TOPOLOGY AND GROUPS

Contents

Introduction

Topology and group theory are strongly intertwined, in ways that are interesting and unexpected when one first meets them. The main interface is the concept of the 'fundamental group', which is a recipe that assigns to each topological space a group. This provides a lot of useful information about the space. For example, one can often show that two spaces are not homeomorphic, by establishing that their fundamental groups are not isomorphic. Chapters II and III are mainly devoted to defining the fundamental group of a space and establishing some of its basic properties. A beautiful application is a proof of the Fundamental Theorem of Algebra.

When computing the fundamental group of a space, the answer is typically given in terms of a list of generators, together with the relations that they satisfy. This is a method for defining groups that occurs in many first courses on group theory. However, to make the notion precise requires some work, and in particular necessitates the introduction of a particularly basic type of group, known as a 'free group'. Free groups were introduced in the second year course on Group Theory, but familiarity with that course will not be assumed here. Chapters IV and V deal with these aspects of group theory.

In Chapters V and VI, the two themes of the course, topology and groups, are brought together. At the end of Chapter V, a central result, the Seifert - van Kampen theorem, is proved. This theorem allows us to compute the fundamental group of almost any topological space. In Chapter VI, covering spaces are introduced, which again form a key interface between algebra and topology. One of the centrepieces of this chapter, and indeed the course, is the Nielsen - Schreier theorem, which asserts that any subgroup of a (finitely generated) free group is free. This purely algebraic result has a purely topological proof.

The structure of the course owes a great deal to the book *Classical Topology and Combinatorial Group Theory* by John Stillwell [7]. This is one of the few books on the subject that gives almost equal weight to both the algebra and the topology, and comes highly recommended. Other suggestions for further reading are included at the end of these notes.

Marc Lackenby
October 2016

Convention

When X and Y are topological spaces, we will only consider *continuous* functions between X and Y, unless we explicitly state otherwise.

Notation

- The composition of two maps $f: X \to Y$ and $g: Y \to Z$ is denoted by gf or $g \circ f$.
- When X and Y are spaces, $X \cong Y$ means that X and Y are homeomorphic.
- When G and H are groups, $G \cong H$ means that G and H are isomorphic.
- The identity element of a group is usually denoted e.
- The identity map on a space X is denoted id_X.
- When $f: X \to Y$ is a function and A is a subset of X, then $f|A$ denotes the restriction of f to A.
- A *based space* (X, b) is a space X and a point $b \in X$. A map $f: (X, b) \to (X', b')$ between based spaces is a map $f: X \to X'$ such that $f(b) = b'$.
- The interval $[0, 1]$ is denoted by I, and ∂I is the set $\{0, 1\}$.
- D^n is the set $\{x \in \mathbb{R}^n : |x| \leq 1\}$, and S^{n-1} is $\{x \in \mathbb{R}^n : |x| = 1\}$.
- We also view S^1 as $\{x \in \mathbb{C} : |x| = 1\}$.
- A *path* in a space X is a map $f: I \to X$.
- The space with one point is often denoted $\{*\}$.
- For spaces X and Y, and a point $y \in Y$, we denote the map $X \to Y$ that sends everything to y by c_y.

Chapter I: Constructing spaces

Many topological spaces are, on a local scale, quite simple. For example, a 'surface' is a topological space with the property that each point has a neighbourhood homeomorphic to $\mathbb{R}^2$. However, although these spaces are locally very simple, their global topology can quite complicated.

In this introductory section, we give three different ways of building spaces, using only simple building blocks. One typically starts with a collection of standard pieces, and one glues them together to create the desired space. More formally, one describes an equivalence relation on the pieces, and the new space is the set of equivalence classes, given the quotient topology.

In Section I.1, we will describe graphs, where the building blocks are vertices (which are just points) and edges (which are intervals). We will also introduce a way of building graphs using groups. The resulting graphs are known as Cayley graphs, and they descibe many properties of the group in a topological way.

In Section I.2, we will define simplicial complexes, where the building blocks are 'simplices', which are generalisation of triangles. Simplicial complexes are technically very useful, but are often slightly unwieldy. Therefore, in Section I.3, we will describe a generalisation known as a 'cell complex' which is often a much more efficient way of building a topological space.

I.1: Cayley graphs of groups

Groups are algebraic objects defined by a short list of axioms. They arise naturally when studying the symmetries of objects. This gives the first hint of their strong links with geometry. In this introductory section, we will show how any finitely generated group can be visualised topologically. This is an initial indication of the fruitful link between group theory, geometry and topology.

We start with the definition of a (countable) graph. Intuitively, this is a finite or countable collection of points, known as 'vertices' or 'nodes' joined by a finite or countable collection of arcs, known as 'edges'. An example, which gives the general idea, is shown in Figure I.2. Unlike some authors, we allow an edge to have both its endpoints in the same vertex, and also allow multiple edges between a pair of vertices. The formal definition is as follows.

Definition I.1. A *(countable) graph* Γ is specified by the following data:

- a finite or countable set V, known as its *vertices*;
- a finite or countable set E, known as its *edges*;
- a function δ which sends each edge e to a subset of V with either 1 or 2 elements. The set $\delta(e)$ is known as the *endpoints* of e.

From this, one constructs the associated topological space, also known as the graph Γ, as follows. Start with a disjoint union of points, one for each vertex, and a disjoint union of copies of the interval I, one for each edge. For each $e \in E$, identify 0 in the associated copy of I with one vertex in $\delta(e)$, and identify 1 in the copy of I with the other vertex in $\delta(e)$.

$$V = \{v_1, v_2, v_3\} \qquad E = \{e_1, e_2, e_3\}$$

$$e_1 \overset{\delta}{\mapsto} \{v_1, v_2\} \qquad e_2 \mapsto \{v_2\} \qquad e_3 \mapsto \{v_2, v_3\}$$

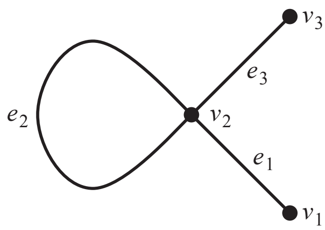

Figure I.2.

Definition I.3. An *orientation* on the graph Γ is a choice of functions $\iota\colon E \to V$ and $\tau\colon E \to V$ such that, for each $e \in E$, $\delta(e) = \{\iota(e), \tau(e)\}$. We say that $\iota(e)$ and $\tau(e)$ are, respectively, the *initial* and *terminal* vertices of the edge e, and we view the edge as running from initial vertex to the terminal vertex.

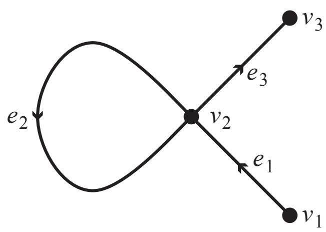

$$V = \{v_1, v_2, v_3\} \qquad E = \{e_1, e_2, e_3\}$$

$E \overset{\iota}{\to} V$	$E \overset{\tau}{\to} V$
$e_1 \mapsto v_1$	$e_1 \mapsto v_2$
$e_2 \mapsto v_2$	$e_2 \mapsto v_2$
$e_3 \mapsto v_2$	$e_3 \mapsto v_3$

Figure I.4.

Definition I.5. Let G be a group and let S be a set of generators for G. The associated *Cayley graph* is an oriented graph with vertex set G and edge set $G \times S$. That is, each edge is associated with a pair (g, s), where $g \in G$ and $s \in S$. The functions ι and τ are specified by:

$$G \times S \xrightarrow{\iota} G \qquad G \times S \xrightarrow{\tau} G$$
$$(g, s) \mapsto g \qquad (g, s) \mapsto gs$$

Thus the edge associated with the pair (g, s) runs from g to gs. We say that this edge is *labelled* by the generator s.

Example I.6. The Cayley graph of $\mathbb{Z}$ with respect to the generator 1 is shown in Figure I.7. The Cayley graph of $\mathbb{Z}$ with respect to the generators $\{2, 3\}$ is shown in Figure I.8.

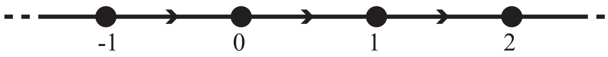

Figure I.7.

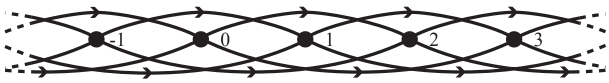

Figure I.8.

Example I.9. The Cayley graph of $\mathbb{Z} \times \mathbb{Z}$ with respect to the generators $x = (1, 0)$ and $y = (0, 1)$ is given in Figure I.10.

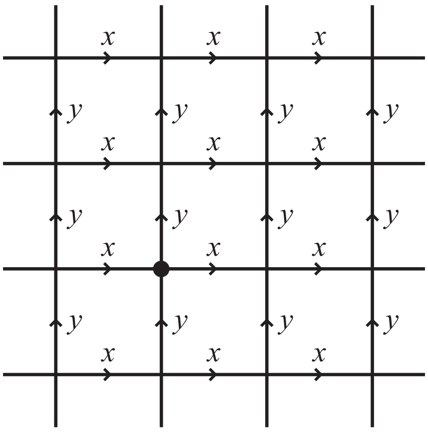

Figure I.10.

Example I.11. The finite permutation group S_3 is generated by $x = (1\ 2\ 3)$ and $y = (1\ 2)$. Its Cayley graph with respect to $\{x, y\}$ is given in Figure I.12.

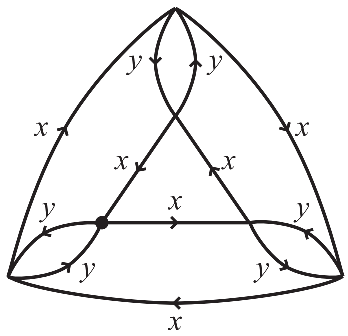

Figure I.12.

Note from Example I.6 that the Cayley graph of a group depends on the choice of generators. However, the graphs in Figures I.7 and I.8 look 'similar'. One can make this precise, and examine geometric and topological properties of a Cayley graph that do not depend on the choice of generators. This is known as geometric group theory, which is an important area of modern mathematical research. Here, we will merely content ourselves with an observation: that one can read off some algebraic properties of the group from topological properties of the Cayley graph, and *vice versa*. For example, suppose that a group element g can be written as a product $s_1^{\epsilon_1} s_2^{\epsilon_2} \dots s_n^{\epsilon_n}$ where each $s_i \in S$ and $\epsilon_i \in \{-1, 1\}$. It is always possible to express g in this way because S is a generating set. This specifies a path, starting at the identity vertex, and running along the edge labelled s_1 (in the forwards direction if $\epsilon_1 = 1$, backwards if $\epsilon_1 = -1$), then along the edge labelled s_2, and so on. Then we end at the vertex labelled g. Thus, we have proved the following.

Proposition I.13. *Any two points in a Cayley graph can be joined by a path.*

Conversely, pick any path from the identity vertex to the g vertex. Then this specifies a way of expressing g as a product of the generators and their inverses.

Thus, there is a correspondence between certain equations in the group and closed loops in the Cayley graph that start and end at the identity vertex. More specifically, the equality $s_1^{\epsilon_1} s_2^{\epsilon_2} \dots s_n^{\epsilon_n} = e$ holds, where each $s_i \in S$ and $\epsilon_i \in \{-1, 1\}$ and e is the identity element of the group, if and only if the corresponding path starting at the identity vertex is in fact a closed loop. Conversely, any such loop gives an equality of the above form.

We will return to groups later in the course, but now we introduce some important topological concepts.

I.2: Simplicial complexes

It turns out that many spaces can be constructed using building blocks that are generalisations of a triangle.

Definition I.14. The *standard n-simplex* is the set

$$\Delta^n = \{(x_0, \dots, x_n) \in \mathbb{R}^{n+1} : x_i \geq 0 \ \forall i \text{ and } \sum_i x_i = 1\}.$$

The non-negative integer n is the *dimension* of this simplex. Its *vertices*, denoted $V(\Delta^n)$, are those points $(x_0, \dots, x_n)$ in Δ^n where $x_i = 1$ for some i (and hence $x_j = 0$ for all $j \neq i$). For each non-empty subset A of $\{0, \dots, n\}$ there is a corresponding *face* of Δ^n, which is

$$\{(x_0, \dots, x_n) \in \Delta^n : x_i = 0 \ \forall i \notin A\}.$$

Note that Δ^n is a face of itself (setting $A = \{0, \dots, n\}$). The *inside* of Δ^n is

$$\text{inside}(\Delta^n) = \{(x_0, \dots, x_n) \in \Delta^n : x_i > 0 \ \forall i\}.$$

Note that the inside of Δ^0 is Δ^0.

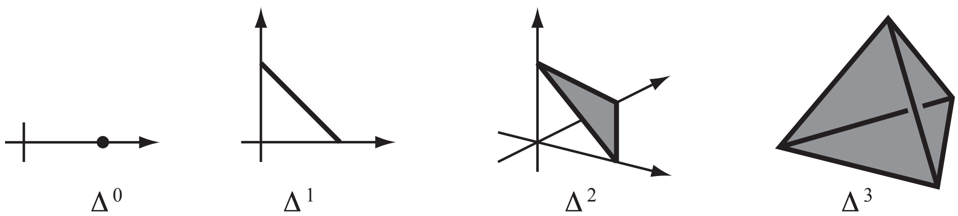

Figure I.15.

Note that $V(\Delta^n)$ is a basis for $\mathbb{R}^{n+1}$. Hence, any function $f \colon V(\Delta^n) \to \mathbb{R}^m$ extends to a unique linear map $\mathbb{R}^{n+1} \to \mathbb{R}^m$. The restriction of this to Δ^n is known as the *affine extension* of f, or is just called *affine*.

Definition I.16. A *face inclusion* of a standard m-simplex into a standard n-simplex (where $m < n$) is the affine extension of an injection $V(\Delta^m) \to V(\Delta^n)$.

For example, there are six face inclusions $\Delta^1 \to \Delta^2$, corresponding to the six injections $\{1, 2\} \to \{1, 2, 3\}$.

We start with the data used to build our spaces:

Definition I.17. An *abstract simplicial complex* is a pair (V, Σ), where V is a set (whose elements are called *vertices*) and Σ is a set of non-empty finite subsets of V (called *simplices*) such that

(i) for each $v \in V$, the 1-element set $\{v\}$ is in Σ;

(ii) if σ is an element of Σ, so is any non-empty subset of σ.

We say that (V, Σ) is *finite* if V is a finite set.

We now give a method of constructing topological spaces:

Definition I.18. The *topological realisation* $|K|$ of an abstract simplicial complex $K = (V, \Sigma)$ is the space obtained by the following procedure:

(i) For each $\sigma \in \Sigma$, take a copy of the standard n-simplex, where $n + 1$ is the number of elements of σ. Denote this simplex by Δ_σ. Label its vertices with the elements of σ.

(ii) Whenever $\sigma \subset \tau \in \Sigma$, identify Δ_σ with a subset of Δ_τ, via the face inclusion which sends the elements of σ to the corresponding elements of τ.

In other words, $|K|$ is a quotient space, obtained by starting with the disjoint union of the simplices in (i), and then imposing the equivalence relation that is described in (ii).

Example I.19. Let $V = \{1, 2, 3\}$ and let Σ be the following collection:

$$\{1\}, \quad \{2\}, \quad \{3\}, \quad \{1,2\}, \quad \{2,3\}, \quad \{3,1\}.$$

Then, to build the topological realisation of (V, Σ), we start with three 0-simplices and three 1-simplices. We identify the 0-simplex $\{1\}$ with a vertex of each of the two 1-simplices $\{1, 2\}$ and $\{1, 3\}$. In this way, we glue the simplices $\{1, 2\}$ and $\{1, 3\}$ together. We perform a similar operation for the other two 0-simplices. We end with the three edges of a triangle (but not its interior, since $\{1, 2, 3\} \notin \Sigma$).

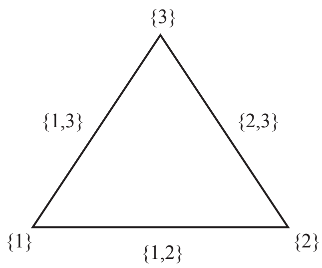

Figure I.20.

Let $K = (V, \Sigma)$ be an abstract simplicial complex. Note that, for each $\sigma \in \Sigma$, there is a homeomorphic copy of Δ_σ inside $|K|$. Note also that $|K|$ is the union of the insides of the simplices, and that the insides of any two simplices are disjoint.

We introduce a quick method for referring to points in $|K|$. Any point $x \in |K|$ lies in the inside of a unique simplex $\sigma = (v_0, \ldots, v_n)$. It is therefore expressed as

$$x = \sum_{i=0}^{n} \lambda_i v_i,$$

for unique positive numbers $\lambda_0, \ldots, \lambda_n$ which sum to one. If $V = \{w_0, \ldots, w_m\}$, we also write $x = \sum \mu_i w_i$, with the understanding that $\mu_i = 0$ if $w_i \notin \{v_0, \ldots, v_n\}$.

If $|K|$ is the topological realisation of an abstract simplicial complex K, we denote the images of the vertices in $|K|$ by $V(|K|)$.

Whenever we refer to a *simplicial complex*, we will mean either an abstract simplicial complex or its topological realisation.

Definition I.21. A *triangulation* of a space X is a simplicial complex K together with a choice of homeomorphism $|K| \to X$.

Example I.22. The torus $S^1 \times S^1$ has a triangulation using nine vertices, as shown in Figure I.23. For reasons of clarity, we have omitted from the middle diagram the edges that are diagonal in the left diagram.

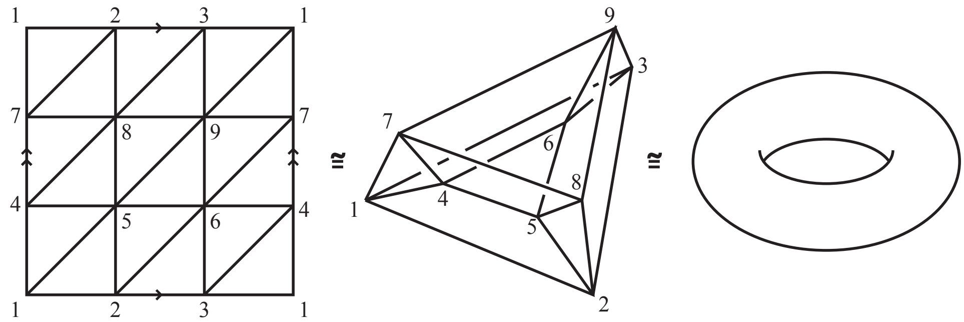

Figure I.23.

Definition I.24. A *subcomplex* of a simplicial complex (V, Σ) is a simplicial complex (V', Σ') such that $V' \subset V$ and $\Sigma' \subset \Sigma$.

Definition I.25. A *simplicial map* between abstract simplicial complexes (V_1, Σ_1) and (V_2, Σ_2) is a function $f\colon V_1 \to V_2$, such that, for all $\sigma_1 \in \Sigma_1$, $f(\sigma_1) = \sigma_2$ for some $\sigma_2 \in \Sigma_2$. It is a *simplicial isomorphism* if it has a simplicial inverse.

Note that a simplicial map need not be injective, and so may decrease the dimension of a simplex.

A simplicial map f between abstract simplicial complexes K_1 and K_2 induces a continuous map $|f|\colon |K_1| \to |K_2|$ as follows. Define $|f|$ on $V(|K_1|)$ according the recipe given by f, and then extend over each simplex using the unique affine extension. Such a map is also known as a *simplicial map*. Note that a simplicial map between finite simplicial complexes is specified by a finite amount of data: one only needs to know the images of the vertices, and then the map is uniquely determined.

Typically, we will not be concerned with an explicit triangulation of a space, and will be quite willing to change the underlying simplicial complex. One way to do this is by 'subdividing' the simplicial complex, which is defined as follows.

Definition I.26. A *subdivision* of a simplicial complex K is a simplicial complex K' together with a homeomorphism $h\colon |K'| \to |K|$ such that, for any simplex σ' of K', $h(\sigma')$ lies entirely in a simplex of $|K|$ and the restriction of h to σ' is affine.

Example I.27. Let K be the triangulation of $I \times I$ shown on the left in Figure I.28. For any positive integer r, let K' be the triangulation of $I \times I$ obtained by dividing $I \times I$ into a lattice of r^2 congruent squares, and then dividing each of these along the diagonal that runs from bottom-right to top-left, as shown in Figure I.28. Then K' is a subdivision of K. We denote it by $(I \times I)_{(r)}$.

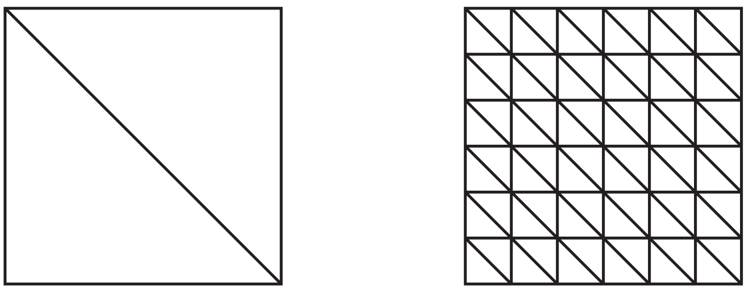

Figure I.28.

I.3: Cell complexes

Simplicial complexes are a very useful technical tool. However, they are somewhat awkward, because even simple spaces such as the torus typically require many simplices in any triangulation. In this section, we introduce a useful generalisation of simplicial complexes, which are a more efficient ways of building topological spaces.

Definition I.29. Let X be a space, and let $f\colon S^{n-1} \to X$ be a map. Then the space *obtained by attaching an n-cell to X along f* is defined to be the quotient of the disjoint union $X \sqcup D^n$, such that, for each point $x \in X$, $f^{-1}(x)$ and x are all identified to a point. It is denoted by $X \cup_f D^n$.

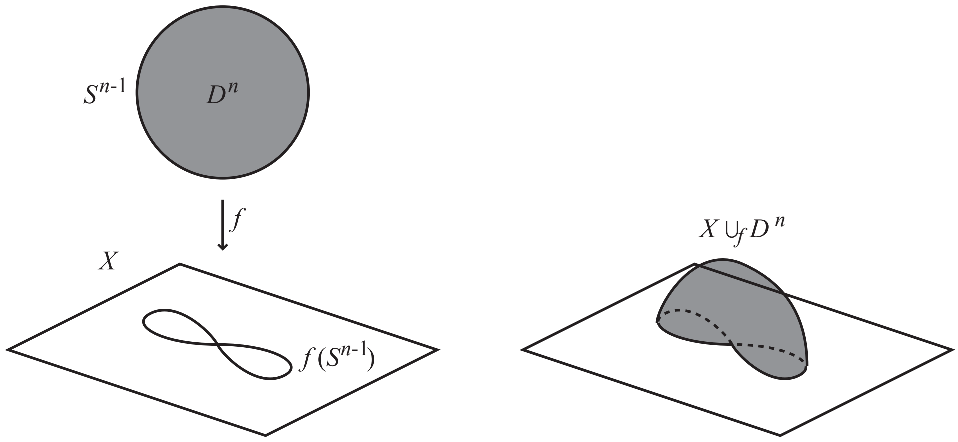

Figure I.30.

Remark I.31. There is a homeomorphic image of X and the interior of D^n in $X \cup_f D^n$. There is also an induced map $D^n \to X \cup_f D^n$, but this need not be injective, since points in the boundary of D^n may be identified.

Definition I.32. A *(finite) cell complex* is a space X decomposed as

$$K^0 \subset K^1 \subset \ldots \subset K^n = X$$

where

(i) K^0 is a finite set of points, and

(ii) K^i is obtained from K^{i-1} by attaching a finite collection of i-cells.

Example I.33. A finite graph is precisely a finite cell complex that consists only of 0-cells and 1-cells.

Remark I.34. Any finite simplicial complex is a finite cell complex, in a natural way, by letting each n-simplex be an n-cell.

Example I.35. The torus $S^1 \times S^1$ has a cell structure, consisting of one 0-cell, two 1-cells, and a single 2-cell. Viewing K^1 as a graph, give its two edges an orientation, and label them a and b. The attaching map $f\colon S^1 \to K^1$ of the 2-cell sends the circle along the path $aba^{-1}b^{-1}$.

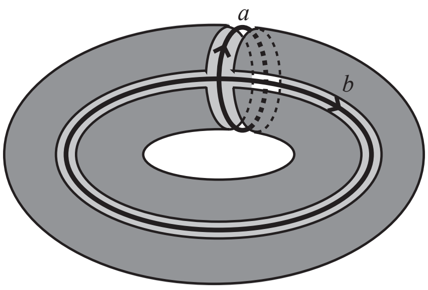

Figure I.36.

Chapter II: Homotopy

II.1: Definition and basic properties

In this section, we introduce a central concept in topology.

Let X and Y be topological spaces.

Definition II.1. A *homotopy* between two maps $f\colon X \to Y$ and $g\colon X \to Y$ is a map $H\colon X \times I \to Y$ such that $H(x,0) = f(x)$ and $H(x,1) = g(x)$ for all $x \in X$. We then say that f and g are *homotopic*, and write $f \simeq g$ or $H\colon f \simeq g$ or $f \overset{H}{\simeq} g$.

One can view I as specifying a 'time' parameter $t \in I$. Thus, as t increases from 0 to 1, the functions $x \mapsto H(x,t)$ specify a 1-parameter family of maps $X \to Y$ which 'interpolate' between f and g. It can be helpful to consider the images of X in Y, as t varies. This resembles a moving picture.

Example II.2. Suppose that Y is a subset of $\mathbb{R}^n$ that is *convex*; this means that for any two points y_1 and y_2 in Y, the line between them, consisting of points $(1-t)y_1 + ty_2$ for $t \in [0,1]$, also lies in Y. Then any two maps $f\colon X \to Y$ and $g\colon X \to Y$ are homotopic, via the homotopy

$$\begin{aligned} H\colon X \times I &\to Y \\ (x,t) &\mapsto (1-t)f(x) + tg(x). \end{aligned}$$

This is known as a *straight-line homotopy*.

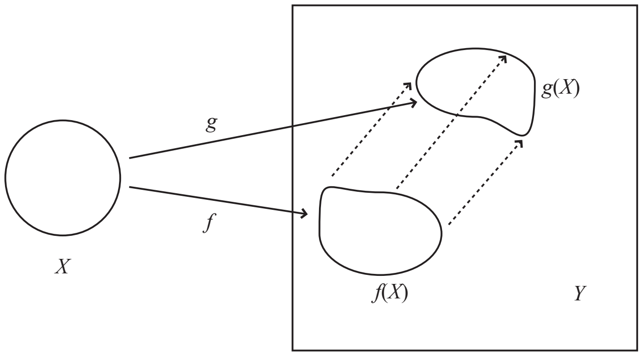

Figure II.3.

Lemma II.4. *For any two spaces X and Y, homotopy is an equivalence relation on the set of continuous maps $X \to Y$.*

Proof. We check the three parts of the definition of an equivalence relation.

Reflexive: For any map $f\colon X \to Y$, $H\colon f \simeq f$, where $H(x,t) = f(x)$ for all t.

Symmetric: If $H\colon f \simeq g$, then $\overline{H}\colon g \simeq f$, where $\overline{H}(x,t) = H(x, 1-t)$.

Transitive: Suppose that $H\colon f \simeq g$ and $K\colon g \simeq h$. Then $L\colon f \simeq h$, where

$$L(x,t) = \begin{cases} H(x,2t) & \text{if } 0 \leq t \leq \frac{1}{2}; \\ K(x,2t-1) & \text{if } \frac{1}{2} \leq t \leq 1. \end{cases}$$

To verify that L is continuous, we need the following lemma. □

Lemma II.5. (Gluing lemma) *If $\{C_1, \dots, C_n\}$ is a finite covering of a space X by closed subsets and $f\colon X \to Y$ is a function, whose restriction to each C_i is continuous, then f is continuous.*

Proof. The map f is continuous if and only if $f^{-1}(C)$ is closed for each closed subset C of Y. But $f^{-1}(C) = \bigcup_{i=1}^{n} f^{-1}(C) \cap C_i$, which is a finite union of closed sets, and hence closed. □

A special case of II.4 is where X is a single point. Then continuous maps $X \to Y$ 'are' points of Y, and homotopies between them are paths. So, the relation of being connected by a path is an equivalence relation on Y. Equivalence classes are known as *path-components* of Y. If Y has a single path-component, it is known as *path-connected*.

Lemma II.6. *Consider the following continuous maps:*

$$W \xrightarrow{f} X \underset{h}{\overset{g}{\rightrightarrows}} Y \xrightarrow{k} Z$$

If $g \simeq h$, then $gf \simeq hf$ and $kg \simeq kh$.

Proof. Let H be the homotopy between g and h. Then $k \circ H: X \times I \to Y \to Z$ is a homotopy between kg and kh. Similarly $H \circ (f \times \mathrm{id}_I): W \times I \to X \times I \to Y$ is a homotopy between gf and hf. □

It is very natural to study homotopies between maps. However, rather less natural is the following notion.

Definition II.7. Two spaces X and Y are *homotopy equivalent*, written $X \simeq Y$, if there are maps

$$X \underset{g}{\overset{f}{\rightleftarrows}} Y$$

such that $gf \simeq \mathrm{id}_X$ and $fg \simeq \mathrm{id}_Y$.

The point here is that we are currently studying maps 'up to homotopy'. So, if gf and fg are both homotopic to the identity, then for the present purposes we view them as 'almost' being the identity. Hence, f and g behave as though they are 'isomorphisms'. So, we should view X and Y as somehow equivalent. All this is rather vague, and is intended to motivate the definition. The fact is, though, it is not in general at all obvious whether two spaces are homotopy equivalent.

Lemma II.8. *Homotopy equivalence is an equivalence relation on spaces.*

Proof. Reflexivity and symmetry are obvious. For transitivity, suppose that we have the following maps

$$X \underset{g}{\overset{f}{\rightleftarrows}} Y \underset{k}{\overset{h}{\rightleftarrows}} Z$$

where fg, gf, hk and kh are all homotopic to the relevant identity map. Then $gkhf \simeq g(\mathrm{id}_Y)f$, by Lemma II.6, which equals gf, which is homotopic to id_X. So, $(gk)(hf) \simeq \mathrm{id}_X$. Similarly, $(hf)(gk) \simeq \mathrm{id}_Z$. □

Definition II.9. A space X is *contractible* if it is homotopy equivalent to the space with one point.

There is a unique map $X \to \{*\}$, and any map $\{*\} \to X$ sends $*$ to some point $x \in X$. Note that $\{*\} \to X \to \{*\}$ is always the identity, and that $X \to \{*\} \to X$ is the constant map c_x. Hence, X is contractible if and only if $\mathrm{id}_X \simeq c_x$, for some $x \in X$.

Example II.10. If X is a convex subspace of $\mathbb{R}^n$, then, for any $x \in X$, $c_x \simeq \mathrm{id}_X$. Hence, X is contractible. In particular, $\mathbb{R}^n$ is contractible, as is D^n.

Definition II.11. When A is a subspace of a space X and $i\colon A \to X$ is the inclusion map, we say that a map $r\colon X \to A$ such that $ri = \mathrm{id}_A$ and $ir \simeq \mathrm{id}_X$ is a *homotopy retract*. Under these circumstances, A and X are homotopy equivalent.

Example II.12. Let $i\colon S^{n-1} \to \mathbb{R}^n - \{0\}$ be the inclusion map, and define

$$\begin{aligned} r\colon \mathbb{R}^n - \{0\} &\to S^{n-1} \\ x &\mapsto x/|x|. \end{aligned}$$

(See Figure II.13.) Then $ri = \mathrm{id}_{S^{n-1}}$ and $H\colon ir \simeq \mathrm{id}_{\mathbb{R}^n - \{0\}}$ via

$$\begin{aligned} H\colon \mathbb{R}^n - \{0\} \times I &\to \mathbb{R}^n - \{0\} \\ (x,t) &\mapsto tx + (1-t)x/|x|. \end{aligned}$$

This is well-defined because the straight line between x and $x/|x|$ does not go through the origin. Hence r is a homotopy retract and $S^{n-1} \simeq \mathbb{R}^n - \{0\}$.

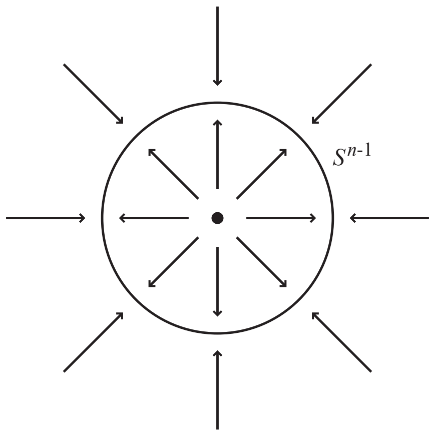

Figure II.13.

Example II.14. Let M denote the Möbius band. Recall that this is the space obtained from $I \times I$ by identifying $(0, y)$ with $(1, 1-y)$ for each $y \in I$. (See Figure II.15.)

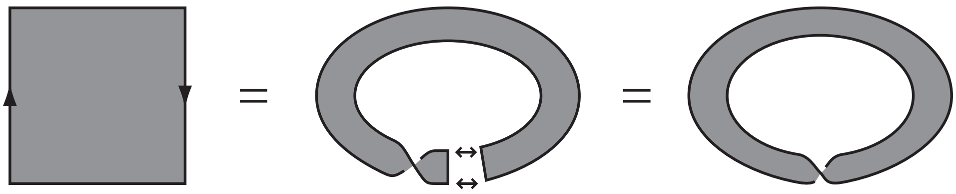

Figure II.15

There is an inclusion map $i\colon S^1 \to M$, sending $e^{2\pi ix}$ to $(x, \frac{1}{2})$. And there is a retraction map $r\colon M \to i(S^1)$, sending $(x, y) \mapsto (x, \frac{1}{2})$. Then r is a homotopy retract, since the composition of r and the inclusion of $i(S^1)$ into M is homotopic to id_M via

$$\begin{aligned} H\colon M \times I &\to M \\ (x, y, t) &\mapsto (x, (1-t)/2 + ty). \end{aligned}$$

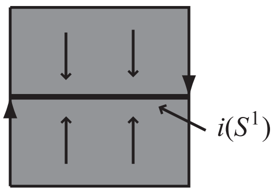

Figure II.16.

Similarly, $S^1 \times \{\frac{1}{2}\}$ is a homotopy retract of $S^1 \times I$. Hence, $M \simeq S^1 \simeq S^1 \times I$.

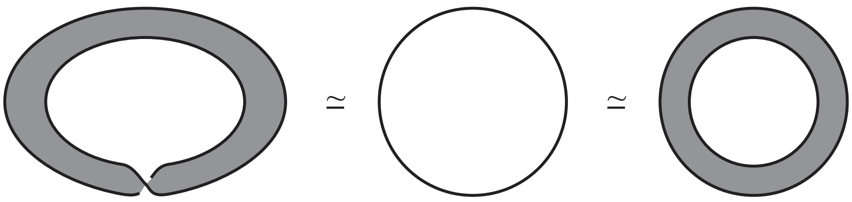

Figure II.17.

The following is an important variant of the notion of a homotopy.

Definition II.18. Let X and Y be spaces and let A be a subspace of X. Then two maps $f, g\colon X \to Y$ are *homotopic relative to* A if $f|A = g|A$ and there is a homotopy $H\colon f \simeq g$ such that $H(x, t) = f(x) = g(x)$ for all $x \in A$ and all $t \in I$.

Remark II.19. There are versions of Lemmas II.4 and II.6 for homotopies relative to subspaces. We will feel free to use these as necessary.

II.2. The Simplicial Approximation Theorem

Almost all of the spaces considered in this course will admit the structure of a simplicial complex. If K and L are simplicial complexes, and $f\colon |K| \to |L|$ is some continuous map, it will be very useful to be able to homotope f to a simplicial map. In general, this is not always possible. But it is if we pass to a 'sufficiently fine' subdivision of K.

Theorem II.20. (Simplicial Approximation Theorem) *Let K and L be simplicial complexes, where K is finite, and let $f\colon |K| \to |L|$ be a continuous map. Then there is some subdivision K' of K and a simplicial map $g\colon K' \to L$ such that $|g|$ is homotopic to f.*

The remainder of this section will be devoted to proving this result. In fact, we prove a slightly weaker version of the theorem, and will indicate how the proof may be extended to give the above result.

We start by describing a canonical way of choosing neighbourhoods in a simplicial complex.

Definition II.21. Let K be a simplicial complex, and let x be a point in $|K|$. The *star* of x in $|K|$ is the following subset of $|K|$:

$$\mathrm{st}_K(x) = \bigcup\{\mathrm{inside}(\sigma) : \sigma \text{ is a simplex of } |K| \text{ and } x \in \sigma\}.$$

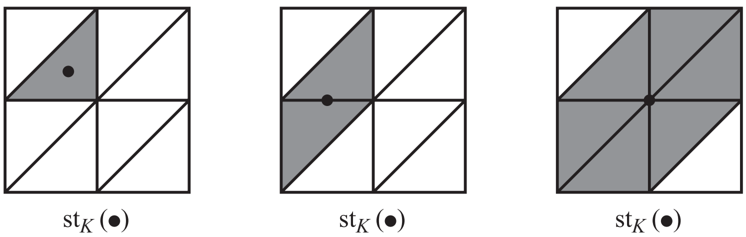

Figure II.22.

Lemma II.23. *For any $x \in |K|$, $\mathrm{st}_K(x)$ is open in $|K|$.*

Proof. Consider

$$|K| - \mathrm{st}_K(x) = \bigcup\{\mathrm{inside}(\sigma) : \sigma \text{ is a simplex of } |K| \text{ and } x \notin \sigma\}$$
$$= \bigcup\{\sigma : \sigma \text{ is a simplex of } |K| \text{ and } x \notin \sigma\}.$$

The latter equality holds because any point in a simplex σ lies in the inside of some face τ of σ, and $x \notin \sigma \Rightarrow x \notin \tau$. Now, $\bigcup\{\sigma : \sigma \text{ is a simplex of } |K| \text{ and } x \notin \sigma\}$ is clearly a subcomplex of K. Hence, it is closed, and so $\mathrm{st}_K(x)$ is open. □

The following gives a method for constructing simplicial maps.

Proposition II.24. *Let K and L be simplicial complexes, and let $f\colon |K| \to |L|$ be a continuous map. Suppose that, for each vertex v of K, there is a vertex $g(v)$ of L such that $f(\mathrm{st}_K(v)) \subset \mathrm{st}_L(g(v))$. Then g is a simplicial map $V(K) \to V(L)$, and $|g| \simeq f$.*

Example II.25. Triangulate the interval I using a simplicial complex K with two vertices v_0 and v_1 and a single 1-simplex between them. Consider the following map of $|K|$ into the topological realisation of a simplicial complex L:

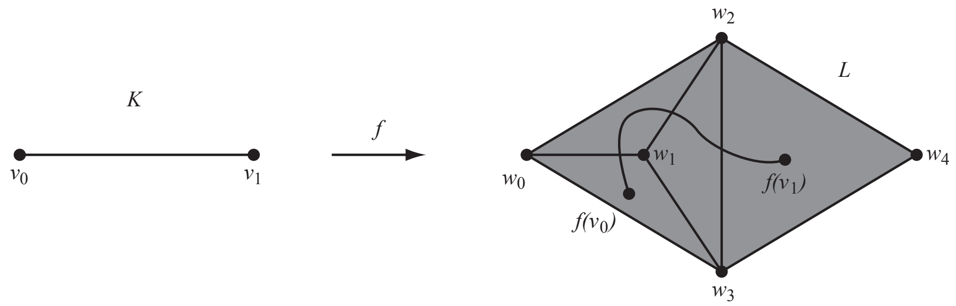

Figure II.26.

Then it is easy to check that $f(\mathrm{st}_K(v_0))$ is not contained within the star of any vertex of L, and so there is no choice of $g(v_0)$ for which the condition $f(\mathrm{st}_K(v_0)) \subset \mathrm{st}_L(g(v_0))$ holds. However, suppose that we subdivide K to the following simplicial complex K', and (by slight abuse of terminology) consider $f\colon |K'| \to |L|$.

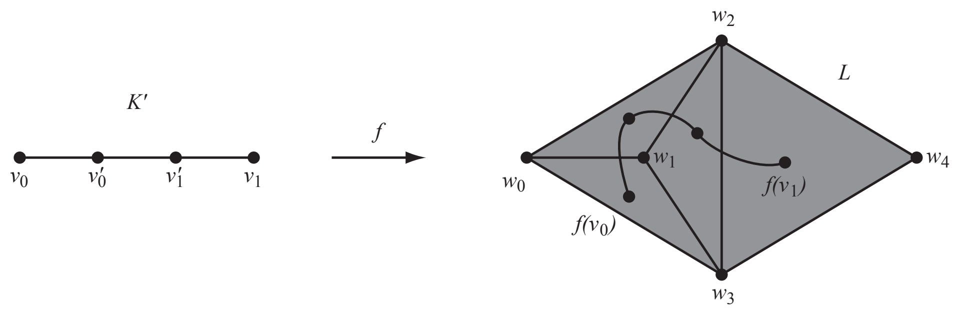

Figure II.27.

Then we have

$$
\begin{aligned}
&f(\mathrm{st}_{K'}(v_0)) \subset \mathrm{st}_L(w_0) \cap \mathrm{st}_L(w_1)\\
&f(\mathrm{st}_{K'}(v_0')) \subset \mathrm{st}_L(w_1)\\
&f(\mathrm{st}_{K'}(v_1')) \subset \mathrm{st}_L(w_2)\\
&f(\mathrm{st}_{K'}(v_1)) \subset \mathrm{st}_L(w_2) \cap \mathrm{st}_L(w_3).
\end{aligned}
$$

So, there are four possible choices for g:

$$\begin{array}{llll}
v_0 \mapsto w_0 & v_0 \mapsto w_1 & v_0 \mapsto w_0 & v_0 \mapsto w_1 \\
v_0' \mapsto w_1 & v_0' \mapsto w_1 & v_0' \mapsto w_1 & v_0' \mapsto w_1 \\
v_1' \mapsto w_2 & v_1' \mapsto w_2 & v_1' \mapsto w_2 & v_1' \mapsto w_2 \\
v_1 \mapsto w_3 & v_1 \mapsto w_3 & v_1 \mapsto w_2 & v_1 \mapsto w_2
\end{array}$$

Proof of II.24. Let $\sigma = (v_0, \dots, v_n)$ be a simplex of K, and let $x \in \mathrm{inside}(\sigma)$. Then $x \in \mathrm{st}_K(v_0) \cap \mathrm{st}_K(v_1) \cap \dots \cap \mathrm{st}_K(v_n)$. So, $f(x) \in f(\mathrm{st}_K(v_0)) \cap \dots \cap f(\mathrm{st}_K(v_n)) \subset \mathrm{st}_L(g(v_0)) \cap \dots \cap \mathrm{st}_L(g(v_n))$. So, if τ is the simplex of L such that $f(x)$ lies in the inside of τ, then all of $g(v_0), \dots, g(v_n)$ must be vertices of τ. So, they span a simplex which is a face of τ and hence a member of L. Thus, g is a simplicial map. Now, x is $\sum_{i=0}^n \lambda_i v_i$, for some $\lambda_i > 0$ with $\sum_{i=0}^n \lambda_i = 1$. So, $|g|(x) = \sum_{i=0}^n \lambda_i g(v_i)$. So, the straight line segment joining $f(x)$ to $|g|(x)$ lies in τ. Thus, there is a well-defined straight-line homotopy between f and $|g|$. □

We will also need the following extension of Proposition II.24, which deals with subcomplexes.

Addendum II.28. *Let K, L, f and g be as in II.24. Let A be any subcomplex of K, and let B be a subcomplex of L such that $f(|A|) \subset |B|$. Then g also maps A into B and the homotopy between $|g|$ and f sends $|A|$ to $|B|$ throughout.*

Proof. Let v be any vertex of A. Let τ be the simplex of L such that $f(v)$ lies in the inside of τ. Then, as argued above, $g(v)$ is a vertex of τ. Since $f(v) \in |B|$, we deduce that τ lies in $|B|$ and hence that $g(v)$ is a vertex of B.

Now consider any point x in $|A|$. Let $(v_0, \dots, v_n)$ be the simplex of K containing x in its inside. Let τ' be the simplex of L such that $f(x)$ lies in the inside of τ'. Then τ' lies in B because $f(x)$ lies in $|B|$. As argued above, all of $g(v_0), \dots, g(v_n)$ must be vertices of τ' and hence vertices of B. The straight-line homotopy between f and $|g|$ sends x into τ' throughout, and hence the image of x remains in $|B|$. □

We may use Proposition II.24 to prove the Simplicial Approximation Theorem, by showing that, given a map $f\colon |K| \to |L|$, one may subdivide K to a simplicial complex K' such that, for all vertices v of K', $f(\mathrm{st}_{K'}(v)) \subset \mathrm{st}_L(g(v))$, for some function $g\colon V(K') \to V(L)$. We shall see that if we take a 'sufficiently fine' subdivision K' of K, then such a function g can always be found. To make this precise, we need to introduce a metric on any finite simplicial complex K.

Definition II.29. The *standard metric* d on a finite simplicial complex $|K|$ is defined to be

$$d\left(\sum_i \lambda_i v_i, \sum_i \lambda'_i v_i\right) = \sum_i |\lambda_i - \lambda'_i|.$$

It is easy to check that this is a metric on $|K|$.

Definition II.30. Let K' be a subdivision of K, and let d be the standard metric on $|K|$. The *coarseness* of the subdivision is

$$\sup\{d(x,y) : x \text{ and } y \text{ belong to the star of the same vertex of } K'\}.$$

Thus, when the coarseness of a subdivision is small, the subdivision is 'fine'.

Example II.31. The subdivision $(I \times I)_{(r)}$ given in Example I.27 has coarseness $4/r$.

We now state a new version of the Simplicial Approximation Theorem, which is in some ways stronger and in some ways weaker than Theorem II.20.

Theorem II.32. (Simplicial Approximation Theorem) *Let K and L be simplicial complexes, where K is finite, and let $f\colon |K| \to |L|$ be a continuous map. Then there is a constant $\delta > 0$ with the following property. If K' is a subdivision of K with coarseness less than δ, then there is a simplicial map $g\colon K' \to L$ such that $|g| \simeq f$.*

We will now start the proof of this result. Recall that the *diameter* of a subset A of a metric space is defined to be

$$\operatorname{diam}(A) = \sup\{d(x,y) : x, y \in A\}.$$

We will need the following elementary fact from basic topology.

Theorem II.33. (Lebesgue Covering Theorem) *Let X be a compact metric space, and let $\mathcal{U}$ be an open covering of X. Then there is a constant $\delta > 0$ such that every subset of X with diameter less than δ is entirely contained within some member of $\mathcal{U}$.*

Proof of II.32. The sets $\{\mathrm{st}_L(w) : w \text{ a vertex of } L\}$ form an open covering of $|L|$, and so the sets $\{f^{-1}(\mathrm{st}_L(w))\}$ form an open covering of $|K|$. Let $\delta > 0$ be the constant from the Lebesgue Covering Theorem for this covering, and let K' be a subdivision of K with coarseness less than δ. Then, for any vertex v of K', $\operatorname{diam}(\mathrm{st}_{K'}(v)) \leq \delta$. So, there is a vertex w of L such that $\mathrm{st}_{K'}(v) \subset f^{-1}(\mathrm{st}_L(w))$. Hence, $f(\mathrm{st}_{K'}(v)) \subset \mathrm{st}_L(w)$. Thus, we set $g(v) = w$ and apply Proposition II.24. □

Later, we will need the following add-on to Theorem II.32.

Addendum II.34. *Let $A_1, \dots, A_n$ be subcomplexes of K, and let $B_1, \dots, B_n$ be subcomplexes of L, such that $f(A_i) \subset B_i$ for each i. Then, for the simplicial map $g: V(K') \to V(L)$ given by Theorem II.32 and for each i, $|g|$ sends A_i to B_i, and the homotopy between f and $|g|$ sends A_i to B_i throughout.*

Proof. We apply Addendum II.28. □

Our original version of the Simplicial Approximation Theorem (II.20) follows immediately from Theorem II.32 and the following result.

Proposition II.35. *A finite simplicial complex K has subdivisions $K^{(r)}$ such that the coarseness of $K^{(r)}$ tends to 0 as $r \to \infty$.*

The basic idea of the proof is simple, but the details are rather technical and not particularly enlightening. Since we will not in fact have much need for Theorem II.20, we only sketch its proof. We start with a recipe for subdividing a single n-simplex.

Let Δ^n be the standard n-simplex. For each face F of Δ^n with vertices $v_1, \dots, v_r$, the *barycentre* of F is $(v_1 + \dots + v_r)/r$. Define a new simplicial complex K' with vertices precisely the barycentres of each of the faces. A set of vertices $w_1, \dots, w_s$ of K', corresponding to faces $F_1, \dots, F_s$ of Δ^n, span a simplex of K' if and only if (after re-ordering $F_1, \dots, F_s$) there are inclusions $F_1 \subset F_2 \subset \dots \subset F_s$. It is possible to show that this is indeed a subdivision of Δ^n, known as the *barycentric subdivision*.

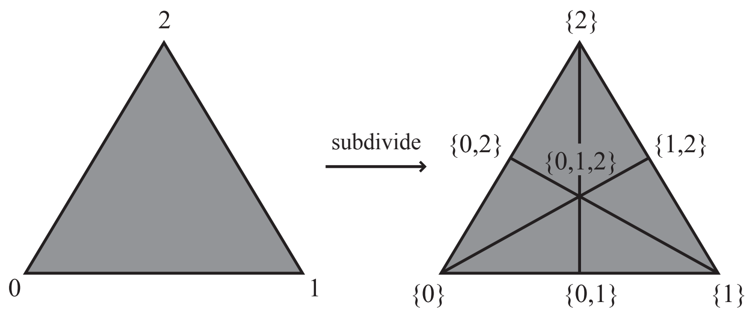

Figure II.36.

If one has a simplicial complex K, it is clear that one can perform this procedure on all simplices of K simultaneously. The result is known as the 'barycentric subdivision' of K. At the level of abstract simplicial complexes, the result is as follows.

Definition II.37. Let $K = (V, \Sigma)$ be an abstract simplicial complex. Then its *barycentric subdivision* $K^{(1)} = (V', \Sigma')$ is an abstract simplicial complex $K^{(1)}$ with vertex set $V' = \Sigma$ and with simplices Σ', specified by the following rule: $(\sigma_0, \ldots, \sigma_n) \in \Sigma'$ if and only if (after possibly re-ordering $\sigma_0, \ldots, \sigma_n$) there are inclusions $\sigma_0 \subset \sigma_1 \subset \ldots \subset \sigma_n$.

We define, for each integer $r \geq 2$, the subdivision $K^{(r)}$ by setting $K^{(r)} = (K^{(r-1)})^{(1)}$. The proof of II.35 proceeds by showing that the coarseness of $K^{(r)}$ tends to 0 as $r \to \infty$. This is intuitively fairly obvious but the details are slightly messy, and are omitted. Theorem II.20 follows immediately.

Chapter III: The fundamental group

III.1: The definition

We now introduce one of the central concepts of the course. It intertwines topology and algebra, by assigning a group to each topological space, known as its fundamental group.

Definition III.1. Let X be a space, and let u and v be paths in X such that $u(1) = v(0)$. Then the *composite path* $u.v$ is given by

$$u.v(t) = \begin{cases} u(2t) & \text{if } 0 \leq t \leq \frac{1}{2}; \\ v(2t-1) & \text{if } \frac{1}{2} \leq t \leq 1. \end{cases}$$

Thus, $u.v$ runs along u, then along v.

Definition III.2. A *loop based at a point* $b \in X$ is a path $\ell\colon I \to X$ such that $\ell(0) = \ell(1) = b$.

For the rest of this section, we will fix a particular point b in X, known as its *basepoint*. We will consider loops based at b, and homotopies *relative to* ∂I between them. In other words, we may vary loops via a homotopy, but only if we keep the endpoints of the loops fixed at b throughout the homotopy.

Definition / Theorem III.3. *The homotopy classes relative to ∂I of loops based at b form a group, called the fundamental group of (X, b), denoted $\pi_1(X, b)$. If ℓ and ℓ' are loops based at b, and $[\ell]$ and $[\ell']$ are their homotopy classes relative to ∂I, their composition $[\ell].[\ell']$ in the group is defined to be $[\ell.\ell']$.*

To prove this, we need to check that composition in $\pi_1(X, b)$ is well-defined and associative, and has an identity and inverses.

Remark III.4. A basepoint is needed to ensure that two loops can be composed.

Remark III.5. If we do not demand that homotopies are relative to ∂I, then it is possible to show that any two paths in the same path-component of X are homotopic, using the fact that I is contractible. Hence, without the restriction that homotopies are relative to ∂I, the fundamental group would be uninteresting.

Well-defined

That composition is well-defined is a consequence of the following lemma. For later applications, we state this lemma not just for loops but also for paths.

Lemma III.6. *Suppose that u and v are paths in X such that $u(1) = v(0)$. Suppose also that u' (respectively, v') is a path with the same endpoints as u (respectively, v). If $u \simeq u'$ relative to ∂I and $v \simeq v'$ relative to ∂I, then $u.v \simeq u'.v'$ relative to ∂I.*

Proof. Let $H\colon u \simeq u'$ and $K\colon v \simeq v'$ be the given homotopies. Define $L\colon I \times I \to X$ by

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

This is continuous by the Gluing Lemma (II.5). So, $L\colon u.v \simeq u'.v'$ relative to ∂I. □

The following diagram explains how L is defined. It is a picture of $I \times I$, and the labelling specifies where each point of $I \times I$ is sent under L.

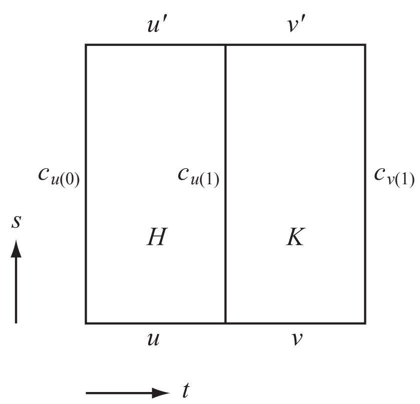

Figure III.7.

Associative

Note that . is not associative. This is because for paths u, v and w such that $u(1) = v(0)$ and $v(1) = w(0)$, $(u.v).w$ and $u.(v.w)$ have the same image but traverse it at different speeds.

Lemma III.8. *Let u, v and w be paths in X such that $u(1) = v(0)$ and $v(1) = w(0)$. Then $u.(v.w) \simeq (u.v).w$ relative to ∂I.*

Proof. An explicit homotopy $H\colon I \times I \to X$ is given by

$$H(t,s) = \begin{cases} u\left(\frac{4t}{2-s}\right) & \text{if } 0 \le t \le \frac{1}{2} - \frac{1}{4}s \\ v\left(4t-2+s\right) & \text{if } \frac{1}{2} - \frac{1}{4}s \le t \le \frac{3}{4} - \frac{1}{4}s \\ w\left(\frac{4t-3+s}{1+s}\right) & \text{if } \frac{3}{4} - \frac{1}{4}s \le t \le 1. \end{cases}$$

This is continuous by the Gluing Lemma (II.5). However, a clearer description of H is as follows. Divide $I \times I$ into the regions shown in Figure III.9. In each region, define H so that each horizontal line traverses u, then v, then w, each at a constant speed. This speed is chosen so that, when one moves from one region of the diagram to another, the path changes from u to v, or from v to w, as appropriate. □

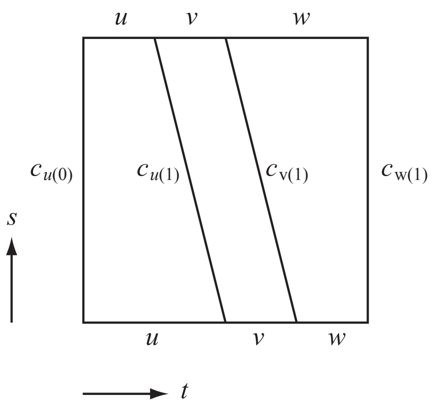

Figure III.9.

IDENTITY

Lemma III.10. *Let u be a path in X with $u(0) = x$ and $u(1) = y$. Then $c_x.u \simeq u$ relative to ∂I, and $u.c_y \simeq u$ relative to ∂I.*

Proof. The required homotopies are given in Figure III.11. □

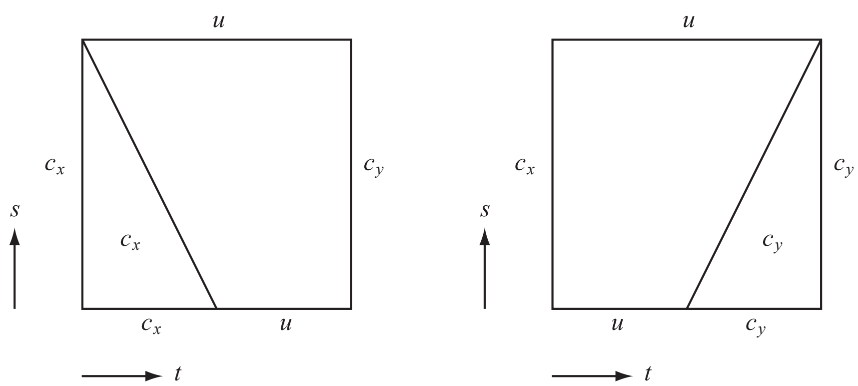

Figure III.11.

Recall that we denote the loop $I \to X$ that sends all of I to b by c_b. Then, by the above lemma, $[c_b]$ is the identity element in $\pi_1(X, b)$.

INVERSES

Lemma III.12. *Let u be a path in X with $u(0) = x$ and $u(1) = y$, and let u^{-1} be the path where $u^{-1}(t) = u(1-t)$. Then $u.u^{-1} \simeq c_x$ relative to ∂I, and $u^{-1}.u \simeq c_y$ relative to ∂I.*

Proof. The required homotopy between $u.u^{-1}$ and c_x is given by

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

The homotopy between $u^{-1}.u$ and c_y is defined similarly. □

This completes the proof of Theorem III.3. □

The terminology u^{-1} in the statement of Lemma III.12 could possibly lead to some confusion, since u^{-1} is not an inverse under composition of maps. Indeed, it is not an inverse for composition of paths, but it does become one when we are are allowed to homotope relative to ∂I.

Examples III.13. Let b be the origin in $\mathbb{R}^2$. Then $\pi_1(\mathbb{R}^2, b)$ is the trivial group. This is because every loop based at b is homotopic relative to ∂I to the constant loop c_b, via the straight-line homotopy.

An example of a space with non-trivial fundamental group is S^1. We let $1 \in S^1$ be the basepoint. (Here, we are viewing S^1 as the unit circle in $\mathbb{C}$, for notational convenience.) In fact, $\pi_1(S^1, 1) \cong \mathbb{Z}$, with a generator represented by the loop ℓ which goes once around the circle:

$$\begin{aligned} \ell\colon I &\to S^1 \\ t &\mapsto e^{2\pi i t}. \end{aligned}$$

However, the proof of this fact is distinctly non-trivial. One must show that any loop based at 1 is homotopic relative to ∂I to ℓ^n, for some $n \in \mathbb{Z}$. And one must show that ℓ^n and ℓ^m are homotopic relative to ∂I only if $n = m$. It will not be until Section III.3 that we have the tools to prove these facts.

We now investigate the dependence of $\pi_1(X, b)$ on the choice of basepoint b. Note that if X_0 is the path-component of X containing the basepoint b, then $\pi_1(X, b) = \pi_1(X_0, b)$, since any loop in X based at b must lie in X_0, and any homotopy between two such loops must also lie in X_0.

Thus, moving the basepoint to another path-component of X may have a significant effect on the fundamental group. But the following proposition implies that moving it within the same path-component has no major effect. In particular, if X is path-connected, then $\pi_1(X, b)$ does not depend (up to isomorphism) on the choice of basepoint, and so we sometimes denote it by $\pi_1(X)$.

Proposition III.14. *If b and b' lie in the same path-component of X, then $\pi_1(X, b) \cong \pi_1(X, b')$.*

Proof. Let w be a path from b to b' in X. If ℓ is a loop based at b, then $w^{-1}.\ell.w$ is a loop based at b', and the function

$$\begin{aligned} w_\sharp \colon \pi_1(X, b) &\to \pi_1(X, b') \\ [\ell] &\mapsto [w^{-1}.\ell.w] \end{aligned}$$

is well-defined by Lemma III.6. (See Figure III.15.) We have

$$\begin{aligned} w_\sharp([\ell]) w_\sharp([\ell']) &= [w^{-1}.\ell.w][w^{-1}.\ell'.w] \\ &= [w^{-1}.\ell.(w.w^{-1}).\ell'.w] \\ &= [w^{-1}.\ell.c_b.\ell'.w] \\ &= [w^{-1}.(\ell.\ell').w] \\ &= w_\sharp([\ell][\ell']) \end{aligned}$$

and so $w_\sharp$ is a homomorphism. Also, $w_\sharp$ has an inverse $(w^{-1})_\sharp$, since

$$(w^{-1})_\sharp(w_\sharp([\ell])) = (w^{-1})_\sharp([w^{-1}.\ell.w]) = [w.w^{-1}.\ell.w.w^{-1}] = [\ell].$$ □

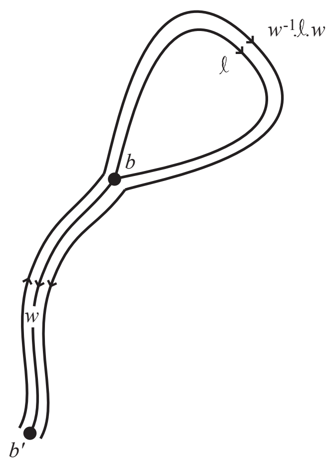

Figure III.15.

Remark III.16. The isomorphism $w_\sharp$ does depend on the choice of w. If u is another path from b to b', then $u_\sharp^{-1} w_\sharp$ is the mapping $[\ell] \mapsto [u.w^{-1}.\ell.w.u^{-1}]$, which is the operation of conjugation by the element $[w.u^{-1}]$ of $\pi_1(X, b)$. Since $\pi_1(X, b)$ need not be abelian, this need not be the identity.

Remark III.17. The fundamental group deals with based loops, but it does yield information about unbased loops $\ell\colon S^1 \to X$. This is because ℓ lies inside a well-defined conjugacy class in $\pi_1(X, b)$, as follows. We may pick an arbitrary path w from b to $\ell(1)$. Then the loop that traverses w, then runs around ℓ and then returns to b along w^{-1} is a loop in X based at b. Applying a homotopy to ℓ does not change the homotopy class relative to ∂I of this loop. Changing the choice to path w would alter this element of $\pi_1(X, b)$ by a conjugacy. Hence, we obtain a well-defined conjugacy class in $\pi_1(X, b)$ from any homotopy class of loop in X.

Proposition III.18. *Let (X, x) and (Y, y) be spaces with basepoints. Then any continuous map $f\colon (X, x) \to (Y, y)$ induces a homomorphism $f_*\colon \pi_1(X, x) \to \pi_1(Y, y)$. Moreover,*

(i) $(\mathrm{id}_X)_* = \mathrm{id}_{\pi_1(X,x)}$

(ii) if $g\colon (Y, y) \to (Z, z)$ is some map, then $(gf)_* = g_* f_*$,

(iii) if $f \simeq f'$ relative to $\{x\}$, then $f_* = f'_*$.

Proof. Define $f_*([\ell]) = [f \circ \ell]$. This is well-defined, by the version of Lemma II.6 for homotopies relative to ∂I. Also, $f \circ (\ell.\ell') = (f \circ \ell).(f \circ \ell')$, and so f_* is a homomorphism. Claims (i) and (ii) are obvious, and claim (iii) follows from the version of Lemma II.6 for homotopies relative to a subspace. □

Proposition III.19. *Let X and Y be path-connected spaces such that $X \simeq Y$. Then $\pi_1(X) \cong \pi_1(Y)$.*

Proof. Let $f\colon X \to Y$ and $g\colon Y \to X$ be homotopy equivalences. If we knew that f and g respected some choice of basepoints and that the homotopies $gf \simeq \mathrm{id}_X$ and $fg \simeq \mathrm{id}_Y$ fixed these basepoints throughout, then it would follow from Proposition III.18 that f_* is an isomorphism with inverse g_*. However, we do not have this information about basepoints, and so another argument is required.

Choose $x_0 \in X$, and let $y_0 = f(x_0)$ and $x_1 = g(y_0)$, so that we have homomorphisms

$$\pi_1(X, x_0) \xrightarrow{f_*} \pi_1(Y, y_0) \xrightarrow{g_*} \pi_1(X, x_1).$$

Let H be the homotopy between gf and id_X. Then $w(t) = H(x_0, t)$ is a path from x_1 to x_0. Let ℓ be a loop in X based at x_0, and consider $K = H \circ (\ell \times \mathrm{id}_I)\colon I \times I \to X$.

Then K is a homotopy fitting into the left-hand square of Figure III.20. Rescale K to fit into the middle trapezium of the right-hand square. Fill in the triangles with maps that are constant on the first variable. This yields a homotopy relative to ∂I between $w^{-1}.(g \circ f \circ \ell).w$ and ℓ. In other words, $w_\sharp g_* f_*[\ell] = [\ell]$. So, $w_\sharp g_* f_* = \mathrm{id}_{\pi_1(X,x_0)}$. In particular, f_* is injective, and, since $w_\sharp$ is an isomorphism, g_* is surjective. By composing f and g the other way round, we deduce that g_* is injective. So, g_* is an isomorphism. □

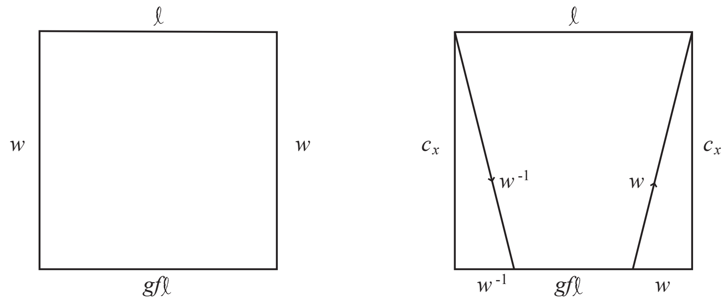

Figure III.20.

Note that, as a consequence, if X is a contractible space, then $\pi_1(X)$ is the trivial group. We give this property a name.

Definition III.21. A space is *simply-connected* if it is path-connected and has trivial fundamental group.

However, it is not the case that simply-connected spaces are necessarily contracible. A counter-example is the 2-sphere, although the proof of this fact is beyond this course.

III.2: A simplicial version

Definition III.22. Let K be a simplicial complex. An *edge path* is a finite sequence $(a_0, \ldots, a_n)$ of vertices of K, such that for each i, $\{a_{i-1}, a_i\}$ spans a simplex of K. (Note: we allow $a_{i-1} = a_i$.) Its *length* is n. An *edge loop* is an edge path with $a_n = a_0$. If $\alpha = (a_0, \ldots, a_n)$ and $\beta = (b_0, \ldots, b_m)$ are edge paths such that $a_n = b_0$, we define $\alpha.\beta$ to be $(a_0, \ldots, a_n, b_1, \ldots, b_m)$.

Definition III.23. Let α be an edge path. An *elementary contraction* of α is an edge path obtained from α by performing one of the following moves:

(0) replacing $\ldots, a_{i-1}, a_i, \ldots$ by $\ldots, a_{i-1}, \ldots$ provided $a_{i-1} = a_i$;

(1) replacing $\ldots, a_{i-1}, a_i, a_{i+1}, \ldots$ by $\ldots, a_{i-1}, \ldots$, provided $a_{i-1} = a_{i+1}$;

(2) replacing $\ldots, a_{i-1}, a_i, a_{i+1}, \ldots$ by $\ldots, a_{i-1}, a_{i+1}, \ldots$, provided $\{a_{i-1}, a_i, a_{i+1}\}$ span a 2-simplex of K.

(See Figure III.24.) We say that α is an *elementary expansion* of β if β is an elementary contraction of α. We say that α and β are *equivalent*, written $\alpha \sim \beta$, if we can pass from α to β by a finite sequence of elementary contractions and expansions. This is clearly an equivalence relation on edge paths.

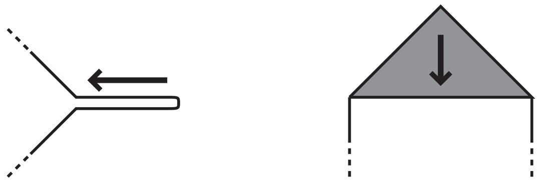

Figure III.24.

Remark III.25. Note that if two edge paths are equivalent, then they have the same initial vertices and the same terminal vertices.

Definition / Theorem III.26. *Let K be a simplicial complex, and let b be a vertex of K. The equivalence classes of edge loops in K based at b form a group denoted $E(K, b)$, called the edge loop group.*

Proof. The product is induced by the product of edge loops. It is easy to check that this respects the equivalence relation $\sim$. It is associative because the product of edge loops is associative. The identity is the equivalence class of (b). The inverse of $(b, b_1, \ldots, b_{n-1}, b)$ is $(b, b_{n-1}, \ldots, b_1, b)$. □

Theorem III.27. *For a simplicial complex K and vertex b, $E(K, b)$ is isomorphic to $\pi_1(|K|, b)$.*

Proof. Let $I_{(n)}$ be the triangulation of $I = [0, 1]$ with n 1-simplices, each of length $\frac{1}{n}$. Then we can regard an edge path of length n as a simplicial map $I_{(n)} \to K$. This defines a mapping

$$\{\text{edge loops in } K \text{ based at } b\} \xrightarrow{\theta} \{\text{loops in } |K| \text{ based at } b\}.$$

If α is obtained from β by an elementary contraction, it is clear that $\theta(\alpha)$ and $\theta(\beta)$ are homotopic relative to ∂I. So θ gives a well-defined mapping $E(K, b) \to \pi_1(|K|, b)$. We must show that it is an isomorphism.

Homomorphism. For edge loops α and β, $\theta(\alpha.\beta) \simeq \theta(\alpha).\theta(\beta)$ relative to ∂I, and so θ is a homomorphism.

Surjective. Let $\ell: I \to |K|$ be any loop in $|K|$ based at b. Give I the triangulation $I_{(1)}$ and view $I_{(n)}$ as a subdivision. Then the coarseness of $I_{(n)}$ is $4/n$, which tends to 0 as $n \to \infty$. So, by the Simplicial Approximation Theorem (II.32) and II.34, we can find a simplicial map $\alpha: I_{(n)} \to K$ for some n, such that $\ell \simeq \theta(\alpha)$ relative to ∂I. So $\theta([\alpha]) = [\ell]$.

Injective. Let $\alpha = (b_0, b_1, \dots, b_{n-1}, b_n)$ be an edge loop based at b. Suppose that $\theta([\alpha])$ is the identity in $\pi_1(|K|, b)$. Then $\theta(\alpha) \simeq c_b$ relative to ∂I, via a homotopy $H: I \times I \to |K|$. Triangulate $I \times I$ using the triangulation $(I \times I)_{(r)}$ of Example I.27. The Simplicial Approximation Theorem gives, for sufficiently large r, a simplicial map $G: (I \times I)_{(r)} \to K$ with $G \simeq H$. Moreover, we can ensure that G sends $\partial I \times I$ and $I \times \{1\}$ to b, by II.34. Also using II.34, we can ensure, when r is a multiple of n, that $G(i/n, 0) = b_i$, and that G sends the 1-simplices between $(i/n, 0)$ and $((i+1)/n, 0)$ to the 1-simplex (b_i, b_{i+1}). So, the restriction of G to $I \times \{0\}$ is an edge path which contracts to α. And, by applying the sequence of moves shown in Figure III.28, we obtain a sequence of elementary contractions and expansions taking this edge path to an edge path, every vertex of which is b. This is equivalent to (b). Hence, $[\alpha]$ is the identity element of $E(K, b)$. □

Theorem III.27 is significant for two reasons. Firstly, $E(K, b)$ is a concrete object that is often computable, and so this gives a method for computing $\pi_1(|K|, b)$. Secondly, it shows that $E(K, b)$ does not depend (up to isomorphism) on the choice of triangulation for $|K|$.

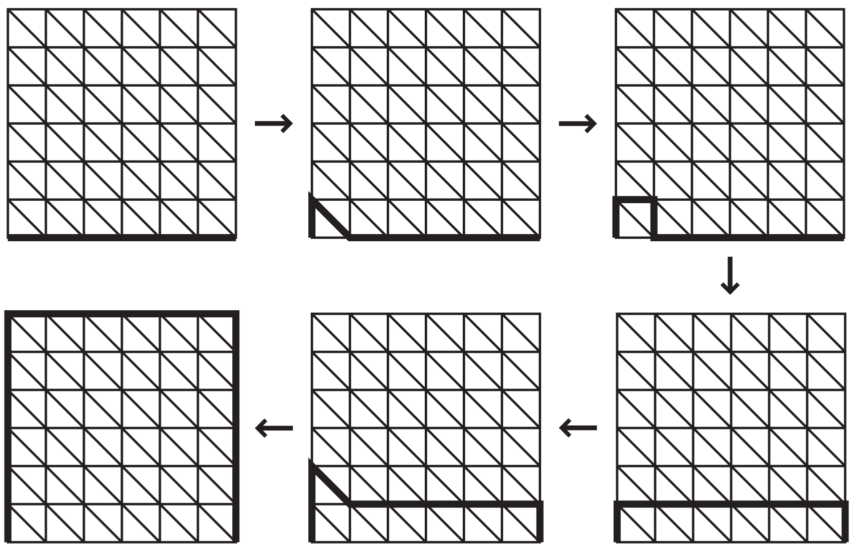

Figure III.28.

Definition III.29. For a simplicial complex K and any non-negative integer n, the *n-skeleton* of K, denoted $\text{skel}^n(K)$, is the subcomplex consisting of the simplices with dimension at most n.

Corollary III.30. *For any simplicial complex K and vertex b, $\pi_1(|K|, b)$ is isomorphic to $\pi_1(|\text{skel}^2(K)|, b)$.*

Proof. The definition of $E(K, b)$ involves only the simplices of dimension at most two. $\square$

Corollary III.31. *For $n \geq 2$, $\pi_1(S^n)$ is trivial.*

Proof. Impose a triangulation on S^n, coming from the n-skeleton of Δ^{n+1}. Then, S^n and Δ^{n+1} have the same 2-skeleton. But Δ^{n+1} is contractible, and so has trivial fundamental group. Hence, so does S^n. $\square$

However, S^1 does not have trivial fundamental group, as we shall prove in the next section.

III.3. The fundamental group of the circle

Theorem III.32. *The fundamental group of the circle is isomorphic to the additive group of integers: $\pi_1(S^1) \cong \mathbb{Z}$.*

As in Examples III.13, we set $1 \in S^1$ to be the basepoint, and let ℓ be the loop which goes once around the circle: $\ell(t) = e^{2\pi i t}$. We will show that a generator for $\pi_1(S^1, 1)$ is represented by ℓ.

Proof of Theorem III.32. Impose a triangulation K on S^1 using three vertices and three 1-simplices. We will show that $E(K, 1)$ is isomorphic to $\mathbb{Z}$. Define an orientation on each 1-simplex, as shown in Figure III.33.

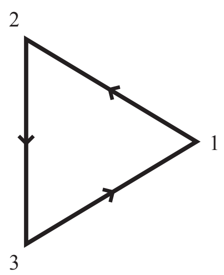

Figure III.33.

Consider a simplicial loop $\alpha = (b_0, b_1, \ldots, b_{n-1}, b_n)$ based at 1. If $b_i = b_{i+1}$ for some i, then we may perform an elementary contraction. Similarly, if the loop traverses some 1-simplex, and then the same 1-simplex in reverse, we may perform an elementary contraction. Hence, a shortest loop equivalent to α traverses all the simplices with same orientation. It is therefore ℓ^n for some $n \in \mathbb{Z}$.

We must show that if ℓ^n and ℓ^m are equivalent, then $n = m$. Define the *winding number* of a simplicial path α to be the number of times it traverses $(1, 2)$ in the forwards direction minus the number of times it traverses it in the backwards direction. Then the winding number of ℓ^n is n. It is clear that an elementary contraction or expansion leaves the winding number unchanged.

Hence, any edge loop is equivalent to ℓ^n, for a unique integer n. This sets up a bijection $E(K, 1) \to \mathbb{Z}$. This is an isomorphism, since $\ell^n.\ell^m = \ell^{n+m}$. □

III.4. The Fundamental Theorem of Algebra

A beautiful application of the above result is the following key theorem.

Theorem III.34. (Fundamental Theorem of Algebra). *Any non-constant polynomial with complex coefficients has at least one root in* $\mathbb{C}$.

Proof. Let $p(x) = a_n x^n + \ldots + a_0$ be the polynomial, where $a_n \neq 0$ and $n > 0$. Let $C_r = \{x \in \mathbb{C} : |x| = r\}$, where r will be a large real number, yet to be chosen. Let $k = p(r)/r^n$ and let $q(x) = kx^n$. Then $p(r) = q(r)$.

Claim. If r is sufficiently large, then $p|C_r$ and $q|C_r$ and the straight line homotopy between them all miss 0.

If not, then for some $x \in C_r$ and some $t \in [0, 1]$,

$$(1 - t)p(x) + tq(x) = 0.$$

This is equivalent to

$$(1 - t)(a_n x^n + \ldots + a_0) + t\left(\frac{a_n|x|^n + \ldots + a_0}{|x|^n}\right)x^n = 0$$

which is equivalent to

$$a_n x^n + \ldots + a_0 = t\left(a_{n-1}x^{n-1} + \ldots + a_0 - a_{n-1}\frac{x^n}{|x|} - \ldots - a_0\frac{x^n}{|x|^n}\right).$$

The left-hand side has order x^n, whereas the right-hand side has order at most x^{n-1}. Hence, $|t| \to \infty$ as $|x| \to \infty$. In particular, there is no solution in the range $t \in [0,1]$ when r is sufficiently large. This proves the claim.

So, $p|C_r : C_r \to \mathbb{C} - \{0\}$ and $q|C_r : C_r \to \mathbb{C} - \{0\}$ are homotopic relative to $\{r\}$. Suppose that $p(x)$ has no root in $\mathbb{C}$. Then p is a function $\mathbb{C} \to \mathbb{C} - \{0\}$. Hence, the following diagram commutes:

$$\begin{array}{ccc} \mathbb{C} & \xrightarrow{p} & \mathbb{C} - \{0\} \\ \uparrow i & \nearrow p|C_r & \\ C_r & & \end{array}$$

where i is the inclusion map. So, at the level of fundamental groups, the following commutes:

$$\begin{array}{ccc} 0 = \pi_1(\mathbb{C}, r) & \xrightarrow{p_*} & \pi_1(\mathbb{C} - \{0\}, r) \cong \mathbb{Z} \\ \uparrow i_* & \nearrow (p|C_r)_* & \\ \pi_1(C_r, r) \cong \mathbb{Z} & & \end{array}$$

So, $(p|C_r)_*$ is the zero homomorphism. But $(p|C_r)_* = (q|C_r)_*$ which sends a generator of $\pi_1(C_r)$ to n times a generator of $\pi_1(\mathbb{C} - \{0\})$, which is a contradiction. □

Chapter IV: Free Groups

IV.1: Definitions

In this chapter, we will embark on the part of the course that deals with groups. Our focus will initially be a special type of group, known as a free group.

For any set S, we will define a group $F(S)$, known as the *free group on* S. Informally, one should view S as an 'alphabet', and elements of $F(S)$ as 'words' in this alphabet. So, for example, if $S = \{a, b\}$, then ab and ba are elements of $F(S)$. The group operation is 'concatenation'; in other words, to compose two words in $F(S)$, we simply write one down and then follow it by the other. For example, the product of ab and ba is $abba$.

The above discussion is slightly oversimplified, because it does not take account of the fact that groups have inverses. So, whenever a is an element of S, we must allow not only a but a^{-1} to appear in the words. But then, $aa^{-1}b$ and b should represent the same element of the group. So, in fact, elements of $F(S)$ are not words in the alphabet S, but are equivalence classes of words.

We are now ready to give the formal definitions. Throughout, S is some set, known as the *alphabet*. From this set, create a new set S^{-1}. This is a copy of the set S, but for each element x of S, we denote the corresponding element of S^{-1} by x^{-1}. We insist that $S \cap S^{-1} = \emptyset$. When $x^{-1} \in S^{-1}$, we say that $(x^{-1})^{-1} = x$.

Definition IV.1. A *word* w is a finite sequence $x_1, \dots, x_m$, where $m \in \mathbb{Z}_{\geq 0}$ and each $x_i \in S \cup S^{-1}$. We write w as $x_1 x_2 \dots x_m$. Note that the empty sequence, where $m = 0$, is allowed as a word. We denote it by $\emptyset$.

Definition IV.2. The *concatenation* of two words $x_1 x_2 \dots x_m$ and $y_1 y_2 \dots y_n$ is the word $x_1 x_2 \dots x_m y_1 y_2 \dots y_n$.

Definition IV.3. A word w' is an *elementary contraction* of a word w, written $w \searrow w'$, if $w = y_1 x x^{-1} y_2$ and $w' = y_1 y_2$, for words y_1 and y_2, and some $x \in S \cup S^{-1}$. We also write $w' \nearrow w$, and say that w is an *elementary expansion* of w'.

Definition IV.4. Two words w' and w are *equivalent*, written $w \sim w'$, if there are words $w_1, \dots, w_n$, where $w = w_1$ and $w' = w_n$, and for each i, $w_i \nearrow w_{i+1}$ or $w_i \searrow w_{i+1}$. The equivalence class of a word w is denoted $[w]$.

Definition IV.5. The *free group on the set* S, denoted $F(S)$, consists of equivalence classes of words in the alphabet S. The composition of two elements $[w]$ and $[w']$ is the class $[ww']$. The identity element is $[\emptyset]$, and is denoted e. The inverse of an element $[x_1 x_2 \dots x_n]$ is $[x_n^{-1} \dots x_2^{-1} x_1^{-1}]$.

One should check that composition is well-defined: if $w_1 \sim w_1'$ and $w_2 \sim w_2'$, then $w_1 w_2 \sim w_1' w_2'$. But this is obvious from the definitions.

Definition IV.6. If a group G is isomorphic to $F(S)$, for some set S, then the copy of S in G is known as a *free generating set*.

IV.2. Reduced representatives

Definition IV.7. A word is *reduced* if it does not admit an elementary contraction.

Proposition IV.8. *Any element of the free group $F(S)$ is represented by a unique reduced word.*

Lemma IV.9. *Let w_1, w_2 and w_3 be words, such that $w_1 \nearrow w_2 \searrow w_3$. Then either there is a word w_2' such that $w_1 \searrow w_2' \nearrow w_3$ or $w_1 = w_3$.*

Proof. Since $w_1 \nearrow w_2$, we can write $w_1 = ab$, and $w_2 = axx^{-1}b$, for some $x \in S \cup S^{-1}$ and some words a and b. As $w_2 \searrow w_3$, w_3 is obtained from w_2 by removing yy^{-1}, for some $y \in S \cup S^{-1}$. The letters xx^{-1} and yy^{-1} intersect in either zero, one or two letters. We will consider these three possibilities in turn. If they do not intersect, then it is possible to remove yy^{-1} from w_1 before inserting xx^{-1}. Hence, if we denote by w_2' the word obtained by removing yy^{-1} from w_1, then $w_1 \searrow w_2' \nearrow w_3$, as required. Suppose now that xx^{-1} and yy^{-1} intersect in a single letter. Then $x = y^{-1}$, and so in w_2, there is chain of letters $xx^{-1}x$ or $x^{-1}xx^{-1}$, and w_1 and w_3 are obtained from w_2 by the two possible ways of performing an elementary contraction on these three letters. In particular, $w_1 = w_3$, as required. Finally, if xx^{-1} and yy^{-1} intersect in two letters, then clearly, all we have done in the sequence $w_1 \nearrow w_2 \searrow w_3$ is to insert a pair of letters and then remove it again, and so $w_1 = w_3$. □

Proof of Proposition IV.8. An elementary contraction to a word reduces its length by two. Hence, a shortest representative for an element of $F(S)$ must be reduced. We must show that there is a unique such representative. Suppose that, on the contrary, there are distinct reduced words w and w' that are equivalent. Then, by definition, there is a sequence of words $w_1, w_2, \dots, w_n$ such that $w = w_1$ and $w' = w_n$ and, for each i, $w_i \nearrow w_{i+1}$ or $w_i \searrow w_{i+1}$. Consider a shortest such sequence. This implies that $w_i \neq w_j$ for any $i \neq j$. This is because if w_i did equal w_j, then we could miss out all the words in the sequence between them, creating a shorter sequence of words joining w to w'. Suppose that, at some point, $w_i \nearrow w_{i+1} \searrow w_{i+2}$. Then, by Lemma IV.9, there is another word w'_{i+1}, such that $w_i \searrow w'_{i+1} \nearrow w_{i+2}$. In this way, we may perform all $\searrow$ moves before all $\nearrow$ ones. Hence, the sequence starts with $w_1 \searrow w_2$ or ends with $w_{n-1} \nearrow w_n$. But, this implies that either w or w' was not reduced, which is a contradiction. □

IV.3. The universal property

Given a set S, there is a function $i: S \to F(S)$, known as the *canonical inclusion*, sending each element of S to the corresponding generator of $F(S)$. The following is known as the 'universal property' of free groups.

Theorem IV.10. *Given any set S, any group G and any function $f: S \to G$, there is a unique homomorphism $\phi: F(S) \to G$ such that the following diagram commutes:*

$$\begin{array}{ccc} S & \xrightarrow{f} & G \\ \downarrow i & \nearrow \phi & \\ F(S) & & \end{array}$$

where $i: S \to F(S)$ is the canonical inclusion.

Proof. We first show the existence of ϕ. Consider any word $w = x_1^{\epsilon_1} \dots x_n^{\epsilon_n}$, where each $x_i \in S$ and each $\epsilon_i \in \{-1, 1\}$. Define $\phi(w)$ to be $f(x_1)^{\epsilon_1} \dots f(x_n)^{\epsilon_n}$. We need to show that this descends to a well-defined function $\phi\colon F(S) \to G$; in other words, that if two words w and w' are equivalent, then they have the same image under ϕ. It suffices to show this when w' is an elementary contraction of w, where $w = w_1 x x^{-1} w_2$ or $w = w_1 x^{-1} x w_2$, and $w' = w_1 w_2$. Then, in the case where $w = w_1 x x^{-1} w_2$

$$\phi(w) = \phi(w_1) f(x) f(x)^{-1} \phi(w_2) = \phi(w_1)\phi(w_2) = \phi(w').$$

Similarly, when $w = w_1 x^{-1} x w_2$, $\phi(w) = \phi(w')$. It is clear that ϕ is a homomorphism. Finally, ϕ is the unique such homomorphism for which the diagram commutes. This is because for each $x \in S$, $\phi(x) = f(x)$, and a homomorphism between groups is determined by what it does to a set of generators. □

We say that $f\colon S \to G$ *induces* the homomorphism $\phi\colon F(S) \to G$.

IV.4: The fundamental group of a graph

Recall the following definition from Section I.1:

Definition I.1. A *graph* Γ is specified by the following data:

- a finite or countable set V, known as its *vertices*;
- a finite or countable set E, known as its *edges*;
- a function δ which sends an edge e to a subset of V with either 1 or 2 elements. The set $\delta(e)$ is known as the *endpoints* of e.

From this, one constructs the associated topological space, also known as the graph Γ, as follows. Start with a disjoint union of points, one for each vertex, and a disjoint union of copies of I, one for each edge. For each $e \in E$, identify 0 in the associated copy of I with one vertex in $\delta(e)$, and identify 1 in the copy of I with the other vertex in $\delta(e)$.

Our goal in this section is to prove the following key result.

Theorem IV.11. *The fundamental group of a connected graph is a free group.*

We need some terminology and basic results from graph theory.

Definition IV.12. Let Γ be a graph with vertex set V, edge set E, and endpoint function δ. A *subgraph* of Γ is a graph with vertex set $V' \subset V$ and edge set $E' \subset E$ and with endpoint function being the restriction of δ. For this to be defined, it is necessary that, for each $e \in E'$, $\delta(e) \subset V'$. Clearly, if Γ is oriented, the subgraph inherits an orientation.

Definition IV.13. An *edge path* in a graph Γ is a concatenation $u_1 \dots u_n$, where each u_i is either a path running along a single edge at unit speed or a constant path based at a vertex. An *edge loop* is an edge path $u: I \to \Gamma$, where $u(0) = u(1)$. An edge path (respectively, edge loop) $u: I \to \Gamma$ is *embedded* if u is injective (respectively, the only two points in I with the same image under u are 0 and 1).

Definition IV.14. A *tree* is a connected graph that contains no embedded edge loops.

Lemma IV.15. *In a tree, there is a unique embedded edge path between distinct vertices.*

Proof. Any two distinct vertices are certainly connected by an edge path, since a tree is assumed to be connected. A shortest such edge path is embedded. We need to show that that it is unique. Suppose that, on the contrary, there are two distinct embedded edge paths $p = u_1 \dots u_n$ and $p' = u'_1 \dots u'_{n'}$ between a distinct pair of vertices. Let $u_i(0)$ be the point on p where the paths first diverge. Let $u_j(1)$ be the next point on p which lies in the image of p'. Then the concatenation of $u_i \dots u_j$ with the sub-arc of p' between $u_j(1)$ and $u_i(0)$ is an embedded edge loop. This contradicts the hypothesis that this is a tree. □

Definition IV.16. A *maximal tree* in a connected graph Γ is a subgraph T that is a tree, but where the addition of any edge of $E(\Gamma) - E(T)$ to T gives a graph that is not a tree.

Lemma IV.17. *Let Γ be a connected graph and let T be a subgraph that is a tree. Then the following are equivalent:*

(i) $V(T) = V(\Gamma)$;

(ii) T *is maximal.*

Proof. (i) $\Rightarrow$ (ii): Let e be an edge of $E(\Gamma) - E(T)$. If the endpoints of e are the same vertex, then adding e to T certainly results in a subgraph that is not a tree, since it contains an embedded edge loop. So, we may assume that the endpoints of e are distinct. They lie in T, since $V(T) = V(\Gamma)$. They are connected by an embedded edge path p in T, by Lemma IV.15. Then $p \cup e$ is an embedded loop in $T \cup e$, and so this is not a tree.

(ii) $\Rightarrow$ (i): Suppose that T is a maximal tree. Suppose that there is a vertex v of Γ that is not in $V(T)$. Pick a shortest edge path from T to v, which exists because Γ is connected. The first edge of this path starts in $V(T)$ but cannot end in $V(T)$. We may therefore add it to T and create a larger tree, which contradicts maximality. □

Lemma IV.18. *Any connected graph* Γ *contains a maximal tree.*

Proof. Recall from the definition of a graph that we are assuming that $V(\Gamma)$ is finite or countable. We may therefore pick a total ordering on $V(\Gamma)$. We may assume that, for each integer $i \geq 2$, the i^{th} vertex shares an edge with one of the earlier vertices. We will construct a nested sequence of subgraphs $T_1 \subset T_2 \subset \dots$, each of which is a tree, and where $V(T_i)$ is the first i vertices of the ordering. Set T_1 to be the first vertex. By assumption, there is an edge e joining the i^{th} vertex to one of the previous vertices. Then set T_i to be $T_{i-1} \cup e$. We have not created any embedded edge loops, and so, inductively, each T_i is a tree. We claim that $T = \bigcup_i T_i$ is a tree. Suppose that it contained an embedded edge loop ℓ. Then, since ℓ consists of only finitely many edges, these edges must all appear in some T_i. But T_i would then not be a tree, which is a contradiction. Since T contains all the vertices of Γ, it is maximal, by Lemma IV.17. □

Proof of Theorem IV.11. Let T be a maximal tree in Γ. Let b be a vertex of Γ, which we take as the basepoint. For any vertex v of Γ, let $\theta(v)$ be the unique embedded edge path from b to v in T. This exists because $V(T) = V(\Gamma)$, by Lemma IV.17. Set $E(\Gamma)$ and $E(T)$ to be the edges of Γ and T respectively. Assign an orientation to each edge e of $E(\Gamma) - E(T)$, and let $\iota(e)$ and $\tau(e)$ be its initial and terminal vertices. We claim that the elements $\{\theta(\iota(e)).e.\theta(\tau(e))^{-1} : e \in E(\Gamma) - E(T)\}$ form a free generating set for $\pi_1(\Gamma, b)$.

We will use the edge loop group, defined in III.26. However, Γ need not be a simplicial complex, since graphs are allowed to have edges with both endpoints at the same vertex and to have multiple edges running between two vertices. But there is an easy way to rectify this problem. Simply subdivide each edge of Γ into three edges. The resulting graph Γ' is clearly homeomorphic to Γ and it is a simplicial complex. For each edge e of Γ, we assign the label e to the middle of the three corresponding edges of Γ'.

We will set up an isomorphism $\phi\colon F(E(\Gamma) - E(T)) \to E(\Gamma', b)$. This is induced, using Theorem IV.10, by the function $E(\Gamma) - E(T) \to E(\Gamma', b)$ that sends each edge e of $E(\Gamma) - E(T)$ to $\theta(\iota(e)).e.\theta(\tau(e))^{-1}$.

To show that this is an isomorphism, we set up an inverse ψ. Any edge loop ℓ in Γ' defines a word in the alphabet $E(\Gamma')$: whenever the path traverses an edge e in the forwards direction (respectively, the backwards direction), write down e (respectively, e^{-1}). Remove all letters in the word that correspond to edges lying in T. Also, for each edge of Γ, remove the letters corresponding to the outer two of the three edges of Γ'. Define $\psi(\ell)$ to be the resulting word in the alphabet $E(\Gamma) - E(T)$. We must show that ψ is well-defined, by establishing that the resulting element of $F(E(\Gamma) - E(T))$ is unchanged if the loop is modified by an elementary expansion or contraction. Crucially, Γ' has no 2-simplices, and so we need only worry about moves (0) and (1) in Definition III.23.

If move (0) is applied to a loop $\ell = (b, b_1, \dots, b_{n-1}, b)$, removing a repeated vertex b_i, then the corresponding word in the alphabet $E(\Gamma')$ is unchanged, and hence so is $\psi(\ell)$.

If move (1) is applied to ℓ, removing an edge e followed by its reverse, this has the effect of an elementary contraction on the word $\psi(\ell)$ (if e is in the middle of an edge in $E(\Gamma) - E(T)$) or no effect at all (otherwise).

Hence, ψ is a well-defined function $E(\Gamma', b) \to F(E(\Gamma) - E(T))$. It is clearly a homomorphism since concatenation of edge loops results in concatenation of words in $F(E(\Gamma) - E(T))$. It is trivial to check that ψ and ϕ are mutual inverses. □

Remark IV.19. The above proof not only establishes that $\pi_1(\Gamma, b)$ is free but also gives an explicit free generating set, as the following examples demonstrate.

Example IV.20. Let Γ be the graph with a single vertex b and four edges e_1, e_2, e_3 and e_4. Then a maximal tree T in Γ consists of just the vertex b. Thus, the four edges do not lie in T, and so $\pi_1(\Gamma, b)$ is a free group on four generators, where the i^{th} generator is a loop going once around the edge e_i.

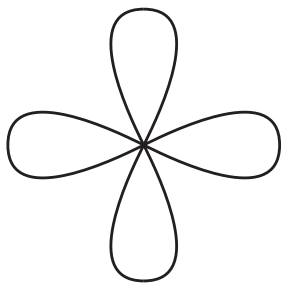

Figure IV.21.

Example IV.22. Consider the graph Γ in the plane with a vertex at each lattice point (x, y), where $x, y \in \mathbb{Z}$, and where (x, y) is joined to $(x-1, y)$, $(x+1, y)$, $(x, y-1)$ and $(x, y+1)$. For ease of reference, label the horizontal edges with a and orient them from left to right, and label the vertical edges b and orient them upwards. The union of all the vertical edges and the x-axis forms a maximal tree T in Γ. Place the basepoint b at the origin. Then, there is a free generator of $\pi_1(\Gamma, b)$ for each edge running from (x, y) to $(x+1, y)$, for each $x, y \in \mathbb{Z}$ where $y \neq 0$. The corresponding loop in Γ is the path $a^x b^y a b^{-y} a^{-x-1}$.

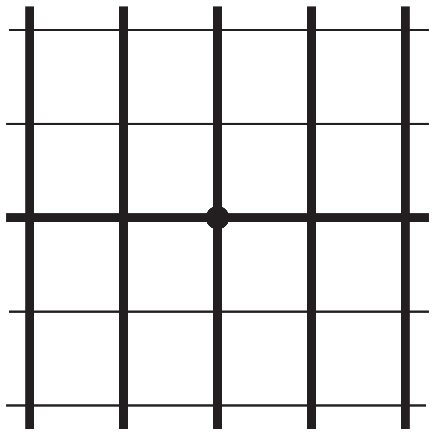

Figure IV.22.

Chapter V: Group presentations

V.1: Generators and relations

Most undergraduates have come across groups defined using 'generators and relations'. A common example is the dihedral group D_{2n}, which is 'defined' as

$$\langle \sigma, \tau \mid \sigma^n = e, \tau^2 = e, \tau\sigma\tau = \sigma^{-1} \rangle.$$

The idea is that σ and τ generate the group, and the 'relations' of the group are given by the equalities on the right-hand side. These relations should be, in some sense, the 'only' ones that hold. However, very rarely is it explained precisely what this means! Of course, one is allowed to 'deduce' relations from the given ones. For example, if $\sigma^n = e$ and $\tau^2 = e$, then $\sigma^n\tau^2 = e$. However, there are slightly more subtle relations that also follow. For example, $\tau\sigma^n\tau$ is also the identity, since $\tau\sigma^n\tau = \tau e \tau = e$. So, it is clear that some more work must be done before one can specify a group using generations and relation, *with complete rigour*. It turns out that free groups play a central rôle in this process.

Definition V.1. Let B be a subset of a group G. The *normal subgroup generated by* B is the intersection of all normal subgroups of G that contain B. It is denoted $\langle\langle B \rangle\rangle$.

Remark V.2. The intersection of a collection of normal subgroups is again a normal subgroup. Hence, $\langle\langle B\rangle\rangle$ is normal in G. It is therefore the smallest normal subgroup of G that contains B, in the sense that any other normal subgroup that contains B also contains $\langle\langle B\rangle\rangle$.

It can be specified quite precisely, as follows.

Proposition V.3. *The subgroup $\langle\langle B\rangle\rangle$ consists of all expressions of the form*

$$\prod_{i=1}^{n} g_i b_i^{\epsilon_i} g_i^{-1},$$

where $n \in \mathbb{Z}_{\geq 0}$, $g_i \in G$, $b_i \in B$ and $\epsilon_i = \pm 1$, for all i.

Proof. Any normal subgroup containing B must contain all elements of the form gbg^{-1} and $gb^{-1}g^{-1}$ ($b \in B$, $g \in G$). Hence, it must contain all finite products of these:

$$\prod_{i=1}^{n} g_i b_i^{\epsilon_i} g_i^{-1}.$$

Let N be the set of all these finite products. We have therefore shown that $N \subset \langle\langle B\rangle\rangle$. We will show that N is in fact a normal subgroup. It clearly contains B, and so, by Remark V.2, we must have $\langle\langle B\rangle\rangle \subset N$. This will prove the proposition. To show that N is a normal subgroup, we check the various conditions:

Identity: N clearly contains e.

Inverses: The inverse of $\prod_{i=1}^{n} g_i b_i^{\epsilon_i} g_i^{-1}$ is $\prod_{i=n}^{1} g_i b_i^{-\epsilon_i} g_i^{-1}$, which also lies in N.

Closure: The product of two elements in N clearly lies in N.

Normality: For $\prod_{i=1}^{n} g_i b_i^{\epsilon_i} g_i^{-1}$ in N and $g \in G$,

$$g\left(\prod_{i=1}^{n} g_i b_i^{\epsilon_i} g_i^{-1}\right) g^{-1} = \prod_{i=1}^{n} (gg_i) b_i^{\epsilon_i} (g_i^{-1} g^{-1}) = \prod_{i=1}^{n} (gg_i) b_i^{\epsilon_i} (gg_i)^{-1},$$

which lies in N. □

We can now specify what it means to define a group via generators and relations. The generators will come from a set X. The relations will be words in X, which we can view as instructions that force these words to be the identity in the group. The precise definition is as follows.

Definition V.4. Let X be a set, and let R be a collection of elements of $F(X)$. The group with presentation $\langle X|R\rangle$ is defined to be $F(X)/\langle\langle R\rangle\rangle$.

We sometimes slightly abuse notation by allowing relations of the form '$w_1 = w_2$', where this is shorthand for the relation $w_1 w_2^{-1}$.

Example V.5. We can now genuinely define the dihedral group D_{2n} to be

$$\langle \sigma, \tau \mid \sigma^n, \tau^2, \tau\sigma\tau\sigma \rangle.$$

Note that $\tau\sigma^n\tau$ does indeed represent the identity element of D_{2n} because it lies in the subgroup of $F(\{\sigma, \tau\})$ normally generated by $\{\sigma^n, \tau^2, \tau\sigma\tau\sigma\}$. To see this, note that in $F(\{\sigma, \tau\})$,

$$\tau\sigma^n\tau = (\tau\sigma^n\tau^{-1})\tau^2,$$

which is in the form specified by Proposition V.3.

One can now answer the question: which relations hold in $G = \langle X|R \rangle$? Equivalently, when do two words w and w' in the alphabet X represent the same element of G? The answer is: precisely when there is an element y of $\langle\langle R \rangle\rangle$ such that $w' = wy$, where equality holds in the free group $F(X)$. However, the following gives an alternative way of deciding whether w and w' represent the same element of G.

Proposition V.6. *Let $G = \langle X|R \rangle$. Then two words w and w' in the alphabet X represent the same element of G if and only if they differ by a finite sequence of the following moves:*

(1) perform an elementary contraction or expansion;

(2) insert somewhere into the word one of the relations in R or its inverse.

Proof. Certainly, applying moves (1) and (2) to a word does not change the element of G that it represents. We must show that if two words w and w' represent the same element of G, then they differ by a sequence of moves (1) and (2). We know, that as elements of $F(X)$, we have the equality $w' = wy$, where $y \in \langle\langle R \rangle\rangle$. So, by Proposition V.3, there is an integer $n \in \mathbb{Z}_{\geq 0}$, and elements $g_i \in F(X)$, and elements $r_i \in R$, and $\epsilon_i = \pm 1$, such that

$$w' = w \prod_{i=1}^{n} g_i r_i^{\epsilon_i} g_i^{-1}.$$

We can obtain $wg_1g_1^{-1}$ from w by a sequence of (1) moves, and then obtain $wg_1r_1^{\epsilon_1}g_1^{-1}$ from this by move (2). Continuing in this way, we obtain w' from w by a sequence of moves (1) and (2). □

Example V.7. We have already seen that in the dihedral group D_{2n}, $\tau\sigma^n\tau$ represents the identity element. Hence, it can be turned into e by moves (1) and (2), as follows:

$$\tau\sigma^n\tau \xrightarrow{(2)} \tau\sigma^n\sigma^{-n}\tau \xrightarrow{n\times(1)} \tau\tau \xrightarrow{(2)} \tau^2\tau^{-2} \xrightarrow{2\times(1)} e.$$

We now show that any group G has a presentation. Let $F(G)$ be the free group on the generating set G. Thus, $F(G)$ consists of all (equivalence classes of) words in the alphabet G. Hence, if x_1 and x_2 are non-trivial elements of G and $x_3 = x_1x_2$ in G, then x_3 and x_1x_2 represent *distinct* elements of $F(G)$, because they are non-equivalent words in the alphabet G. There is a canonical homomorphism $F(G) \to G$ sending each generator of $F(G)$ to the corresponding element of G, which is clearly surjective. Let $R(G)$ be the kernel of this homomorphism. Thus, for example, $x_3x_2^{-1}x_1^{-1}$ lies in $R(G)$. Then, by the first isomorphism theorem for groups, G is isomorphic to $F(G)/R(G)$. Hence, G has presentation $\langle G|R(G)\rangle$.

Definition V.8. The *canonical presentation* for G is $\langle G|R(G)\rangle$.

The canonical presentation of a group is extremely inefficient. As soon as G is infinite, the canonical presentation has infinitely many generators. Its main rôle comes from the fact that it depends only on the group G and involves no arbitrary choices. But we shall mostly be interested in more efficient presentations of a group, given by the following definition.

Definition V.9. A presentation of $\langle X|R\rangle$ is *finite* if X and R are both finite sets. A group is *finitely presented* if it has a finite presentation.

Remark V.10. We will establish in Chapter V the following nice characterisation of finitely presented groups: a group is finitely presented if and only if it is isomorphic to the fundamental group of a finite simplicial complex.

The following result allows us to check whether a function from a group $\langle X|R\rangle$ to another group is a homomorphism.

Lemma V.11. *Let $\langle X|R\rangle$ and H be groups. Let a function $f\colon X \to H$ induce a homomorphism $\phi\colon F(X) \to H$. This descends to a homomorphism $\langle X|R\rangle \to H$ if and only if $\phi(r) = e$ for all $r \in R$.*

Proof. Clearly, the condition that $\phi(r) = e$ is necessary for ϕ to give a homomorphism, since any $r \in R$ represents the identity element of $\langle X|R\rangle$. Conversely, suppose that $\phi(r) = e$ for all $r \in R$. By Proposition V.3, any element w of $\langle\langle R\rangle\rangle$ can be written as

$$\prod_{i=1}^{n} w_i r_i^{\epsilon_i} w_i^{-1},$$

where $n \in \mathbb{Z}_{\geq 0}$, $w_i \in F(X)$, $r_i \in R$ and $\epsilon_i = \pm 1$, for all i. Since $\phi(r) = e$ for all $r \in R$, $\phi(w)$ is also e. Hence, ϕ descends to a homomorphism $F(X)/\langle\langle R\rangle\rangle \to H$, as required. $\square$

V.2: Tietze transformations

One can ask: when do two finite presentations represent the same group? The answer is: if and only if they differ by a sequence of so-called Tietze transformations.

Definition V.12. A *Tietze transformation* is one of the following moves applied to a finite presentation $\langle x_1, \ldots, x_m \mid r_1, \ldots, r_n \rangle$:

(T1) Re-order the generators or relations.

(T2) Add or remove the relation e.

(T3) Perform an elementary contraction or expansion to a relation r_i.

(T4) Insert a relation r_i or its inverse into one of the other r_j, or remove it.

(T5) Add a generator x_{m+1} together with a relation $w(x_1, \ldots, x_m)x_{m+1}^{-1}$, which defines it as a word in the old generators, or perform the reverse of this operation.

It is clear that none of the above transformations alters the group.

Example V.13. We claim that $\langle a, b | abab^{-1} \rangle$ and $\langle b, c | cbbc \rangle$ are presentations of the same group. We establish this via Tietze transformations:

$$
\begin{aligned}
\langle a, b | abab^{-1} \rangle &\xrightarrow{(T5)} \langle a, b, c | abab^{-1}, ab^{-1}c^{-1} \rangle \\
&\xrightarrow{(T4)} \langle a, b, c | (ab^{-1}c^{-1})^{-1}abab^{-1}, ab^{-1}c^{-1} \rangle \\
&\xrightarrow{(T4)} \langle a, b, c | (ab^{-1}c^{-1})^{-1}ab(ab^{-1}c^{-1})^{-1}ab^{-1}, ab^{-1}c^{-1} \rangle \\
&\xrightarrow{(T3)} \langle a, b, c | cbbc, ab^{-1}c^{-1} \rangle \\
&\xrightarrow{(T2)} \langle a, b, c | cbbc, ab^{-1}c^{-1}, e \rangle \\
&\xrightarrow{(T4)} \langle a, b, c | cbbc, ab^{-1}c^{-1}, (ab^{-1}c^{-1})^{-1} \rangle \\
&\xrightarrow{(T4)} \langle a, b, c | cbbc, ab^{-1}c^{-1}(ab^{-1}c^{-1})^{-1}, (ab^{-1}c^{-1})^{-1} \rangle \\
&\xrightarrow{(T3)} \langle a, b, c | cbbc, e, (ab^{-1}c^{-1})^{-1} \rangle \\
&\xrightarrow{(T2)} \langle a, b, c | cbbc, cba^{-1} \rangle \\
&\xrightarrow{(T5)} \langle b, c | cbbc \rangle
\end{aligned}
$$

Theorem V.14. (Tietze) *Any two finite presentations of a group G are convertible into each other by a finite sequence of Tietze transformations.*

Lemma V.15. *Let $\langle x_1, \dots, x_m | r_1, \dots, r_n \rangle$ be a presentation for a group G. Suppose that a word w in the generators $x_1, \dots, x_m$ is trivial in G. Then, there is a sequence of (T2), (T3) and (T4) moves taking the presentation to $\langle x_1, \dots, x_m | r_1, \dots, r_n, w \rangle$.*

Proof. We first apply move (T2), adding the relation e. Since w is trivial in G, it can be obtained from e by a sequence of moves (1) and (2) in Proposition V.6. So, there is a sequence of (T3) and (T4) moves taking the presentation to $\langle x_1, \dots, x_m | r_1, \dots, r_n, w \rangle$, as required. □

Proof of Theorem V.14. Let $\langle x_1, \dots, x_m | r_1, \dots, r_n \rangle$ and $\langle x'_1, \dots, x'_{m'} | r'_1, \dots, r'_{n'} \rangle$ be two presentations of G. Since each x'_i is an element of $\langle x_1, \dots, x_m | r_1, \dots, r_n \rangle$, it can be written as a word χ'_i in the generators $x_1, \dots, x_m$. Similarly, each x_i can be written as a word χ_i in the generators $x'_1, \dots, x'_{m'}$.

We start by applying the move (T5) m' times to the first presentation to obtain

$$\langle x_1, \dots, x_m, x'_1, \dots, x'_{m'} | r_1, \dots, r_n, \chi'_1 (x'_1)^{-1}, \dots, \chi'_{m'} (x'_{m'})^{-1} \rangle.$$

The relation $x_i = \chi_i$ holds in the group. Hence, by Lemma V.15, there is a sequence of (T2), (T3) and (T4) moves taking the presentation to

$$\left\langle x_1, \dots, x_m, x'_1, \dots, x'_{m'} \,\middle|\, \begin{array}{c} r_1, \dots, r_n, \\ \chi'_1 (x'_1)^{-1}, \dots, \chi'_{m'} (x'_{m'})^{-1}, \\ \chi_1 x_1^{-1}, \dots, \chi_m x_m^{-1} \end{array} \right\rangle.$$

The relations $r'_1, \dots, r'_{n'}$ represent trivial words in the group. Hence, by Lemma V.15, there is a sequence of (T2), (T3) and (T4) moves taking the presentation to

$$\left\langle x_1, \dots, x_m, x'_1, \dots, x'_{m'} \,\middle|\, \begin{array}{c} r_1, \dots, r_n, r'_1, \dots, r'_{n'}, \\ \chi'_1 (x'_1)^{-1}, \dots, \chi'_{m'} (x'_{m'})^{-1}, \\ \chi_1 x_1^{-1}, \dots, \chi_m x_m^{-1} \end{array} \right\rangle.$$

This presentation is symmetric with respect to primed and unprimed symbols, except they occur in a different order. Hence, we can first apply (T1) moves, and then reversing the above derivation, we can obtain the presentation $\langle x'_1, \dots, x'_{m'} | r'_1, \dots, r'_{n'} \rangle$, as required. □

V.3: Push-outs

In this section, we use presentations to define the following construction which is important in group theory.

Definition V.16. Let G_0, G_1 and G_2 be groups, and let $\phi_1 \colon G_0 \to G_1$ and $\phi_2 \colon G_0 \to G_2$ be homomorphisms. Let $\langle X_1 | R_1 \rangle$ and $\langle X_2 | R_2 \rangle$ be the canonical presentations of G_1 and G_2, where $X_1 \cap X_2 = \emptyset$. Then the *push-out* $G_1 *_{G_0} G_2$ of

$$G_1 \xleftarrow{\phi_1} G_0 \xrightarrow{\phi_2} G_2$$

is the group

$$\langle X_1 \cup X_2 \mid R_1 \cup R_2 \cup \{\phi_1(g) = \phi_2(g) : g \in G_0\} \rangle.$$

Remark V.17. The notation $G_1 *_{G_0} G_2$ is ambiguous, since the resulting group depends not just on G_0, G_1 and G_2, but on the homomorphisms ϕ_1 and ϕ_2. However, this terminology has become standard, at least when ϕ_1 and ϕ_2 are injective.

The canonical presentations of G_1 and G_2 were used to ensure that the resulting answer was clearly independent of the choice of presentation. However, we will see later (in Lemma V.20) that one may substitute other presentations for G_1 and G_2 in the definition and obtain the same group.

Remark V.18. By Lemma V.11, the inclusions $X_1 \to X_1 \cup X_2$ and $X_2 \to X_1 \cup X_2$ induce *canonical homomorphisms* $\alpha_1 \colon G_1 \to G_1 *_{G_0} G_2$ and $\alpha_2 \colon G_2 \to G_1 *_{G_0} G_2$. The following diagram commutes:

$$\begin{array}{ccc} G_0 & \xrightarrow{\phi_1} & G_1 \\ \downarrow{\scriptstyle \phi_2} & & \downarrow{\scriptstyle \alpha_1} \\ G_2 & \xrightarrow{\alpha_2} & G_1 *_{G_0} G_2 \end{array}$$

This is because the relation $\phi_1(g) = \phi_2(g)$, for each $g \in G_0$, holds in $G_1 *_{G_0} G_2$.

Proposition V.19. (Universal property of push-outs) *Let $G_1 *_{G_0} G_2$ be the push-out of*

$$G_1 \xleftarrow{\phi_1} G_0 \xrightarrow{\phi_2} G_2.$$

Let H be a group and let $\beta_1 \colon G_1 \to H$ and $\beta_2 \colon G_2 \to H$ be homomorphisms such that the following diagram commutes:

$$\begin{array}{ccc} G_0 & \xrightarrow{\phi_1} & G_1 \\ \downarrow{\scriptstyle \phi_2} & & \downarrow{\scriptstyle \beta_1} \\ G_2 & \xrightarrow{\beta_2} & H \end{array}$$

Then there is a unique homomorphism $\beta \colon G_1 *_{G_0} G_2 \to H$ such that the following diagram *commutes:*

$$\begin{array}{ccc} & & G_1 \\ & \swarrow \scriptstyle{\alpha_1} & \downarrow \scriptstyle{\beta_1} \\ G_1 *_{G_0} G_2 & \xrightarrow{\beta} & H \\ & \nwarrow \scriptstyle{\alpha_2} & \uparrow \scriptstyle{\beta_2} \\ & & G_2 \end{array}$$

Proof. The push-out $G_1 *_{G_0} G_2$ has generators $G_1 \cup G_2$. Define β on these generators by $\beta(g_i) = \beta_i(g_i)$, for $g_i \in G_i$. This is forced upon us by the commutativity of the final diagram in V.19. Thus, if the homomorphism β exists, then it is unique. But we must check that it is well-defined, by verifying that $\beta(r) = e$ for any relation r of $G_1 *_{G_0} G_2$. This is certainly true for the relations in G_1 and G_2, since β_1 and β_2 are well-defined. The other type of relation is $\phi_1(g)(\phi_2(g))^{-1}$ for $g \in G_0$. But $(\beta\phi_1(g))(\beta\phi_2(g))^{-1} = e$ by the commutativity of the square in V.19. □

Lemma V.20. *Let G_0, G_1, G_2, ϕ_1 and ϕ_2 be as in Definition V.16. Let $\langle X_1' | R_1' \rangle$ and $\langle X_2' | R_2' \rangle$ be any presentations for G_1 and G_2, where $X_1' \cap X_2' = \emptyset$. Then the push-out is isomorphic to*

$$\langle X_1' \cup X_2' \mid R_1' \cup R_2' \cup \{\phi_1(g) = \phi_2(g) : g \in G_0\} \rangle.$$

Proof. Let G be the push-out, and let H be the group with presentation as in V.20. Let $G_1 \to \langle X_1' | R_1' \rangle$ and $G_2 \to \langle X_2' | R_2' \rangle$ be the 'identity' maps. By Lemma V.11, these induce homomorphisms $\beta_1 \colon G_1 \to H$ and $\beta_2 \colon G_2 \to H$. The square of V.19 commutes because the relations $\phi_1(g) = \phi_2(g)$ hold for each $g \in G_0$. Thus, by VI.19, there is a homomorphism $\beta \colon G \to H$ such that the final diagram in V.19 commutes. We will define an inverse $\phi \colon H \to G$. There is a function $X_i' \to G_i$ sending each generator to the corresponding element of G_i. Compose this with α_i to give a function $f \colon X_1' \cup X_2' \to G$. This induces a homomorphism $\phi \colon F(X_1' \cup X_2') \to G$. By Lemma V.11, this descends to a homomorphism $\phi \colon H \to G$ because $\phi(r) = e$ for each relation r in the presentation of H. It is clear that this is an inverse for β. Hence, G is isomorphic to H. □

Definition V.21. When G_0 is the trivial group, then the push-out $G_1 *_{G_0} G_2$ depends only on G_1 and G_2. It is known as the *free product* $G_1 * G_2$.

Example V.22. The free product $\mathbb{Z} * \mathbb{Z}$ is isomorphic to the free group on two generators. This is because we may take presentations $\langle x | \ \rangle$ and $\langle y | \ \rangle$ for the first and second copies of $\mathbb{Z}$. Lemma V.20 gives that $\mathbb{Z} * \mathbb{Z}$ has presentation $\langle x, y | \ \rangle$.

Definition V.23. When $\phi_1 \colon G_0 \to G_1$ and $\phi_2 \colon G_0 \to G_2$ are injective, the push-out $G_1 *_{G_0} G_2$ is known as the *amalgamated free product* of G_1 and G_2 along G_0.

V.4: The Seifert - Van Kampen Theorem

In this section, the two main themes of this course – topology and group theory – are woven together. We will prove a result, the Seifert - van Kampen theorem, which will allow us to compute the fundamental group of any finite simplicial complex. The answer will be given as a group presentation.

Theorem V.24. (Seifert - van Kampen) *Let K be a space, which is a union of two path-connected open sets K_1 and K_2, where $K_1 \cap K_2$ is also path-connected. Let b be a point in $K_1 \cap K_2$, and let $i_1 \colon K_1 \cap K_2 \to K_1$ and $i_2 \colon K_1 \cap K_2 \to K_2$ be the inclusion maps. Then $\pi_1(K, b)$ is isomorphic to the push-out of*

$$\pi_1(K_1, b) \xleftarrow{i_{1*}} \pi_1(K_1 \cap K_2, b) \xrightarrow{i_{2*}} \pi_1(K_2, b).$$

Moreover, the homomorphisms $\pi_1(K_1, b) \to \pi_1(K, b)$ and $\pi_1(K_2, b) \to \pi_1(K, b)$, which are the composition of the canonical homomorphisms to the push-out and the isomorphism to $\pi_1(K, b)$, are the maps induced by inclusion.

An alternative formulation is the following.

Theorem V.25. *Let K, K_1, K_2, b, i_1 and i_2 be as in V.24. Let $\langle X_1 | R_1 \rangle$ and $\langle X_2 | R_2 \rangle$ be presentations for $\pi_1(K_1, b)$ and $\pi_1(K_2, b)$, with $X_1 \cap X_2 = \emptyset$. Then a presentation of $\pi_1(K, b)$ is given by*

$$\langle X_1 \cup X_2 \mid R_1 \cup R_2 \cup \{ i_{1*}(g) = i_{2*}(g) : \text{ for each } g \in \pi_1(X_1 \cap X_2, b) \} \rangle.$$

Moreover, the homomorphism $\langle X_i | R_i \rangle \to \pi_1(K, b)$ (where $i = 1$ or 2) arising from the inclusion of generators $X_i \to X_1 \cup X_2$ is the map induced by the inclusion $K_i \to K$.

Proof. Let G be the push-out of V.24. Then G has a presentation as in V.25. Now, the following diagram commutes:

$$\begin{array}{ccc} \pi_1(K_1 \cap K_2, b) & \xrightarrow{i_{1*}} & \pi_1(K_1, b) \\ \downarrow{\scriptstyle i_{2*}} & & \downarrow \\ \pi_1(K_2, b) & \longrightarrow & \pi_1(K, b) \end{array}$$

where the unlabelled maps are the homomorphisms induced by inclusion. Hence, the universal property of push-outs (V.19) gives a homomorphism $\beta \colon G \to \pi_1(K, b)$. We will show that this is an isomorphism, by showing that it is surjective and injective.

Surjective. Suppose we are given a loop ℓ in K based at b. We need to show that ℓ is homotopic, relative to ∂I, to a composition of loops based at b, each of which lies entirely in K_1 or entirely in K_2. The inverse images $\{\ell^{-1}K_1, \ell^{-1}K_2\}$ form an open covering of I. By the Lebesgue Covering Theorem (II.33) applied to I, there is some subdivision $I_{(n)}$ of I such that each simplex of $I_{(n)}$ maps entirely into K_1 or K_2. This then expresses ℓ as a composition of paths $u_1 \dots u_n$, each of which lies entirely in K_1 or K_2. Pick a path $\theta(x)$ for each point x in K from b to x. We may insist that if $x \in K_i$, then $\theta(x)$ lies in K_i, since K_1, K_2 and $K_1 \cap K_2$ are path-connected. We also let $\theta(b)$ be the constant path c_b. Then $u_1 \dots u_n$ is homotopic relative to ∂I to $\big(\theta(u_1(0)).u_1.\theta(u_1(1))^{-1}\big) . \big(\theta(u_2(0)).u_2.\theta(u_2(1))^{-1}\big) \dots \big(\theta(u_n(0)).u_n.\theta(u_n(1))^{-1}\big)$. This is a composition of loops based at b, each of which lies entirely in K_1 or entirely in K_2, as required.

Injective. Suppose that g is an element of G such that $\beta(g)$ is the identity. Then g is a composition of generators, each of which lies in $\pi_1(K_1, b)$ or $\pi_1(K_2, b)$, and so $\beta(g)$ is represented by a corresponding composition of loops $\ell_1 \dots \ell_n$. Suppose that $\ell_1 \dots \ell_n$ is homotopic in K to the constant loop relative to ∂I. We need to show that g is the identity element of G. We will do this establishing that there is a sequence of moves (1) and (2) of V.6 taking g to the trivial word in G.

Let $H \colon I \times I \to K$ be the homotopy relative to ∂I between $\ell_1 \dots \ell_n$ and c_b. Then $\{H^{-1}K_1, H^{-1}K_2\}$ is an open covering of $I \times I$. Hence, by the Lebesgue Covering Theorem (II.33), there is some subdivision $(I \times I)_{(r)}$ (as in I.27) such that each simplex of $(I \times I)_{(r)}$ maps entirely into K_1 or entirely into K_2. We may realise the homotopy between $\ell_1 \dots \ell_n$ and c_b using the moves in Figure III.28. Each of these has the effect of homotoping a sub-path, relative to its endpoints, in such a way that during the homotopy, this sub-path remains entirely in K_1 or K_2. Let $\lambda_1, \dots, \lambda_N$ be this sequence of based loops, where $\lambda_1 = \ell_1 \dots \ell_n$, and $\lambda_N = c_b$. Each λ_i is naturally a composition of paths, each path arising from an edge of $(I \times I)_{(r)}$. Assign each of these paths a label in the set $\{1, 2\}$, with the restriction that if the path has label j, then it lies entirely in K_j. We now modify each λ_i to give a new loop λ'_i. We have expressed λ_i as a composition of paths. Replace each of these paths u with the path $\theta(u(0)).u.(\theta(u(1))^{-1}$. The composition of these new paths is λ'_i. It is expressed as a composition of loops, each based at b, and each inheriting a label in $\{1, 2\}$. Clearly, λ'_i is homotopic, relative to ∂I, to λ_i. Now, the homotopy from λ_i to λ_{i+1} induces a homotopy between λ'_i and λ'_{i+1}. This is relative to ∂I and is supported entirely in K_1 or in K_2. Now, the expression of λ'_i as a composition of labelled loops determines a word w_i in the generators $X_1 \cup X_2$. When we pass from λ'_i to λ'_{i+1}, we claim that w_{i+1} is obtained from w_i by a sequence of moves (1) and (2) of V.6. Suppose that the homotopy is supported in K_i. Now, the sub-loops may or not have the label i. But if one of the loops does not have the label i, then it lies in $K_1 \cap K_2$ and so represents an element $g \in \pi_1(K_1 \cap K_2, b)$. So, we may first apply a relation $i_{1*}(g) = i_{2*}(g)$, which has the effect of making a label change. Since the homotopy is supported in K_i, the start and end loops represent the same element of $\pi_1(K_i, b)$, and so we can get from one to the other by applying moves (1) and (2), using the relations R_i. Thus, the element of G is unchanged in the sequence of moves taking w_1 to w_N. Since $w_1 = g$ and w_N clearly represents the identity in G, g must also have been the identity. Hence, β is injective.

The fact that the homomorphisms $\pi_1(K_i, b) \to \pi_1(K, b)$ arising from the push-out construction are the maps induced by inclusion, is clear. □

One might wonder whether, in the Seifert - van Kampen theorem, it is necessary to require that K_1 and K_2 are open and path-connected.

Certainly, the hypothesis that K_1, K_2 and $K_1 \cap K_2$ be path-connected is necessary. For example, the circle can be decomposed into two open sets $K_1 = S^1 - \{1\}$ and $K_2 = S^1 - \{-1\}$. So, K_1, K_2 and both components of $K_1 \cap K_2$ are open intervals and hence have trivial fundamental group. However, $\pi_1(S^1)$ is non-trivial.

The assumption that K_1 and K_2 are are open also cannot be simply dropped: there are examples of spaces K decomposed into path-connected closed subsets K_1 and K_2, and where $K_1 \cap K_2$ is path-connected, but where the conclusion of the theorem does not hold. However, there is a version of the theorem where K is a path-connected simplicial complex, and K_1, K_2 and $K_1 \cap K_2$ are path-connected subcomplexes (and hence closed subsets of K).

V.5: Topological applications

Definition V.26. The *wedge* $(X, x) \vee (Y, y)$ of two spaces with basepoints is the quotient of the disjoint union $X \sqcup Y$, under the identification $x \sim y$. Its basepoint is the image of x and y in this quotient.

By picking an arbitrary basepoint b in S^1, and wedging n copies of (S^1, b) together, we obtain the space $\bigvee^n S^1$, which is known as a *bouquet of circles*, which is shown below. An application of the Seifert - van Kampen theorem immediately gives its fundamental group.

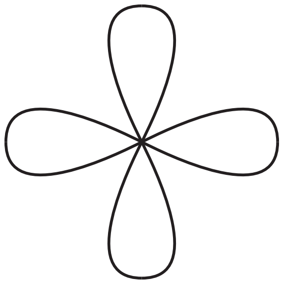

Figure V.27.

Corollary V.28. *The fundamental group of $\bigvee^n S^1$ is isomorphic to the free group on n generators.*

Proof. We apply induction on n. The induction starts with $n = 1$, where $\pi_1(S^1) \cong \mathbb{Z}$, by Theorem III.32. For the inductive step, suppose that $\pi_1(\bigvee^{n-1} S^1)$ is the free group on $n-1$ generators. Let b be the vertex of the wedge, which we take to be the basepoint. Let N be a small open neighbourhood of b. Decompose $\bigvee^n S^1$ as $K_1 = N \cup \bigvee^{n-1} S^1$ and $K_2 = N \cup S^1$. Then $\bigvee^{n-1} S^1$ is a homotopy retract of K_1, and S^1 is a homotopy retract of K_2. The intersection $K_1 \cap K_2$ is N, which is clearly contractible. So, V.25 implies that $\pi_1(\bigvee^n S^1)$ has a presentation with n generators and no relations. □

Note that $\bigvee^n S^1$ is a graph, and so Corollary V.28 can also be proved using Theorem IV.11.

Another important application of the Seifert - van Kampen theorem is that it allows us to compute the fundamental group of any cell complex. Recall the following definitions from Chapter I.

Definition I.29. Let X be a space, and let $f\colon S^{n-1} \to X$ be a map. Then the space *obtained by attaching an n-cell to X along f* is defined to be the quotient of the disjoint union $X \sqcup D^n$, such that, for each point $x \in X$, $f^{-1}(x)$ and x are all identified to a point. It is denoted by $X \cup_f D^n$.

Definition I.32. A *(finite) cell complex* is a space X decomposed as

$$K^0 \subset K^1 \subset \ldots \subset K^n = X$$

where

(i) K^0 is a finite set of points, and

(ii) K^i is obtained from K^{i-1} by attaching a finite collection of i-cells.

Theorem V.29. *Let K be a connected cell complex, and let $\ell_i\colon S^1 \to K^1$ be the attaching maps of its 2-cells, where $1 \leq i \leq n$. Let b be a basepoint in K^0. Let $[\ell_i]$ be the conjugacy class of the loop ℓ_i in $\pi_1(K^1, b)$. Then $\pi_1(K, b)$ is isomorphic to $\pi_1(K^1, b)/\langle\langle[\ell_1], \dots, [\ell_n]\rangle\rangle$.*

Remark V.30. The loops ℓ_i are not necessarily based at the basepoint b, and hence do not give well-defined elements of $\pi_1(K^1, b)$. However, they do give well-defined conjugacy classes, by Remark III.17. Let w_i be a path from b to $\ell_i(1)$, and let ℓ'_i be $w_i.\ell_i.w_i^{-1}$. Then ℓ'_i is a loop based at b, and $\langle\langle \ell'_1, \dots, \ell'_n\rangle\rangle = \langle\langle[\ell_1], \dots, [\ell_n]\rangle\rangle$. Since $\pi_1(K^1, b)$ is free, by Theorem IV.11, this therefore gives a presentation for $\pi_1(K, b)$.

Example V.31. The cell decomposition for the torus given in Example I.35 gives the presentation $\langle a, b \mid aba^{-1}b^{-1}\rangle$ for its fundamental group. This is clearly isomorphic to $\mathbb{Z} \times \mathbb{Z}$. An alternative method for constructing this cell structure of the torus arises from its description as a square with side identifications. Its four corners are all identified to a single point, giving a 0-cell. The four edges are identified to two 1-cells. The square itself is a 2-cell. One can also see easily why the loops ab and ba are homotopic relative to ∂I.

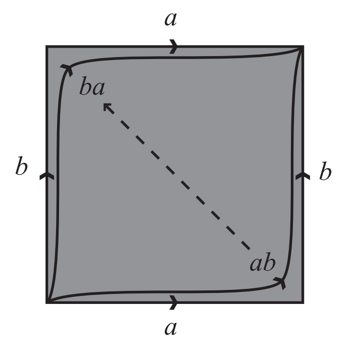

Figure V.32.

Proof of Theorem V.29. We will show how the fundamental group behaves when an n-cell is attached to a space, when $n \geq 2$. Therefore, let X be a path-connected space, and let $f\colon S^{n-1} \to X$ be the attaching map of an n-cell. Let $Y = X \cup_f D^n$. Decompose Y into the open sets

$$K_1 = \{z \in D^n : |z| < \tfrac{2}{3}\}$$
$$K_2 = \{z \in D^n : |z| > \tfrac{1}{3}\} \sqcup X/\sim.$$

Figure V.33.

Then K_1 is homeomorphic to an open n-ball, and $K_1 \cap K_2$ is homeomorphic to $S^{n-1} \times (\frac{1}{3}, \frac{2}{3})$, which is homotopy equivalent to S^{n-1}. Finally, K_2 is also path-connected, and it is homotopy equivalent to X. This is realised by the inclusion map $i\colon X \to K_2$ and a retraction map $r\colon K_2 \to X$. This is defined to be the identity map on X, and on $\{z \in D^n : |z| > \frac{1}{3}\}$, it is the map

$$\{z \in D^n : |z| > \tfrac{1}{3}\} \to S^{n-1} \xrightarrow{f} X,$$

where the first map is radial projection from the origin. It is clear that ri is the identity on X and that ir is homotopic to the identity on K_2. Now apply the Seifert - van Kampen theorem (V.25). When $n > 2$, $\pi_1(K_1 \cap K_2)$ and $\pi_1(K_1)$ are both trivial, and so attaching an n-cell has no effect on the fundamental group. When $n = 2$, $\pi_1(K_1 \cap K_2) \cong \mathbb{Z}$, and $\pi_1(K_1)$ is trivial, and so attaching a 2-cell has the effect of adding a relation to $\pi_1(X)$ that represents the (conjugacy class of) the loop $[f]$. □

Corollary V.34. *Any finitely presented group can be realised as the fundamental group of a finite connected cell complex. Moreover, this may be given a triangulation.*

Proof. Let $\langle x_1, \dots, x_m \mid r_1, \dots, r_n \rangle$ be a finite presentation of a group. Let K^0 be a single vertex. Let K^1 be a bouquet of m circles. This may be given a triangulation in which each circle consists of three 1-simplices. Then $\pi_1(K^1)$ is a free group on m generators, where each generator consists of a loop that goes round one of the circles. Now attach 2-cells along the words r_j. The resulting space has the required fundamental group, by Theorem V.29. Giving the 2-cells a simplicial structure as shown in Figure V.35 (in the case where the attaching word has length 2 in the generators), we give the whole space a triangulation. □

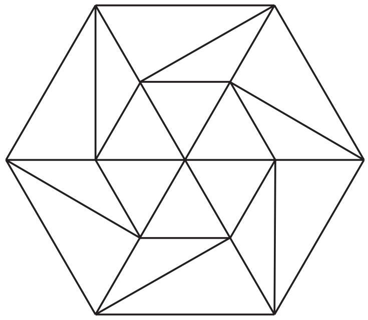

Figure V.35.

Corollary V.36. *The following are equivalent for a group G:*

(i) G is finitely presented;

(ii) G is isomorphic to the fundamental group of a finite connected simplicial complex;

(iii) G is isomorphic to the fundamental group of a finite connected cell complex.

Proof. (i) $\Rightarrow$ (ii): This is V.34.

(ii) $\Rightarrow$ (iii): This is because any finite simplicial complex is a finite cell complex (I.34).

(iii) $\Rightarrow$ (i): This follows from Remark V.30. □

Chapter VI: Covering spaces

VI.1: Definitions and basic properties

In this chapter, we will introduce an important topological construction: a covering space of a given space X. Covering spaces relate to the subgroups of the fundamental group of X. They also give an alternative method for computing fundamental groups. In addition, they relate to Cayley graphs, which were introduced in Section I.1 as a way to visualise abstract finitely generated groups.

Definition VI.1. A (continuous) map $p: \tilde{X} \to X$ is a *covering map* if X and $\tilde{X}$ are non-empty path-connected spaces and, given any $x \in X$, there exists some open set U_x containing x, such that $p^{-1}U_x$ is a disjoint union of open sets V_j (where j lies in some indexing set J) and $p|V_j: V_j \to U_x$ is a homeomorphism for all $j \in J$. The open sets U_x are known as *elementary open sets*. We say that $\tilde{X}$ is a *covering space* of X. Sometimes $\tilde{X}$ and X are given basepoints $\tilde{b}$ and b such that $p(\tilde{b}) = b$. Then $p: (\tilde{X}, \tilde{b}) \to (X, b)$ is a *based covering map*.

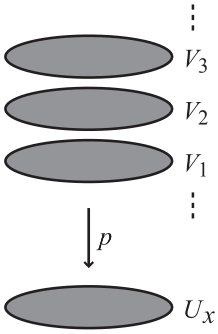

Figure VI.2.

Remark VI.3. The indexing set J *a priori* depends on x. However, we will see below that the cardinality of J is independent of x.

Example VI.4. The following is a covering map

$$\begin{aligned} \mathbb{R} &\xrightarrow{p} S^1 \\ t &\longmapsto e^{2\pi i t}. \end{aligned}$$

Given x in S^1, take U_x to be the open semi-circle with x as its midpoint. For example, $p^{-1}U_1 = \bigcup_{n\in\mathbb{Z}}(n - \frac{1}{4}, n + \frac{1}{4})$.

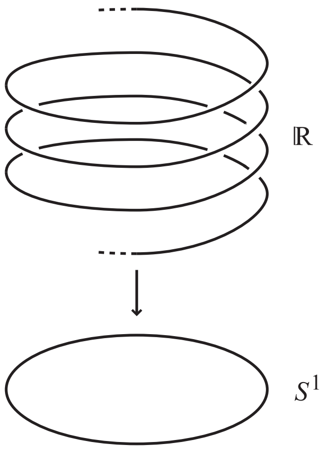

Figure VI.5.

Example VI.6. For any non-zero integer n, the map

$$\begin{aligned} S^1 &\longrightarrow S^1 \\ z &\longmapsto z^n. \end{aligned}$$

is a covering.

Example VI.7. Let $\mathbb{R}P^n$ be the set of 1-dimensional subspaces of $\mathbb{R}^{n+1}$. Define $p\colon S^n \to \mathbb{R}P^n$ to be the map that sends a point $y \in S^n \subset \mathbb{R}^{n+1}$ to the 1-dimensional subspace through y. For each point x in $\mathbb{R}P^n$, $p^{-1}(x)$ is two points. We give $\mathbb{R}P^n$ the quotient topology from the map p. Then, provided U_x is a sufficiently small open set around $x \in \mathbb{R}P^n$, $p^{-1}(U_x)$ is two copies of U_x, and the restriction of p to each of these copies is a homeomorphism onto U_x. Hence, p is a covering map.

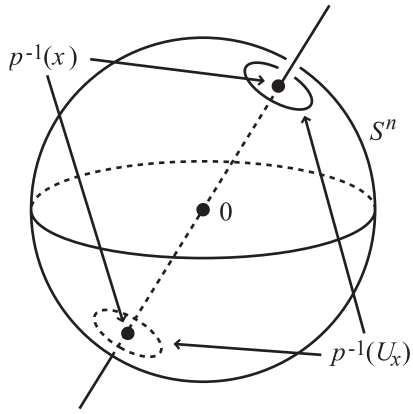

Figure VI.8.

Proposition VI.9. *Let $p\colon \tilde{X} \to X$ be a covering map. Then*

1. *p sends open sets to open sets;*
2. *for x_1 and x_2 in X, $p^{-1}(x_1)$ and $p^{-1}(x_2)$ have the same cardinality;*
3. *p is surjective;*
4. *p is a quotient map.*

Proof. (1) Let U be an open set in $\tilde{X}$. For each $y \in U$, we will find an open set containing $p(y)$ contained in $p(U)$. This will show that $p(U)$ is open, as required. Let V_j be the copy of $U_{p(y)}$ in $p^{-1}U_{p(y)}$ that contains y. Since the restriction of p to V_j is a homeomorphism, $p(V_j \cap U)$ is open in X. This is an open set containing $p(y)$ lying in $p(U)$.

(2) The cardinality of $p^{-1}(x)$ is clearly locally constant on X. Since X is connected, it is therefore globally constant.

(3) Since $\tilde{X}$ is non-empty, $p^{-1}(x)$ is non-empty for some $x \in X$. By (2), $p^{-1}(x)$ is therefore non-empty for each $x \in X$, and so p is surjective.

(4) A surjective open mapping is a quotient map. □

Definition VI.10. The *degree* of a covering map $p: \tilde{X} \to X$ is the cardinality of $p^{-1}(x)$, for any $x \in X$.

Definition VI.11. If $p: \tilde{X} \to X$ is a covering map and $f: Y \to X$ is a map, then a *lift* of f is a map $\tilde{f}: Y \to \tilde{X}$ such that $p\tilde{f} = f$:

$$\begin{array}{ccc} & & \tilde{X} \\ & \tilde{f}\nearrow & \downarrow p \\ Y & \xrightarrow{f} & X \end{array}$$

Example VI.12. Given the covering map $p: \mathbb{R} \to S^1$ of VI.4, the map $f: I \to S^1$ sending $t \mapsto e^{2\pi it}$ lifts to $\tilde{f}: I \to \mathbb{R}$, where $\tilde{f}(t) = t$. However, the identity map $S^1 \to S^1$ does not lift. One way to see this is to suppose that a lift $\tilde{f}: S^1 \to \mathbb{R}$ does exist. Then $\tilde{f}(1) = n$, for some $n \in \mathbb{Z}$. From the commutative diagram in VI.11, we obtain the following:

$$\begin{array}{ccc} & & \pi_1(\mathbb{R}, n) \\ & \nearrow \tilde{f}_* & \downarrow p_* \\ \pi_1(S^1, 1) & \xrightarrow{\text{id}} & \pi_1(S^1, 1) \end{array}$$

However, this is clearly impossible, since $\pi_1(\mathbb{R})$ is trivial, whereas $\pi_1(S^1)$ is non-trivial.

Theorem VI.13. (Uniqueness of lifts) *Let $p: \tilde{X} \to X$ be a covering map, and let $f: Y \to X$ be a map, where Y is connected. Suppose that g and h are lifts of f and that $g(y_0) = h(y_0)$ for some $y_0 \in Y$. Then $g = h$.*

Proof. Let $C = \{y \in Y : g(y) = h(y)\}$. By hypothesis, $y_0 \in C$, and so C is non-empty. We will show that C is closed, open and hence all of Y, since Y is connected.

Since p is a covering map, there is an elementary open set $U_{f(y)}$ containing $f(y)$, for any $y \in Y$, and open sets V_1 and V_2 in $\tilde{X}$ such that $p|V_1$ and $p|V_2$ are homeomorphisms from V_1 and V_2 to $U_{f(y)}$, with $g(y) \in V_1$ and $h(y) \in V_2$.

We now show that C is both closed and open.

Closed. Let y be a point in $Y - C$. Then $V_1 \cap V_2 = \emptyset$. Therefore, $g^{-1}(V_1) \cap h^{-1}(V_2)$ is contained in $Y - C$. This is an open set containing y. Hence $Y - C$ is open.

Open. Suppose that y is in C. Then $V_1 = V_2$. Consider $g^{-1}(V_1) \cap h^{-1}(V_2)$. On this set, $p \circ g = p \circ h$. Since $p|V_1$ is an injection, $g = h$ on this set. This is therefore in C. It is an open set containing y, lying in C. So, C is open. □

Theorem VI.14. (Path lifting) *Let $p: \tilde{X} \to X$ be a covering map. Let $\alpha: I \to X$ be a path with $\alpha(0) = x$. Given $\tilde{x} \in p^{-1}(x)$, α has a lift $\tilde{\alpha}: I \to \tilde{X}$ such that $\tilde{\alpha}(0) = \tilde{x}$.*

Remark VI.15. By Theorem VI.13, $\tilde{\alpha}$ is the unique lift of α such that $\tilde{\alpha}(0) = \tilde{x}$.

Proof of Theorem VI.14. Let A be the set

$$A = \{t \in I : \text{ there exists a lift of } \alpha|[0,t] \text{ starting at } \tilde{x}\}.$$

Then A is non-empty, since it contains 0. Let T be the supremum of A. Pick an elementary open set $U_{\alpha(T)}$ around $\alpha(T)$. Pick an $\epsilon > 0$ such that $(T-\epsilon, T+\epsilon) \cap [0,1]$ is mapped into $U_{\alpha(T)}$ by α. Let $t = \max\{0, T - \frac{\epsilon}{2}\}$. Let $\tilde{\alpha}: [0,t] \to \tilde{X}$ be a lift of $\alpha|[0,t]$ starting at $\tilde{x}$. Let V_j be the copy of $U_{\alpha(T)}$ in $p^{-1}U_{\alpha(T)}$ that contains $\tilde{\alpha}(t)$. Then the homeomorphism $U_{\alpha(T)} \cong V_j$ specifies a way of extending $\tilde{\alpha}$ to a lift of $\alpha|([0, T+\epsilon) \cap [0,1])$. This contradicts the definition of T unless $T = 1$. Moreover it implies that A is all of I. But $\tilde{\alpha}$ has then been defined on all of $[0,1]$ as required. □

Theorem VI.16. (Homotopy lifting) *Let $p: \tilde{X} \to X$ be a covering map. Let Y be a space, and let $H: Y \times I \to X$ be a map. If h is a lift of $H|Y \times \{0\}$, then H has a unique lift $\tilde{H}: Y \times I \to \tilde{X}$ such that $\tilde{H}|Y \times \{0\} = h$.*

Remark VI.17. When $Y = I$, such a map h always exists, by Theorem VI.14.

Proof [not examinable]. Fix $y \in Y$. Then $t \mapsto H(y,t)$ is path in X starting at $H(y,0)$. So, by Theorems VI.14 and VI.13, there is a unique lift to $\tilde{X}$, starting at $h(y)$. Define this to be $\tilde{H}(y,t)$. We must show that $\tilde{H}$ is continuous. We will show this to be the case at $(y,t) \in Y \times I$ for every $t \in I$. Since H is continuous, (y,t) has a product neighbourhood $N_t \times (a_t, b_t)$ such that $H(N_t \times (a_t, b_t)) \subset U_{H(y,t)}$. The intervals (a_t, b_t) cover I, and so by the Lebesgue Covering Theorem (I.58), there is some positive integer n such that $[i/n, (i+1)/n]$ lies in one of these intervals (a_{t_i}, b_{t_i}), for each integer i between 0 and $n-1$. Set N to be $\bigcap_{i=0}^{n-1} N_{t_i}$. We will show, by induction on i, that $\tilde{H}$ is continuous on $N \times [i/n, (i+1)/n]$, although we will feel free to shrink N to a smaller neighbourhood of y, if necessary. Now, for each i, $H(N \times [i/n, (i+1)/n])$ lies inside an elementary open neighbourhood $U_{H(y,t_i)}$. Hence, there is an open set V_j (depending on i) in $\tilde{X}$ that projects homeomorphically onto $U_{H(y,t_i)}$ and that contains $\tilde{H}(y, i/n)$. Now, by induction, $\tilde{H}|N \times \{i/n\}$ is continuous. Hence, by replacing N by a smaller neighbourhood of y if necessary, we may assume that $\tilde{H}|N \times \{i/n\}$ lies in V_j. Hence, $\tilde{H}|N \times [i/n, (i+1)/n]$ is the composition of H and the homeomorphism $U_{H(y,t_i)} \cong V_j$. In particular, it is continuous on $N \times [i/n, (i+1)/n]$. So, $\tilde{H}$ is continuous on $N \times I$, and is, in particular, continuous at (y,t) as required. □

Corollary VI.18. *If $p: (\tilde{X}, \tilde{b}) \to (X, b)$ is a based covering map, then $p_*: \pi_1(\tilde{X}, \tilde{b}) \to \pi_1(X, b)$ is an injection.*

Proof. If ℓ is a loop in $\tilde{X}$ based at $\tilde{b}$, then $p \circ \ell$ is a loop in X based at b. Suppose that $p_*[\ell] = [p \circ \ell]$ is trivial in $\pi_1(X, b)$, and let $H: I \times I \to X$ be the homotopy relative to ∂I between $p \circ \ell$ and c_b. Now, ℓ is a lift of $H|I \times \{0\}$. So, by Theorem VI.16, there is a lift $\tilde{H}: I \times I \to \tilde{X}$ of H such that $\tilde{H}|I \times \{0\} = \ell$. Note that $\tilde{H}|\{0\} \times I$, $\tilde{H}|\{1\} \times I$ and $\tilde{H}|I \times \{1\}$ are all constant maps, since the lift of a constant map from a path-connected space is constant. They must all map to $\tilde{b}$, since this is where ℓ sends ∂I. Hence, $\tilde{H}$ is a homotopy relative to ∂I between ℓ and $c_{\tilde{b}}$. Thus, $[\ell]$ is trivial in $\pi_1(X, b)$ and so p_* is an injection. □

VI.2: The inverse image of the basepoint

Fix a based covering map $p: (\tilde{X}, \tilde{b}) \to (X, b)$.

If two loops ℓ and ℓ' based at b are homotopic relative to ∂I, they can be lifted to paths $\tilde{\ell}$ and $\tilde{\ell}'$ starting at $\tilde{b}$. By Theorem VI.16, they are homotopic relative to ∂I. In particular, $\tilde{\ell}(1) = \tilde{\ell}'(1)$. So, we can define a function

$$\pi_1(X, b) \xrightarrow{\lambda} p^{-1}(b)$$
$$[\ell] \mapsto \tilde{\ell}(1).$$

Proposition VI.19. *For elements g_1 and g_2 of $\pi_1(X, b)$, $\lambda(g_1) = \lambda(g_2)$ if and only if g_1 and g_2 belong to the same right coset of $p_*\pi_1(\tilde{X}, \tilde{b})$. This induces a bijection between right cosets of $p_*\pi_1(\tilde{X}, \tilde{b})$ and points of $p^{-1}(b)$.*

Proof. Let ℓ_1 and ℓ_2 be loops based at b such that $[\ell_i] = g_i$. Suppose that $\tilde{\ell}_1(1) = \tilde{\ell}_2(1)$. Then $\tilde{\ell}_1.\tilde{\ell}_2^{-1}$ is a loop based at $\tilde{b}$. The map p sends this to $\ell_1.\ell_2^{-1}$, and so $[\ell_1][\ell_2]^{-1} = p_*[\tilde{\ell}_1.\tilde{\ell}_2^{-1}] \in p_*\pi_1(\tilde{X}, \tilde{b})$. Therefore, g_1 and g_2 belong to the same right coset of $p_*\pi_1(\tilde{X}, \tilde{b})$.

Conversely, suppose that $[\ell_1]$ and $[\ell_2]$ belong to the same right coset of $p_*\pi_1(\tilde{X}, \tilde{b})$ and so $[\ell_1][\ell_2]^{-1} \in p_*\pi_1(\tilde{X}, \tilde{b})$. Then $\ell_1.\ell_2^{-1}$ is homotopic relative to ∂I to $p \circ \ell$, for some loop ℓ in $\tilde{X}$ based at $\tilde{b}$. This homotopy lifts to a homotopy relative to ∂I between ℓ and a lift of $\ell_1.\ell_2^{-1}$. In particular, $\ell_1.\ell_2^{-1}$ lifts to a loop based at $\tilde{b}$. Now, the first half of this loop is clearly $\tilde{\ell}_1$. The second half is $\tilde{\ell}_2^{-1}$ because its reverse is the lift of ℓ_2 starting at $\tilde{b}$. Hence, $\tilde{\ell}_1(1) = \tilde{\ell}_2(1)$.

So, λ induces an injection from right cosets of $p_*\pi_1(\tilde{X}, \tilde{b})$ to $p^{-1}(b)$. We only need to show that λ is surjective. Since $\tilde{X}$ is path-connected, there is a path u from $\tilde{b}$ to any other point x of $p^{-1}(b)$. Then $p \circ u$ is a loop in X that lifts to u. So, $\lambda([p \circ u]) = x$. □

Corollary VI.20. *A loop ℓ in X based at b lifts to a loop based at $\tilde{b}$ if and only if $[\ell] \in p_*\pi_1(\tilde{X}, \tilde{b})$.*

Proof. By definition, ℓ lifts to a loop based at $\tilde{b}$ if and only $\lambda[\ell] = \tilde{b}$. But $\tilde{b}$ corresponds to the identity coset of $p_*\pi_1(\tilde{X}, \tilde{b})$. So, this is equivalent to $[\ell] \in p_*\pi_1(\tilde{X}, \tilde{b})$. □

When $p_*\pi_1(\tilde{X}, \tilde{b})$ is a normal subgroup of $\pi_1(X, b)$, the right cosets of $p_*\pi_1(\tilde{X}, \tilde{b})$ form a group, namely the quotient group $\pi_1(X, b)/p_*\pi_1(\tilde{X}, \tilde{b})$. By Proposition VI.19, $p^{-1}(b)$ is in one-one correspondence with this group. Thus, one can 'see' the group $\pi_1(X, b)/p_*\pi_1(\tilde{X}, \tilde{b})$ in the points $p^{-1}(b)$. The following procedure allows us to read off the group structure from these points.

Procedure VI.21. Suppose that $p_*\pi_1(\tilde{X}, \tilde{b})$ is a normal subgroup of $\pi_1(X, b)$. Let b_1 and b_2 be points in $p^{-1}(b)$. These correspond to elements of $\pi_1(X, b)/p_*\pi_1(\tilde{X}, \tilde{b})$. We wish to find the point $p^{-1}(b)$ (called $b_1.b_2$) corresponding to the product of these two elements. Let $\tilde{\ell}_1$ and $\tilde{\ell}_2$ be paths from $\tilde{b}$ to b_1 and b_2 respectively. Then $\ell_1 = p \circ \tilde{\ell}_1$ and $\ell_2 = p \circ \tilde{\ell}_2$ are loops in X based at b such that $\lambda([\ell_i]) = b_i$. To compute $\lambda([\ell_1].[\ell_2])$, we lift $\ell_1.\ell_2$ to a path based at b, and then $b_1.b_2$ is its endpoint. Alternatively, we can note that the second half of this path is the lift to ℓ_2 that starts at b_1.

Definition VI.22. When $\tilde{X}$ is simply-connected, a based covering map $p\colon (\tilde{X}, \tilde{b}) \to (X, b)$ is known as the *universal cover* of X.

In this case, $p^{-1}(b)$ is in one-one correspondence with $\pi_1(X, b)$. This is a useful method for computing the fundamental group of spaces. For example, we can compute the fundamental group of the circle, without using the method of Chapter III, which relied on the Simplicial Approximation Theorem.

Theorem VI.23. *The fundamental group of the circle is isomorphic to $\mathbb{Z}$.*

Proof. We know that

$$\mathbb{R} \xrightarrow{p} S^1$$
$$t \mapsto e^{2\pi i t}$$

is a covering map, and that $\mathbb{R}$ is simply-connected. So, $\pi_1(S^1, 1)$ is in one-one correspondence with $p^{-1}(1) = \mathbb{Z}$. We must verify that the group structure on $\pi_1(S^1, 1)$ coincides with the additive group structure on $\mathbb{Z}$. Let n_1 and n_2 be points in $\mathbb{Z}$. We wish to compute the composition $n_1.n_2$, using Procedure VI.20. Let $\tilde{\ell}_2$ be a path from 0 to n_2, and let $\ell_2 = p \circ \tilde{\ell}_2$. Then the lift of ℓ_2 that starts at n_1 ends at $n_1 + n_2$. Hence, λ induces an isomorphism $\pi_1(S^1, 1) \to \mathbb{Z}$. □

Example VI.24. The torus $S^1 \times S^1$ has universal cover

$$\mathbb{R} \times \mathbb{R} \xrightarrow{p} S^1 \times S^1$$
$$(x, y) \mapsto (e^{2\pi ix}, e^{2\pi iy}).$$

Set $(1, 1)$ to be the basepoint b in $S^1 \times S^1$, and let $(0, 0)$ be the basepoint in $\mathbb{R} \times \mathbb{R}$. Note that $p^{-1}(b)$ is $\mathbb{Z} \times \mathbb{Z} \subset \mathbb{R} \times \mathbb{R}$. Hence, $\pi_1(S^1 \times S^1)$ is in one-one correspondence with $\mathbb{Z} \times \mathbb{Z}$. It is trivial to check, as in VI.23, that the additive group structure on $\mathbb{Z} \times \mathbb{Z}$ is the same as that on $\pi_1(S^1 \times S^1)$. This computation of the fundamental group of the torus is an alternative to the method of V.17.

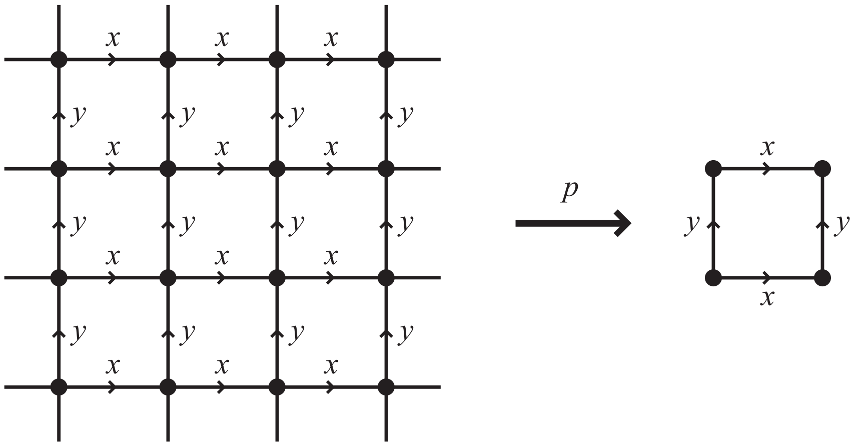

Figure VI.25.

Theorem VI.26. *The fundamental group of* $\mathbb{R}P^n$ *is* $\mathbb{Z}/2\mathbb{Z}$ *if* $n > 1$*, and is* $\mathbb{Z}$ *if* $n = 1$*.*

Proof. It is clear that $\mathbb{R}P^1$ is homeomorphic to S^1, since it is obtained by identifying the two endpoints of the upper hemisphere of S^1. Hence, $\pi_1(\mathbb{R}P^1) \cong \mathbb{Z}$. So, suppose that $n > 1$. By Example VI.7, there is a covering map $p: S^n \to \mathbb{R}P^n$, which has degree 2. Since S^n is simply-connected, the inverse image of the basepoint is in one-one correspondence with $\pi_1(\mathbb{R}P^n)$. Hence, $\pi_1(\mathbb{R}P^n)$ is the unique group of order 2, namely $\mathbb{Z}/2\mathbb{Z}$. □

VI.3: Uniqueness of coverings

Definition VI.27. A space Y is *locally path-connected* if, for each point y of Y and each neighbourhood V of y, there is an open neighbourhood of y contained in V that is path-connected.

Remark VI.28. Any simplicial complex is locally path-connected.

Theorem VI.29. (Existence of lifts) *Let $p\colon (\tilde{X}, \tilde{b}) \to (X, b)$ be a based covering map. Let Y be a path-connected, locally path-connected space and let $f\colon (Y, y_0) \to (X, b)$ be some map. Then f has a lift $\tilde{f}\colon (Y, y_0) \to (\tilde{X}, \tilde{b})$ if and only if $f_*\pi_1(Y, y_0) \subset p_*\pi_1(\tilde{X}, \tilde{b})$.*

Proof. The necessity of the condition $f_*\pi_1(Y, y_0) \subset p_*\pi_1(\tilde{X}, \tilde{b})$ follows immediately from the commutativity of the diagram in VI.11.

Suppose now that this condition holds. We shall construct the required lift $\tilde{f}$. Given a point $y \in Y$, let α be a path from y_0 to y. Then $\beta = f \circ \alpha$ is a path in X. Lift this to a path $\tilde{\beta}$ in $\tilde{X}$ starting at $\tilde{b}$. Define $\tilde{f}(y)$ to be $\tilde{\beta}(1)$. Then, $p\tilde{f}(y) = p\tilde{\beta}(1) = \beta(1) = f\alpha(1) = f(y)$, and so the commutativity of VI.11 is established.

However, we must show that $\tilde{f}$ is independent of the choice of path α, and hence well-defined. Let α' be another path from y_0 to y, and let β' and $\tilde{\beta}'$ be the corresponding paths in X and $\tilde{X}$. Then $\alpha'.\alpha^{-1}$ is a loop based at y_0. So, $(f \circ \alpha').(f \circ \alpha^{-1})$ is a loop based at b. Since $[(f \circ \alpha').(f \circ \alpha^{-1})] = f_*[\alpha'.\alpha^{-1}] \in \mathrm{Im}(f_*) \subset \mathrm{Im}(p_*)$, $(f \circ \alpha').(f \circ \alpha^{-1})$ lifts to a loop in $\tilde{X}$ based at $\tilde{b}$. This is therefore $\tilde{\beta}'.\tilde{\beta}^{-1}$. In particular, $\tilde{\beta}(1) = \tilde{\beta}'(1)$, and so $\tilde{f}(y)$ is well-defined.

We now show that $\tilde{f}$ is continuous. Let y be a point of Y, and let α be a path from y_0 to y. Let $U_{f(y)}$ be an elementary open neighbourhood of $f(y)$, and let V_j be the homeomorphic copy of $U_{f(y)}$ in $\tilde{X}$ containing $\tilde{f}(y)$. Since Y is locally path-connected, there is a path-connected open neighbourhood W of y that lies in $f^{-1}U_{f(y)}$. Hence, for any point y' in W, there is a path α' from y to y' that lies in W. We may use $\alpha.\alpha'$ to define $\tilde{f}(y')$. Since $f(W)$ is contained in an elementary open set, the lift of $f\alpha'$ is simply $f\alpha'$ composed with the homeomorphism $U_{f(y)} \cong V_j$. It follows that $\tilde{f}$ is continuous at y. □

Definition VI.30. Two based covering spaces $p\colon (\tilde{X}, \tilde{b}) \to (X, b)$ and $p'\colon (\tilde{X}', \tilde{b}') \to (X, b)$ are *equivalent* if there is a homeomorphism $f\colon (\tilde{X}, \tilde{b}) \to (\tilde{X}', \tilde{b}')$ such that $p = p'f$:

$$\begin{array}{ccc} (\tilde{X}, \tilde{b}) & \xrightarrow{f} & (\tilde{X}', \tilde{b}') \\ & \searrow^{p} \quad \swarrow^{p'} & \\ & (X, b) & \end{array}$$

Theorem VI.31. (Uniqueness of covering spaces) *Let X be a path-connected, locally path-connected space, and let b be a basepoint in X. Then, for any subgroup H of $\pi_1(X, b)$, there is at most one based covering space $p\colon (\tilde{X}, \tilde{b}) \to (X, b)$, up to equivalence, such that $p_*\pi_1(\tilde{X}, \tilde{b}) = H$.*

Proof. Let $p' \colon (\tilde{X}', \tilde{b}') \to (X, b)$ be another covering such that $p'_* \pi_1(\tilde{X}', \tilde{b}') = H$. Then, by Theorem VI.29, p' admits a lift $\tilde{p}' \colon (\tilde{X}', \tilde{b}') \to (\tilde{X}, \tilde{b})$. Similarly, p admits a lift $\tilde{p} \colon (\tilde{X}, \tilde{b}) \to (\tilde{X}', \tilde{b}')$. The following diagram therefore commutes:

$$\begin{array}{ccccc} (\tilde{X}', \tilde{b}') & \xrightarrow{\tilde{p}'} & (\tilde{X}, \tilde{b}) & \xrightarrow{\tilde{p}} & (\tilde{X}', \tilde{b}') \\ & \searrow p' & \downarrow p & \swarrow p' & \\ & & (X, b) & & \end{array}$$

Now, the composition of the mappings on the top line is a lift of p'. But $\mathrm{id}_{\tilde{X}'}$ is also. Hence, $\tilde{p}\tilde{p}' = \mathrm{id}_{\tilde{X}'}$, by VI.13. Similarly, $\tilde{p}'\tilde{p} = \mathrm{id}_{\tilde{X}}$. Therefore, $\tilde{p}'$ is a homeomorphism, and so the coverings are equivalent. □

VI.4: Construction of covering spaces

Theorem VI.32. *Let K be a path-connected simplicial complex, and let b be a vertex of K. Then, for any subgroup H of $\pi_1(K, b)$, there is a based covering $p \colon (\tilde{K}, \tilde{b}) \to (K, b)$ such that $p_* \pi_1(\tilde{K}, \tilde{b}) = H$. Moreover, $\tilde{K}$ is a simplicial complex and p is a simplicial map.*

The idea behind this construction has its roots in Section VI.2. There, we showed that $p^{-1}(b)$ is in one-one correspondence with the right cosets of $p_* \pi_1(\tilde{K}, \tilde{b})$. Thus, a point of $p^{-1}(b)$ corresponds to an equivalence class of loops based at b, where two loops ℓ and ℓ' are 'equivalent' if $[\ell][\ell']^{-1} \in p_* \pi_1(\tilde{K}, \tilde{b})$. So, if we only know the subgroup $p_* \pi_1(\tilde{K}, \tilde{b})$, but have no further information about p or $\tilde{K}$, we can still reconstruct $p^{-1}(b)$. How do we reconstruct the whole of $\tilde{K}$? The answer is that, instead of considering *loops* based at b, we use *paths* that start at b but do not necessarily end at b. The precise construction is as follows.

Proof of VI.32 [not examinable]. We first define the vertices $V(\tilde{K})$ of $\tilde{K}$. Each such vertex will be an equivalence class of edge paths. We will consider only edge paths that start at b, but not necessarily ending at b. We call a path with these properties *based at b*. Two such paths α and β are *H-equivalent* if there is some loop ℓ based at b with $[\ell] \in H$ such that $\alpha \sim \ell.\beta$, where $\sim$ is as defined in III.23 using elementary contractions and expansions. It is straightforward to check that this is an equivalence relation. Then $V(\tilde{K})$ is defined to be the set of H-equivalence classes.

We now define the simplices of $\tilde{K}$. We declare that vertices in $V(\tilde{K})$ span a simplex if and only if they have representative edge paths $\alpha.(v_0, v_0), \alpha.(v_0, v_1), \ldots, \alpha.(v_0, v_n)$, where α is some edge path based at b and $(v_0, \ldots, v_n)$ are the vertices of a simplex of K.

Define $\tilde{b}$ to be the equivalence class of the path (b) of length zero.

Claim. $\tilde{K}$ is an abstract simplicial complex.

We must check the two conditions in the definition (I.35). Firstly, if $v \in V(\tilde{K})$, then v is represented by an edge path α, ending at v_0, say. This is equivalent to $\alpha.(v_0, v_0)$, and so (v) is a simplex of $\tilde{K}$.

Suppose that $[\alpha.(v_0, v_0)], [\alpha.(v_0, v_1)], \ldots, [\alpha.(v_0, v_n)]$ span a simplex of $\tilde{K}$, and let $[\alpha.(v_0, v_{i_0})], \ldots, [\alpha.(v_0, v_{i_m})]$ be a non-empty subset of these vertices. Then, changing the representative paths of these vertices, we may write them as $[\alpha.(v_0, v_{i_0}).(v_{i_0}, v_{i_0})], \ldots,$ $[\alpha.(v_0, v_{i_0}).(v_{i_0}, v_{i_m})]$, which clearly span a simplex of $\tilde{K}$, as required. This proves the claim.

Claim. $\tilde{K}$ is path-connected.

It suffices to show that any vertex of $\tilde{K}$ is connected to $\tilde{b}$ by an edge path in $\tilde{K}$. Let $[\alpha]$ be a vertex of $\tilde{K}$, where $\alpha = (b, b_1, \ldots, b_n)$ is an edge path. Then $[(b)], [(b, b_1)], \ldots,$ $[(b, b_1, \ldots, b_n)]$ is an edge path in $\tilde{K}$ joining $\tilde{b}$ to $[\alpha]$. This proves the claim.

Define $p \colon V(\tilde{K}) \to V(K)$ be sending an H-equivalence class of edge paths to their terminal point. This is well-defined, since H-equivalent edge paths have the same terminal points.

Claim. p is a simplicial map, and the restriction of p to any simplex of $\tilde{K}$ is injective.

If $(w_0, \ldots w_n)$ span a simplex of $\tilde{K}$, then, by definition, they have representatives $\alpha.(v_0, v_0), \alpha.(v_0, v_1), \ldots, \alpha.(v_0, v_n)$, where $(v_0, \ldots, v_n)$ span a simplex of K. Then $p(w_i) = v_i$, for each i, and so $(p(w_0), \ldots, p(w_n))$ span a simplex of K. Note that this has not decreased the dimension of the simplex, as required.

We need to show that p is a covering map, and to do this, we must specify, for each $x \in K$, an elementary open set U_x. We set this to be $\mathrm{st}_K(x)$. Let $\mathrm{cl}(\mathrm{st}_K(x))$ denote the closure of the star of x. Note that this is a subcomplex of K.

Claim. For each vertex $w \in \tilde{K}$, the restriction of p to $\mathrm{cl}(\mathrm{st}_{\tilde{K}}(w))$ is a simplicial isomorphism onto $\mathrm{cl}(\mathrm{st}_K(p(w)))$. Hence, the restriction of p to $\mathrm{st}_{\tilde{K}}(w)$ is a homeomorphism onto $\mathrm{st}_K(p(w))$.

Let $v = p(w)$. We first check that the restriction of p to the vertices of $\mathrm{cl}(\mathrm{st}_{\tilde{K}}(w))$ is a bijection onto the vertices of $\mathrm{cl}(\mathrm{st}_K(v))$.

Injection. Let w_1 and w_2 be distinct vertices of $\mathrm{cl}(\mathrm{st}_{\tilde{K}}(w))$. Then w and w_i span a simplex, and so they have representative edge paths $\alpha_i.(v,v)$ and $\alpha_i.(v,v_i)$, where $p(w_i)=v_i$ and (v,v_i) span a simplex of K. Since $\alpha_1.(v,v)$ and $\alpha_2.(v,v)$ both represent the same vertex w, they must be H-equivalent, and hence so are α_1 and α_2. Therefore, w_2 is represented by $\alpha_1.(v,v_2)$. Since we are assuming that w_1 and w_2 are distinct, we must have $v_1 \neq v_2$. In other words, $p(w_1) \neq p(w_2)$, which proves that this is an injection.

Surjection. Let v_1 be a vertex in $\mathrm{cl}(\mathrm{st}_K(v))$. Then, v and v_1 span a simplex in K. Let α be an edge path that represents w. Then $\alpha.(v,v)$ and $\alpha.(v,v_1)$ span a simplex of $\tilde{K}$. Hence, $\alpha.(v,v_1)$ is a vertex in $\mathrm{cl}(\mathrm{st}_{\tilde{K}}(w))$ that maps to v_1.

Suppose now that $(v_0,\dots,v_n)$ span a simplex in $\mathrm{cl}(\mathrm{st}_K(v))$. Let $w_i = \mathrm{cl}(\mathrm{st}_{\tilde{K}}(w))) \cap p^{-1}(v_i)$. We must show that $(w_0,\dots,w_n)$ span a simplex of $\tilde{K}$. Suppose first that $v \in \{v_0,\dots,v_n\}$. Then we may set $v_0 = v$, by relabelling the vertices. Now, $w=[\alpha]$ for some edge path α based at b. Then $[\alpha.(v_0,v_0)],\dots,[\alpha.(v_0,v_n)]$, which equals $(w_0,\dots,w_n)$, spans the required simplex of $\tilde{K}$. When $v \notin \{v_0,\dots,v_n\}$, then $(v,v_0,\dots,v_n)$ span a simplex of K, and so $(w,w_0,\dots,w_n)$ span a simplex of $\tilde{K}$, and therefore so do $(w_0,\dots,w_n)$. This proves the claim.

Claim. For each point $y \in \tilde{K}$, $p|\mathrm{st}_{\tilde{K}}(y)$ is a homeomorphism onto $\mathrm{st}_K(p(y))$.

The previous claim established this when y is a vertex. The point y lies in the inside of a simplex σ. Let w be a vertex of σ. Then $\mathrm{st}_{\tilde{K}}(w)$ contains $\mathrm{st}_{\tilde{K}}(y)$. Since the restriction of p to $\mathrm{cl}(\mathrm{st}_{\tilde{K}}(w))$ is a simplicial isomorphism onto its image, it preserves stars. Therefore, the $p|\mathrm{st}_{\tilde{K}}(y)$ is a homeomorphism onto $\mathrm{st}_{\tilde{K}}(p(y))$, as required.

We define, for a point $x \in X$, the indexing set J in the definition of a covering map to be $p^{-1}(x)$. For $y \in J = p^{-1}(x)$, we set V_y to be $\mathrm{st}_{\tilde{K}}(y)$. We have shown that the restriction of p to any V_y is a homeomorphism onto U_x. Thus, to verify that p is a covering map, all we now need to do is prove the following claim.

Claim. Let y and y' be distinct points of $p^{-1}(x)$. Then $V_y \cap V_{y'} = \emptyset$.

Suppose that V_y and $V_{y'}$ intersect in a point z. Now, by the definition of $\mathrm{st}_{\tilde{K}}(y)$, there is a simplex σ that contains y, which also contains z in its inside. Hence, y lies in $\mathrm{cl}(V_z)$. The same holds for y'. But the restriction of p to $\mathrm{cl}(V_z)$ is a homeomorphism onto its image. In particular, it is injective. It is therefore impossible for the distinct points y and y' in $\mathrm{cl}(V_z)$ to have the same image x under p.

Our final claim, to complete the proof of the theorem, is the following.

Claim. $p_*\pi_1(\tilde{K}, \tilde{b}) = H$.

We must show that a loop ℓ based at b lifts to a loop if and only if $[\ell] \in H$. It suffices to consider the case where ℓ is an edge loop $(b, b_1, \ldots, b_{n-1}, b)$. The lift of this is an edge path $[(b)], [(b, b_1)], \ldots, [(b, b_1, \ldots, b_{n-1}, b)]$. This is a loop if and only if the final vertex $[\ell]$ is equal to $[(b)]$. This occurs if and only if ℓ is H-equivalent to (b), which means that $[\ell]$ is in H. □

Theorems VI.31 and VI.32 and Remark VI.28 give the following classification of covering spaces.

Theorem VI.33. *Let K be a path-connected simplicial complex, and let b be a vertex of K. Then, there is precisely one based covering space, up to equivalence, for each subgroup of $\pi_1(K, b)$.*

The following theorem is a very striking and beautiful consequence. But more striking still is the fact that, although its statement is purely algebraic, its proof is topological. There do exist algebraic proofs, but none are as elegant as this.

Theorem VI.34. (Nielsen-Schreier) *Any subgroup of a finitely generated free group is free.*

Proof. Let F be the free group on n generators. Let X be the wedge of n circles, and let b be the central vertex. Then $\pi_1(X, b) \cong F$, by Corollary V.5. Let H be any subgroup of F. Then, by Theorem VI.32, there is a based covering $p\colon (\tilde{X}, \tilde{b}) \to (X, b)$ such that $p_*\pi_1(\tilde{X}, \tilde{b}) = H$. By Corollary VI.18, p_* is injective and so $\pi_1(\tilde{X}, \tilde{b}) \cong H$. Theorem VI.32 states that $\tilde{X}$ is a simplicial complex, and that p is a simplicial map. Since p is a local homeomorphism, $\tilde{X}$ can contain only zero and one-dimensional simplices. Therefore, $\tilde{X}$ is a graph, and so by Theorem IV.11, its fundamental group is free. □

By examining the above proof and using Remark IV.19, we can obtain an explicit free generating set for the subgroup, as the following example demonstrates.

Example VI.35. Let F be the free group on two generators x and y. Let H be the kernel of the homomorphism $F \to (\mathbb{Z}/2\mathbb{Z}) \times (\mathbb{Z}/2\mathbb{Z})$ sending x to $(1, 0)$ and y to $(0, 1)$. Let X be the wedge of two circles with central vertex b. Then $\pi_1(X, b)$ is a free group on two generators x and y, represented by loops going round one of the circles once. The based covering space $p\colon (\tilde{X}, \tilde{b}) \to (X, b)$ such that $p_*\pi_1(\tilde{X}, \tilde{b}) = H$ is shown in Figure VI.36.

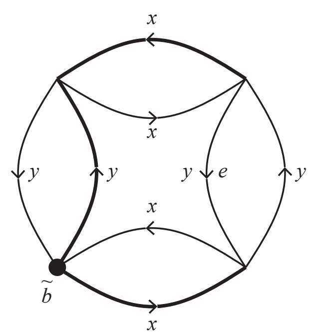

Figure VI.36.

To see that the equality $p_*\pi_1(\tilde{X},\tilde{b}) = H$ does indeed hold, note that a loop $[\ell]$ based at b lifts to a loop in $\tilde{X}$ if and only if it runs over edges labelled x an even number of times and edges labelled y an even number of times, if and only if $[\ell] \in H$. A maximal tree T in $\tilde{X}$ is shown in bold. The proof of Theorem IV.11 gives that there is a free generator of $\pi_1(\tilde{X},\tilde{b})$ for each (oriented) edge e of $E(\tilde{X}) - E(T)$. It is represented by the loop that starts at $\tilde{b}$, runs along T as far as $\iota(e)$, then runs over e, and then returns to b by a path in T. For example, corresponding to oriented edge e shown in the figure, there is the loop $xyx^{-1}y^{-1}$. Thus, a free generating set for H is as follows:

$$xyx^{-1}y^{-1}, \qquad x^2, \qquad y^2, \qquad yxyx^{-1}, \qquad yx^2y^{-1}.$$

Another interesting application of the covering space theory we have developed is the following algebraic result. Again, its proof is purely topological.

Theorem VI.37. *Let G be a finitely presented group, and let H be a finite index subgroup. Then, H is finitely presented.*

Proof. Let K be a finite simplicial complex, with basepoint b, such that $\pi_1(K,b) \cong G$, which exists by Corollary V.19. Let $p\colon (\tilde{K},\tilde{b}) \to (K,b)$ be the based covering such that $p_*\pi_1(\tilde{K},\tilde{b}) = H$. Then $p^{-1}(b)$ is in one-one correspondence with the right cosets of H in G. Since we are assuming H has finite index, $p^{-1}(b)$ is finite, and hence the degree of p is finite. Therefore, $\tilde{K}$ is also a finite simplicial complex. By Corollary V.21, its fundamental group is finitely presented. □

Example VI.38. Let G be the group $(\mathbb{Z}/2\mathbb{Z}) * \mathbb{Z}$, which has presentation $\langle x, y | x^2 \rangle$. Let $\phi\colon G \to \mathbb{Z}/2\mathbb{Z}$ be the homomorphism that sends x to 0 and y to 1, and let H be its kernel. We realize G as $\pi_1(K,b)$ for the 2-complex K as in V.19, and we let $p\colon (\tilde{K},\tilde{b}) \to (K,b)$ be the based covering space corresponding to H. The 0-cells and 1-cells of $\tilde{K}$ are shown in Figure VI.39. A maximal tree is the single edge shown in bold, and there is one free generator of the fundamental group of this graph for each of the edges c, d and f. The space $\tilde{K}$ has two 2-cells, one running over c twice, and one running over d twice. So, a presentation for H is $\langle c, d, f | c^2, d^2 \rangle$. Thus, H is isomorphic to $(\mathbb{Z}/2\mathbb{Z}) * (\mathbb{Z}/2\mathbb{Z}) * \mathbb{Z}$. Note that the generators c, d and e are respectively x, yxy^{-1} y^2.

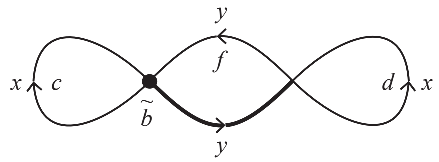

Figure VI.39.

VI.5: The universal cover and Cayley graphs

A corollary of Theorem VI.32 is that a connected simplicial complex K always has a universal cover $p\colon (\tilde{K}, \tilde{b}) \to (K, b)$. By VI.19, the points in $p^{-1}(b)$ are in one-one correspondence with the elements of $\pi_1(K, b)$. Thus, it is possible to 'see' this group in $\tilde{K}$. This is particularly useful when G is a finitely presented group, and K is the 2-complex constructed in V.19 with $\pi_1(K, b) \cong G$. In this section, we will investigate $\tilde{K}$ in this case.

Recall that G is given as a finite presentation $\langle x_1, \dots, x_m \mid r_1, \dots, r_n \rangle$, and that K is a 2-complex with a single 0-cell b, a 1-cell for each generator and a 2-cell for each relation. Assign each 1-cell of K an orientation, so that a loop running forwards along this edge represents one of the generators $x_1, \dots, x_m$. This cell complex has a triangulation and so has a universal cover $p\colon (\tilde{K}, \tilde{b}) \to (K, b)$. The cell structure on K induces a cell structure on $\tilde{K}$, as follows. The inverse image of b is a discrete collection of points, which we take to be 0-cells of $\tilde{K}$. Each 1-cell of K has an induced map $D^1 \to K$. One may lift this to $D^1 \to \tilde{K}$, starting at any point in $p^{-1}(b)$. The union of these lifts we take to be the 1-cells of $\tilde{K}$. Each 1-cell of $\tilde{K}$ is labelled by one of the generators $x_1, \dots, x_m$ and inherits an orientation from the corresponding 1-cell of K. Let Γ be the 1-skeleton of $\tilde{K}$. We will show in Proposition VI.40 that this is, in fact, the Cayley graph of G with respect to the generators $x_1, \dots, x_m$. Each 2-cell of K has interior homeomorphic to a disc. The inverse image of this in $\tilde{K}$ is a union of open discs. These therefore yield 2-cells attached to Γ. Hence, $\tilde{K}$ inherits the structure of a cell complex. This is known as the *Cayley 2-complex* associated with the presentation $\langle x_1, \dots, x_m \mid r_1, \dots, r_n \rangle$ of G.

Proposition VI.40. *The 1-skeleton of $\tilde{K}$ is the Cayley graph of G with respect to the generators $\{x_1, \dots, x_m\}$.*

Proof. Recall from I.5 that this is defined as follows. It has a vertex for each element of G. For each element $g \in G$ and each generator x_i, there is an edge running from the vertex labelled g to the vertex labelled gx_i. We need to check that Γ does indeed have this form. Let b be the 0-cell of K. By VI.19, the vertices $p^{-1}(b)$ are in one-one correspondence with the group G. Now, there are edges pointing of out b with labels $x_1, \dots, x_m$, and there are also edges pointing into b with these labels. Since p is a local homeomorphism, we have the same picture near each point v of $p^{-1}(b)$. Let v correspond to the group element g. We need to verify that the edge labelled x_i starting at v ends at the edge labelled gx_i. Procedure VI.21 states that, if one lifts the loop x_i in K to a path, starting at v, the endpoint of this path is the vertex labelled gx_i. Thus, the edge labelled x_i starting at g ends at gx_i, as required. □

Example VI.41. Let $\langle x, y | xyx^{-1}y^{-1} \rangle$ be a presentation of $\mathbb{Z} \times \mathbb{Z}$. The resulting 2-complex K with fundamental group $\mathbb{Z} \times \mathbb{Z}$ is the torus. The Cayley 2-complex has 1-skeleton shown in Figure VI.42. The 2-cells of this complex fill in each of the squares, and so the Cayley 2-complex is a copy of the plane $\mathbb{R} \times \mathbb{R}$.

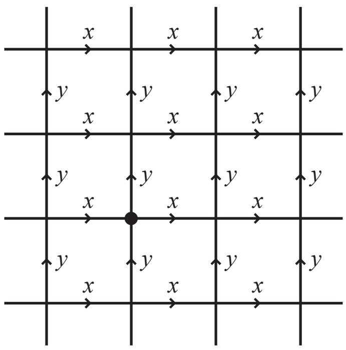

Figure VI.42.

Example VI.43. Start with the standard presentation of the free group F on two generators, with generating set $\{x, y\}$ and no relations. The associated 2-complex K is $S^1 \vee S^1$. Its universal cover $\tilde{K}$ is shown in Figure VI.44. The labelling on the edges specifies the map $\tilde{K} \to K$, which is indeed a covering map. Since $\tilde{K}$ is a tree, its fundamental group is trivial, by Theorem IV.11 and Remark IV.9, and so it is the universal cover of K. It is therefore the Cayley graph of F with respect to $\{x, y\}$.

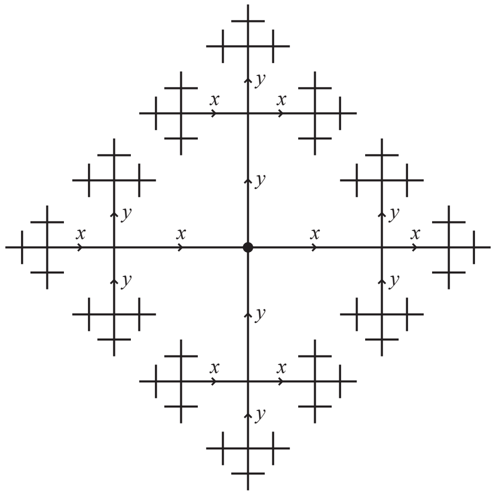

Figure VI.44.

VI.6: Normal subgroups and regular covering spaces

One can see that from Figures VI.42 and VI.44 that where one places the basepoint in a Cayley graph seems to be somewhat arbitrary. In other words, all points in the inverse image of the basepoint seem to look 'the same'. We can make this precise by the following definition.

Definition VI.45. Let $p: \tilde{X} \to X$ be a covering map. Then a *covering transformation* is a homeomorphism $t: \tilde{X} \to \tilde{X}$ such that the following diagram commutes:

$$\begin{array}{ccc} \tilde{X} & \xrightarrow{t} & \tilde{X} \\ & \searrow^{p} \quad \swarrow^{p} & \\ & X & \end{array}$$

Definition VI.46. A covering map $p: (\tilde{X}, \tilde{b}) \to (X, b)$ is *regular* if any two points of $p^{-1}(b)$ differ by a covering transformation.

It is clear that the covering maps in Examples VI.41 and VI.43 are regular. However, this need not always be the case, as the following example shows.

Example VI.47. Let $\tilde{X}$ be the space shown in Figure VI.48. There is a covering map $p: \tilde{X} \to S^1 \vee S^1$, sending the vertices of $\tilde{X}$ to the basepoint of $S^1 \vee S^1$, and mapping in the edges of $\tilde{X}$ using the recipe given by the edge labels. That this is a cover follows from the fact that each vertex of $\tilde{X}$ has four edges emanating from it, two labelled x

and y pointing in, and two labelled x and y pointing out. However, it is clear that there is no covering transformation taking the left-hand vertex of $\tilde{X}$ to the central one. This is because there is a loop labelled x emanating from the former, but there is no such loop based at the latter.

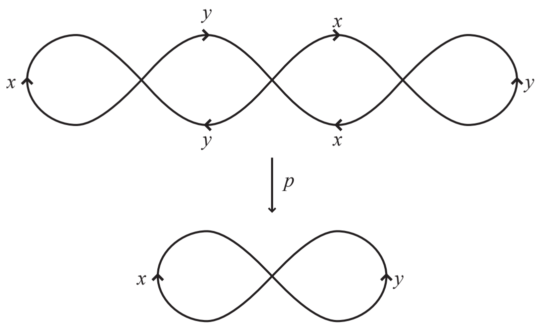

Figure VI.48.

It turns out that the distinction between this example and the earlier ones is due to whether $p_*\pi_1(\tilde{X}, \tilde{b})$ is a normal subgroup of $\pi_1(X, b)$.

Theorem VI.49. *Let $p\colon (\tilde{X}, \tilde{b}) \to (X, b)$ be a regular covering map. Then $p_*\pi_1(\tilde{X}, \tilde{b})$ is a normal subgroup of $\pi_1(X, b)$.*

Proof. Let h be an element of $p_*\pi_1(\tilde{X}, \tilde{b})$, and let g be some element of $\pi_1(X, b)$. We must show that $ghg^{-1} \in p_*\pi_1(\tilde{X}, \tilde{b})$. There is a loop ℓ in $\tilde{X}$ based at $\tilde{b}$ such that $[p \circ \ell] = h$. Also, there is a loop α in X based at b such that $[\alpha] = g$. Then α lifts to a path $\tilde{\alpha}$ based at $\tilde{b}$. By assumption, there is a covering transformation t taking $\tilde{b}$ to $\tilde{\alpha}(1)$. This takes ℓ to the loop $t\ell$. Hence, $\alpha.(p\ell).\alpha^{-1}$ lifts to $\tilde{\alpha}.(t\ell).\tilde{\alpha}^{-1}$, which is a loop. Therefore, $ghg^{-1} = [\alpha.(p\ell).\alpha^{-1}]$ is in $p_*\pi_1(\tilde{X}, \tilde{b})$. □

Theorem VI.50. *Let $p\colon (\tilde{X}, \tilde{b}) \to (X, b)$ be a covering map, where X is locally path-connected. Suppose that $p_*\pi_1(\tilde{X}, \tilde{b})$ is a normal subgroup of $\pi_1(X, b)$. Then p is regular.*

Proof. Let $\tilde{b}'$ be a point in $p^{-1}(b)$. It suffices to give a covering transformation taking $\tilde{b}$ to $\tilde{b}'$. Consider the following diagram

$$\begin{array}{ccc} & & (\tilde{X}, \tilde{b}) \\ & & \downarrow p \\ (\tilde{X}, \tilde{b}') & \xrightarrow{p} & (X, b) \end{array}$$

We claim that there is a lift $t\colon (\tilde{X}, \tilde{b}') \to (X, b)$. To prove this, we will use Theorem VI.29. We need to know that $p_*\pi_1(\tilde{X}, \tilde{b}') \subset p_*\pi_1(\tilde{X}, \tilde{b})$. Let α be a loop in X that lifts to a path $\tilde{\alpha}$ from $\tilde{b}$ to $\tilde{b}'$. Given any loop ℓ based at $\tilde{b}'$, $\tilde{\alpha}.\ell.\tilde{\alpha}^{-1}$ is a loop based at $\tilde{b}$. So, $[\alpha].[(p\ell)].[\alpha^{-1}]$ lies in $p_*\pi_1(\tilde{X}, \tilde{b})$, by Corollary VI.20. Since $p_*\pi_1(\tilde{X}, \tilde{b})$ is a normal subgroup, $[\alpha]^{-1}.[\alpha].[p\ell].[\alpha^{-1}].[\alpha]$ also lies in $p_*\pi_1(\tilde{X},\tilde{b})$. So, $[p\ell] \in p_*\pi_1(\tilde{X},\tilde{b})$, as required. So, the lift t exists. Repeat the argument, with the roles of $(\tilde{X},\tilde{b})$ and $(\tilde{X},\tilde{b}')$ reversed to get a diagram

$$\begin{array}{ccccc} (\tilde{X}',\tilde{b}') & \xrightarrow{t'} & (\tilde{X},\tilde{b}) & \xrightarrow{t} & (\tilde{X}',\tilde{b}') \\ & \searrow p' & \downarrow p & \swarrow p' & \\ & & (X,b) & & \end{array}$$

By uniqueness of lifts, the top line tt' equals the identity. Similarly, $t't$ is the identity. So, t is a homeomorphism, and therefore a covering transformation. □

This is most significant when $\tilde{X}$ is the universal cover of X. We saw in Section VI.5 that in this case, the group $\pi_1(X,b)$ can be 'seen' in the inverse image of b. But also, as a consequence of Theorem VI.48, $\pi_1(X,b)$ is realised a group of homeomorphisms of X.

Theorems VI.49 and VI.50 are just some initial steps in a large theory which relates the algebraic properties of a group with topological properties of covering spaces, especially the universal cover. This theory, known as geometric group theory, is a very active and exciting field of mathematical research. See [2] for more information about this subject.

Further reading

1. D. Cohen, *Combinatorial group theory: a topological approach*, LMS Student Texts 14 (1989), Chs 1-7.

2. P. de la Harpe, *Topics in Geometric Group Theory*, Chicago Lectures in Mathematics, University of Chicago Press (2000).

3. M. Hall, Jr, *The Theory of Groups*, Macmillan (1959), Chs. 1-7, 12, 17

4. A. Hatcher, *Algebraic Topology*, Cambridge University Press (2001), Ch. 1.

5. D.L. Johnson, *Presentations of groups*, London Mathematical Society, Student Texts 15, Second Edition, Cambridge University Press, (1997). Chs. 1-5, 10,13

6. W. Magnus, A. Karrass, and D. Solitar, *Combinatorial Group Theory*, Dover Publications, (1976). Chs. 1-4

7. J. Stillwell, *Classical Topology and Combinatorial Group Theory* (Springer-Verlag, 1993)