need to check that

$$\bar{\alpha} \cdot (f \circ \gamma) \cdot \alpha \simeq g \circ \gamma$$

as paths, for any loop γ in X based at x_0. The map $G : I \times I \to Y$ defined by

$$G(s,t) = F(\gamma(s),t)$$

is a map from the square $I \times I$ to Y. We consider two paths from $(0,1)$ to $(1,1)$ in the square:

$$a(t) = (t,1)$$

and $b = b_1 \cdot b_2 \cdot b_3$, where

$$b_1(t) = (0,1-t),\ b_2(t) = (t,0),\ b_3(t) = (1,t)$$

for all $t \in I$. Note that

$$G \circ a(t) = F(\gamma(t),1) = g \circ \gamma(t)\ ,$$

and

$$G \circ b = \bar{\alpha} \cdot (f \circ \gamma) \cdot \alpha\ .$$

But $I \times I$ is a convex domain in $\mathbb{R}^2$, so the straight-line homotopy H shows that a and b are homotopic as paths. Therefore, $G \circ H$ is the homotopy we need. □

We can now prove that homotopy equivalences induce isomorphisms.

Theorem 1.24. *If $f : X \to Y$ is a homotopy equivalence then*

$$f_* : \pi_1(X,x_0) \to \pi_1(Y,f(x_0))$$

is an isomorphism for any $x_0 \in X$.

Proof. We need to prove that f_* is bijective. Let g be a homotopy inverse to f, and F a homotopy from $g \circ f$ to id_X. Let $\alpha(t) = F(x_0,t)$. Then, by several lemmas above,

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

which is an isomorphism, so f_* is injective and g_* is surjective. Now let G be a homotopy from $f \circ g$ to id_Y and let $\beta(t) = G(f(x_0),t)$. Then, similarly,

$$f_* \circ g_* = (f \circ g)_* = \beta_\# \circ (\mathrm{id}_Y)_* = \beta_\#$$

and, since this is an isomorphism, g_* is surjective, so an isomorphism. (Care is needed here, since the f_* in the second equation is not the same as the f_* in the first equation – the base points are different. However, the two g_*'s are the same.) Hence

$$f_* = (g_*)^{-1} \circ \alpha_\#$$

is an isomorphism, as required. □

Corollary 1.25. *Contractible spaces are simply connected.*

Proof. If X is contractible then there is $x_0 \in X$ and a homotopy F between id_X and constant map $X \to x_0$. In particular, $F(x, \cdot)$ is a path from arbitrary $x \in X$ to x_0, so X is path connected. Now X is simply connected by Theorem 1.24. □

We therefore have a tool for proving that spaces are simply connected. However, next we need some examples of spaces with *non-trivial* fundamental groups.

2 Covering spaces

2.1 Definition and first examples

Informally, a covering space is a way of 'unwrapping' the loops in a space X.

Definition 2.1. Let $p : \widehat{X} \to X$ be a continuous map. An open set $U \subseteq X$ is *evenly covered* if there is a discrete set Δ_U and an identification

$$p^{-1}(U) \equiv \Delta_U \times U$$

so that, on $p^{-1}(U)$, p coincides with projection onto U. We often write $U_\delta := \{\delta\} \times U$ and $p_\delta := p|_{U_\delta}$. Note that, if $x \in U$, then each $\{\delta\} \times U$ contains exactly one element of $p^{-1}(x)$, and so we may canonically identify Δ_U with $p^{-1}(x)$. In particular, $p^{-1}(U) \equiv \coprod_{\delta \in \Delta_U} U_\delta$.

If every point in X has an evenly covered neighbourhood then *p is a covering map*, and *$\widehat{X}$ is a covering space of X*.

Clearly homeomorphisms are covering maps, and it is also easy to see that products of covering maps are covering maps. Here's the first really interesting example.

Example 2.2. Let $\widehat{X} = \mathbb{R}$, $X = S^1 \subseteq \mathbb{C}$, and $p : \mathbb{R} \to S^1$ defined by

$$p(t) = e^{2\pi i t}$$

for all t. Let $U \subseteq S^1$ be any proper open subset of S^1 that contains 1, and let $z_1 \notin U$. Then we may choose a branch of log that's well defined on $S^1 - z_1$, with $\log(1) = 0$. Now every point $\hat{z} \in p^{-1}(U)$ can be written uniquely as

$$\hat{z} = k + \frac{\log(z)}{2\pi i}$$

for some $k \in \mathbb{Z}$ and $z = p(\hat{z})$. Therefore, $p^{-1}(U) \equiv \mathbb{Z} \times U$, and U is evenly covered with $\Delta_U = \mathbb{Z}$. Since every $z \in S^1$ is in some such U, it follows that p is a covering map.

Example 2.3. Let $\widehat{X} = X = S^1 \subseteq \mathbb{C}$, and let $p_n : \widehat{X} \to X$ be defined by

$$z \mapsto z^n$$

for all z. Again, let $U \subseteq S^1$ be any proper open subset of S^1 that contains 1, and let $z_1 \notin U$. We may choose a branch of $\sqrt[n]{}$ that's well defined on $S^1 - z_1$, with $\sqrt[n]{1} = 1$. Now every point $\hat{z} \in p^{-1}(U)$ can be written uniquely as

$$\hat{z} = e^{2\pi i \frac{k}{n}} \sqrt[n]{z}$$

for $z = p_n(\hat{z})$. As in the previous example, we obtain that U is evenly covered with Δ_U equal to the set of nth roots of unity, and since every $z \in S^1$ is in some such U, it follows that p is a covering map.

Example 2.4. Let $\widehat{X} = S^2$, let $G = \mathbb{Z}/2\mathbb{Z}$ act on S^2 via the antipodal map $x \mapsto -x$, and let X be the the quotient $G\backslash S^2$ – that is, the set of G-orbits on S^2, equipped with the quotient topology. Let $p : S^2 \to X$ be the quotient map.

A point in X is a pair of antipodal points $\{x, -x\}$, or equivalently the intersection of S^2 with a line $l \subseteq \mathbb{R}^3$ through the origin. Let $U \subseteq X$ be the set of lines that make an angle strictly less than $\pi/2$ with l. Then

$$p^{-1}(U) = S^2 - C_l$$

where $C_l \subseteq S^2$ is the great circle of points in S^2 perpendicular to l. In particular,

$$p^{-1}(U) = U_+ \sqcup U_-$$

where $U_\pm$ is the set of points at spherical distance less than $\pi/2$ from $\pm x$. Furthermore, it follows from the definition of the quotient topology that each $p|_{U_\pm} : U_\pm \to U$ is a homeomorphism, and $U_+ \sqcup U_- \cong \{\pm 1\} \times U$. Therefore, p is a covering map.

The space X is called the *real projective plane*, and is denoted by $\mathbb{R}P^2$.

Remark 2.5. In the same way, we can construct *real projective n-space*, $\mathbb{R}P^n$, as the quotient of S^n by the antipodal action of $\mathbb{Z}/2\mathbb{Z}$.

All of the covering spaces that we have seen above have an interesting property: every point in the base has the same number of pre-images in the cover.

Definition 2.6. Let $n \in \mathbb{N} \cup \{\infty\}$. A covering map $p : \widehat{X} \to X$ is n-sheeted if $\#p^{-1}(x) = n$ for any $x \in X$. In this case, we call n the *degree* of p.

2.2 Lifting properties

How are covering spaces related to the fundamental group? Informally, we can think of a covering space $\widehat{X}$ as *unwrapping* X. This also has the effect of *unwrapping* loops in X, which enables us to see that they are homotopically non-trivial. The next definition makes this precise.

Definition 2.7. Let $p : \widehat{X} \to X$ be a covering map and $f : Y \to X$ a continuous map. A *lift of f to $\widetilde{X}$* is a continuous map $\hat{f} : Y \to \widehat{X}$ so that $f = p \circ \hat{f}$. That is, the following diagram commutes.

$$\begin{tikzcd} & \widehat{X} \arrow[d, "p"] \\ Y \arrow[ur, "\hat{f}"] \arrow[r, "f"] & X \end{tikzcd}$$

The first important fact about lifts is that they are unique, in a suitable sense. More precisely, if Y is connected then a lift of f is determined by its value at a point. We will also prove our results in this subsection under the hypotheses that some of the spaces we consider are locally path-connected.

Definition 2.8. A space X is *locally path-connected* if, for every $x \in X$ and every open neighbourhood U of X, there is a path-connected neighbourhood V of x with $x \in V \subseteq U$.

In general these lemmas can be proved without the hypothesis of local path-connectedness – see, for instance, Hatcher – but all the spaces we consider in this course will be locally path-connected, so this hypothesis is fine.

Lemma 2.9. *Let $p : \widetilde{X} \to X$ be a covering map and $\hat{f}_1, \hat{f}_2 : Y \to \widehat{X}$ lifts of a map f from a connected, locally path-connected space Y to X. If there is $y_0 \in Y$ such that $\hat{f}_1(y_0) = \hat{f}_2(y_0)$ then $\hat{f}_1(y) = \hat{f}_2(y)$ for all $y \in Y$.*

Proof. The lemma follows from the claim that the set

$$S = \{y \in Y \mid \hat{f}_1(y) = \hat{f}_2(y)\}$$

is both open and closed; since Y is connected, it follows that S is either empty or the whole of Y, which proves the lemma.

Let $y_0 \in Y$ be arbitrary, let U be an evenly covered neighbourhood of $f(y_0)$, and let $V \subseteq f^{-1}(U)$ be a path-connected neighbourhood of y_0. We will show that $y_0 \in S$ if and only if $V \subseteq S$, which proves the claim. To this end, let $y \in V$ be arbitrary, and let α be a path in V from y_0 to y. Then $\hat{f}_i \circ \alpha$ is a path in $\widetilde{X}$ from $\hat{f}_i(y_0)$ to $\hat{f}_i(y)$; furthermore, since

$$p \circ \hat{f}_i \circ \alpha(t) = f \circ \alpha(t) \subseteq f(V) \subseteq U$$

we see that $\hat{f}_i \circ \alpha$ is a path in $p^{-1}(U)$. Therefore, for each i, $\hat{f}_i(y)$ and $\hat{f}_i(y_0)$ lie in the same path component of $p^{-1}(U)$; in particular, they lie in U_{δ_i} for some $\delta_i \in \Delta_U$.

Suppose now that $y_0 \in S$. Then $\hat{f}_1(y_0) = \hat{f}_2(y_0)$, so $\delta_1 = \delta_2$, since distinct U_{δ_i} are disjoint. Therefore,

$$\hat{f}_1(y) = p_{\delta_1}^{-1} \circ f(y) = p_{\delta_2}^{-1} \circ f(y) = \hat{f}_2(y)$$

and so $y \in S$.

Likewise, suppose that $y_0 \notin S$. Then $\hat{f}_1(y_0) \neq \hat{f}_2(y_0)$, but each U_δ contains a unique point of $p^{-1}(y_0)$, so $\delta_1 \neq \delta_2$. Since $\hat{f}_1(y)$ and $\hat{f}_2(y)$ lie in disjoint subsets of $p^{-1}(U)$, it follows that $\hat{f}_1(y) \neq \hat{f}_2(y)$ and so $y \notin S$. This completes the proof of the claim, and hence the lemma. □

In particular, lifts of paths are determined by their initial points.

Definition 2.10. Let $\gamma : I \to X$ be a path from x_0 and $p : \widehat{X} \to X$ a covering map. A *lift of* γ *at* $\hat{x}_0$ is a lift $\hat{\gamma}$ of γ with $\hat{x}_0 = \hat{\gamma}(0)$. Note that $\hat{x}_0$ necessarily is in $p^{-1}(x_0)$.

Having proved uniqueness, we next move on to existence. Our first lifting lemma tells us that lifts of paths always exist.

Lemma 2.11 (Path-lifting lemma). *Let $p : \widehat{X} \to X$ be a covering map and let $\gamma : I \to X$ be a path from x_0. For any $\hat{x}_0 \in p^{-1}(x_0)$ there is a unique lift $\hat{\gamma}$ of γ at $\hat{x}_0$.*

Proof. Uniqueness follows immediately from Lemma 2.9, so we just need to prove existence. We will prove that the set

$$S = \{t \in I \mid \gamma|_{[0,t]} \text{ lifts at } \hat{x}_0 \text{ to } \widehat{X}\}$$

is open and closed. Since $0 \in S$ and I is connected, it follows that $S = I$, which proves the theorem.

Let $t_0 \in I$ be arbitrary. Let U be an evenly covered neighbourhood of $\gamma(t_0)$ and let $V \subseteq \gamma^{-1}(U)$ be a path-connected neighbourhood of t_0. We will prove that $t_0 \in S$ if and only if $V \subseteq S$, from which it follows that S is open and closed, as required. Let $t \in V$ be arbitrary.

Suppose first that $t_0 \in S$ but $t \notin S$. Note in particular that $t > t_0$. Then $\hat{\gamma}$ makes sense at t_0, so $\hat{\gamma}(t_0) \in U_\delta$ for some δ. Since $[t_0, t] \subseteq V$, we have that $\gamma([t_0, t]) \subseteq U$. Therefore, the path

$$s \mapsto \begin{cases} \hat{\gamma}(s) & s \leq t_0 \\ p_\delta^{-1} \circ \gamma(s) & s \in [t_0, t] \end{cases}$$

is a lift of $\gamma|_{[0,t]}$ at $\hat{x}_0$, so $t \in S$, which is a contradiction. Therefore, S is open.

If $t_0 \notin S$ but $t \in S$, then the same argument with the roles of t_0 and t reversed leads to a contradiction. Therefore, S is closed, and the proof is complete. □

As a first simple application, we can now prove that the n-sheeted property of all our examples is not a coincidence.

Lemma 2.12. *If $p : \widehat{X} \to X$ is a covering map and X is path-connected then p is n-sheeted for some $n \in \mathbb{N} \cup \infty$.*

Proof. Let γ be a path in X from x to y. For any $\hat{x} \in p^{-1}(x)$, let $\hat{\gamma}_{\hat{x}}$ be the unique lift of γ at $\hat{x}$. We can now define a map

$$p^{-1}(x) \to p^{-1}(y)$$

by

$$\hat{x} \mapsto \hat{\gamma}_{\hat{x}}(1)\ .$$

Furthermore, the map $p^{-1}(y) \to p^{-1}(x)$ defined by $\bar{\gamma}$ provides an inverse, so this is a bijection. $\square$

Having lifted paths, we next lift homotopies.

Lemma 2.13 (Homotopy-lifting lemma). *Let $p : \widehat{X} \to X$ be a covering map and let $f_0 : Y \to X$ be a map from a locally path-connected space Y. Let $F : Y \times I \to X$ be a homotopy with $F(y,0) = f_0(y)$ for all y, and let $\hat{f}_0 : Y \to \widehat{X}$ be a lift of f_0. There is a unique lift $\widehat{F}$ of F to $\widehat{X}$ so that $\widehat{F}(\cdot, 0) = \hat{f}_0$.*

Proof. For each $y \in Y$, the homotopy F defines a path

$$\gamma_y(t) = F(y,t)$$

from $f_0(y)$. By the path-lifting lemma, each γ_y lifts at $\hat{f}_0(y)$ to a path $\hat{\gamma}_y$ in $\widehat{X}$. By the uniqueness of lifts, we must have

$$\widehat{F}(y,t) = \hat{\gamma}_y(t)$$

for all $y \in Y$ and $t \in I$. It remains to prove that $\widehat{F}$ is continuous. To do this, we define a different lift $\widetilde{F}$ of F which is continuous by definition, and prove that the two lifts agree.

Consider $y_0 \in Y$. For any t, $F(y_0, t)$ has an evenly covered neighbourhood U_t in X. By compactness of $\{y_0\} \times I$, we may take finitely many intervals $\{J_i\}$ that cover I and a path-connected neighbourhood V of y_0 so that, for each i, $F(V \times J_i)$ is contained in some evenly covered set U_i. Let U_{δ_i} be the unique slice of $p^{-1}(U_i)$ such that $\widehat{F}(\{y_0\} \times J_i) \subseteq U_{\delta_i}$.

For any $(y,t) \in V \times I$, we now define

$$\widetilde{F}(y,t) = p_{\delta_i}^{-1} \circ F(y,t)$$

whenever $t \in J_i$. We need to check that this is well defined. Suppose, therefore, that $t \in J_i \cap J_j$. Let α be a path in V from y_0 to y and let

$$\alpha_t(s) = F(\alpha(s), t)\ .$$

Then $p_{\delta_i}^{-1} \circ \alpha_t$ is a lift of α_t at $p_{\delta_i}^{-1} \circ \alpha_t(0)$ and, likewise, $p_{\delta_j}^{-1} \circ \alpha_t$ is a lift of α_t at $p_{\delta_j}^{-1} \circ \alpha_t(0)$. But

$$p_{\delta_i}^{-1} \circ \alpha_t(0) = \widehat{F}(y_0, t) = p_{\delta_j}^{-1} \circ \alpha_t(0)$$

so, by uniqueness of lifts, $p_{\delta_i}^{-1} \circ \alpha_t(1) = p_{\delta_j}^{-1} \circ \alpha_t(1)$. Therefore, $p_{\delta_i}^{-1} \circ F(y,t) = p_{\delta_j}^{-1} \circ F(y,t)$, which proves that $\widetilde{F}$ is well defined.

Since V is connected and $\widetilde{F}(\cdot, 0)$ is a lift of f_0 that agrees with $\hat{f}_0$ at y_0, we have that $\widetilde{F}(y,0) = \hat{f}_0(y)$ for all $y \in V$, by uniqueness of lifts. Now, for each $y \in V$, $\widetilde{F}(y,\cdot)$ is a lift of γ_y at $\hat{f}_0(y)$, and so $\widetilde{F}(y,t) = \hat{\gamma}_y(t)$ by uniqueness of lifts. Therefore, $\widetilde{F}$ and $\widehat{F}$ agree on $V \times I$. But $\widetilde{F}$ is continuous by construction, so $\widehat{F}$ is too. □

Since we are frequently interested in homotopies of paths, we also need to know that a homotopy *of paths* lifts to a homotopy *of paths*.

Lemma 2.14. *Let $p : \widehat{X} \to X$ be a covering map and let $F : I \times I \to X$ be a homotopy of paths. Then any lift $\widehat{F}$ of F to $\widehat{X}$ is also a homotopy of paths.*

Proof. Since F is a homotopy of paths, $F(\cdot, 0)$ and $F(\cdot, 1)$ are constant paths, at say x_0 and x_1 respectively. Now $\widehat{F}(\cdot, 0)$ is a path in the discrete set $p^{-1}(x_0)$, hence is constant. Similarly, $\widehat{F}(\cdot, 1)$ is constant. Therefore, $\widehat{F}$ is a homotopy of paths. □

2.3 Applications to calculations of fundamental groups

The lifting properties turn covering spaces into a tool for understanding fundamental groups.

Lemma 2.15. *Let $p : \widehat{X} \to X$ be a covering map with $p(\hat{x}) = x$. The induced map*

$$p_* : \pi_1(\widehat{X}, \hat{x}) \to \pi_1(X, x)$$

is injective.

Proof. We need to prove that $\ker p_*$ is trivial. Suppose therefore that $[\hat{\gamma}] \in \ker p_*$. That is to say, $\hat{\gamma}$ is a loop in $\widehat{X}$ based at $\hat{x}$ such that$\gamma = p \circ \hat{\gamma}$ is homotopic to c_x. By Lemma 2.14, the homotopy of paths F between γ and c_x lifts to a homotopy of paths between $\hat{\gamma}$ and $c_{\hat{x}}$. Therefore, $[\gamma]$ is trivial, as required. □

We may therefore identify $\pi_1(\widehat{X}, \hat{x}_0)$ with the subgroup $p_*\pi_1(\widehat{X}, \hat{x}) \leq \pi_1(X, x)$. We next combine this with an elaboration of the idea of the proof of Lemma 2.12.

Let $[\gamma] \in \pi_1(X,x)$ and let $p : \widehat{X} \to X$ be a covering map. A loop γ based at x naturally defines a map $p^{-1}(x) \to p^{-1}(x)$ by mapping

$$\hat{x} \mapsto \hat{\gamma}_{\hat{x}}(1) \ ,$$

the endpoint of the lift of γ at $\hat{x}$. Lemma 2.14 now implies that the endpoint $\hat{\gamma}_{\hat{x}}(1)$ only depends on the homotopy-class of the loop γ and therefore this map defines an *action* of $\pi_1(X,x)$ on $p^{-1}(x)$.

However, care is needed! Because of the order in which we concatenate paths, this action is a *right* action.[1] We will write $\hat{x}.\gamma := \hat{\gamma}_{\hat{x}}(1)$, and it is easy to check from the uniqueness of lifts that $(\hat{x}.\gamma).\delta = \hat{x}.(\gamma.\delta)$.

Covering spaces can be used to compute fundamental groups because this action has a natural algebraic interpretation.

Lemma 2.16. *Let $p : \widehat{X} \to X$ be a covering map and suppose that $\widehat{X}$ is path connected. Let $x \in X$. The map*

$$\begin{array}{rcl} p_*\pi_1(\widehat{X},\hat{x})\backslash\pi_1(X,x_0) & \to & p^{-1}(x) \\ p_*\pi_1(\widehat{X},\hat{x})[\gamma] & \mapsto & \hat{x}.\gamma \ . \end{array}$$

is a bijection for any choice $\hat{x} \in p^{-1}(x)$. Furthermore, this bijection is equivariant*, in the sense that*

$$p_*\pi_1(\widehat{X},\hat{x})[\gamma][\delta] \mapsto \hat{x}.(\gamma \cdot \delta)$$

for all loops γ, δ based at x.

Proof. The lemma is just the Orbit-Stabiliser theorem, applied to the right action of $\pi_1(X,x)$ on $p^{-1}(x)$. To conclude, we need to show that the action is transitive (so $p^{-1}(x)$ is the whole orbit) and that the stabiliser of $\hat{x}$ is $p_*\pi_1(\hat{X},\hat{x})$.

Transitivity follows immediately from the hypothesis that $\widehat{X}$ is path-connected: any $\hat{y} \in p^{-1}(x)$ is end point of a path $\hat{\gamma}$ from $\hat{x}$, whose image is a loop γ based at x. By uniqueness of lifts, $\hat{\gamma} = \hat{\gamma}_{\hat{x}}$, and so $\hat{x}.\gamma = \hat{y}$.

The stabiliser of $\hat{x}$ consists of the homotopy classes of those loops γ in X based at x whose lift at $\hat{x}$ is a loop. This is the definition of $p_*\pi_1(\widehat{X},\hat{x})$. $\square$

[1] Recall that a *right action* of a group G on a set X is a map $X \times G \to X$, denoted by $(x,g) \mapsto x.g$, so that $(x.g).h = x.(gh)$.

One easy consequence is that the *degree* of the covering map is equal to the index of the subgroup. So already the existence of a non-trivial connected covering space implies that the fundamental group is non-trivial.

Corollary 2.17. *Let $p : \widehat{X} \to X$ be a covering map. If $\widehat{X}$ and X are both path connected then*

$$\deg(p) = |\pi_1(X, x) : p_*\pi_1(\hat{X}, \hat{x})|$$

for any choice of $x \in X$ and $\hat{x} \in p^{-1}(x)$.

So Examples 2.2 and 2.3 show that $\pi_1(S^1, 1)$ is an infinite group, with subgroups of every index. Lemma 2.16 is particularly useful when $\widehat{X}$ is simply connected.

Definition 2.18. If $\widehat{X} \to X$ is a covering map and $\widehat{X}$ is simply connected then $\widehat{X}$ is called a *universal cover* of X.

In particular, we can completely understand the group structure of $\pi_1(X, x)$ by understanding the action loops by path-lifting on the preimage of the base point.

Corollary 2.19. *If $p : \widehat{X} \to X$ is a universal cover then any choice of $\hat{x} \in p^{-1}(x)$ defines a bijection*

$$\pi_1(X, x) \to p^{-1}(x) \ ,$$

and furthermore the group structure is determined by $\hat{x}.(\gamma.\delta) = (\hat{x}.\gamma).\delta$.

We can use this result to compute the fundamental group of the circle.

Example 2.20. Example 2.2 tells us that

$$\begin{array}{rcl} p : \mathbb{R} & \to & S^1 \\ t & \mapsto & e^{2\pi i t} \end{array}$$

is a covering map. Since $\mathbb{R}$ is contractible, this is a universal cover. By Corollary 2.19, path lifting at 0 defines a bijection $\pi_1(S^1, 1) \to p^{-1}(1) = \mathbb{Z}$. We next give an example of a loop in each homotopy class, and use the right action to compute the group operation.

For $n \in \mathbb{Z}$, consider

$$\tilde{\gamma}_n(t) = nt \ ,$$

which is a path in $\mathbb{R}$ from 0 to n. Then $\gamma_n = p \circ \tilde{\gamma}_n$ is a loop in S^1 based at 1, every element of $\pi_1(S^1, 1)$ is represented by such a loop, and the bijection with $\mathbb{Z}$ is given by $[\gamma_n] \mapsto n$. Furthermore, for any $m \in \mathbb{Z}$, the path $m + \tilde{\gamma}_n$ is the lift of γ_n based at $m = \gamma_m(1)$, so

$$0.(\gamma_m \cdot \gamma_n) = (0.\gamma_m).\gamma_n = m.\gamma_n = m + \tilde{\gamma}_n(1) = m + n = 0.\gamma_{m+n}$$

by the properties of right actions. Therefore, $[\gamma_m \cdot \gamma_n] = [\gamma_{m+n}]$, and so the bijection $\pi_1(S^1, 1) \to \mathbb{Z}$ is an isomorphism of groups.

2.4 The fundamental group of the circle

In the last section, we saw that $\pi_1(S^1, 1) \cong \mathbb{Z}$. This is our first example of a non-trivial fundamental group, and it has some attractive applications. The first is the *extension problem* of the kind we saw in the introduction.

Theorem 2.21. *The identity* id_{S^1} *does not extend to a map from the disc* $D^2 \to S^1$. *In other words,* S^1 *is not a retract of* D^2.

Proof. Suppose that $\mathrm{id}_{S^1} = r \circ \iota$ where $\iota : S^1 \to D^2$ is the natural inclusion and $r : D^2 \to S^1$ is a retraction. Since D^2 is contractible and hence simply connected, we have a factorisation

$$(\mathrm{id}_{S^1})_* = r_* \circ \iota_*$$

through the trivial group. Therefore $\mathrm{id}_{\mathbb{Z}}$ maps 1 to 0, which is a contradiction. $\square$

One attractive corollary is the *Brouwer fixed point theorem.*

Corollary 2.22. *Every continuous map* $f : D^2 \to D^2$ *has a fixed point.*

Proof. Let f be a continuous map $D^2 \to D^2$ without a fixed point. Let $g : D^2 \to S^1$ be defined by projecting $f(x)$ through x onto S^1. This is then a continuous retraction of D^2 onto S^1, which contradicts Theorem 2.21. $\square$

Another famous application is Gauss's proof of the fundamental theorem of algebra.

Theorem 2.23. *Every non-constant polynomial* $p : \mathbb{C} \to \mathbb{C}$ *has a root.*

Sketch proof of Theorem 2.23. Let $r : \mathbb{C} - 0 \to S^1$ be the usual retraction, and let $\lambda_R : S^1 \to \mathbb{C}$ be multiplication by R. If p has no root then

$$f_R = r \circ p \circ \lambda_R$$

defines a map $S^1 \to S^1$ for any $R > 0$. In particular, for any R_1, R_2, the two maps f_{R_1} and f_{R_2} are homotopic, so they induce the same homomorphism p_* on fundamental groups

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

which is equal to multiplication by some constant l.

The theorem now follows by noticing that, for very small R, $d = 0$, whereas for very large R, $d = \deg(p)$. $\square$

2.5 Existence of universal covers

We have seen that universal covers are a useful tool for computing fundamental groups. It's therefore an important fact that they can be guaranteed to exist in a wide variety of settings. The next theorem records this fact. The proof is non-examinable, and we only sketch it here.

Theorem 2.24. *If X is a path-connected, locally simply connected[2] topological space, then there exists a universal cover $p : \widetilde{X} \to X$.*

Sketch proof (non-examinable): The idea of the construction is a natural consequence of the lifting properties we have already proved. Fix a basepoint $x_0 \in X$, and consider the set of *all* paths γ starting at x_0.

$$\mathcal{X} = \{\gamma : I \to X \mid \gamma(0) = x_0\}$$

We define

$$\widetilde{X} \coloneqq \mathcal{X} / \simeq$$

the quotient of $\mathcal{X}$ by homotopy *of paths*. The map $p : \widetilde{X} \to X$ is defined by

$$p([\gamma]) = \gamma(1) \ ,$$

so it sends (homotopy classes of) paths to their endpoints.

Now the difficulty of the proof is to define a suitable topology on $\widetilde{X}$ and prove that p then has the desired properties. $\square$

[2]Recall that this means that, for every neighbourhood U of every point x, there exists a neighbourhood $V \subseteq U$ which is simply connected.

2.6 The Galois correspondence

Now we know that universal covers exist, we can go on and understand *all* covering spaces in terms of the algebra of the fundamental group. Roughly, we would like to say that path-connected covering spaces correspond to subgroups of the fundamental group. To make this precise, we need an appropriate notion of equivalence for covering spaces.

Definition 2.25. Let X be a path-connected space and $p_1 : \widehat{X}_1 \to X$, $p_2 : \widehat{X} \to X$ be covering maps. An *isomorphism of covering spaces* is a homeomorphism $\phi : \widehat{X}_1 \to \widehat{X}_2$ so that the following diagram commutes.

$$\begin{tikzcd} \widehat{X}_1 \arrow[rr, "\phi"] \arrow[dr, "p_1"'] & & \widehat{X}_2 \arrow[dl, "p_2"] \\ & X & \end{tikzcd}$$

That is, $p_2 \circ \phi = p_1$. Note that ϕ^{-1} is also an isomorphism of covering spaces. If $\widehat{X}_i$ are equipped with base points $\hat{x}_i$ and $\phi(\hat{x}_1) = \hat{x}_2$, then ϕ is said to be *based*.

Remark 2.26. The isomorphism ϕ is a lift of p_1 to $\widehat{X}_2$, so we can apply the results we already have about lifts to covering transformations. In particular, *based* isomorphisms are uniquely determined by where they send their base points.

We can now state a correspondence between covering spaces and subgroups of the fundamental group. This correspondence is often called the *Galois correspondence*, since it is analogous to (and, in the context of Riemann surfaces, equivalent to) the fundamental correspondence of Galois theory.

Theorem 2.27 (Galois correspondence with base points)**.** *Let X be a path-connected, locally simply connected space with base point x_0. The map that sends a covering space $p : \widehat{X} \to X$ equipped with a base point $\hat{x}_0 \in p^{-1}(x_0)$ to the subgroup $p_*\pi_1(\widehat{X}, \hat{x}_0)$ induces a bijection between the set of based-isomorphism-classes of path-connected covering spaces with base point and the set of subgroups of $\pi_1(X, x_0)$.*

The proof of this theorem is non-examinable, and hence omitted. But it's easy to illustrate the theorem in the example that we have studied most intently so far: the circle.

Example 2.28. Example 2.20 tells us that $\pi_1(S^1, 1) \cong \mathbb{Z}$. Therefore, every subgroup can be written uniquely as

$$H_n = \langle n \rangle$$

where $n \in \mathbb{N}$. The universal cover $p : \mathbb{R} \to S^1$ corresponds to the trivial subgroup $\langle 0 \rangle$, while the n-sheeted cover p_n (seen in Example 2.3) corresponds to H_n. Therefore, every connected covering space of S^1 is based isomorphic to exactly one of p_n or p.

One nice consequence is that universal covers are unique up to (based) isomorphism.

Corollary 2.29. *Let X be a path-connected, locally simply connected space. Any two universal covers $p_1 : \widetilde{X}_1 \to X$ and $p_2 : \widetilde{X}_2 \to X$ are isomorphic.*

Proof. Choose base points $x \in X$ and $\hat{x}_i \in p_i^{-1}(x)$. Since both $\widehat{X}_i$ are simply connected, they both map under the Galois correspondence to the trivial subgroup $1 \leq \pi_1(X, x)$. Therefore, by Theorem 2.27, they are isomorphic. □

Theorem 2.27 is beautiful, but the dependence on base points is annoying. Fortunately, we can deduce a base-point-free version.

Corollary 2.30 (Galois correspondence without base points)**.** *Let X be a path-connected, locally simply connected space with base point x_0. The map that sends a covering space $p : \widehat{X} \to X$ equipped with a base point $\hat{x}_0 \in p^{-1}(x_0)$ to the subgroup $p_*\pi_1(\widehat{X}, \hat{x}_0)$ induces a bijection between the set of isomorphism classes of path-connected covering spaces (*without distinguished base point*) and* conjugacy classes *of subgroups of $\pi_1(X, x_0)$.*

Proof. It follows immediately from Theorem 2.27 that the map is surjective. To see that it is injective, we need to show that if $p_{1*}\pi_1(\widehat{X}_1, \hat{x}_1)$ and $p_{2*}\pi_1(\widehat{X}_2, \hat{x}_2)$ are conjugate then there is an isomorphism of covering spaces $\phi : \widehat{X}_1 \to \widehat{X}_2$ (not necessarily respecting base points).

Suppose therefore that

$$p_{1*}\pi_1(\widehat{X}_1, \hat{x}_1) = [\gamma] p_{2*}\pi_1(\widehat{X}_2, \hat{x}_2)[\bar{\gamma}]$$

for some $[\gamma] \in \pi_1(X, x)$. Let $\hat{\bar{\gamma}}$ be the lift of $\bar{\gamma}$ at $\hat{x}_2$, and let $\hat{x}_2'$ be the end point of $\hat{\bar{\gamma}}$. By Lemma 1.20, we see that

$$[\gamma] p_{2*}\pi_1(\widehat{X}_2, \hat{x}_2)[\bar{\gamma}] = p_{2*}\hat{\gamma}_{\#}\pi_1(\widehat{X}_2, \hat{x}_2) = p_{2*}\pi_1(\widehat{X}_2, \hat{x}_2') \ .$$

Therefore,

$$p_{1*}\pi_1(\widehat{X}_1, \hat{x}_1) = p_{2*}\pi_1(\widehat{X}_2, \hat{x}_2') \ ,$$

so by Theorem 2.27, $p_1 : \widehat{X}_1 \to X$ and $p_2 : \widehat{X}_2 \to X$ are isomorphic. □

Note that we don't see the difference between Theorem 2.27 and Corollary 2.30 in our favourite example of $X = S^1$, because $\pi_1(S_1) \cong \mathbb{Z}$ is abelian, so every subgroup is normal and every conjugacy class of subgroups is a singleton.

3 The Seifert–van Kampen theorem

Our next task is to compute more non-trivial fundamental groups. Our main tool is the Seifert–van Kampen theorem, which is a 'gluing theorem': it computes the fundamental group of a space X which is the result of 'gluing together' two subspaces A and B. It tells us that $\pi_1(X)$ is constructed by 'gluing together' $\pi_1(A)$ and $\pi_1(B)$. To state the theorem, we need to know how to 'glue together' two groups.

3.1 Free groups and group presentations

You may have informally seen the notion of a group presentation already, if you have seen expressions like the following definition of a dihedral group.

$$D_{2n} = \langle r, s \mid s^2 = r^n = 1,\ srs = s^{-1} \rangle$$

To state the Seifert–van Kampen theorem clearly, we will need to define group presentations formally. We start with a *free groups*, which are the groups with no relations at all: that is, these will be the groups with presentations of the following form.

$$\langle a_1, a_2, \dots \mid \rangle$$

The set of generators $A = \{a_i\}$ is called an *alphabet*. It can be any set (though we will usually only need it to be countable), and the free group with those generators will be called $F(A)$. We now give the formal definition of a free group.

Definition 3.1. Let A be a set, let $F(A)$ be a group, and let $A \to F(A)$ be a map of sets. We say that $F(A)$ is the *free group on A* if it satisfies the following *universal property*. For any group G and any set map $A \to G$ there is a unique homomorphism of groups $f : F(A) \to G$ so that the following diagram commutes.

$$\begin{tikzcd} F(A) \arrow[dr, "f"] & \\ A \arrow[u] \arrow[r] & G \end{tikzcd}$$

We call the resulting homomorphism f the *canonical homomorphism.*

Roughly speaking, this definition says that $F(A)$ maps onto anything that can happen to A in a group.

Example 3.2. Let $A = \{\alpha\}$ contain one element. Let $A \to \mathbb{Z}$ be the map $\alpha \mapsto 1$. Given any assignment

$$\alpha \mapsto g \in G$$

for G a group, the homomorphism $f(n) = g^n$ satisfies the universal property.

Remark 3.3. We make some remarks on what Definition 3.1 means for existence and uniqueness.

(i) Definition 3.1 really is a definition! It implies that $F(A)$ is unique up to canonical isomorphism. Indeed, suppose that $A \to F'(A)$ also satisfies the universal property. Putting $G = F'(A)$ in the universal property for $F(A)$, we get a canonical homomorphism $f : F(A) \to F'(A)$. Putting $G = F(A)$ in the universal property for $F'(A)$, we get a canonical homomorphism $f' : F'(A) \to F(A)$. Next put $G = F(A)$ in the universal property of $F(A)$. Then $f' \circ f$ and $\mathrm{id}_{F(A)}$ both satisfy the properties for the canonical homomorphism, so by uniqueness, $f' \circ f = \mathrm{id}_{F(A)}$. Similarly, $f \circ f' = \mathrm{id}_{F'(A)}$, and so we obtain that $F(A)$ and $F'(A)$ are (canonically) isomorphic.

(ii) What Definition 3.1 does *not* make clear is that $F(A)$ exists. We shall defer that question until after the Seifert–van Kampen theorem. However, it is a fact that $F(A)$ always exists.

(iii) If A is a finite set of cardinality r (such as $\{a_1, \dots a_r\}$ or $A = \{a, b, \dots\}$) then we will write $F_r \equiv F(A)$ and call this *the free group of rank* r.

We identify elements $a \in A$ with their images in $F(A)$, so we can write expressions like

$$w = aba^{-1}b^{-1}$$

for elements of $F(A)$ if $a, b \in A$. Note that the subgroup $\langle A\rangle \leq F(A)$ also satisfies the universal property, so it must in fact be the whole of $F(A)$. Therefore, A generates $F(A)$.

Now that we have free groups, it is a simple matter to define presentations.

Definition 3.4. A *(group) presentation* consists of a set A and a subset of *relations* $R \subseteq F(A)$. It *presents* the group

$$\langle A \mid R\rangle := F(A)/\langle\!\langle R\rangle\!\rangle$$

(where $\langle\!\langle R\rangle\!\rangle$ denotes the normal closure of R, i.e. the smallest normal subgroup to contain R). The presentation is called *finite* if A and R are both finite sets.

Group presentations also have a universal property. Just as in Remark 3.3(i), this property defines $\langle A \mid R\rangle$ up to canonical isomorphism.

Lemma 3.5. *Consider the group $\langle A \mid R\rangle$ and the quotient map $q : F(A) \to \langle A \mid R\rangle$. Whenever G is a group and $f : F(A) \to G$ is a homomorphism such that $R \subseteq \ker f$, there is a unique homomorphism $g : \langle A \mid R\rangle \to G$ that makes the following diagram commute.*

$$\begin{tikzcd}
\langle A \mid R\rangle \arrow[dr, "g"] & \\
F(A) \arrow[u, "q"] \arrow[r, "f"'] & G
\end{tikzcd}$$

Proof. It is easy to see that

$$\langle\!\langle R\rangle\!\rangle = \{\prod_{i=1}^{n} \gamma_i r_i^{\pm 1} \gamma_i^{-1} \mid n \in \mathbb{N},\ r_i \in R,\ \gamma_i \in \Gamma\}$$

where $n = 0$ corresponds to the identity. In particular, since $f(r) = 1$ for all $r \in R$, it follows that $f(w) = 1$ for all $w \in \langle\!\langle R\rangle\!\rangle$.

An arbitrary element of $\langle A \mid R\rangle$ is of the form

$$q(w) = w\langle\!\langle R\rangle\!\rangle$$

for $w \in F(A)$. By hypothesis, we must define

$$g \circ q(w) = f(w)$$

so g, if it exists, is unique. Furthermore, if $q(w) = q(v)$ then $v^{-1}w \in \langle\!\langle R\rangle\!\rangle$, so $f(w) = f(v)$ and $g \circ q(w) = g \circ q(v)$. This shows that g is well defined. $\square$

Example 3.6. In Example 3.2 we saw that $F(\{a\}) \cong \mathbb{Z}$. Since $\mathbb{Z}$ is abelian, every subgroup is normal and so $\langle\langle a^n \rangle\rangle = \langle a^n \rangle$. Therefore

$$\langle a \mid a^n \rangle = \langle a \rangle / \langle a^n \rangle \cong \mathbb{Z}/n\mathbb{Z}$$

as you would expect.

Example 3.7. Define the dihedral group D_{2n} to be the group of symmetries of the regular n-gon. Standard arguments (as seen in *1a Groups*) show that this is generated by a rotation ρ through angle $2\pi/n$ and a reflection σ, that these satisfy the relations $\rho^n = 1$, $\sigma^2 = 1$ and $\sigma\rho\sigma = \rho^{-1}$, and that D_{2n} has elements.

Consider now the alphabet $A = \{r, s\}$ and the relations $R = \{r^n, s^2, rsrs\}$. The universal properties of free groups and presentations together imply that the assignment $r \mapsto \rho$ and $s \mapsto \sigma$ extends uniquely to a homomorphism

$$\langle r, s \mid r^n,\ s^2,\ rsrs \rangle \to D_{2n}$$

which is surjective since ρ, σ generate D_{2n}; in particular, $\langle r, s \mid r^n,\ s^2,\ rsrs \rangle$ is of cardinality at least $2n$. Using the relations R, we can show that any element of $\langle r, s \mid r^n,\ s^2,\ rsrs \rangle$ is an image of an element of $F(\{r, s\})$ of one of the forms

$$1, r, \ldots, r^{n-1}, s, sr, sr^2, \ldots, sr^{n-1}\ .$$

Therefore, the cardinality of $\langle r, s \mid r^n,\ s^2,\ rsrs \rangle$ is at most $2n$, and so the canonical map $\langle r, s \mid r^n,\ s^2,\ rsrs \rangle \to D_{2n}$ is an isomorphism.

Example 3.8. Let G be a group. Consider the identity set map $G \to G$. By the universal property of free groups, there is a canonical homomorphism $F(G) \to G$. Let R be the kernel of this map. It follows immediately that

$$G \cong \langle G \mid R \rangle$$

by the first isomorphism theorem. Thus, every group has a presentation – the *tautological* presentation. However, these presentations are very inefficient – indeed, they are never finite unless G is finite – and so they are not usually of much use.

We can now give a precise definition of what it means to 'glue two groups A and B along a group C'. The definition is again via a universal property.

Definition 3.9. Consider a commutative square of group homomorphisms.

$$\begin{array}{ccc} C & \xrightarrow{i} & A \\ \downarrow{\scriptstyle j} & & \downarrow{\scriptstyle k} \\ B & \xrightarrow{l} & \Gamma \end{array}$$

(That is, we assume that $k \circ i = l \circ j$.) This diagram is called a *pushout* if it satisfies the following universal property. Whenever G is a group and there are homomorphisms $f : A \to G$, $g : B \to G$ such that $f \circ i = g \circ j$, there exists a unique homomorphism $\phi : \Gamma \to G$ such that $f = \phi \circ k$ and $g = \phi \circ l$. This is illustrated in the following commutative diagram.

$$\begin{tikzcd} C \arrow[r, "i"] \arrow[d, "j"] & A \arrow[d, "k"] \arrow[ddr, bend left, "f"] & \\ B \arrow[r, "l"] \arrow[drr, bend right, "g"'] & \Gamma \arrow[dr, "\phi"] & \\ & & G \end{tikzcd}$$

As in Remark 3.3(i), this definition defines Γ, together with the maps k and l, uniquely up to canonical isomorphism. We may write $\Gamma = A \coprod_C B$ (suppressing the dependence on i, j).

This has a number of important special cases.

Definition 3.10. If C is trivial then $A \coprod_C B$ is called the *free product* of A and B, and is denoted by $A * B$. More generally, if i, j are injective then $A \coprod_C B$ is called the *free product with amalgamation* and denoted by $A *_C B$.

Example 3.11. The free product $\mathbb{Z} * \mathbb{Z}$ satisfies the universal property of the free group F_2, so $F_2 \cong \mathbb{Z} * \mathbb{Z}$. This argument generalises to show that

$$F(S_1 \sqcup S_2) \cong F(S_1) * F(S_2) \ ,$$

and so by induction we can see that

$$\underbrace{\mathbb{Z} * \ldots * \mathbb{Z}}_{r \text{ times}} \cong F_r$$

for any $r \in \mathbb{N}$.

At the other end of the spectrum, we have the situation where one of the groups A or B is trivial. In this case, we see that pushouts also generalise quotients.

Lemma 3.12. *For a map of groups $i : C \to A$ and the trivial map $j : C \to 1$, we have*

$$A \coprod_C 1 \cong A/\langle\!\langle i(C) \rangle\!\rangle \ .$$

Proof. We just need to check that the quotient $A/\langle\!\langle i(C) \rangle\!\rangle$ satisfies the desired universal property. Indeed, the hypotheses of the pushout diagram tell us to consider a homomorphism $f : A \to G$ such that, for all $c \in C$, $f \circ i(c) = 1$. But this implies that f descends to a unique map $\phi : A/\langle\!\langle i(C) \rangle\!\rangle \to G$, as required. □

More generally, if we have presentations for the groups A and B and generators for the group C then we can write down a presentation for the pushout $A \coprod_C B$.

Lemma 3.13. *Suppose that $A = \langle S_1 \mid R_1 \rangle$ and $B = \langle S_2 \mid R_2 \rangle$. Suppose furthermore that T is a generating set for C, and let $\tilde{i} : T \to F(S_1)$ be a lift of i and $\tilde{j} : T \to F(S_2)$ a lift of j. Then*

$$\Gamma = \langle S_1 \sqcup S_2 \mid R_1 \cup R_2 \cup \{\tilde{i}(t)^{-1}\tilde{j}(t) \mid t \in T\} \rangle$$

is a presentation for $A \coprod_C B$.

Proof. We need to check that the presentation satisfies the universal property of pushouts. Suppose we are given homomorphisms $f : A \to G$ and $g : B \to G$. Consider the following monstrous commutative diagram.

$$\begin{tikzcd}
T \arrow[rr, "\tilde{i}"] \arrow[dr] \arrow[dd, "\tilde{j}"] & & F(S_1) \arrow[d, "q_1"] \arrow[ddr, bend left, "\tilde{f}"] & \\
& C \arrow[r, "i"] \arrow[d, "j"] & A \arrow[d, "k"] \arrow[dr, "f"] & \\
F(S_2) \arrow[r, "q_2"] \arrow[rrr, bend right, "\tilde{g}"] & B \arrow[r, "l"] \arrow[rr, "g"] & \Gamma & G
\end{tikzcd}$$

where $q_1 : F(S_1) \to A$ and $q_2 : F(S_2) \to B$ are the natural quotient maps, $k : A \to \Gamma$ and $l : B \to \Gamma$ are the canonical homomorphisms induced by the natural maps $S_1 \to \Gamma$ and $S_2 \to \Gamma$, and $\tilde{f} = f \circ q_1$, $\tilde{g} = g \circ q_2$. In particular, note that $\tilde{f}(R_1) = \tilde{g}(R_2) = 1$ and $\tilde{f} \circ \tilde{i}(t) = \tilde{g} \circ \tilde{j}(t)$ for any $t \in T$. The universal property of free products now gives us a canonical homomorphism $\tilde{\phi} : F(S_1) * F(S_2) \to G$ that agrees with $\tilde{f}$ on $F(S_1)$ and $\tilde{g}$ on $F(S_2)$; in particular, $\tilde{\phi}(R_1) = \tilde{\phi}(R_2) = 1$. Furthermore, for any $t \in T$,

$$\tilde{\phi} \circ \tilde{i}(t) = \tilde{f} \circ \tilde{i}(t) = \tilde{g} \circ \tilde{j}(t) = \tilde{\phi} \circ \tilde{j}(t)$$

so $\tilde{\imath}(t)^{-1}\tilde{j}(t) \in \ker \tilde{\phi}$. Therefore, by the universal property of presentations, $\tilde{\phi}$ induces a unique homomorphism $\phi : \Gamma \to G$. It remains to check that ϕ makes the diagram commute.

For an arbitrary element $a = q_1(w) \in A$ (with $w \in F(S_1)$), we have that

$$\phi \circ k(a) = \tilde{\phi}(w) = \tilde{f}(w) = f \circ q_1(w) = f(a)$$

as required, and similarly $\phi \circ l(b) = g(b)$ for $b \in B$. So ϕ does indeed make the diagram commute, as required. □

3.2 The Seifert–van Kampen theorem for wedges

We'll start with a simple version of the theorem.

Definition 3.14. Let X, Y be spaces equipped with base points x_0, y_0 respectively. The *wedge* of X and Y is defined to be

$$X \vee Y := X \sqcup Y / \sim$$

where $\sim$ is the smallest equivalence relation such that $x_0 \sim y_0$. Note that as long as X, Y are path-connected, the wedge only depends on the choices of x_0, y_0 up to homotopy equivalence, so the fact that the notation suppresses the base points isn't a big problem.

We can now state our first version of the Seifert–van Kampen theorem.

Theorem 3.15 (Seifert–van Kampen theorem for wedges)**.** *Suppose that $X = Y_1 \vee Y_2$ with Y_1, Y_2 path-connected and x_0 the wedge point. Suppose, furthermore, that Y_1 and Y_2 each contain an open neighbourhood that deformation retracts to the wedge point. Then*

$$\pi_1(X, x_0) = \pi_1(Y_1, x_0) * \pi_1(Y_2, x_0) \ .$$

Sketch proof (non-examinable): Consider homomorphisms $f_i : \pi_1(Y_i, x_0) \to G$. We need to construct the canonical homomorphism $\phi : \pi_1(X, x_0) \to G$.

First, we modify X by a homotopy equivalence to make it easier to analyse. Consider

$$X' = Y_1 \sqcup I \sqcup Y_2 / \sim$$

so that $0 \in I$ is glued to the base point of Y_1 and $1 \in I$ is glued to the base point of Y_2. Then it is not difficult to prove that $X' \simeq X$ (cf. Example Sheet 1, question 9). Let's take $x_0 = 1/2 \in I$, and write $Y_1' = Y_1 \cup [0, 1/2]$ and $Y_2' = Y_2 \cup [1/2, 1]$.

Consider a loop γ in X' based at x_0. We can 'straighten' γ so that it can be written as a concatenation

$$\gamma = \alpha_1 \cdot \beta_1 \cdot \alpha_2 \cdot \ldots \cdot \beta_{n-1} \cdot \alpha_n \cdot \beta_n$$

where the α_i are loops in Y_1' based at x_0 and the β_i' are loopsin Y_2' based at x_0. Now we are forced to set

$$\phi(\gamma) = f_1(\alpha_1) f_2(\beta_1) \ldots f_1(\alpha_n) f_2(\beta_n) \ .$$

It follows easily that ϕ, if well defined, is unique and a homomorphism, as required, but we still need to check that it is well defined.

To check that it is well defined, we need to show that ϕ was independent of the homotopy-representative γ. Let

$$\gamma' = \alpha_1' \cdot \beta_1' \cdot \alpha_2' \cdot \ldots \cdot \beta_{m-1}' \cdot \alpha_m' \cdot \beta_m'$$

be another 'straightened' loop in the same homotopy-class as γ, and let F be a homotopy of paths from γ to γ'. We can also 'straighten' the homotopy F so that it is always *transverse* to x_0. This means that $F^{-1}(x_0)$ is a finite embedded union of arcs and circles in the square $I \times I$, where the endpoints of the arcs lie in the boundary of $I \times I$.

There are now various cases to consider. If there is a circle in $F^{-1}(x_0)$ then the disc that it bounds can be cut out and removed (since the whole boundary is sent to x_0). If there is an arc of $F^{-1}(0)$ starts and ends on γ then that shows that a sub-path of γ is homotopy as a path to x_0. This shows that γ can be modified by a homotopy to reduce n without changing $\phi(\gamma)$. Likewise, if there is an arc in $F^{-1}(x_0)$ with endpoints on γ' thenwe can modify γ' by a homotopy to reduce m. A similar argument applies if we have an arc in $F^{-1}(0)$ with one endpoint on γ or γ' and one endpoint on the constant sides of the square.

At the end of this process, $F^{-1}(0)$ consists of a finite number of arcs, each with one endpoint on γ and the other on γ'. This shows that $m = n$, and also that $\alpha_i \simeq \alpha_i'$ and $\beta_i \simeq \beta_i'$ as paths, for each i. It follows that $\phi(\gamma) = \phi(\gamma')$, so ϕ is well defined. □

Let's give some applications immediately.

Example 3.16. Theorem 3.15 combined with Example 2.20 and Lemma 3.11 show us that

$$\pi_1(S^1 \vee S^1) \cong \mathbb{Z} * \mathbb{Z} \cong F_2 \ .$$

In particular, the free group of rank 2 exists.

Similar considerations tell us that all free groups exist.

Example 3.17. Let A be any finite set and consider the wedge product

$$\bigvee_A S^1 := (A \times S^1)/\sim$$

where $\sim$ is the smallest equivalence relation that identifies $(a, 1)$ for all $a \in A$. It now follows that

$$\pi_1(\bigvee_A S^1) \cong F(A)$$

by induction, just as in Example 3.16. Thus all free groups of finite rank F_r exists.

Remark 3.18. In fact, it is not difficult to show that $\pi_1(\bigvee_A S^1) \cong F(A)$ for any set A, so in fact all free groups exist.

3.3 The full theorem

In fact, the theorem holds in much greater generality. We are interested in the situation where our topological space X is covered by sets Y_1 and Y_2, with intersection Z. We'll write $X = Y_1 \cup_Z Y_2$ for this situation.

Theorem 3.19 (Seifert–van Kampen theorem)**.** *Suppose that $Y_1, Y_2 \subseteq X$ are open subsets and $X = Y_1 \cup_Z Y_2$ with Y_1, Y_2, Z all path-connected. Let $x_0 \in Z$, and let $i_k : Z \to Y_i$, $j_k : Y_i \to X$ be the inclusion maps (for $k = 1, 2$). Then the diagram*

$$\begin{tikzcd}
\pi_1(Z, x_0) \arrow[r, "i_1"] \arrow[d, "i_2"] & \pi_1(Y_1, x_0) \arrow[d, "j_1"] \\
\pi_1(Y_2, x_0) \arrow[r, "j_2"] & \pi_1(X, x_0)
\end{tikzcd}$$

is a pushout.

The proof is a more elaborate version of the proof of Theorem 3.15, and we won't give it here.

The conclusion of the theorem looks abstract, but in fact Lemma 3.13 shows us that if we have presentations of $\pi_1(Y_1)$ and $\pi_1(Y_2)$ and a good understanding of the maps i_1, i_2 then we can write down a presentation for $\pi_1(X)$. Let's see some important examples.

Example 3.20. Let S^n be the n-sphere, let $x_\pm = (\pm 1, 0, \dots, 0)$, let $U_\pm = S^n - \{x_\mp\}$ and let $V = S^n - \{x_+, x_-\}$. Then we see that

$$S^n = U_+ \cup_V U_-$$

so we can apply the Seifert–van Kampen theorem as long as V is path-connected.

Now, stereographic projection provides homeomorphisms $U_\pm \cong \mathbb{R}^n$, so these are contractible and hence simply connected. We may also easily write down a *horizontal* projection π of V to a cylinder $(-1,1) \times S^{n-1}$: for $x = (x_1, x_2, \dots x_n) \in V$ and $x' = (x_2, \dots, x_n) \in \mathbb{R}^{n-1}$, the map

$$\begin{aligned} \pi : V &\to (-1,1) \times S^{n-1} \\ x &\mapsto (x_1, x'/\|x'\|) \end{aligned}$$

is a homeomorphism. Thus we see that $V \simeq S^{n-1}$ which is path-connected as long as $n > 1$. Choosing a base point in V, the Seifert–van Kampen theorem now tells us that the commutative square

$$\begin{tikzcd}
\pi_1(V) \arrow[r] \arrow[d] & 1 \arrow[d] \\
1 \arrow[r] & \pi_1(S^n)
\end{tikzcd}$$

is a pushout, so it follows that $\pi_1(S^n)$ is trivial by Lemma 3.12. Therefore S^n is simply connected for $n \geq 2$.

An inconvenient feature of Theorem 3.19 is that the pieces Y_1, Y_2 are required to be open. In fact, we can work with closed pieces as long as they are sufficiently nice.

Definition 3.21. A subset $Y \subseteq X$ is called a *neighbourhood deformation retract* if Y has an open neighbourhood

$$Y \subseteq V \subseteq X$$

such that Y is a deformation retract of V.

It's not hard to see that the Seifert–van Kampen theorem works equally well when Y_1, Y_2 are closed and the intersection Z is a neighbourhood deformation retract.

Corollary 3.22. *Suppose that $Y_1, Y_2 \subseteq X$ are closed subsets and $X = Y_1 \cup_Z Y_2$ with Y_1, Y_2, Z all path-connected. Let $x_0 \in Z$, and let $i_k : Z \to Y_i$, $j_k : Y_i \to X$ be the inclusion maps (for $k = 1, 2$). Suppose furthermore that Z is a neighbourhood deformation retract in both Y_1 and Y_2. Then the diagram*

$$\begin{tikzcd}
\pi_1(Z, x_0) \arrow[r, "i_1"] \arrow[d, "i_2"] & \pi_1(Y_1, x_0) \arrow[d, "j_1"] \\
\pi_1(Y_2, x_0) \arrow[r, "j_2"] & \pi_1(X, x_0)
\end{tikzcd}$$

is a pushout.

Proof. Let

$$Z \subseteq V_2 \subseteq Y_2$$

be an open neighbourhood of Z so that Z is a deformation retract of V_2. Then $U_1 = Y_1 \cup V_2$ is open in X, and it is easy to see that Y_1 is a deformation retract of U_1. Similarly, Y_2 is a deformation retract of an open set $U_2 = Y_2 \cup V_1$, and Z is a deformation retract of the intersection $W = U_1 \cap U_2 = V_1 \cup V_2$. This allows us to rewrite the commutative square in the statement of the corollary as

$$\begin{tikzcd}
\pi_1(W, x_0) \arrow[r, "i_1"] \arrow[d, "i_2"] & \pi_1(U_1, x_0) \arrow[d, "j_1"] \\
\pi_1(U_2, x_0) \arrow[r, "j_2"] & \pi_1(X, x_0)
\end{tikzcd}$$

and the result now follows by applying Theorem 3.19. $\square$

3.4 Attaching cells

The Seifert–van Kampen theorem allows us to write down presentations for a very large class of 'reasonable' spaces.

Definition 3.23. Let X be a topological space and let $\alpha : S^{n-1} \to X$ be a continuous map. The space

$$X \cup_\alpha D^n := X \sqcup D^n / \sim$$

where $\sim$ is the smallest equivalence relation that identifies $\theta \in S^{n-1} = \partial D^n$ with $\alpha(\theta)$ is said to be the result of *attaching an n-cell to X*.

The Seifert–van Kampen theorem enables us to precisely compute the effect on the fundamental group of attaching an n-cell. We first start with the case when $n \geq 3$.

Lemma 3.24. *If $n \geq 3$ then the inclusion map $i : X \to X \cup_\alpha D^n$ induces an isomorphism on fundamental groups.*

Proof. The *mapping cylinder* of α is the space

$$M_\alpha := (X \sqcup S^{n-1} \times I) / \sim$$

where $(\theta, 0) \sim \alpha(\theta)$ for all $\theta \in S^{n-1}$. Identify S^{n-1} with $S^{n-1} \times \{1\} \subseteq M_\alpha$, and note that S^{n-1} is now a neighbourhood deformation retract in both M_α and D^n. Note also that

$$X \cup_\alpha D^n \cong M_\alpha \cup_{S^{n-1}} D^n \ .$$

We can now apply Corollary 3.22 with a basepoint in S^{n-1} to see that the diagram

$$\begin{array}{ccc} \pi_1(S^{n-1}, X_0) & \longrightarrow & \pi_1(M_\alpha, x_0) \\ \downarrow & & \downarrow \\ \pi_1(D^n, x_0) & \longrightarrow & \pi_1(X \cup_\alpha D^n, x_0) \end{array}$$

is a pushout. Since $\pi_1(S^{n-1}, X_0) \cong \pi_1(D^n, x_0) \cong 1$, it follows easily that the map $\pi_1(M_\alpha, x_0) \to \pi_1(X \cup_\alpha D^n, x_0)$ induced by inclusion is an isomorphism. Changing base points, this also applies to any base point in X, and since X is a deformation retract of M_α, the result follows. $\square$

In this sense, the fundamental group doesn't 'see' any n-dimensional information about the space if $n \geq 3$. Something more interesting happens when we attach a 2-cell. In this case, given a choice of base point on S^1, the map α defines a based loop in X, corresponding to an element $[\alpha]$ of the fundamental group of X. Attaching a 2-cell along this loop precisely corresponds to killing it in the fundamental group.

Lemma 3.25. *Let $\alpha : S^1 \to X$ be a continuous map, and choose a base point $x_0 = \alpha(\theta_0)$ for $\theta_0 \in S^1$. Then*

$$\pi_1(X \cup_\alpha D^2, x_0) \cong \pi_1(X, x_0)/\langle\langle [\alpha] \rangle\rangle$$

and the inclusion map i induces the quotient map on fundamental groups.

Proof. The proof proceeds in exactly the same way as Lemma 3.24, and we obtain the fact that, for a base point in $x_0 \in S^1$, the following diagram is a pushout.

$$\begin{tikzcd}
\pi_1(S^1, \theta_0) \arrow[r] \arrow[d] & \pi_1(X, x_0) \arrow[d] \\
\pi_1(D^2, \theta_0) \arrow[r] & \pi_1(X \cup_\alpha D^2, x_0)
\end{tikzcd}$$

Since D_2 is simply connected and $\alpha_* \pi_1(S^1, \theta_0) = [\alpha]$, the result follows from Lemma 3.12. □

Another way to state Lemma 3.25 is that if $\langle A \mid R \rangle$ is a presentation for $\pi_1 X$ and we attach a 2-cell to X along a loop α then the fundamental group of the resulting space has presentation $\langle A \mid R, [\gamma] \rangle$. This enables us to characterise very precisely the fundamental groups of spaces constructed in this way. In particular, we can construct spaces with a very wide range of fundamental groups!

Theorem 3.26. *For any finitely presented group G there is a compact space X such that*

$$\pi_1(X, x_0) \cong G$$

for any base point x_0.

Proof. Suppose that $G = \langle A \mid R \rangle$ for A and R finite. By example 3.17 there is a path-connected space $X^{(1)}$ such that $X^{(1) \cong F(A)}$. Applying Lemma 3.25 inductively, we obtain a path-connected space X with $\pi_1(X, x_0) \cong G$. □

In fact, using Example 3.8 and a little more care, *every* group is the fundamental group of a space X. However, we cannot in general assume that the space is compact.

3.5 The classification of surfaces

We'll finish our discussion of the fundamental group by computing the fundamental groups of another important class of examples.

Definition 3.27. An *n-dimensional (topological) manifold* is a Hausdorff topological space M such that every point $x \in M$ has a neighbourhood U homeomorphic to an open neighbourhood in $\mathbb{R}^n$. A *2-dimensional manifold* is called a *surface*.

Manifolds are the basic objects of study in topology. We can build a large number of examples by attaching 2-cells to wedges of circles. We start with the most basic example.

Example 3.28. Let $\alpha : S^1 \to *$ be the constant map and consider

$$X = * \cup_\alpha D^2 \ .$$

If we identify $\mathbb{R}^2$ with the interior of D^2 then stereographic projection defines a homeomorphism between the interior of D^2 and $S^2 - \{x_0\}$ for any point x_0. This homeomorphism extends to a continuous bijection

$$X \to S^2$$

which is a homeomorphism since X is compact and S^2 is Hausdorff.

This example extends naturally to a large family of examples.

Example 3.29. Let $g \in \mathbb{N}$ and consider the wedge

$$\Gamma_{2g} := \bigvee_{i=1}^{2g} S^1_i$$

where each S^1_i is a circle. For each $1 \le i \le g$, let $\alpha_i : I \to S^1_i$ and $\beta_i : I \to S^1_{i+g}$ be paths that restrict to homeomorphisms on the interior of I. We now consider the path

$$\rho_g := \alpha_1 \cdot \beta_1 \cdot \bar{\alpha}_1 \cdot \bar{\beta}_1 \cdot \alpha_2 \cdot \ldots \cdot \bar{\beta}_{g-1} \cdot \alpha_g \cdot \beta_g \cdot \bar{\alpha}_g \cdot \bar{\beta}_g$$

and let $\Sigma_g := \Gamma_{2g} \cup_{\rho_g} D^2$.

The space Σ_g is actually a compact surface. To see this, we need to check that every point in Σ_g has a neighbourhood homeomorphic to an open disc in $\mathbb{R}^2$. This is clear for points in the interior of D^2. For points $x \in \Gamma_{2g}$ which are not the wedge point, it's also easy to see: x is in the image of exactly two points of ∂D^2, each of which has a neighbourhood which is an open half-disc; these two half-discs glue together to give a disc.

Finally, we need to consider the cone point x_0. This is the image of $2g$ points in ∂D^2, which we should think of as the vertices of a regular $2g$-gon. Each of these has a small neighbourhood which we can think of as an open sector of angle π/g. These neighbourhoods necessarily glue together in Σ_g to make a wedge of discs, and in fact we can check from the combinatorics of the concatenation ρ_g that they glue together to make a single disc.

It's not hard to see that this corresponds to identifying sides of a $2g$-gon in a certain pattern of pairs. The surface Σ_g is called the *orientable surface of genus g*. Lemma 3.25 implies that

$$\pi_1 \Sigma_g \cong \langle a_1, \ldots, a_g, b_1, \ldots, b_g \mid a_1 b_1 a_1^{-1} b_1^{-1} \ldots a_g b_g a_g^{-1} b_g^{-1} \rangle \ .$$

Note that $\Sigma_0 \cong S^2$ and Σ_1 is homeomorphic to the 2-torus $T^2 := S^1 \times S^1$.

Example 3.30. Let $g \in \mathbb{N}$ and consider the wedge

$$\Gamma_{g+1} := \bigvee_{i=0}^{g} S^1_i$$

where each S^1_i is a circle. For each $0 \le i \le g$, let $\alpha_i : I \to S^1_i$ be a path that restricts to a homeomorphism on the interior of I. We now consider the path

$$\sigma_g := \alpha_0 \cdot \alpha_0 \cdot \alpha_1 \cdot \alpha_1 \cdot \ldots \alpha_h \cdot \alpha_g$$

and let $S_g := \Gamma_{g+1} \cup_{\sigma_g} D^2$. As in Example 3.29, S_g is a compact surface, called the *non-orientable surface of genus g*. Again, Lemma 3.25 implies that

$$\pi_1 S_g \cong \langle a_0, a_1, \ldots, a_g \mid a_0^2 a_1^2 \ldots a_g^2 \rangle \ .$$

It's a nice exercise to check that $S_0 \cong \mathbb{R}P^2$. The surface S_1 is the *Klein bottle*.

Amazingly, in dimension 2 we have a complete classification of compact manifolds.

Theorem 3.31 (Classification of compact surfaces)**.** *Any compact surface X is homeomorphic to Σ_g or S_g, or some g.*

The proof of this theorem is beyond the scope of this course, and so is omitted here. But Theorem 3.31 leaves open the following important question: are the surfaces Σ_g and S_g pairwise non-homeomorphic, or even non-homotopy-equivalent? We can answer this using our calculations of the fundamental group. The following lemma characterises the groups $\pi_1\Sigma_g$ and $\pi_1 S_g$ via their abelian quotient groups.

Lemma 3.32. *Let $g \in \mathbb{N}$. The group $\pi_1\Sigma_g$ surjects $\mathbb{Z}^{2g}$ but not $\mathbb{Z}^{2g} \oplus (\mathbb{Z}/2\mathbb{Z})$. The group $\pi_1 S_g$ surjects $\mathbb{Z}^g \oplus (\mathbb{Z}/2\mathbb{Z})$ but not $\mathbb{Z}^{g+1}$.*

Proof. We start with

$$\pi_1\Sigma_g \cong \langle a_1, \dots, a_g, b_1, \dots, b_g \mid a_1 b_1 a_1^{-1} b_1^{-1} \dots a_g b_g a_g^{-1} b_g^{-1} \rangle \ .$$

Let $\{\bar{a}_i, \bar{b}_i \mid 1 \le i \le g\}$ be a basis for $\mathbb{Z}^{2g}$. The assignment $a_i \mapsto \bar{a}_i$, $b_i \mapsto \bar{b}_i$ sends ρ_g to 0, so by Lemma 3.5, it extends to a surjective homomorphism.

On the other hand, suppose $f : \pi_1\Sigma_g \to \mathbb{Z}^{2g} \oplus (\mathbb{Z}/2\mathbb{Z})$ is a surjection. Composing with reduction modulo 2, we get a surjection $\pi_1\Sigma_g \to (\mathbb{Z}/2\mathbb{Z})^{2g+1}$, and so the latter is generated by the $2g$ elements $\{f(a_i), f(b_i) \mid 1 \le i \le g\}$, which is absurd.

Now consider

$$\pi_1 S_g \cong \langle a_0, a_1, \dots, a_g \mid a_0^2 a_1^2 \dots a_g^2 \rangle \ .$$

Let $\{\bar{a}_i \mid 1 \le i \le g\}$ be a basis for the $\mathbb{Z}^g$ factor of $\mathbb{Z}^g \oplus (\mathbb{Z}/2\mathbb{Z})$, and let $\bar{c}_0$ generate the $\mathbb{Z}/2\mathbb{Z}$ factor. The assignment $a_i \mapsto \bar{a}_i$ for $1 \le i \le g$ and

$$a_0 \mapsto \bar{c}_0 - \sum_{i=1}^{g} \bar{a}_i$$

sends σ_g to 0, so extends to a surjective homomorphism $\pi_1 S_g \to \mathbb{Z}^g \oplus (\mathbb{Z}/2\mathbb{Z})$.

On the other hand, suppose $f : \pi_1 S_g \to \mathbb{Z}^{g+1}$ is a surjection. Then the set $\{f(a_i) \mid 0 \le i \le g\}$ generates $\mathbb{Z}^{g+1}$, but

$$\sum_{i=1}^{g} 2f(a_i) = 0 \ ,$$

which is absurd. ∎

4 Simplicial complexes

We have used the fundamental group to prove some very nice theorems (such as the Brouwer fixed-point theorem), but its usefulness is limited by the fact that it doesn't tell us anything interesting about the higher-dimensional spheres, for instance. There are various kinds of 'higher-dimensional' invariants that one can define. In this course, we are going to use *simplicial homology*. But before we do that, we need to define the class of spaces for which it works. These are the *simplicial complexes* of the title of this section.

4.1 Simplices and stuff

Simplicial complexes form a large, but very well behaved, class of spaces. The basic building block is a *simplex* (plural *simplices*).

Definition 4.1. A finite set $V = \{v_0, \ldots, v_n\} \subseteq \mathbb{R}^m$ is said to be *in general position* if the smallest affine subspace that contains A is of dimension n. There are various other formulations that are easily seen to be equivalent:

(a) the set of vectors $\{v_1 - v_0, \ldots, v_n - v_0\}$ is linearly independent;

(b) for any scalars $s_1, \ldots, s_n$, if

$$\sum_{i=1}^{n} s_i(v_i - v_0) = 0$$

then every $s_i = 0$;

(c) for any scalars $t_0, \ldots, t_n$ such that $\sum_i t_0 = 0$, if

$$\sum_{i=0}^{n} t_i v_i = 0$$

then every $t_i = 0$.

In particular, $m \geq n$.

Simplices are the convex hulls of finite sets of points in general position.

Definition 4.2. Let $n \geq 0$. If $V = \{v_0, \ldots, v_n\} \subseteq \mathbb{R}^m$ then the set

$$\langle V \rangle := \left\{ \sum_{i=0}^{n} t_i v_i \;\middle|\; \sum_{i=0}^{n} t_i = 1\ ,\ t_i \geq 0 \right\}$$

is called the *span* of V. If V is in general position then $\langle V\rangle$ is said to be an n-*simplex*, and n is the *dimension* of $\langle V\rangle$.

Each simplex comes with a natural collection of subsimplices.

Definition 4.3. If $V \subseteq \mathbb{R}^n$ is in general position and $U \subseteq V$ then $\langle U\rangle \subseteq \langle V\rangle$ is said to be a *face* of the simplex $\langle V\rangle$, and we write $\langle U\rangle \leq \langle V\rangle$. If $U \neq V$ then $\langle U\rangle$ is a *proper* face of $\langle V\rangle$.

Remark 4.4. Since the $\varnothing \subseteq V$, every simplex σ has an empty face, which is the (empty) span of the empty set.

A *simplicial complex* is now defined to be a finite collection of simplices that can be glued together along their faces.

Definition 4.5. A *(geometric) simplicial complex* is a finite set K of simplices in $\mathbb{R}^m$ (for some suitably large m) such that:

(a) if $\sigma \in K$ and $\tau \leq \sigma$ then $\tau \in K$;

(b) if $\sigma, \tau \in K$ then the intersection $\sigma \cap \tau$ is a face of both σ and τ.

The *dimension* of K is the largest n such that K contains an n-simplex. The d-*skeleton* of K is defined to be

$$K_{(d)} := \{\sigma \in K \mid \dim(\sigma) \leq d\} \ .$$

Example 4.6. The set of faces of a simplex σ is clearly a simplicial complex.

Example 4.7. The set of proper faces of an n-simplex σ is also a simplicial complex, namely the $(n-1)$-skeleton of σ. It is called the *boundary* of σ and denoted by $\partial\sigma$. The set of points that are in σ but not in a simplex of $\partial\sigma$ are called the *interior* of σ, denoted by $\mathring{\sigma}$.

Remark 4.8. For $n > 0$, the interior $\mathring{\sigma}$ coincides with the usual notion. However, a 0-simplex σ has no non-empty proper faces, so the boundary is empty and $\mathring{\sigma} = \sigma$!

Simplicial complexes provide us with a convenient way of building spaces.

Definition 4.9. If K is a simplicial complex in $\mathbb{R}^m$, the *polyhedron* or *realisation* of K is

$$|K| := \bigcup_{\sigma \in K} \sigma \ ,$$

the union of the simplices in K. More generally, if X is any topological space, a *triangulation* of X is a homeomorphism $h : |K| \to X$, where K is some simplicial complex.

Example 4.10. Consider $\mathbb{R}^{n+1}$ equipped with a basis $e_0, \ldots, e_n$. The *standard n-simplex* is the span of the basis vectors, $\sigma_n := \langle e_0, \ldots, e_n \rangle$. It's easy to see that there exists a triangulation $h : \sigma_n \to D^n$.

Example 4.11. The *standard simplicial $(n-1)$-sphere* is defined to be $\partial\sigma_n$, where σ_n is the standard n-simplex. It's so-called because the triangulation $h : \sigma_n \to D^n$ restricts to a triangulation $|\partial\sigma_n| \to S^{n-1}$.

Example 4.12. As above, let $e_0, \ldots, e_n$ be the standard basis for $\mathbb{R}^{n+1}$. Let $E := \{\pm e_0, \ldots, e_n\}$ and let E_0 be the set of subsets $S \subseteq E$ such that, for each $i \in I$, we do not have both $e_i \in S$ and $-e_i \in S$. Each $S \in E_0$ is in general position, and so

$$K := \{\langle S \rangle \mid S \in E_0\}$$

is a set of simplices. If $T \subseteq S \in E_0$ then $T \in E_0$, from which it follows that K is closed under passing to faces. Furthermore, it is easy to check directly that, for any $S, T \in E_0$,

$$\langle S \rangle \cap \langle T \rangle = \langle S \cap T \rangle$$

from which it follows that K is closed under taking intersections. Therefore, K is a simplicial complex. For any vector $v \in \mathbb{R}^{n+1} - \{0\}$, it's easy to see that the ray from 0 through v passes through a unique point of $|K|$. It follows that the map

$$\begin{array}{rcl} h : |K| & \to & S^n \\ v & \mapsto & \dfrac{v}{\|v\|} \end{array}$$

is a triangulation.

Having defined our objects of study, we of course next need to define the maps between them.

Definition 4.13. Let K, L be simplicial complexes. A *simplicial map* is a map $f : K \to L$ such that:

(i) each 0-simplex $\langle v \rangle \in K$ is sent to a 0-simplex $\langle f(v) \rangle \in L$; and

(ii) $f(\langle v_0, \ldots, v_n \rangle) = \langle f(v_0), \ldots, f(v_n) \rangle$ for any $\langle v_0, \ldots, v_n \rangle \in K$.

Note that the images are *not* required to be in general position, just to span a simplex.

The *realisation* of $f : K \to L$ is the continuous map $|f| : |K| \to |L|$ defined to be equal to

$$f_\sigma\left(\sum_{i=0}^{m} t_i v_i\right) = \sum_{i=0}^{m} t_i f(v_i)$$

on $\sigma = \langle v_0, \dots, v_m \rangle$ Note that $|f|$ is well defined and continuous, since f_σ is continuous and $f_\sigma = f_\tau$ when $\tau \leq \sigma$.

Remark 4.14. The realisation of the 0-skeleton, $|K_{(0)}|$, is called the set of *vertices* of K, and denoted by V_K. A very convenient feature of simplicial maps is that they are determined by their values on the vertices.

4.2 Barycentric subdivision

Since we need to study arbitrary continuous maps, we would like to prove that every continuous map between realisations is homotopic to a realisation of a simplicial map. However, we quickly run into a problem.

Example 4.15. Let $K = \partial\sigma_2$, the standard simplicial circle. On the one hand, $|K| \cong S^1$, so there are infinitely many homotopy classes of maps $|K| \to |K|$. On the other hand, a simplicial map $K \to K$ is determined by the images of the three vertices, so there are at most $3^3 = 27$ simplicial maps $K \to K$. In particular, it is certainly not the case that every continuous map is homotopic to a realisation of a simplicial map.

We fix this problem by *subdividing*.

Definition 4.16. Let $V = \{v_0, \dots, v_n\}$ be in general position. The point

$$\hat{\sigma} \coloneqq \frac{1}{n+1}\sum_{i=0}^{n} v_n$$

is called the *barycentre* of the simplex $\sigma \coloneqq \langle V \rangle$.

Barycentres fit nicely together, and from this we obtain a way to subdivide a simplicial complex. The idea is to that the vertices of the subdivided complex should be the barycentres of the simplices of the original complex.

Definition 4.17. Let K be a simplicial complex. The *barycentric subdivision* K' is the simplicial complex defined as follows. The vertices are the barycentres of the simplices of K.

$$V_K \coloneqq \{\hat{\sigma} \mid \sigma \in K\}$$

These vertices span a simplex precisely when the corresponding collection of simplices is a nested collection of faces. That is

$$\langle \hat{\sigma}_1, \hat{\sigma}_2, \ldots, \hat{\sigma}_n \rangle \in K'$$

if and only if $\sigma_1 \leq \sigma_2 \leq \ldots \leq \sigma_n$.

It's not completely obvious that K' has the properties we expect, but this is checked in the following result.

Lemma 4.18. *Let K be a simplicial complex. Then K' is a simplicial complex, and $|K'| = |K|$.*

Proof. We first check that each of the spans in the definition really is a simplex. By removing redundancies, we may suppose that $\sigma_1 \nleq \sigma_2 \nleq \ldots \nleq \sigma_n$. Now suppose that $\sum_{i=0}^{n} t_i = 0$ and $\sum_{i=0}^{n} t_i \hat{\sigma}_i = 0$. Let j be maximal such that $t_j \neq 0$. It follows that

$$\hat{\sigma}_j = -\sum_{i=0}^{j-1} \frac{t_i}{t_j} \hat{\sigma}_i$$

which implies that $\hat{\sigma}_j$ is contained in a proper face of σ_j, a contradiction. Therefore the set of barycentres $\{\hat{\sigma}_1, \hat{\sigma}_2, \ldots, \hat{\sigma}_n\}$ is in general position, and so $\langle \hat{\sigma}_1, \hat{\sigma}_2, \ldots, \hat{\sigma}_n \rangle$ is a simplex, as claimed.

We next check that K' really is a simplicial complex. We do this by induction on $\dim K$. It's immediate from the definition that K' is closed under passing to faces, so it remains to check that K' is closed under taking intersections. The simplex $\langle \hat{\sigma}_1, \hat{\sigma}_2, \ldots, \hat{\sigma}_n \rangle$ of K' is contained in σ_n, so we may reduce to the case of two simplices $\sigma', \tau' \in K'$ both contained in some common simplex $\delta \in K$. If either σ' or τ' doesn't contain $\hat{\delta}$, then their intersection is contained in $\partial\delta$. If both σ', τ' contain $\hat{\delta}$, then their intersection is equal to the span of $\hat{\delta}$ and the intersection of $(\sigma' \cap \partial\delta) \cap (\tau' \cap \partial\delta)$. In either case, we reduce to the case of $\hat{\delta}$ which is of lower dimension.

Finally, it remains to check that $|K'| = |K|$. By definition, each simplex $\langle \hat{\sigma}_1, \hat{\sigma}_2, \ldots, \hat{\sigma}_n \rangle \in K'$ is contained in $\sigma_n \in K$, so $|K'| \subseteq |K|$. The reverse inclusionis again proved by induction on $\dim K$. Consider $\sigma = \langle v_0, \ldots, v_m \rangle \in K$ and $x \in \sigma$, with a view to showing that $x \in |K'|$. If $x = \hat{\sigma}$ there is nothing to prove. Otherwise, we may consider the projection $\pi(x)$ of x from $\hat{\sigma}$ onto the boundary $|\partial\sigma|$. By induction, $\pi(x)$ is in a simplex

$$\langle \hat{\sigma}_1, \hat{\sigma}_2, \ldots, \hat{\sigma}_{n-1} \rangle \in K'$$

where each σ_i is a proper face of σ. It now follows that $x \in \langle \hat{\sigma}_1, \hat{\sigma}_2, \ldots, \hat{\sigma}_{n-1}, \hat{\sigma} \rangle \in K'$, as required. $\square$

Our aim is to subdivide repeatedly, to make the complex finer.

Definition 4.19. We inductively define $K^{(0)} := K$ and $K^{(r)} := (K^{(r-1)})'$ for $r > 0$. The complex $K^{(r)}$ is called the *rth barycentric subdivision of K*.

We next need to formalise the intuitive notion of how 'fine' a complex is.

Definition 4.20. For K a simplicial complex, the quantity

$$\mathrm{mesh}(K) := \max_{\langle u,v \rangle \in K} \|u - v\|$$

is called the *mesh* of K. That is, it is the ℓ_2-norm of the longest 1-simplex in K.

Remark 4.21. Note that $\mathrm{mesh}(K)$ is also the maximal diameter of any simplex of K.

The next lemma now makes precise the idea that $K^{(r)}$ becomes increasingly fine as r increases.

Lemma 4.22. *If* $\dim K = n$ *then*

$$\mathrm{mesh}(K^{(r)}) \le \left(\frac{n}{n+1} \right)^r \mathrm{mesh}(K)$$

for any r. In particular, $\mathrm{mesh}(K^{(r)}) \to 0$ *as $r \to \infty$.*

Proof. It suffices to handle the case $r = 1$. If $\langle u, v \rangle \in K'$ then $u = \hat{\tau}$ and $v = \hat{\sigma}$, for $\tau \le \sigma$ without loss of generality. Since distance to a simplex is maximised on a vertex, we may further assume that τ is a 0-simplex, so $\hat{\tau}$ is a vertex of σ. Let $\sigma = \langle v_0, \ldots, v_m \rangle$ with $\hat{\tau} = v_0$. Now

$$\begin{aligned}
\|\hat{\tau} - \hat{\sigma}\| &= \left\| v_0 - \frac{1}{m+1} \sum_{i=0}^{m} v_i \right\| \\
&= \left\| \frac{m}{m+1} v_0 - \frac{1}{m+1} \sum_{i=1}^{m} v_i \right\| \\
&= \frac{1}{m+1} \left\| \sum_{i=1}^{m} v_0 - v_i \right\| \\
&\le \frac{m}{m+1} \mathrm{mesh}(K) \\
&\le \frac{n}{n+1} \mathrm{mesh}(K)
\end{aligned}$$

as required. ☐

4.3 The simplicial approximation theorem

The main theorem of this section will tell us that every continuous map between realisations of simplicial complexes is homotopic to the realisation of a simplicial map, after subdividing often enough.

Definition 4.23. Let K be a simplicial complex. The *(open) star* of a vertex v of K is the union of the interiors of the simplices that contain v.

$$\mathrm{St}_K(v) \coloneqq \bigcup_{v\in\sigma} \mathring{\sigma}$$

Definition 4.24. Let K and L be simplicial complexes, and $\phi : |K| \to |L|$ a continuous map. A simplicial map $f : K \to L$ is a *simplicial approximation to* ϕ if

$$\phi(\mathrm{St}_K(v)) \subseteq \mathrm{St}_L(f(v))$$

for every vertex v of K.

For simplicial approximations to do the job we want them to, we need them to be homotopic to the original map. The next lemma guarantees this.

Lemma 4.25. *If $f : K \to L$ is a simplicial approximation to $\phi : |K| \to |L|$ then $|f| \simeq \phi$.*

Proof. Let $|L|$ be contained in the ambient vector space $\mathbb{R}^m$. We will show that straight-line homotopy between $|f|$ and ϕ has image contained in $|L|$.

Consider $x \in |K|$, contained in the interior of a unique simplex σ, and let $\phi(x)$ be contained in the interior of the simplex $\tau \in L$. We will show that $f(\sigma)$ is a face of τ.

Let v_i be a vertex of σ. Then $x \in \mathrm{St}_K(v_i)$ and so

$$\phi(x) \in \phi(\mathrm{St}_K(v_i)) \subseteq \mathrm{St}_L(f(v_i))$$

since f is a simplicial approximation to ϕ. Therefore, $\mathring{\tau} \subseteq \mathrm{St}_L(f(v_i))$ and so $f(v_i)$ is a vertex of τ. So every vertex of σ is sent by f to a vertex of τ, and $f(\sigma)$ is a face of τ, as claimed.

Since τ is convex in $\mathbb{R}^m$, it follows that the straight line between $|f|(x)$ and $\phi(x)$ is contained τ, and the result follows. ☐

We are now ready to prove the main theorem.

Theorem 4.26 (Simplicial approximation theorem)**.** *Let K, L be simplicial complexes and $\phi : |K| \to |L|$ a continuous map. For some $r \in \mathbb{N}$, there exists a simplicial approximation $f : K^{(r)} \to L$ to ϕ.*

Proof. The collection of open sets $\mathcal{U} = \{\phi^{-1}\mathrm{St}_L(u) \mid u \in V_L\}$ is an open cover of $|K|$. We next prove that there is a $\delta > 0$ such that, for every $x \in |K|$, the ball $B(x, \delta)$ is contained in some element of $\mathcal{U}$. (This statement about open covers is called the *Lebesgue number lemma*.)

If not then, for each n, there is $x_n \in |K|$ so that $B(x_n, 1/n)$ is not contained in any U. Taking a subsequence, we may assume that the sequence (x_n) converges to a limit x. But then $x \in U$ for some $U \in \mathcal{U}$, so $B(x, \epsilon) \subseteq U$ for some $\epsilon > 0$. Now $d(x_n, x) < \epsilon/2$ for sufficiently large n, and so $B(x_n, \epsilon/2) \subseteq U$. When $1/n < \epsilon/2$, this contradicts the assumption that $B(x_n, 1/n)$ is not contained in any U .

By Lemma 4.22, we may choose r sufficiently large that $\mathrm{mesh}(K^{(r)}) < \delta$. For each vertex v of $K^{(r)}$, we now have that

$$\mathrm{St}_{K^{(r)}}(v) \subseteq B(v, \delta) \subseteq \phi^{-1}(\mathrm{St}_L(u))$$

for some u a vertex of L. Set $f(v) = u$. This is now constructed to be a simplicial approximation, but we still need to check that it is a simplicial map: i.e. that f sends simplices of $K^{(r)}$ to simplices of L.

As in the proof of Lemma 4.25, if σ is a simplex of K and $x \in \mathring{\sigma}$ then $f(\sigma)$ is a face of τ, the unique simplex of L containing $\phi(x)$ in its interior. In particular, $f(\sigma)$ is a simplex, as required. □

We have already seen that it is sometimes useful to be able to construct *relative* homotopies. This is easy to arrange, at least relative to vertices.

Remark 4.27. If S is a subset of V_K such that $\phi(S) \subseteq V_L$ then we may take $f(v) = \phi(v)$ for all $v \in S$ (by choosing δ in the proof a little smaller if necessary). In this case, the straight-line homotopy between $|f|$ and ϕ is relative to S.

5 Homology

We are now ready to define a new algebraic invariant of a topological space – *homology*. We will see that it has the advantage that, unlike the fundamental group, it can detect higher-dimensional phenomena.

5.1 Simplicial homology

The fundamental group was defined using paths. The corresponding notion for homology is *chains*. These are formal sums of simplices. It's easy to interpret 2σ as two copies of σ, but it's less clear how to interpret the expression $-\sigma$. To do this we define *oriented* simplices.

Definition 5.1. Let $V = (v_0, \dots, v_n)$ be an *ordered* set of points in general position in $\mathbb{R}^m$. We consider the natural action of the symmetric group S_{n+1} on V. As long as $n \geq 1$, there are two orbits of the alternating group A_{n+1} under this action. An *orientation* on the simplex $\langle V \rangle$ is a choice of one of these orbits. An *oriented simplex* σ consists of the simplex σ together with a choice of orientation on σ. We will abuse notation and write $\langle v_0, \dots, v_n \rangle$ to mean the oriented simplex determined by the orbit of $(v_0, \dots, v_n)$. For an oriented simplex σ, the notation $\bar{\sigma}$ denotes the same simplex with the opposite orientation.

Example 5.2. Consider an unoriented 1-simplex $\sigma = \langle v_0, v_1 \rangle$. The two orientations on σ correspond to the two orderings $\langle v_0, v_1 \rangle$ and $\langle v_1, v_0 \rangle$. We can represent these pictorially by drawing a small arrow on σ from the first vertex to the second vertex. So the two orientation corresponds to the intuitive idea that we can travel along σ in two different directions.

Example 5.3. Consider an unoriented 2-simplex $\sigma = \langle v_0, v_1, v_2 \rangle$. The two orientations on σ correspond to the two cyclic orderings $\langle v_0, v_1, v_2 \rangle$ and $\langle v_2, v_1, v_0 \rangle$. On a drawing of σ, these correspond to travelling either clockwise or anticlockwise round the boundary of σ. We can represent these pictorially by drawing a small rotating arrow on σ.

We are now ready to define the chains associated to a simplicial complex.

Definition 5.4. Let K be a simplicial complex. For each $n \in \mathbb{N}$, the group of *n-chains on K* is the free abelian group

$$C_n(K) := \bigoplus_{\dim \sigma = n} \langle \sigma \rangle$$

formally generated by the n-simplices of K. That is, elements of $C_n(K)$ are expressions of the form $a\sigma + b\tau + c\upsilon + \dots$ where σ, τ, υ are n-simplices of K and $a, b, c \in \mathbb{Z}$. In particular, if K has no n-simplices (because $n > \dim K$ or $n < 0$) then $C_n(K) = 0$.

Choose an orientation arbitrarily on each simplex σ of K (of dimension > 0). We then identify $-\sigma$ with $\bar{\sigma}$, the simplex with the opposite orientation. Note that the arbitrary choice of orientation did not change $C_n(K)$ dramatically, since changing the orientation of a simplex is reflected in the automorphisms of $C_n(K)$ that changes the sign of one of the generators.

So far the connection between $C_n(K)$ and the topology of K is rather slight. However, we next define homomorphisms that are intimately connected with the way that K is constructed. Informally, these homomorphisms send simplices to their boundaries. To give a formal definition, we need to be careful about orientations.

Definition 5.5. The *boundary homomorphism* $\partial \equiv \partial_n : C_n(K) \to C_{n-1}(K)$ is the homomorphism defined on a basis element $\sigma = \langle v_0, \dots, v_n \rangle$ as follows.

$$\partial\sigma := \sum_{i=0}^{n} (-1)^i \langle v_0, \dots, \hat{v}_i, \dots, v_n \rangle$$

Here, the notation $\langle v_0, \dots, \hat{v}_i, \dots, v_n \rangle$ means the oriented $(n-1)$-simplex defined by omitting v_i from the list. Note that that this is well defined. Indeed, if we apply the transposition $(j, j+1)$ then $d_n(\sigma)$ becomes

$$\sum_{i=0}^{j} (-1)^i \langle v_0, \dots, \hat{v}_i, \dots, v_{j+1}, v_j, \dots, v_n \rangle + (-1)^j \langle v_0, \dots, \hat{v}_{j+1}, v_j, \dots, v_n \rangle$$
$$+ (-1)^{j+1} \langle v_0, \dots, v_{j+1}, \hat{v}_j, \dots, v_n \rangle + \sum_{i=j+2}^{n} \langle v_0, \dots, v_{j+1}, v_j, \dots, \hat{v}_i, \dots, v_n \rangle$$

which equals $-d_n(\sigma)$. Since consecutive transpositions generate S_{n+1}, it follows that permutations act by multiplying by their sign, and so A_{n+1} acts trivially, as required.

Example 5.6. Let $\sigma = \langle v_0, v_1 \rangle$. Then $\partial\sigma = \langle v_1 \rangle - \langle v_0 \rangle$. Thus, the boundary can be interpreted as the difference between the end vertex and the start vertex.

Example 5.7. Let $\sigma = \langle v_0, v_1, v_2 \rangle$. Then $\partial\sigma = \langle v_1, v_2 \rangle - \langle v_0, v_2 \rangle + \langle v_0, v_1 \rangle = \langle v_0, v_1 \rangle + \langle v_1, v_2 \rangle + \langle v_2, v_0 \rangle$. Thus, $\partial\sigma$ can be thought of as the 1-chain that goes round the boundary, in the same direction as the orientation on σ.

This map allows us to define two important subgroups of chains.

Definition 5.8. Let $n \in \mathbb{Z}$. The group of *n-cycles* is the group

$$Z_n(K) := \ker(\partial_n : C_n(K) \to C_{n-1}K)) \ .$$

The group of *n-boundaries* is the group

$$B_n(K) := \operatorname{im}(\partial_n : C_{n+1}(K) \to C_n(K)) \ .$$

So the n-cycles are the n-chains with no boundaries, and the n-boundaries are the n-chains that are boundaries.

Continuing the analogy with the fundamental group, if chains are analogous to paths, then cycles are analogous to loops, and boundaries are analogous to homotopies. The fact that makes it possible to define homology is that every boundary is a cycle.

Lemma 5.9. *For each n, $B_n(K) \subseteq Z_n(K)$, that is, the composition $\partial \circ \partial = 0$ in every dimension.*

Proof. We only need to check that $\partial \circ \partial(\sigma) = 0$ for an arbitrary oriented simplex $\sigma = \langle v_0, \ldots, v_n \rangle$. By definition,

$$\partial\sigma = \sum_{i=0}^{n} (-1)^i \langle v_0, \ldots, \hat{v}_i, \ldots, v_n \rangle \ .$$

When working out $\partial \circ \partial(\sigma)$, we take out another vertex v_j, and the sum breaks into two parts depending on whether $j < i$ or $i < j$. We compute:

$$\begin{aligned} \partial \circ \partial(\sigma) &= \sum_{j<i} (-1)^j (-1)^i \langle v_0, \ldots, \hat{v}_j, \ldots, \hat{v}_i, \ldots, v_n \rangle + \\ &\quad \sum_{j>i} (-1)^{j-1} (-1)^i \langle v_0, \ldots, \hat{v}_i, \ldots, \hat{v}_j, \ldots, v_n \rangle \\ &= \sum_{j<i} (-1)^{i+j} \langle v_0, \ldots, \hat{v}_j, \ldots, \hat{v}_i, \ldots, v_n \rangle - \\ &\quad \sum_{j>i} (-1)^{i+j} \langle v_0, \ldots, \hat{v}_i, \ldots, \hat{v}_j, \ldots, v_n \rangle \ . \end{aligned}$$

Reversing the roles of i and j, we see that the two terms cancel, and the result follows. □

We can now define simplicial homology, which measures how many cycles fail to be boundaries.

Definition 5.10. The *nth (simplicial) homology group of a simplicial complex* K is defined to be the quotient

$$H_n(K) := Z_n(K)/B_n(K) \ .$$

Remark 5.11. Note the happy fact that, since $H_n(K)$ is defined in terms of linear maps of abelian groups, we can *in principle* always compute it. However, this is very tedious in practice, and we will shortly see some very useful techniques for simplifying computations.

It takes a while to understand this definition intuitively. The idea is that the boundaries 'fill in' some of the cycles, so the homology groups measure how many holes remain in K. This becomes a bit clearer if we compute a couple of examples.

Example 5.12. Let K be the standard simplicial circle (i.e. 1-sphere). The 0-simplices are $\langle e_0 \rangle$, $\langle e_1 \rangle$ and $\langle e_2 \rangle$, and the 1-simplices are $\langle e_0, e_1 \rangle$, $\langle e_1, e_2 \rangle$ and $\langle e_2, e_0 \rangle$. We take these oriented simplices to be bases for $C_0(K)$ and $C_1(K)$. So $C_0(K) \cong C_1(K) \cong \mathbb{Z}^3$, and all the other chain groups are trivial. In particular, $H_k(K) = 0$ when $k < 0$ or $k > 2$.

The only non-zero boundary map is ∂_1; when written in these bases it has the following matrix.

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

To compute the homology groups, we put it in Smith normal form.

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

From this we see that $\operatorname{im} \partial_1$ is a $\mathbb{Z}^2$ direct factor of $C_0(K)$ and that $\ker \partial_1 \cong \mathbb{Z}$. We can now finish the computation.

$$\begin{aligned} H_0(K) &= Z_0(K)/B_0(K) = C_0(K)/B_0(K) \cong \mathbb{Z}^2/\mathbb{Z}^2 \cong \mathbb{Z} \\ H_1(K) &= Z_1(K)/B_1(K) \cong \mathbb{Z}/0 \cong \mathbb{Z} \end{aligned}$$

The interpretation here is that $H_0(K)$ tells us that $|K|$ is path-connected (see Lemma 5.14 below), while $H_1(K)$ tells us that $|K|$ has a one-dimensional 'hole'.

Example 5.13. Let L be the standard 2-simplex σ_2, together with its faces. The 0- and 1-simplices are as in Example 5.12. There is also the unique 2-simplex $\langle e_0, e_1, e_2 \rangle$. Therefore, $C_i(L) = C_i(K)$ for $i = 0, 1$, and ∂_1 is unchanged. But we also have $C_2(K) = \langle \sigma \rangle$ and another non-trivial boundary map ∂_2, represented by the following matrix.

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

Since the image of ∂_2 is primitive and contained in $\ker \partial_1$, in fact $\text{im } \partial_2 = \ker \partial_1$. We also see that $\ker \partial_2 = 0$. We can now compute the interesting homology groups of L.

$$\begin{aligned} H_0(L) &= Z_0(L)/B_0(L) = C_0(L)/B_0(L) \cong \mathbb{Z}^2/\mathbb{Z}^2 \cong \mathbb{Z} \\ H_1(L) &= Z_1(L)/B_1(L) = \ker \partial_1 / \text{im } \partial_2 \cong 0 \\ H_2(L) &= Z_2(L)/B_2(L) = 0/0 \cong 0 \end{aligned}$$

The interpretation here is that the 2-simplex added a boundary that 'filled in' the 1-dimensional hole in $|K|$. Note that this is the same as the homology of a point. This shouldn't surprise us, since $|L|$ is contractible.

The next lemma makes formal the idea that we've already seen, that $H_0(K) \cong \mathbb{Z}$ tells us that $|K|$ is path-connected.

Lemma 5.14. *Let K be a simplicial complex. If d is the number of path components of $|K|$ then $H_0(K) \cong \mathbb{Z}^d$.*

Proof. Let $\pi_0(K)$ be the set of path components of $|K|$.[3] Let $\mathbb{Z}[\pi_0(K)]$ be the free abelian group with basis $\pi_0(K)$. There is a natural map

$$q : C_0(K) \to \mathbb{Z}[\pi_0(K)]$$

that sends 0-simplex $\langle v \rangle$ of K to the path component $[v]$. Since every path component of $|K|$ contains a vertex, this map is surjective. Since $Z_0(K) = C_0(K)$, the result follows if we can show that the $\ker q = B_0(K)$.

[3] Why is this good terminology?

For each 1-simplex $\sigma = \langle v_0, v_1 \rangle \in C_1(K)$,

$$q \circ \partial\sigma = q(\langle v_1 \rangle) - q(\langle v_0 \rangle) = 0$$

since v_0 and v_1 are in the same path component. Therefore, $q \circ \partial_1 = 0$ and so $B_0(K) \subseteq \ker q$.

For the other direction, note that $\ker q$ is the subspace generated by all elements of the form $\langle v \rangle - \langle u \rangle$ where u and v are in the same path component. In this case, there is a continuous path from u, and v, and a simplicial approximation then defines a simplicial path. Summing the oriented simplices that appear in this path gives a 1-chain

$$c = \langle v_0, v_1 \rangle + \langle v_1, v_2 \rangle + \ldots + \langle v_{k-1}, v_k \rangle$$

where $u = v_0$ and $v = v_k$. Now $\partial c = \langle v \rangle - \langle u \rangle$, and so $\langle v \rangle - \langle u \rangle \in B_0(K)$. Therefore ker $q \subseteq B_0(K)$. This completes the proof. □

5.2 Chain maps and chain homotopies

As well as computing more examples, we would love to know that homology groups are invariants of the homeomorphism, or even homotopy, types of simplicial complexes. To achieve this, we'll need some more theory. To do this, we need to introduce a slightly more abstract setting.

Definition 5.15. A *chain complex* $C_\bullet$ is a sequence of abelian groups $\{C_n \mid n \in \mathbb{Z}\}$ with $C_n = 0$ for $n < 0$ and homomorphisms

$$\partial_n : C_n \to C_{n-1}$$

such that $\partial_{n-1} \circ \partial_n = 0$. A *chain map* between two chain complexes $f_\bullet : C_\bullet \to D_\bullet$ is a collection of homomorphisms $f_n : C_n \to D_n$ such that

$$\begin{array}{ccc} C_n & \xrightarrow{\partial_n} & C_{n-1} \\ \downarrow{\scriptstyle f_n} & & \downarrow{\scriptstyle f_{n-1}} \\ C_n & \xrightarrow{\partial_n} & D_{n-1} \end{array}$$

commutes for all n. That is, $\partial \circ f_\bullet = f_\bullet \circ \partial$.

We can make definitions for any chain complex in analogy with the definitions of the last section.

Definition 5.16. Let $C_\bullet$ be a chain complex. The group of *n-boundaries* is $B_n(C_\bullet) \coloneqq \operatorname{im} \partial_{n+1}$. The group of n-cycles is $Z_n(C_\bullet) \coloneqq \ker \partial_n$. The *nth homology group* of $C_\bullet$ is then $H_n(C_\bullet) \coloneqq Z_n(C_\bullet)/B_n(C_\bullet)$.

An automatic consequence of these definitions is that a chain map $f_\bullet : C_\bullet \to D_\bullet$ induces a homomorphism on homology.

Lemma 5.17. *Let $f_\bullet : C_\bullet \to D_\bullet$ be a chain map. The formula*

$$[c] \mapsto [f_n(c)]$$

gives a well-defined homomorphism $f_* : H_n(C_\bullet) \to H_n(D_\bullet)$ for all n.

Proof. Let $c \in Z_n(C_\bullet)$. By the definition of a chain map,

$$\partial_n \circ f_n(c) = f_{n-1} \circ \partial_n(c) = 0$$

so $f_n(c) \in Z_n(D_\bullet)$. Therefore $[f_n(c)]$ is an element of $H_n(D_\bullet)$. To see that this map descends to $H_n(C_\bullet)$, suppose that $c - c' \in B_n(C_\bullet)$. Then, by the definition of $B_n(C)$, there exists $b \in C_{n+1}$ such that $c - c' = \partial_{n+1}(b)$. Now

$$f_n(c - c') = f_n \circ \partial_{n+1}(b) = \partial_{n+1} \circ f_n(b)$$

so $f_n(c - c') \in B_n(D_\bullet)$. Therefore $[f_n(c)] = [f_n(c')]$ in $H_n(D_\bullet)$, so f_* is well defined on homology. It's clear from the definition that it's a homomorphism. $\square$

Lemma 5.9 shows us that the chain groups of a simplicial complex form a chain complex. Simplicial maps also induce chain maps.

Lemma 5.18. *A simplicial map $f : K \to L$ induces a chain map $f_\bullet : C_\bullet(K) \to C_\bullet(L)$ via the following assignment.*

$$f_n : \sigma \mapsto \begin{cases} f(\sigma) & \dim f(\sigma) = n \\ 0 & \text{otherwise} \end{cases}$$

In particular, it also induces a homomorphism $f_* : H_n(K) \to H_n(L)$.

Proof. Let $\sigma = \langle v_0, \ldots, v_n \rangle$ as usual. We need to check that $\partial_n \circ f_n(\sigma) = f_{n-1} \circ \partial_n(\sigma)$. If $f(\sigma)$ is a simplex of dimension n then every face is also sent to a simplex of the same dimension, and the result is clear. Likewise, if $\dim f(\sigma) \leq n-2$ then every face is also sent to a simplex of strictly lower dimension, and the result is also clear.

Therefore, the case of interest is when (without loss of generality) $f(v_0) = f(v_1)$ and $f(\langle v_1, \ldots, v_n \rangle)$ is of dimension $n-1$. In this case, $f_n(\sigma) = 0$ and so $\partial_n \circ f_n(\sigma) = 0$. To evaluate the other side of the equation, consider

$$\partial_n(\sigma) = \sum_{i=0}^{n} \langle v_0, \ldots, \hat{v}_i, \ldots, v_n \rangle$$

as usual. Since $f(v_0) = f(v_1)$, whenever $i \neq 0, 1$, the corresponding face is sent to something of strictly lower dimension and so killed by f_{n-1}. Therefore

$$f_{n-1} \circ \partial_n(\sigma) = \langle f(v_1), f(v_2), \ldots, f(v_n) \rangle - \langle f(v_0), f(v_2), \ldots, f(v_n) \rangle = 0$$

as required. □

Remark 5.19. The maps on homology induced by simplicial maps enjoy the usual 'functorial' properties.

(i) If $f : K \to L$ and $g : L \to M$ are simplicial maps then $(g \circ f)_* = g_* \circ f_*$.

(ii) If K is a simplicial complex then $(\mathrm{id}_K)_* = \mathrm{id}_{H_n(K)}$ (for any n).

This is a useful tool, but to prove homotopy invariance, we will need an analogue of homotopy that works in this context. This is provided by the following definition.

Definition 5.20. Let $C_\bullet$ and $D_\bullet$ be chain complexes, and let $f_\bullet, g_\bullet : C_\bullet \to D_\bullet$ be chain maps. A *chain homotopy* $h_\bullet$ between $f_\bullet$ and $g_\bullet$ is a collection of homomorphisms $h_n : C_n \to D_{n+1}$ so that

$$g_n(c) - f_n(c) = \partial_{n+1} \circ h_n(c) + h_{n-1} \circ \partial_n(c)$$

for all $c \in C_n$. In this case, we say that the chain maps $f_\bullet$ and $g_\bullet$ are *chain homotopic*, and write $f_\bullet \simeq g_\bullet$.

It is a minor miracle of this subject that this entirely algebraic definition captures many of the features of topological homotopies. In particular, chain-homotopic maps induce the same map on homology.

Lemma 5.21. *If $f_\bullet \simeq g_\bullet : C_\bullet \to D_\bullet$ the induced maps $f_*, g_* : H_n(C_\bullet) \to H_n(D_\bullet)$ are equal for every n.*

Proof. Consider $c \in Z_n(C_\bullet)$. Since, $\partial_n(c) = 0$, the definition of chain homotopy gives

$$g_n(c) - f_n(c) = \partial_{n+1} \circ h_n(c) + h_{n-1} \circ \partial_n(c) = \partial_{n+1} \circ h_n(c) \in B_n(D_\bullet)$$

so $[g_n(c)] = [f_n(c)]$ in homology, as required. □

The following example may give a hint of why chain homotopies are like actual homotopies.

Example 5.22. Let K be the 2-complex that consists of the standard 2-simplex $\sigma_2 = \langle e_0, e_1, e_2 \rangle$ and its faces. Let L consist of the face $\langle e_0, e_1 \rangle$ and its faces. Let $i : L \to K$ be the natural inclusion, and $r : K \to L$ the simplicial retraction that fixes e_0 and e_1, and sends $e_2 \mapsto e_0$. Naturally, $r \circ i = \mathrm{id}_L$. We now define a chain homotopy $h_n : C_n(K) \to C_{n+1}(K)$ by sending every simplex to 0 except for the following.

$$\begin{aligned} h_0 &: \langle e_2 \rangle \mapsto \langle e_2, e_0 \rangle \\ h_1 &: \langle e_1, e_2 \rangle \mapsto -\langle e_0, e_1, e_2 \rangle \end{aligned}$$

We can now check directly that this defines a chain homotopy between $i \circ r$ and id_K by evaluating the two sides of the defining equation each of the seven simplices of K. Most of these checks are fairly trivial. The three least trivial are:

$$\begin{aligned} (\partial_1 \circ h_0 + h_{-1} \circ \partial_0)(\langle e_2 \rangle) &= \partial_1(\langle e_2, e_0 \rangle) \\ &= \langle e_0 \rangle - \langle e_2 \rangle \\ &= (i_0 \circ r_0 - \mathrm{id}_{C_0(K)})(\langle e_2 \rangle) \ ; \end{aligned}$$

$$\begin{aligned} (\partial_2 \circ h_1 + h_0 \circ \partial_1)(\langle e_1, e_2 \rangle) &= \partial_2(-\langle e_0, e_1, e_2 \rangle) + h_0(\langle e_2 \rangle - \langle e_1 \rangle) \\ &= -\langle e_0, e_1 \rangle - \langle e_1, e_2 \rangle - \langle e_2, e_0 \rangle + \langle e_2, e_0 \rangle \\ &= \langle e_1, e_0 \rangle - \langle e_1, e_2 \rangle \\ &= (i_1 \circ r_1 - \mathrm{id}_{C_1(K)})(\langle e_1, e_2 \rangle) \ ; \end{aligned}$$

and

$$\begin{aligned} (\partial_3 \circ h_2 + h_1 \circ \partial_2)(\langle e_0, e_1, e_2 \rangle) &= h_1(\langle e_0, e_1 \rangle + \langle e_1, e_2 \rangle + \langle e_2, e_0 \rangle) \\ &= -\langle e_0, e_1, e_2 \rangle \\ &= (i_2 \circ r_2 - \mathrm{id}_{C_2(K)})(\langle e_0, e_1, e_2 \rangle) \end{aligned}$$

as required.

We can use chain homotopies to check that a certain simple class of simplicial complexes, which are obviously contractible, have the homology of a point.

Definition 5.23. A simplicial complex K is a *cone* if there is a vertex x_0 such, for every simplex $\tau \in K$, there exists $\sigma \in K$ such that $x_0 \in \sigma$ and $\tau \leq \sigma$. The vertex x_0 is called a *cone point*.

Lemma 5.24. *If K is a cone then the homology of K is as follows.*

$$H_n(K) \cong \begin{cases} \mathbb{Z} & n = 0 \\ 0 & \text{otherwise} \end{cases}$$

Proof. Let $i : \{\langle x_0 \rangle\} \to K$ be the inclusion and $r : K \to \{\langle x_0 \rangle\}$ the unique retraction. Evidently $r_* \circ i_*$ is the identity. We will exhibit a chain homotopy between $\mathrm{id}_{C_\bullet(K)}$ and $i_\bullet \circ r_\bullet$. By Lemma 5.21, this implies that r_* is an isomorphism, and the result follows.

Let $\sigma = \langle v_0, \ldots, v_n \rangle \in K$. Then

$$h_n(\sigma) = \begin{cases} 0 & x_0 \in \sigma \\ \langle x_0, v_0, \ldots, v_n \rangle & \text{otherwise} \end{cases}$$

We need to check this is the required chain homotopy. The proof divides into several cases. First, assume that $n > 0$. Suppose that $x_0 \notin \sigma$. Then

$$\begin{aligned} (\partial_{n+1} \circ h_n + h_{n-1} \circ \partial_n)(\sigma) &= \langle v_0, \ldots, v_n \rangle - \sum_{i=0}^{n} (-1)^i \langle x_0, v_0, \ldots, \hat{v}_i, \ldots, v_n \rangle \\ &+ \sum_{i=0}^{n} (-1)^i \langle x_0, v_0, \ldots, \hat{v}_i, \ldots, v_n \rangle \\ &= \sigma \\ &= (\mathrm{id}_{C_n(K)} - i_n \circ r_n)(\sigma) \end{aligned}$$

as required.

If $x_0 \in \sigma$ then $x_0 = v_j$ for some j. Now $h_n(\sigma) = 0$ and so

$$\begin{aligned} (\partial_{n+1} \circ h_n + h_{n-1} \circ \partial_n)(\sigma) &= h_{n-1}\left(\sum_{i=0}^{n} (-1)^i \langle v_0, \ldots, \hat{v}_i, \ldots v_n \rangle \right) \\ &= (-1)^j \langle x_0, v_0, \ldots, \hat{v}_j, \ldots, v_n \rangle \\ &= \langle v_0, \ldots, v_{j-1}, x_0, v_{j+1}, \ldots, v_n \rangle \\ &= (\mathrm{id}_{C_n(K)} - i_n \circ r_n)(\sigma) \end{aligned}$$

where the third equality holds because a $(j+1)$-cycle has sign $(-1)^j$.

Now assume that $n=0$, so $\sigma = \langle v_0 \rangle$. If $v_0 \neq x_0$ then

$$(\partial_1 \circ h_0 + h_{-1} \circ \partial_0)(\sigma) = \partial_1 \langle x_0, v_0 \rangle = \langle v_0 \rangle - \langle x_0 \rangle = (\mathrm{id}_{C_0(K)} - i_0 \circ r_0)(\sigma)$$

as required. Finally, if $\sigma = \langle x_0 \rangle$ then

$$(\partial_1 \circ h_0 + h_{-1} \circ \partial_0)(\sigma) = 0 = (\mathrm{id}_{C_0(K)} - i_0 \circ r_0)(\sigma)$$

and the proof is complete. □

5.3 The homology of the simplex and the sphere

This enables us to compute some homology groups quite quickly.

Example 5.25. Let K be the simplicial complex that consists of the standard n-simplex σ_n, together with its faces. Then any vertex is a cone point, so

$$H_k(K) \cong \begin{cases} \mathbb{Z} & k = 0 \\ 0 & \text{otherwise} \end{cases}$$

by Lemma 5.24.

From this, it's not difficult to compute the homology of the (standard simplicial) sphere.

Example 5.26. Let $L = \partial \sigma_n$, the standard simplicial $(n-1)$-sphere, for $n \geq 2$. The chain complex of L is very close to the chain complex of K from Example 5.25.

$$\begin{tikzcd}
0 \arrow[r] & 0 \arrow[r] \arrow[d] & C_{n-1}(L) \arrow[r, "\partial_{n-1}"] \arrow[d, "="] & \cdots \arrow[r, "\partial_1"] & C_0(L) \arrow[r] \arrow[d, "="] & 0 \\
0 \arrow[r] & C_n(K) \arrow[r, "\partial_n"] & C_{n-1}(K) \arrow[r, "\partial_{n-1}"] & \cdots \arrow[r, "\partial_1"] & C_0(K) \arrow[r] & 0
\end{tikzcd}$$

In particular, for $k \leq n-2$, we see that $C_{k+1}(L) = C_{k+1}(K)$ and $C_k(L) = C_k(K)$, and so $H_k(L) = H_k(K)$ as computed in Example 5.25.

The only remaining dimension of interest is $k = n-1$. From $H_{n-1}(K) \cong 0$, it follows that $Z_{n-1}(K) = B_{n-1}(K)$. But $C_n(K) = \langle \sigma_n \rangle$, and so the boundary map ∂_n is clearly injective; in particular, $B_{n-1}(K) \cong C_n(K)$. Therefore

$$Z_{n-1}(L) = Z_{n-1}(K) = B_{n-1}(K) \cong C_n(K) \cong \mathbb{Z} \ .$$

But $C_n(L) \cong 0$, so $B_{n-1}(L) \cong 0$ and $H_{n-1}(L) = Z_{n-1}(L)/B_{n-1}(L) \cong \mathbb{Z}/0 \cong \mathbb{Z}$. In summary, for $n \geq 2$, the homology of the standard $(n-1)$-dimensional sphere is as follows.

$$H_k(K) \cong \begin{cases} \mathbb{Z} & k = 0, n-1 \\ 0 & \text{otherwise} \end{cases}$$

A similar calculation works when $n = 1$.

$$H_k(K) \cong \begin{cases} \mathbb{Z}^2 & k = 0 \\ 0 & \text{otherwise} \end{cases}$$

This is very encouraging, because it shows that, unlike the fundamental group, homology distinguishes the spheres in all dimensions. It seems to confirm our intuition that nth homology should detect n-dimensional holes.

5.4 Continuous maps and homotopies

We would like to know that the groups $H_n(K)$ are invariant of the realisation $|K|$, not just of the complex K. More generally, we would like continuous maps $\phi : |K| \to |L|$ to induce well defined maps on homology. The idea should be clear: let $f : K^{(r)} \to L$ be a simplicial approximation to ϕ, and then set

$$\phi_* := f_* : H_n(K^{(r)}) \to H_n(L)$$

for any $n \in \mathbb{Z}$. However, there are two obvious problems: we made a choice of simplicial approximation, so we don't know that ϕ_* is well defined; and we got a map from $H_n(K^{(r)})$, not from $H_n(L)$. So we need to check that $H_n(K^{(r)})$ is canonically isomorphic to $H_n(K)$, and also that the induced map ϕ_* is independent of the choice fo simplicial approximation f. The key idea is the next definition, which is a *simplicial* version of a homotopy.

Definition 5.27. Two simplicial maps $f, g : K \to L$ are *contiguous* if, for every $\sigma \in K$ there exists $\tau \in L$ such that $f(\sigma)$ and $g(\sigma)$ are both faces of τ.

Remark 5.28. Suppose that $\phi : |K| \to |L|$ is a continuous map, and $f, g : K \to L$ are both simplicial approximations to ϕ. Suppose that $x \in \mathring{\sigma}$ and that $\phi(x) \in \mathring{\tau}$. In the proof of Lemma 4.25, we proved that $f(\sigma) \leq \tau$. Likewise, $g(\sigma) \leq \tau$, and so f and g are contiguous.

Contiguous maps can be thought of as homotopic in a simplicial sense. The next lemma checks that they do in fact induce equal maps on homology.

Lemma 5.29. *If $f, g : K \to L$ are contiguous then*

$$f_* = g_* : H_n(K) \to H_n(L)$$

for all n.

Proof. It suffices to exhibit a chain homotopy between f_* and g_*. We find this chain homotopy in a surprising way. Fix a total order $<$ on the vertices of K. We can then write each $\sigma \in K$ uniquely as

$$\sigma = \langle v_0, \dots, v_n \rangle$$

such that $v_0 < \dots < v_n$. For notational simplicity, we'll adopt the convention that

$$\langle f(v_0), \dots, f(v_i), g(v_i), \dots, g(v_n) \rangle = 0$$

if these $n+1$ vertices are not in general posiiton. Then define $h_n : C_n(K) \to C_n(L)$ by setting

$$h_n(\langle v_0, \dots, v_n \rangle) = \sum_{i=0}^{n} (-1)^i \langle f(v_0), \dots, f(v_i), g(v_i), \dots, g(v_n) \rangle$$

for each n-simplex $\sigma = \langle v_0, \dots, v_n \rangle$ of K. We next check that this is a chain homotopy on the n-simplex σ.

$$\begin{aligned}
\partial \circ h + h \circ \partial(\sigma) &= \partial\left(\sum_{j=0}^{n} (-1)^j \langle f(v_0), \dots, f(v_j), g(v_j), \dots, g(v_n) \rangle \right) \\
&+ h\left(\sum_{i=0}^{n} (-1)^i \langle v_0, \dots, \hat{v}_i, \dots, v_n \rangle \right) \\
&= \sum_{i \le j} (-1)^{i+j} \langle f(v_0), \dots, \widehat{f(v_i)}, \dots, f(v_j), g(v_j), \dots, g(v_n) \rangle \\
&- \sum_{i \ge j} (-1)^{i+j} \langle f(v_0), \dots, f(v_j), g(v_j), \dots, \widehat{g(v_i)}, \dots, g(v_n) \rangle \\
&+ \sum_{j < i} (-1)^{i+j} \langle f(v_0), \dots, f(v_j), g(v_j), \dots, \widehat{g(v_i)}, \dots, g(v_n) \rangle \\
&- \sum_{j > i} (-1)^{i+j} \langle f(v_0), \dots, \widehat{f(v_i)}, \dots, f(v_j), g(v_j), \dots, g(v_n) \rangle \\
&= \sum_{i=0}^{n} \langle f(v_0), \dots, f(v_{i-1}), g(v_i), \dots, g(v_n) \rangle \\
&- \sum_{i=0}^{n} \langle f(v_0), \dots, f(v_i), g(v_{i+1}), \dots, g(v_n) \rangle \\
&= \langle g(v_0), \dots, g(v_n) \rangle - \langle f(v_0), \dots, f(v_n) \rangle
\end{aligned}$$

□

Combining Remark 5.28 with Lemma 5.29, we see that different choices of simplicial approximation will induce the same homomorphism on homology. We next need to deal with the identification of $H_n(K)$ and $H_n(K')$.

Lemma 5.30. *Let K be a simplicial complex. A simplicial map $s : K' \to K$ is a simplicial approximation to the identity $\mathrm{id}_{|K|}$ if and only if $s(\hat{\sigma})$ is a vertex of σ, for all $\sigma \in K$. Furthermore, such an s exists.*

Proof. Suppose that $s : K' \to K$ is a simplicial approximation to $\mathrm{id}_{|K|}$. By the definition of simplicial approximation,

$$\mathring{\sigma} \subseteq \mathrm{id}_{|K|}(\mathrm{St}_{K'}(\hat{\sigma})) \subseteq \mathrm{St}_K(s(\hat{\sigma}))$$

so in particular $s(\hat{\sigma})$ is a vertex of σ.

Conversely, suppose that $s(\hat{\sigma})$ is a vertex of σ for each $\sigma \in K$. Consider a simplex $\tau' \in K'$ is so that $\mathring{\tau}' \subseteq \mathrm{St}_{K'}(\hat{\sigma})$. Then the interior of τ' is contained in the interior of a simplex $\tau \in K$ such that $\sigma \leq \tau$. In particular, $s(\hat{\sigma})$ is also a vertex of τ. Therefore

$$\mathring{\tau}' \subseteq \mathring{\tau} \subseteq \mathrm{St}_K(s(\hat{\sigma})) \ .$$

Since the interiors of such simplices τ' cover $\mathrm{St}_{K'}(\hat{\sigma})$, it follows that $\mathrm{id}_{|K|}(\mathrm{St}_{K'}(\hat{\sigma})) \subseteq \mathrm{St}_K(s(\hat{\sigma}))$, so s is indeed a simplicial approximation to $\mathrm{id}_{|K|}$.

To see that such an s exists, for each vertex $\hat{\sigma}$ of K' choose any $s(\hat{\sigma})$ that is a vertex of σ. A simplex of K' is of the form

$$\langle \hat{\sigma}_0, \ldots, \hat{\sigma}_n \rangle$$

where each $\sigma_0 \leq \sigma_1 \leq \ldots \leq \sigma_n$. In particular, all vertices of any σ_i are vertices of σ_n, and so

$$\langle s(\hat{\sigma}_0), \ldots, s(\hat{\sigma}_n) \rangle$$

is a face of σ_n. Therefore, s defines a simplicial map. □

Finally we need the following fact, whose proof will be postponed to a subsequent section.

Proposition 5.31. *If $s : K' \to K$ is a simplicial approximation to the identity, then the induced map on homology*

$$s_* : H_n(K') \to H_n(K)$$

is an isomorphism for all n.

Combining Proposition 5.31 with Remark 5.28, 5.29 and Lemma 5.30, we obtain a special case of the invariance we have been looking for.

Corollary 5.32. *Let K be a simplicial complex. We may canonically identify $H_n(K)$ with $H_n(K')$.*

In light of Corollary 5.32, we will write $H_n(K^{(r)}) \equiv H_n(K)$. We are at last ready to define the homology of a triangulable space, and the homomorphism induced by a continuous map.

Definition 5.33. Let $\alpha : |K| \to X$ be a triangulation. Then we define

$$H_n(X) := H_n(K)$$

for all n.

Let $\phi : X \to Y$ be a continuous map of triangulable spaces; let $\alpha : |K| \to X$ and $\beta : |L| \to Y$ be triangulations. Let $f : K^{(r)} \to L$ be a simplicial approximation to $\beta^{-1} \circ \phi \circ \alpha$. Then we define

$$\phi_* := f_* : H_n(K) = H_n(X) \to H_n(L) = H_n(Y) \ .$$

for each n. Here, we are using Corollary 5.32 to identify $H_n(K^{(r)})$ with $H_n(K)$. By Remark 5.28 and Lemma 5.29, ϕ_* is independent of the choice of simplicial approximation.

This is a big step forward, but we would also like to know that homotopic maps induce equal maps on homology. This is the content of the next theorem, whose proof we will only sketch.

Theorem 5.34. *Suppose that $\alpha : |K| \to X$ and $\beta : |L| \to Y$ are triangulations. If $\phi \simeq \psi : X \to Y$ are homotopic continuous maps then*

$$\phi_* = \psi_* : H_n(X) \to H_n(Y)$$

for every n.

Sketch proof. By hypothesis, $\beta^{-1} \circ \phi \circ \alpha \simeq \beta^{-1} \circ \psi \circ \alpha$; let $\Phi : |K| \times I \to |L|$ realise this homotopy. The product $|K| \times I$ is the realisation of a simplicial complex M (see Example Sheet 3, Question 9), in which $|K| \times \{0\}$ and $|K| \times \{1\}$ sit naturally as subcomplexes, K_0 and K_1 respectively.

Furthermore, for each subcomplex Ł $\subseteq\in K$, $|L| \times I$ is the realisation of a subcomplex M_L of M. We also notice that there is a simple formula for the boundary of a product:

$$\partial(M_\sigma) = (\sigma \times \{1\}) \cup (\sigma \times \{0\}) \cup M_{\partial\sigma}$$

for any simplex σ.

Let $F : M^{(r)} \to L$ be a simplicial approximation to Φ. By the definition of induced homomorphism, we can take take ϕ_* to be induced by $f = F|_{K_0^{(r)}}$ and ψ_* to be induced by $g = F|_{K_1^{(r)}}$. To prove the theorem, we need to write down a chain homotopy between f and g.

Let $i : K^{(r)} \to K_0^{(r)} \subseteq M^{(r)}$ and $j : K^{(r)} \to K_1^{(r)} \subseteq M^{(r)}$ be the natural inclusions. For each n-simplex $\sigma \in K^{(r)}$, let $h_n(\sigma)$ be the sum of the (suitably oriented) $(n+1)$-simplices that occur in $M_\sigma^{(r)}$. Now, the formula for the boundary of a product translates at the level of chains into the formula

$$\partial_{n+1} \circ h_n(\sigma) = j_n(\sigma) - i_n(\sigma) - h_{n-1} \circ \partial_n(\sigma)$$

which rearranges to show that $h_\bullet$ defines a chain homotopy between $i_\bullet$ and $j_\bullet$. Since $F_\bullet$ is a chain map, it follows that

$$F_n \circ j_n - F_n \circ i_n = \partial_{n+1} \circ F_{n+1} \circ h_n + F_n \circ h_{n-1} \circ \partial_n$$

for all n. Because $f_n = F_n \circ i_n$ and $g_n = F_n \circ j_n$, we have that $F_\bullet \circ h_\bullet$ is the required chain homotopy. □

Finally, for homology to be a useful tool, we also need to check that it has the usual 'functorial' properties.

Lemma 5.35. *Let $X \overset{\phi}{\to} Y \overset{\psi}{\to} Z$ be continuous maps of triangulable spaces. Then*

$$(\psi \circ \phi)_* = \psi_* \circ \phi_* : H_n(X) \to H_n(Z)$$

for all n. Furthermore, $(\mathrm{id}_X)_* = \mathrm{id}_{H_n(X)}$ for all n.

Proof. Let $\alpha : |K| \to X$, $\beta : |L| \to Y$, $\gamma : |M| \to Z$ be triangulations. The second assertion is clear, because we may take id_K as a simplicial approximation to $\alpha^{-1} \circ \mathrm{id}_X \circ \alpha = \mathrm{id}_{|K|}$.

For the first assertion, take $g : L^{(s)} \to M$ to be a simplicial approximation to $\gamma^{-1} \circ \psi \circ \beta$ and $f : K^{(r)} \to L^{(s)}$ a simplicial approximation to $\beta^{-1} \circ \phi \circ \alpha$. For any vertex v of $K^{(r)}$ we now have

$$\gamma^{-1} \circ \psi \circ \phi \circ \alpha(\mathrm{St}_{K^{(r)}}(x)) \subseteq \gamma^{-1} \circ \psi \circ \beta(\mathrm{St}_{L^{(s)}}(f(x))) \subseteq \mathrm{St}_M(g \circ f(x))$$

so $g \circ f$ is a simplicial approximation to $\gamma^{-1} \circ \psi \circ \phi \circ \alpha$. Therefore

$$(\psi \circ \phi)_* = (g \circ f)_* = g_* \circ f_* = \psi_* \circ \phi_*$$

as required. □

As an immediate consequence of Theorem 5.34 and Lemma 5.35, we obtain that homology is an invariant of homotopy equivalence.

Corollary 5.36. *Let X, Y be triangulable spaces. A homotopy equivalence $\phi : X \to Y$ induces isomorphisms $\phi_* : H_n(X) \to H_n(Y)$ for all n.*

6 Homology calculations

With the results of the previous section in hand, homology becomes a powerful tool. We'll start by recording some applications.

6.1 Homology of spheres and applications

Part of our motivation for developing homology groups was to detect the 'higher-dimensional holes' in spheres.

Example 6.1. In Example 5.26, we saw that no sphere S^{n-1} has the same homology groups as a point. Therefore, by Corollary 5.36, spheres are never contractible. We also saw that distinct spheres have different homology groups. This shows that $S^m \simeq S^n$ if and only if $m = n$.

The next result is a classical theorem, called *invariance of domain.* Its proof is similar, but a little more subtle. It's a nice example of how to use homotopy equivalence to prove that spaces are not homeomorphic.

Theorem 6.2. *If $\mathbb{R}^m \cong \mathbb{R}^n$ then $m = n$.*

Proof. Suppose that $\phi : \mathbb{R}^m \to \mathbb{R}^n$ is a homeomorphism. Composing with a translation of $\mathbb{R}^n$, we may assume that $\phi(0) = 0$. Therefore $\mathbb{R}^m - \{0\} \cong \mathbb{R}^n - \{0\}$. In Remark 1.9, we noticed that punctured Euclidean space is homotopy equivalent to the unit sphere, so it follows that $S^{m-1} \simeq S^{n-1}$. By Example 6.1, it follows that $m = n$. □

Another nice application is to a higher-dimensional Brouwer fixed point theorem. The next theorem is now proved in exactly the same way as Corollary 2.22, except using the $(n-1)$st homology group instead of the fundamental group.

Theorem 6.3. *Any continuous map from the closed n-dimensional ball to itself*

$$\phi : D^n \to D^n$$

has a fixed point.

Proof. The proof is left as an easy exercise, following the proofs of Theorem 2.21 and Corollary 2.22. □

6.2 Mayer–Vietoris theorem

For more sophisticated applications, we need to be able to compute more examples. The main theorem of this section is a gluing theorem for homology groups. It plays an analogous role to the role of the Seifert–van Kampen theorem for the fundamental group.

Definition 6.4. A sequence of homomorphisms of abelian groups

$$\cdots \to A_{i+1} \overset{f_i}{\to} A_i \overset{f_{i-1}}{\to} A_{i-1} \to \cdot$$

is said to be *exact at A_i* if $\mathrm{im}\,(f_i) = \ker(f_{i-1})$. The sequence is *exact* if it is exact at every A_i. A *short exact sequence* is sequence

$$0 \to A \to B \to C \to 0$$

which is exact.

The idea behind exact sequences is that, if one know most of the terms in an exact sequence, then one can figure out the remaining terms. We can think of homology groups as measuring the failure of a chain complex to be exact.

Example 6.5. Exactness at B of

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

means that g is surjective. Exactness at A

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

means that f is surjective. In particular, a *very short* sequence

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

is exact if and only if f is an isomorphism.

The Mayer–Vietoris sequence concerns a simplicial complex K which is the union of two subcomplexes L and M. Let N denote the intersection, and write $K = L \cup_N M$. Let $i : N \to L$, $j : N \to M$, $l : L \to K$ and $m : M \to K$ be the inclusion maps.

Theorem 6.6 (Mayer–Vietoris). *Consider simplicial complexes $K = L \cup_N M$. For every $n \in \mathbb{Z}$, there is a homomorphism*

$$\delta_* : H_n(K) \to H_{n-1}(N)$$

that makes the following sequence exact.

$$\begin{aligned}
&\cdots \xrightarrow{i_* \oplus j_*} H_{n+1}(L) \oplus H_{n+1}(M) \xrightarrow{l_* - m_*} H_{n+1}(K) \xrightarrow{\delta_*} \\
&H_n(N) \xrightarrow{i_* \oplus j_*} H_n(L) \oplus H_n(M) \xrightarrow{l_* - m_*} H_n(K) \xrightarrow{\delta_*} \\
&H_{n-1}(N) \xrightarrow{i_* \oplus j_*} H_{n-1}(L) \oplus H_{n-1}(M) \xrightarrow{l_* - m_*} \cdots
\end{aligned}$$

The theorem follows from a purely algebraic fact. A sequence of chain maps

$$A_\bullet \to B_\bullet \to C_\bullet$$

is exact at $B_\bullet$ if $A_n \to B_n \to C_n$ is exact at B_n for every n. Again, the sequence is exact if it is exact at every term.

Lemma 6.7 (Snake lemma)**.** *Let*

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

be a short exact sequence of chain complexes. For every $n \in \mathbb{Z}$, there is a homomorphism

$$\delta_* : H_{n+1}(C_\bullet) \to H_n(A_\bullet)$$

that makes the following sequence exact.

$$\begin{tikzcd}
\cdots \arrow[r, "f_*"] & H_{n+1}(B_\bullet) \arrow[r, "g_*"] & H_{n+1}(C_\bullet) \arrow[dll, "\delta_*"] \\
H_n(A_\bullet) \arrow[r, "f_*"] & H_n(B_\bullet) \arrow[r, "g_*"] & H_n(C_\bullet) \arrow[dll, "\delta_*"] \\
H_{n-1}(A_\bullet) \arrow[r, "f_*"] & H_{n-1}(B_\bullet) \arrow[r, "g_*"] & \cdots
\end{tikzcd}$$

Proof. The proof is a long sequence of trivial checks. The reader may like to refer to the following commutative diagram.

$$\begin{tikzcd}
 & \vdots \arrow[d] & \vdots \arrow[d] & \vdots \arrow[d] & \\
0 \arrow[r] & A_{n+1} \arrow[r, "f_{n+1}"] \arrow[d, "\partial_{n+1}"] & B_{n+1} \arrow[r, "g_{n+1}"] \arrow[d, "\partial_{n+1}"] & C_{n+1} \arrow[r] \arrow[d, "\partial_{n+1}"] & 0 \\
0 \arrow[r] & A_n \arrow[r, "f_n"] \arrow[d, "\partial_n"] & B_n \arrow[r, "g_n"] \arrow[d, "\partial_n"] & C_n \arrow[r] \arrow[d, "\partial_n"] & 0 \\
0 \arrow[r] & A_{n-1} \arrow[r, "f_{n-1}"] \arrow[d] & B_{n-1} \arrow[r, "g_{n-1}"] \arrow[d] & C_{n-1} \arrow[r] \arrow[d] & 0 \\
 & \vdots & \vdots & \vdots &
\end{tikzcd}$$

The first task is to construct the homomorphism $\delta_* : H_{n+1}(C_\bullet) \to H_n(A_\bullet)$. Consider $[x] \in H_{n+1}(C_\bullet)$, defined by an $(n+1)$-cycle $x \in Z_{n+1}(C_\bullet) \subseteq C_{n+1}$. By exactness at C_{n+1}, there is $y \in B_{n+1}$ such that $g_{n+1}(y) = x$. Because x is a cycle,

$$g_n \circ \partial_{n+1}(y) = \partial_n \circ g_{n+1}(y) = \partial_n(x) = 0$$

so $\partial_{n+1}(y) \in \ker(g_n)$, which equals $\mathrm{im}(f_n)$ by exactness at B_n. Therefore, $\partial_{n+1}(y) = f_n(z)$ for some $z \in A_n$. What is more,

$$f_{n-1} \circ \partial_n(z) = \partial_n \circ f_n(z) = \partial_n \circ \partial_n(y) = 0$$

but, exactness at A_{n-1} tells us that f_{n-1} is injective, so $\partial_n(z) = 0$. Therefore, $z \in Z_n(A_\bullet)$, and so defines a class in $H_n(A_\bullet)$. We set $\delta_*([x]) = [z]$.

We next need to check that this is well defined: that is, if we modify x by a boundary, and make different choices of y and z, then the resulting $(n-1)$-chain only changes by a boundary. Consider therefore $x' \in C_{n+1}$ with $x' - x = \partial_{n+2}(w)$, for $w \in C_{n+2}$, and suppose the construction proceeds similarly with $y' \in B_{n+1}$ and $z' \in A_{n-2}$. Exactness at C_{n+1} implies that $w = g_{n+2}(u)$ for some $u \in B_{n+2}$. Now

$$g_{n+1} \circ \partial_{n+2}(u) = \partial_{n+2} \circ g_{n+2}(u) = \partial_{n+2}(w) = x' - x$$

so in particular $y' - y - \partial_{n+2}(u) \in \ker g_{n+1} = \mathrm{im}\, f_{n+1}$, by exactness at B_{n+1}. Therefore, there is $v \in A_{n+1}$ such that $y' - y = f_{n+1}(v) + \partial_{n+2}(u)$. This gives us

$$f_n(z' - z) = \partial_{n+1}(y' - y) = \partial_{n+1} \circ f_{n+1}(v) + \partial_{n+1} \circ \partial_{n+2}(u) = f_n \circ \partial_{n+1}(v)$$

by commutativity of the diagram. Since f_n is injective, it follows that $z' - z$ is a boundary, as required.

Thirdly, we note that the assignment $\delta_*[x] := [z]$ is a homomorphism, since if $x'' = x + x'$, then we can take $y'' = y + y'$ and $z'' = z + z'$. This completes the construction of δ_*.

We next need to check that the sequence is exact. We start at $H_n(B_\bullet)$, where we need to show that $\mathrm{im}\, f_* = \ker g_*$. If $[a] \in H_n(A_\bullet)$ then

$$g_* \circ f_*[a] = [g_n \circ f_n(a)] = 0$$

by exactness at B_n, so $\mathrm{im}\, f_* \subseteq \ker g_*$. Suppose now that $[b] \in H_n(B_\bullet)$ with $g_*[b] = 0$. This means that $g_n(b) = \partial_{n+1}(x)$ for some $x \in C_{n+1}$, but $x = g_{n+1}(y)$ for $y \in B_{n+1}$ since g_{n+1} is surjective. Therefore

$$g_n \circ \partial_{n+1}(y) = \partial_{n+1} \circ g_{n+1}(y) = \partial_{n+1}(x) = g_n(b)$$

so $b - \partial_{n+1}(y) \in \ker g_n = \mathrm{im}\, f_n$ by exactness at $B_\bullet$. Therefore $b - \partial_{n+1}(y) = f_n(a)$ for some $a \in A_n$. Now

$$f_{n-1} \circ \partial_n(a) = \partial_n \circ f_n(a) = \partial_n(b) - \partial_n \circ \partial_{n+1}(y) = 0$$

since $b \in Z_n(B_\bullet)$ so, since f_{n-1} is injective, $\partial_n(a) = 0$. Therefore, $a \in Z_n(A_\bullet)$ and so defines a class $[a] \in H_n(A_\bullet)$ that satisfies

$$f_*[a] = [f_n(a)] = [b - \partial_{n+1}(y)] = [b]$$

which shows that $[b] \in \operatorname{im} f_*$. Therefore $\ker g_* \subseteq \operatorname{im} f_*$, which completes the proof of exactness at $H_n(B_\bullet)$.

We next turn to exactness at $H_n(A_\bullet)$. Consider first $[z] = \delta_*[x]$, for x and z as in the construction of δ_* above. By construction,

$$f_*[z] = [f_n(z)] = [\partial_{n+1}(y)] = 0$$

so $\operatorname{im} \delta_* \subseteq \ker f_*$. For the reverse inclusion, suppose that $f_*[z] = 0$, which means that $f_n(z) = \partial_{n+1}(y)$ for some $y \in B_{n+1}$. Setting $x = g_{n+1}(y)$, we have that $\delta_*[x] = [z]$ by the construction of δ_*. This completes the proof that $\operatorname{im} \delta_* = \ker f_*$, so the sequence is exact at $H_n(A_\bullet)$.

Finally, we need to check exactness at $H_n(C_\bullet)$, meaning that we need to check that $\operatorname{im} g_* = \ker \delta_*$. Suppose first that $[x] \in \operatorname{im} g_*$. This means that, in the construction of δ_*, we can take $x = g_n(y)$ for y an $(n+1)$-cycle, so $\partial_n(y) = 0$ and we can take $z = 0$. Therefore $\delta_*[x] = 0$, so $\operatorname{im} g_* \subseteq \ker \delta_*$ as required. For the reverse inclusion, suppose that $\delta_*[x] = 0$, meaning that, in the construction of δ_*, we can take $z \in A_{n-1}$ to be a boundary. That is, $z = \partial_n(b)$ for some $b \in B_n$. Now

$$\partial_n \circ f_n(b) = f_{n-1} \circ \partial_n(b) = f_{n-1}(z) = \partial_n(y)$$

so $y - f_n(b)$ is a cycle. Also, $g_n(y - f_n(b)) = g_n(y) = x$, so x is the image of a cycle. In particular, $[x] \in \operatorname{im} g_*$. This completes the proof of exactness at $H_n(C_\bullet)$, and hence the proof of the lemma. $\square$

We can now prove Theorem 6.6 as an application of Lemma 6.7.

Proof of Theorem 6.6. It's easy to see that $H_n(C_\bullet \oplus D_\bullet) \cong H_n(C_\bullet) \oplus H_n(D_\bullet)$. Therefore, by Lemma 6.7, it suffices to show that

$$0 \to C_\bullet(N) \overset{i_\bullet \oplus j_\bullet}{\to} C_\bullet(L) \oplus C_\bullet(M) \overset{l_\bullet - m_\bullet}{\to} C_\bullet(K) \to 0$$

is a short exact sequence of chain complexes.

For the first term, note that $C_n(N)$ is in fact a direct summand of $C_n(L)$ and $C_n(M)$, with i_n and j_n the inclusions. In particular, i_n and j_n are injective, and hence so is $i_n \oplus j_n$.

For the last term, consider $c \in C_n(K)$. Since $K = L \cup M$, we may write $c = c_L + c_M$ where c_L is supported on simplices of L and c_M is supported on simplices of M. This means that $c_L = l_n(b_L)$ and $c_M = m_n(b_M)$ for some $b_L \in C_n(L)$ and $b_M \in C_n(M)$. In particular, $c = l_n(b_L) - m_n(b_M)$, so is in the image of $l_n - m_n$ as required.

Finally, we check exactness at $C_n(L) \oplus C_n(M)$. For a pair $(b_L, b_M) \in C_n(L) \oplus C_n(M)$, we have $l_n(b_L) - m_n(b_M)$ if and only if every simplex that occurs in b_L also occurs in b_M with the same coefficient. This means precisely that both chains are supported on the intersection N, and that there is some $a \in C_n(N)$ such that $b_L = i_n(a)$ and $b_M = j_n(a)$. Therefore $\ker(l_n - m_n) = \operatorname{im}(i_n \oplus j_n)$, as required. □

When applying the Mayer–Vietoris theorem, the next result is often useful.

Lemma 6.8 (Five lemma)**.** *Suppose that the rows of the following commutative diagram are exact.*

$$\begin{array}{ccccccccc} A & \longrightarrow & B & \longrightarrow & C & \longrightarrow & D & \longrightarrow & E \\ \downarrow{\scriptstyle\alpha} & & \downarrow{\scriptstyle\beta} & & \downarrow{\scriptstyle\gamma} & & \downarrow{\scriptstyle\delta} & & \downarrow{\scriptstyle\epsilon} \\ A' & \longrightarrow & B' & \longrightarrow & C' & \longrightarrow & D' & \longrightarrow & E' \end{array}$$

If α, β, δ and ϵ are all isomorphisms, then so is γ.

Proof. The proof is an easy diagram-chase. It is Question 3 of Example Sheet 4. □

We are now at last in a position to prove that the canonical homomorphism identifies the homology of a simplicial complex and its subdivision.

Proof of Proposition 5.31. The proof is by induction on the number of simplices of K. Note that the result is trivial if K has one simplex (necessarily a 0-simplex). For the inductive step, let $\sigma \in K$ be a simplex which is not a proper face of any simplex of K – for instance, any simplex of maximal dimension has this property. Let $L = K - \{\sigma\}$, let M be the subcomplex of K that consists of σ and its faces, and let $N = L \cap M = \partial\sigma$. Abusing notation, let s denote the restrictions of s to the barycentric subdivisions of any of these subcomplexes.

Now M and M' are both cones, so $s_* : H_n(M') \to H_n(M)$ is trivially an isomorphism. Also, L and N necessarily have fewer simplices than K, so by induction $s_* : H_n(L') \to H_n(L)$ and $s_* : H_n(N') \to H_n(N)$ are both isomorphisms. Applying the Mayer–Vietoris theorem to $K = L \cup_N M$ and $K' = L' \cup_{N'} M'$, we obtain the following commutative diagram with exact rows.

$$\begin{tikzcd}
H_{n+1}(N') \arrow[r] \arrow[d, "s_*"] & H_n(L') \oplus H_n(M') \arrow[r] \arrow[d, "s_* \oplus s_*"] & H_n(K') \arrow[r] \arrow[d, "s_*"] & H_{n-1}(N') \arrow[r] \arrow[d, "s_*"] & H_{n-1}(L') \oplus H_{n-1}(M') \arrow[d, "s_* \oplus s_*"] \\
H_{n+1}(N) \arrow[r] & H_n(L) \oplus H_n(M) \arrow[r] & H_n(K) \arrow[r] & H_{n-1}(N) \arrow[r] & H_{n-1}(L) \oplus H_{n-1}(M)
\end{tikzcd}$$

Since all the other vertical arrows are isomorphisms, it follows that $s_* : H_n(K') \to H_n(K)$ is an isomorphism for all n. □

6.3 Homology of compact surfaces

The Mayer–Vietoris sequence can also be used to compute the homology groups of compact surfaces. Recall the compact, orientable surfaces Σ_g from Example 3.29. These were constructed by attaching a 2-cell to a graph.

$$\Sigma_g \cong \Gamma_{2g} \cup_{\rho_g} D^2$$

We therefore start by computing the homology of wedges of circles.

Example 6.9. Let $\Gamma_r := \bigvee_{i=1}^r S^1$. It is easy to see that Γ_r is triangulable: simply triangulate the circles S^1 and identify a single vertex. Now $\Gamma_1 \cong S^1$ and

$$\Gamma_r \cong \Gamma_{r-1} \vee S^1$$

so we may apply the Mayer–Vietoris sequence to compute the homology groups by induction. Let Γ_r be the realisation of $K = L \cup_N M$, where L is a triangulation of Γ_{r-1}, M is a triangulation of S^1 and $N = \{*\}$. Then Mayer–Vietoris gives us an exact sequence, and after filling in the terms that we know, we obtain the following.

$$\begin{tikzcd}
0 \arrow[r] & H_1(\Gamma_{r-1}) \oplus \mathbb{Z} \arrow[r] & H_1(\Gamma_r) \arrow[dll, "\delta_*"'] \\
\mathbb{Z} \arrow[r] & \mathbb{Z} \oplus \mathbb{Z} \arrow[r] & H_0(\Gamma_r)
\end{tikzcd}$$

Since Γ_r is connected we also know that $H_0(\Gamma_r) \cong \mathbb{Z}$, and the first arrow of the bottom row is injective just by the general description given in Lemma 5.14. Therefore im $\delta_* \cong 0$, and so we obtain a very short exact sequence.

$$0 \to H_1(\Gamma_{r-1}) \oplus \mathbb{Z} \to H_1(\Gamma_r) \to 0$$

By induction, it follows that:

$$H_n(\Gamma_r) \cong \begin{cases} \mathbb{Z} & n = 0 \\ \mathbb{Z}^r & n = 1 \\ 0 & \text{otherwise} \end{cases}$$

For future reference, it is useful to note that, if $\alpha_1, \ldots, \alpha_r$ represent a standard set of loops in each copy of S^1, then the homology classes $[\alpha_1], \ldots, [\alpha_r]$ form a basis for $H_1(\Gamma_r)$.

Remark 6.10. In fact, our discussion of the groups H_0 was valid whenever K, L, M, N were connected. In this setting, we can always conclude that $\delta_* : H_1(K) \to H_0(N)$ is the zero map.

Rather like in the proof of Lemma 3.24, it's convenient to attach the 2-cell in two stages: first attach an annulus, and then attach a disc to the other boundary of the annulus. To this end, we define

$$\Sigma_g^* := \Gamma_{2g} \cup_{\rho_g} (S^1 \times I)$$

where ρ_g attaches $S_1 \times \{0\}$ to Γ_{2g}. Note that Σ_g^* is homeomorphic to Σ_g with an open disc removed. Shrinking I to a point shows that Σ_g^* deformation retracts to Γ_{2g}. We may then recover Σ_g as

$$\Sigma_g = \Sigma_g^* \cup_i D^2$$

where $i : S^1 \to \Sigma_g^*$ identifies $S^1 = \partial D^2$ with $S^1 \times \{1\} \subseteq \Sigma_g^*$.

Example 6.11. Choose compatible triangulations so that Σ_g is the realisation of $K = L \cup_N M$, with Σ_g^* the realisation of L, D^2 the realisation of M, and S^1 the realisation of N. (It's not hard to see that such triangulations exist.) Since Σ_g^* deformation retracts to Γ_{2g}, we know the homology groups of L, M and N. Mayer–Vietoris now gives us the following exact sequence. (As in

Remark 6.10, we may ignore the 0th row.)

$$
\begin{tikzcd}
 & 0 \arrow[r] & H_2(\Sigma_g) \arrow[dll, "\delta_*"] \\
\mathbb{Z} \arrow[r, "i_*"] & \mathbb{Z}^{2g} \arrow[r] & H_1(\Sigma_g) \arrow[dll, "\delta_*"] \\
0 & &
\end{tikzcd}
$$

To complete the calculation, we need to work out the map i_*. The element $1 \in \mathbb{Z} \cong H_1(S^1)$ is represented by a 1-cycle γ that goes once around the circle. The inclusion i identifies γ with the boundary circle $\partial\Sigma_g^* \subseteq \Sigma_g^*$, and the deformation retraction to Γ_{2g} identifies this boundary with the image of the circle under ρ_g. In other words, i_* is equal to the map induced by ρ_g. From the definition of ρ_g,

$$\rho_{g*}(1) = [\alpha_1]+[\beta_1]-[\alpha_1]-[\beta_1]+[\alpha_2]+\ldots+[\beta_{g-1}]+[\alpha_g]+[\beta_g]-[\alpha_g]-[\beta_g] = 0$$

and we deduce that i_* is the zero map. This tells us that the above exact sequence breaks into two very short exact sequences:

$$0 \to H_2(\Sigma_g) \to \mathbb{Z} \to 0$$

and

$$0 \to \mathbb{Z}^{2g} \to H_1(\Sigma_g) \to 0$$

from which the homology groups of Σ_g follow.

$$H_n(\Sigma_g) \cong \begin{cases} \mathbb{Z} & n = 0, 2 \\ \mathbb{Z}^{2g} & n = 1 \\ 0 & \text{otherwise} \end{cases}$$

This gives an alternative proof that surfaces Σ_g are pairwise not homotopy-equivalent. It also shows us that Σ_g is not homotopy-equivalent to Γ_{2g}, say; this wasn't obvious from our computation of the fundamental group.

We can compute the homology groups of the non-orientable surfaces S_g similarly.

Example 6.12. Consider the non-orientable surfaces S_g of Example 3.30. The calculation of $H_n(S_g)$ proceeds identically to the orientable case, until we obtain the following long exact sequence.

$$0 \longrightarrow H_2(S_g) \xrightarrow{\delta_*} \mathbb{Z} \xrightarrow{i_*} \mathbb{Z}^{g+1} \longrightarrow H_1(S_g) \xrightarrow{\delta_*} 0$$

As in Example 6.11, i_* is induced by the attaching map of the 2-cell, in this case σ_g. Looking at the definition of σ_g, we see that

$$i_*(1) = \sigma_{g*}(1) = 2[\alpha_0] + 2[\alpha_1] + \cdots + 2[\alpha_g]\ .$$

Therefore i_* is injective, so $H_2(S_g) \cong \ker i_* \cong 0$. We can also compute the first homology group.

$$\begin{aligned} H_1(S_g) &\cong \text{im } l_* \\ &\cong \mathbb{Z}^{g+1}/\ker l_* \\ &= \mathbb{Z}^{g+1}/\text{im } i_* \\ &= \langle[\alpha_0]\rangle \oplus \cdots \oplus \langle[\alpha_g]\rangle/\langle 2[\alpha_0] + \cdots + 2[\alpha_g]\rangle \\ &\cong \mathbb{Z}^g \oplus (\mathbb{Z}/2\mathbb{Z}) \end{aligned}$$

In summary:

$$H_n(S_g) \cong \begin{cases} \mathbb{Z} & n = 0 \\ \mathbb{Z}^g \oplus (\mathbb{Z}/2\mathbb{Z}) & n = 1 \\ 0 & \text{otherwise} \end{cases}$$

6.4 Rational homology and Euler characteristic

So far we have defined homology with coefficients in $\mathbb{Z}$. In fact, homology can be defined with much more general coefficients. In this section, we'll briefly discuss what happens if we take coefficients in $\mathbb{Q}$.

Definition 6.13. Let K be a simplicial complex. For each $n \in \mathbb{Z}$, the vector space of *rational n-chains* $C_n(K;\mathbb{Q})$ is the vector space over $\mathbb{Q}$ with basis the n-simplices of K. As before, we arbitrarily choose orientations on the simplices, and identify $-\sigma$ with $\bar{\sigma}$.

The definition of the boundary map and its basic properties go through as before, and we may therefore define homology as before.

Definition 6.14. Let K be a simplicial complex. Then $Z_n(K;\mathbb{Q})$ is defined to be the kernel of $\partial_n : C_n(K;\mathbb{Q}) \to C_{n-1}(K;\mathbb{Q})$ and $B_n(K;\mathbb{Q})$ is defined to be the image of $\partial_{n+1} : C_{n+1}(K;\mathbb{Q}) \to C_n(K;\mathbb{Q})$. Since $\partial_n \circ \partial_{n+1}$, we have $B_n(K;\mathbb{Q}) \subseteq Z_n(K;\mathbb{Q})$, and we set the *nth rational homology vector space* of K to be

$$H_n(K;\mathbb{Q}) := Z_n(K;\mathbb{Q})/B_n(K;\mathbb{Q})$$

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

This is often easier to work with than the usual (integral) homology, since the groups involved are vector spaces. However, it loses some information, as the next result makes clear.

Lemma 6.15. *Let K be a simplicial complex. If*

$$H_n(K) \cong \mathbb{Z}^b \oplus F$$

where F is a finite abelian group, then $H_n(K;\mathbb{Q}) \cong \mathbb{Q}^b$.

Proof. Let $H_n(K;\mathbb{Q}) \cong \mathbb{Q}^{b'}$ (since it's a finite-dimensional vector space). We need to prove that $b' = b$.

There is a natural comparison maps $C_n(K) \to C_n(K;\mathbb{Q})$ obtained by thinking of an n-chain as a rational n-chain. Since these comparison maps are chain maps, they induce comparison maps $B_n(K) \to B_n(K;\mathbb{Q})$, $Z_n(K) \to Z_n(K;\mathbb{Q})$ and $H_n(K) \to H_n(K;\mathbb{Q})$.

If $c \in Z_n(K,\mathbb{Q})$ then, multiplying by the product m of the denominators, we see that mc is in the image of $Z_n(K)$ for some non-zero m. It follows that $b' \leq b$.

Let $[c_1],\ldots,[c_b] \in H_n(K)$ together generate a copy of $\mathbb{Z}^b$, and consider their images in $H_n(K,\mathbb{Q})$. Suppose there are rationals $\lambda_1,\ldots,\lambda_b$ so that

$$\sum_{i=1}^{b} \lambda_i[c_i] = 0$$

in $H_n(K,\mathbb{Q})$. Clearing denominators, we may take the λ_i to be integers. But all the c_i are integral, so this implies that $\sum_i \lambda_i c_i$ is a boundary in $Z_n(K)$, which in turn implies that

$$\sum_{i=1}^{b} \lambda_i[c_i] = 0$$

in $H_n(K)$. Therefore $\lambda_i = 0$ for all i. This implies that $b' \geq b$, and the proof is complete. $\square$

So rational homology by itself does not have the power to distinguish $\mathbb{R}P^2$ from a point, for instance. Its advantage is that it leads to some invariants that are very easy to compute.

Definition 6.16. Let K be a simplicial complex. The *Euler characteristic* of K is

$$\chi(K) := \sum_{n\in\mathbb{Z}} (-1)^n \dim_{\mathbb{Q}} H_n(K;\mathbb{Q}) \ ,$$

which we note is a finite sum. As usual, if $\alpha : |K| \to X$ is a triangulation, we also set $\chi(X) := \chi(K)$.

The reason this is so easy to compute boils down to the rank-nullity formula.

Lemma 6.17. *Let K be a simplicial complex. Then*

$$\chi(K) = \sum_{n\in\mathbb{Z}} (-1)^n \#\{n\text{–simplices in } K\}$$

which we again note is a finite sum.

Proof. Since the number of n-simplices in K is $\dim_{\mathbb{Q}} C_n(K;\mathbb{Q})$, we in fact prove the following, more natural, identity.

$$\chi(K) = \sum_{n\in\mathbb{Z}} (-1)^n \dim_{\mathbb{Q}} C_n(K;\mathbb{Q})$$

The rank-nullity formula applied to the quotient map $Z_n(K;\mathbb{Q}) \to H_n(K;\mathbb{Q})$ tells us that

$$\dim_{\mathbb{Q}} H_n(K;\mathbb{Q}) + \dim_{\mathbb{Q}} B_n(K;\mathbb{Q}) = \dim_{\mathbb{Q}} Z_n(K;\mathbb{Q})$$

while, applied to the boundary map ∂_n, it tells us that

$$\dim_{\mathbb{Q}} Z_n(K;\mathbb{Q}) + \dim_{\mathbb{Q}} B_{n-1}(K;\mathbb{Q}) = \dim_{\mathbb{Q}} C_n(K\ bQ) \ .$$

We therefore compute:

$$\begin{aligned}
\sum_{n\in\mathbb{Z}} (-1)^n \dim_{\mathbb{Q}} C_n(K;\mathbb{Q}) \quad &= \quad \sum_{n\in\mathbb{Z}} (-1)^n \dim_{\mathbb{Q}}\ Z_n(K;\mathbb{Q}) + \sum_{n\in\mathbb{Z}} (-1)^n \dim_{\mathbb{Q}} B_{n-1}(K;\mathbb{Q}) \\
&= \quad \sum_{n\in\mathbb{Z}} (-1)^n \dim_{\mathbb{Q}} Z_n(K;\mathbb{Q}) - \sum_{n\in\mathbb{Z}} (-1)^n \dim_{\mathbb{Q}} B_n(K;\mathbb{Q}) \\
&= \quad \sum_{n\in\mathbb{Z}} (-1)^n \dim_{\mathbb{Q}} H_n(K;\mathbb{Q})
\end{aligned}$$

as required. $\square$

Remark 6.18. In the case of a 2-dimensional simplicial complex with V 0-simplices, E 1-simplices and F 2-simplices, we recover the familiar formula

$$\chi(K) = V - E + F$$

which you have probably seen before. Interpreting it in terms of homology gives a natural reason why the Euler characteristic is a homeomorphism, and even homotopy, invariant.

6.5 The Lefschetz fixed-point theorem

The final theorem of the course is a far-reaching generalisation of the Brouwer fixed-point theorem.

Definition 6.19. Let X be a triangulable space and $\phi : X \to X$ a continuous map. The *Lefschetz number* of ϕ is defined to be

$$L(\phi) := \sum_{n\in\mathbb{Z}} (-1)^n \mathrm{tr}(\phi_* : H_n(X;\mathbb{Q}) \to H_n(X;\mathbb{Q}))$$

which, as usual, we note is a finite sum.

The first thing to notice is that this is a generalisation of Euler characteristic.

Remark 6.20. If X is triangulable then $\chi(X) = L(\mathrm{id}_X)$, since the trace of the identity map on a vector space is the dimension.

Like Euler characteristic, the Lefschetz number can be computed at the level of chains.

Lemma 6.21. *If $f : K \to K$ is a simplicial self-map of a simplicial complex then*

$$L(|f|) := \sum_{n\in\mathbb{Z}} (-1)^n \ \mathrm{tr}(f_n : C_n(K;\mathbb{Q}) \to C_n(K;\mathbb{Q})) \ .$$

Proof. First, consider the following commutative diagram of linear maps of vector spaces, with exact rows.

$$\begin{tikzcd}
0 \arrow[r] & A \arrow[r] \arrow[d, "\alpha"] & B \arrow[r] \arrow[d, "\beta"] & C \arrow[r] \arrow[d, "\gamma"] & 0 \\
0 \arrow[r] & A' \arrow[r] & B' \arrow[r] & C' \arrow[r] & 0
\end{tikzcd}$$

It's an easy exercise in linear algebra to see that $\mathrm{tr}(\beta) = \mathrm{tr}(\alpha) + \mathrm{tr}(\gamma)$. The proof is now identical to the proof of Lemma 6.17. $\square$

We're now ready to state the Lefschetz fixed-point theorem. We only sketch the proof, which is non-examinable.

Theorem 6.22 (Lefschetz fixed-point theorem)**.** *Let $\phi : X \to X$ be a continuous self-map of a triangulable space. If $L(\phi) \neq 0$ then ϕ has a fixed point.*

Sketch proof. We prove the contrapositive: if ϕ has no fixed point then $L(\phi) = 0$. By compactness, if ϕ has no fixed point then there is $\delta > 0$ such that $\|x - \phi(x)\| > \delta$ for all $x \in X$. We now choose a simiplicial complex K, and identify $|K|$ with X, so that $\mathrm{mesh}(K) < \delta/2$. In particular, if $x \in \sigma \in K$ then $\phi(x) \notin \sigma$.

Let $f : K^{(r)} \to K$ is a simplicial approximation to ϕ. If v is a vertex of $K^{(r)}$ contained in a simplex $\sigma \in K$ then $\phi(v) \in \mathrm{St}_{K^{(r)}}(f(v))$ so $\|\phi(v) - f(v)\| < \delta/2$. But $\|\phi(v) - v\| > \delta$ and so $\|v - f(v)\| > \delta/2$. Therefore, $f(v) \notin \sigma$.

Let $\iota_\bullet : C_\bullet(K;\mathbb{Q}) \to C_\bullet(K^{(r)};\mathbb{Q})$ be the map that induces the canonical identification on homology. For each n-simplex $\sigma \in K$, $\iota_n(\sigma)$ is supported on simplices contained in σ. Therefore, since f takes vertices of σ out of σ, it follows that $f_n \circ \iota_n(\sigma)$ is supported on simplices that are disjoint from σ.

Since ϕ_* is induced at the level of chains by $f_n \circ \iota_n$, we now have

$$L(\phi) = \sum_{n \in \mathbb{Z}} (-1)^n \mathrm{tr}(f_n \circ \iota_n)$$

by Lemma 6.21. But $f_n \circ \iota_n$ moves every simplex off itself and so the corresponding matrix has zeroes on the diagonal. Therefore its trace is zero and the result follows. $\square$

We immediately obtain a dramatic generalisation of the Brouwer fixed-point theorem.

Corollary 6.23. *If X is triangulable and contractible then any continuous self-map $\phi : X \to X$ has a fixed point.*

Proof. The only non-zero term in $L(\phi)$ comes from the 0-dimensional homology. The map $\phi_* : H_0(X;\mathbb{Q}) \to H_0(X;\mathbb{Q}) \cong \mathbb{Q}$ is the identity map, and so $L(\phi) = 1$. Therefore, by Theroem 6.22, ϕ has a fixed point. $\square$