COVERING SPACES AND BASS–SERRE THEORY

ANDREJ VUKOVIĆ

Abstract. This article is an introduction to geometric group theory from the perspective of topology. We begin by reviewing covering space theory, culminating in a proof of the Seifert–van Kampen Theorem. We then develop the notion of graphs of groups. Finally, we study Bass–Serre theory, the theory of groups acting on trees.

1. Introduction

Geometric group theory is the field of mathematics which is concerned with applying topological and geometric methods to the study of group-theoretic problems. One important tool in geometric group theory is Bass–Serre theory, the study of group actions on trees, first developed by Hyman Bass and Jean-Pierre Serre in [7]. This work was in some ways a precursor to the ideas of Gromov, which lie at the foundation of contemporary geometric group theory.

This expository paper concludes with some basic definitions and results from Bass–Serre theory. However, that area of mathematics is highly reliant on earlier results from algebraic topology. In particular, we need to be able to calculate (and define!) fundamental groups, which means we need the Seifert–van Kampen theorem at our disposal. There are two main approaches to proving this theorem. The first, as presented for example in [2, Theorem 1.20], is essentially geometric in character and constructs the required map basically from a free product of fundamental groups. The second, as presented in [1, Theorem 14.4] and also in this paper (Theorem 8.1), is algebraic, and instead requires a large amount of covering space theory which then leads to a much shorter proof. In [1], this proof is attributed to Alexander Grothendieck.

Throughout this paper, we have cited the sources of the results we use. To this end, we have included a citation to the bibliography next to or in the preamble preceding each result.

2. Liftings

Definition 2.1. Given topological spaces X and $\tilde{X}$, we say that a continuous map $p : \tilde{X} \to X$ is a *covering map* if for any $x \in X$ there exists an open neighbourhood U of x such that $p^{-1}(U)$ is a disjoint union of (arbitrarily many) open sets, each of which is mapped homeomorphically by p onto U. If there exists such a p, we say that X is (evenly) *covered* by $\tilde{X}$ and that p is an (even) *covering* of X.

Given two coverings $p : \tilde{X} \to X$ and $p' : Y \to X$, a homeomorphism $f : \tilde{X} \to Y$ is said to be an *isomorphism* of coverings if $p' \circ f = p$, i.e., if the following diagram commutes:

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

Example 2.2. We obtain a *trivial cover* by taking n discrete copies of a topological space X. Formally speaking, our covering space is $\tilde{X} = X \times \{1, ..., n\}$ equipped with the product topology, where $\{1, ..., n\}$ is equipped with the discrete topology. Then our covering map is given by $p(x, n) = x$, and every open set is evenly covered. Here is a trivial two-sheeted covering of the circle S^1:

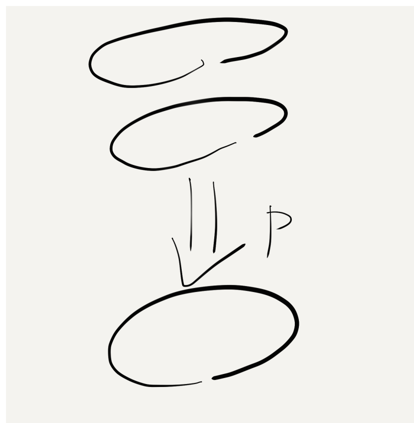

Example 2.3. Consider S^1 as the unit circle in $\mathbb{C}$. Then $p : S^1 \to S^1$ given by $p(z) = z^n$ is a covering map for any $n \in \mathbb{N}$. The following illustration, which is a modification of an image from [2], is a depiction of this map when $n = 3$:

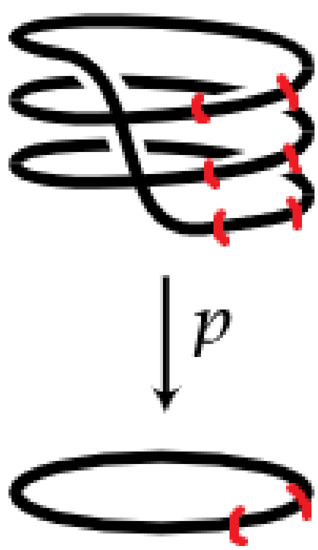

However, this map cannot exactly be depicted in three-dimensional space: one must imagine that the covering space in the picture does not intersect itself. In the picture, I have indicated an open subset of S^1 and its preimage by p.

Definition 2.4. A topological space X is *connected* if it cannot be written as the union of two or more disjoint non-empty open subsets. It is *path-connected* if for any $x, y \in X$ there exists a path, i.e., a continuous map $f : [a, b] \to X$, such that $f(a) = x$ and $f(b) = y$. (We will often use paths with $a = 0$, $b = 1$.)

Lemma 2.5. [1, Lemma 11.5]. *Suppose $p : \tilde{X} \to X$ is a covering, Y is a connected topological space, and $f, g : Y \to \tilde{X}$ are continuous maps such that $p \circ f = p \circ g$. If there exists $y \in Y$ such that $f(y) = g(y)$, then $f = g$.*

Proof. Recall the well-known result that a topological space is connected if and only if it only has two clopen subsets, namely the whole space and the empty set. Applying this to Y, it suffices to show that the set $Z := \{y \in Y | f(y) = g(y)\}$ is open and has open complement. (It is non-empty by assumption.)

Suppose $z \in Z$ and take a neighbourhood N of $p(f(z)) = p(g(z)) \in X$ that is evenly covered by p. Then $p^{-1}(N)$ is a union of disjoint open sets N_i for $i \in I$ where I is some indexing set and each N_i is mapped homeomorphically to N by p. Since f and g are continuous, they must map a neighbourhood U of z into the same N_i. Moreover, since $p \circ f = p \circ g$, f and g must agree on U. We have chosen an arbitrary point $z \in Z$ and shown that it must have an open neighbourhood also contained in Z, so Z is open.

Now suppose $z \notin Z$, i.e., $z \in Z^{\complement}$. By continuity of f and g, we may find a neighbourhood V of z such that $f(V) \cap g(V) = \emptyset$. Thus $V \subseteq Z^{\complement}$, and $Z^{\complement}$ is open as well. □

The following lemma is sufficiently well-known that no proof will be given here, although we have cited one.

Lemma 2.6. (Lebesgue) [1, Lemma A.19]. *Given any covering of a compact metric space K by open sets, there is some $\epsilon > 0$ such that any subset of K of diameter less than ϵ is contained in some open set from the covering.*

The next proposition allows us to lift paths from a space to paths in its covering space.

Proposition 2.7. [1, Proposition 11.6]. *Let $p : \tilde{X} \to X$ be a covering, let $\gamma : [a, b] \to X$ be a continuous path, and let $y \in \tilde{X}$ be such that $p(y) = \gamma(a)$. Then there is a unique continuous lift $\tilde{\gamma} : [a, b] \to \tilde{X}$ such that $\tilde{\gamma}(a) = y$ and $p(\tilde{\gamma}(t)) = \gamma(t)$ for all $t \in [a, b]$.*

Proof. First, note that if such a $\tilde{\gamma}$ exists, then it is necessarily unique because if there exists a continuous $\gamma' : [a, b] \to \tilde{X}$ satisfying $\gamma'(a) = y = \tilde{\gamma}(a)$ and $(p \circ \gamma')(t) = \gamma(t) = (p \circ \tilde{\gamma})(t)$ for all $t \in [a, b]$, then by Lemma 2.5 and the fact that $[a, b]$ is connected, we have $\gamma' = \tilde{\gamma}$.

If p is a trivial covering, then there is a unique component of $\tilde{X}$ containing y and mapping homeomorphically to the component of X containing $\gamma([a, b])$; in this case, we may simply use the inverse homeomorphism to lift the path.

In the general case, apply Lemma 2.6 to the open sets $\gamma^{-1}(N \cap \gamma([a, b]))$ where N varies over open sets in X that are evenly covered by p; the collection of open subsets of X that are evenly covered by p is an open cover of X because every point of X has an open neighbourhood evenly covered by p (by definition of covering maps), so the collection of $\gamma^{-1}(N \cap \gamma([a, b]))$ is an open cover of the compact metric space $[a, b]$. (Note that we use $\gamma^{-1}(N \cap \gamma([a, b]))$

rather than $\gamma^{-1}(N)$; the latter is not well-defined because there may be points of N not in $\gamma([a, b])$.) As N varies, the open sets $\gamma^{-1}(N \cap \gamma([a, b]))$ produce an open cover of $[a, b]$, and by compactness of $[a, b]$ there exists a subdivision of $[a, b]$ into $a = t_0 < ... < t_n = b$ where each $\gamma([t_{i-1}, t_i])$ is contained in some open set that is evenly covered by p. From our solution of the trivial case we may lift the restriction $\gamma([t_0, t_1])$ to a path in $\tilde{X}$ from y to some point y_1. Similarly, we may lift $\gamma([t_1, t_2])$ to a path starting at y_1 and ending at some y_2. Proceeding in this way, we will be done lifting γ after n steps. □

We now present some basic topological terminology that will be useful for us later when we introduce the fundamental group.

Definition 2.8. A *homotopy* of paths in X is a family $f_t : [0, 1] \to X, t \in [0, 1]$ of paths satisfying the following two conditions:
(i) $f_t(0)$ and $f_t(1)$ are independent of t, i.e., the endpoints stay fixed; and
(ii) the map $F : [0, 1]^2 \to X$ given by $F(s, t) = f_t(s)$ is continuous.

When there exist two paths f_0 and f_1 in X and a homotopy that at $t = 0$ gives f_0 and at $t = 1$ gives f_1, we say f_0 and f_1 are *homotopic* and write $f_0 \simeq f_1$.

We want to be able to lift homotopies in a similar way to how we lift paths. We will prove a proposition allowing us to do so, but first we must prove the following lemma, known as the *pasting lemma*, which we will subsequently use.

Lemma 2.9. [4, Theorem 18.3]. *Let X be a topological space, and let $S \subseteq X$ be a subset such that $S = S_1 \cup S_2$ where S_1 and S_2 are closed in X. (In particular, this implies S is also closed in X.) Let Y be a topological space, and let $f : S_1 \to Y$, $g : S_2 \to Y$ be continuous. If $f(x) = g(x)$ for every $x \in S_1 \cap S_2$, then there exists a unique continuous function $h : S \to Y$ defined by $h|_{S_1} = f$ and $h|_{S_2} = g$.*

Proof. First, the map h is well-defined because f and g agree on $S_1 \cap S_2$. Let C be closed in Y. Then $h^{-1}(C) = f^{-1}(C) \cup g^{-1}(C)$ by definition of the restriction. By continuity of f and g, $f^{-1}(C)$ and $g^{-1}(C)$ are closed, so $h^{-1}(C)$ is closed. This proves h is continuous. □

Proposition 2.10. [1, Proposition 11.8]. *Suppose that $p : \tilde{X} \to X$ is a covering and $H : [a, b] \times [0, 1] \to X$ is a continuous map, i.e., H is a path homotopy in X. Suppose that $\gamma_0(t) = H(t, 0), a \leq t \leq b$ is the initial path and that $\tilde{\gamma_0}$ is a lifting of γ_0 to $\tilde{X}$. Then there is a unique lifting $\tilde{H}$ of H to $\tilde{X}$ whose initial path is $\tilde{\gamma_0}$. To be precise, there is a continuous map $\tilde{H} : [a, b] \times [0, 1] \to \tilde{X}$ such that $p \circ \tilde{H} = H$ and $\tilde{H}(t, 0) = \tilde{\gamma_0}(t)$ for all $t \in [a, b]$.*

Proof. Apply Lemma 2.6 to produce subdivisions $a = t_0 < t_1 < ... < t_n = b$ and $0 = 0 = s_0 < s_1 < ... < s_m = 1$ such that if $R_{ij} = [t_{i-1}, t_i] \times [s_{j-1}, s_j]$, then $H(R_{ij})$ is a subset of some evenly covered open set. Thus we can construct $\tilde{H}$ by first lifting the restriction of H to $R_{11}, R_{21}, ..., R_{n1}$, then lifting the restriction of H to $R_{12}, R_{22}, ..., R_{n2}$, and continuing this way until all of H is lifted to $\tilde{H}$. The resulting lift is unique because each $H(R_{ij})$ in the sequence (other than $H(R_{11})$) intersects some previous $H(R_{k\ell})$ in the sequence. (For example, $H(R_{22}) \cap H(R_{12}) \neq \emptyset$ because the sets intersect along their boundaries.) Therefore, the lift of $H(R_{11})$ determines the lifts of all the other $H(R_{ij})$ in succession by Lemma 2.9. □

Next, we prove that being homotopic is an equivalence relation.

Proposition 2.11. [2, Proposition 1.2]. *The relation between two paths of being homotopic in a fixed topological space is an equivalence relation.*

Proof. Reflexivity: Clear because we can use the constant homotopy $f_t = f$ for a given map f.

Symmetry: Clear because if $f_0 \simeq f_1$ by f_t, then $f_1 \simeq f_0$ by f_{1-t}.

Transitivity: Suppose that $f \simeq g$ by f_t and $g \simeq h$ by g_t. Set $h_t = f_{2t}$ for $0 \leq t \leq 1/2$ and $h_t = g_{2t-1}$ for $1/2 \leq t \leq 1$. At $t = 1/2$, we have $f_1 = g_0$, so $h_{1/2}$ is well-defined. Consider the associated map $H(x,t) = h_t(x)$. We have $H(x,t) = F(x, 2t)$ for $0 \leq t \leq 1/2$ and $H(x,t) = G(x, 2t-1)$ for $1/2 \leq t \leq 1$, where F, G are the maps associated to f_t and g_t, respectively. Thus H is continuous on $[0,1] \times [0,1/2]$ and $[0,1] \times [1/2,1]$, and as these sets are closed, we conclude that H is continuous on $[0,1]^2$ by Lemma 2.9. □

Definition 2.12. By Proposition 2.11, we may define the equivalence class of a path $f : [0,1] \to X$ under homotopy and denote it by $[f]$.

Given a path f, define a *reparametrization* of f to be a composition fs where $s : [0,1] \to [0,1]$ is continuous and satisfies $s(0) = 0, s(1) = 1$. Note that $fs \simeq f$ by the homotopy fg_t where $g_t(x) = (1-t)s(x) + tx$.

If we are given paths $f, g : [0,1] \to X$ such that $f(1) = g(0)$, we can define their *path composition* (or *composition* for short, where context makes it clear we are not talking about function composition) $f \cdot g$ as follows. For $0 \leq x \leq 1/2$, we set $(f \cdot g)(x) = f(2x)$ and for $1/2 \leq x \leq 1$, we set $(f \cdot g)(x) = g(2x-1)$. It is clear these agree at $x = 1/2$. Moreover, if $F(x,t)$ is a homotopy $f_0 \simeq f_1$ and $G(x,t)$ is a homotopy $g_0 \simeq g_1$, then $H(x,t) := G(F(x,t),t)$ is a homotopy $g_0 \circ f_0 \simeq g_1 \circ f_1$, so path composition also preserves homotopy classes.

3. Group Actions and G-Coverings

In this section, we follow [1, §11c]. Recall the following definitions from elementary group theory.

Definition 3.1. A *(left) group action* of a group G on a topological space Y is a mapping $G \times Y \to Y$ given by $(g, y) \mapsto g.y$ and satisfying the following:

(i) $g.(h.y) = (gh).y$ for all $g, h \in G$ and $y \in Y$, where gh denotes the multiplication of g with h in G;

(ii) $e.y = y$ for all $y \in Y$, where e is the identity in G; and

(iii) $y \mapsto g.y$ is a homeomorphism of Y for all $g \in G$.

Remark 3.2. Equivalently, a group action can be thought of in the following way. Given a topological space X, let $\mathrm{Hom}(X)$ be the set of homeomorphisms $f : X \to X$. Since the inverse of a homeomorphism and the composition of two homeomorphisms are homeomorphisms and since the identity is a homeomorphism, $\mathrm{Hom}(X)$ is a group under the operation of composition. An action of the group G on X is then a homomorphism $\sigma : G \to \mathrm{Hom}(X)$.

Indeed, for any element $g \in G$, we then have a homeomorphism $f_g : X \to X$. We replace the notation $g.x$ with $f_g(x)$. The fact that f_g is a homeomorphism means condition (iii) of Definition 3.1 is satisfied. The fact that f_e (where e is the identity of G) is the identity map $f_e(x) = x$ means that condition (ii) is satisfied. Finally, letting $\sigma : G \to \text{Hom}(X)$ be the homomorphism, condition (i) follows from

$$f_g(f_h(x)) = (f_g \circ f_h)(x) = (\sigma(g) \circ \sigma(h))(x) = \sigma(gh)(x) = f_{gh}(x).$$

The key point is that in the category of topological spaces, $\text{Hom}(X)$ is the automorphism group of the object X, i.e., the collection of morphisms from X to itself. Similarly, an action of a group G on an object X in some other category (for example, the category of sets) is a group homomorphism $G \to \text{Aut}(X)$. (For example, if X is a finite set, then $\text{Aut}(X)$ is the symmetric group on X.) This observation allows us to define group actions in other categories.

Definition 3.3. Given $y \in Y$ and an action of G on Y, we say the *orbit* of y is the set $G.y := \{gy \in Y | g \in G\}$.

Definition 3.4. An action of a group G on a space X is *free* if for any $x \in X$ $gx = x$ implies g is the identity.

We might suspect that being in the same orbit is an equivalence relation. This is the content of the next proposition, which is stated but not proved in [1, §11c].

Proposition 3.5. *Define the relation R on Y by yRy' if y' is in the orbit $G.y$. Then R is an equivalence relation.*

Proof. Clearly R is reflexive. It is symmetric because if $x, y \in Y$ are in the same orbit, then $x = g.y$, so

$$\begin{aligned} g^{-1}.x &= g^{-1}.(g.y) \\ &= (g^{-1}g).y \\ &= e.y = y. \end{aligned}$$

We have used the existence of inverse elements in groups here. R is also transitive, since if $x, y, z \in Y$ are such that $x = g_1.y$ and $y = g_2.z$, then $x = g_1.(g_2.z) = (g_1g_2).z$, and of course $g_1g_2 \in G$. Therefore, R is an equivalence relation. □

Remark 3.6. In light of Proposition 3.5, we can write $[g]$ for the equivalence class of an element corresponding to its orbit. This is not to be confused with the similar notation from Definition 2.12.

Definition 3.7. Let G act on Y. Then we define the *space of orbits*

$$X = G \backslash Y := \{G.y | y \in Y\}$$

with the quotient topology from the map $p : Y \to X$ that sends a point to the orbit containing it. We may equip X with the quotient topology, so that $U \subseteq X$ is open when $p^{-1}(U)$ is open in Y.

Note that we have not yet said the map $p : Y \to X$ is a covering map. We will come to this point later.

Definition 3.8. We say that G acts *evenly* on Y if for any $y \in Y$ there is a neighbourhood N of y such that $g.N$ and $h.N$ are disjoint for any distinct $g, h \in G$.

Remark 3.9. In calling this an "even" action, we are using Fulton's terminology. The more common terminology is "properly discontinuous". As Fulton notes, this terminology is used inconsistently in the literature; sometimes it is less restrictive, requiring only that each point has a neighbourhood such that at most finitely many translates of that neighbourhood can intersect. In such cases, the correct terminology for what we call an "even" action is a "free and properly discontinuous" action. To avoid misleading a reader who is used to one definition or the other of a "properly discontinuous" action, we have chosen to stick with Fulton's word for it.

Lemma 3.10. [1, Lemma 11.17]. *If G acts evenly on Y, then the projection $p : Y \to G\backslash Y$ is a covering map.*

Proof. First, notice that p is continuous because if U is open in $G\backslash Y$, then $p^{-1}(U)$ is open in Y by definition of the quotient topology, so p is continuous. Next, let us show p is open. By definition, again, of the quotient topology, it suffices to show that $p^{-1} \circ p$ is open, i.e., that if U is open in Y, then $p^{-1}(p(U))$ is open in Y. But we have that

$$p^{-1}(p(U)) = \bigcup_{g \in G} g.U,$$

where every $g.U$ is open since g acts homeomorphically on Y, so $p^{-1}(p(U))$ is open and p is an open map. This will help us prove that p is a covering map.

Moreover, as G acts evenly on Y, for any $y \in Y$ we can find a neighbourhood U of y such that $\bigcup_{g \in G} g.U$ is a disjoint union. In particular, if we pick some representative $y \in Y$ so that $[y] \in G\backslash Y$ and choose such a neighbourhood U of y, we have that $p(U)$ is a neighbourhood of $[y]$, that $p^{-1}(p(U)) = \bigsqcup_{g \in G} g.U$, and that $p|_{g.U}$ is continuous, open, and surjective for each $g \in G$. It only remains to check that $p|_{g.U}$ is injective, and then we will have that p is a covering map.

Suppose we take $y_1, y_2 \in U$ such that $p(gy_1) = p(gy_2)$. Then there exists $h \in G$ such that $h.g.y_1 = g.y_2$. Because G acts evenly on Y it follows that $h = e$, so $y_1 = y_2$. Therefore, p is a covering map. □

If a covering map is induced from an even group action as in Lemma 3.10, we call it a *G-covering*. We restate this more precisely.

Definition 3.11. Suppose $p : Y \to G\backslash Y$ is a covering map, where Y is a topological space on which the group G acts evenly. Then p is called a *G-covering*. If $p : Y \to X$, $p' : Y' \to X$ are G-coverings and $\varphi : Y \to Y'$ is a homeomorphism such that $p' \circ \varphi = p$ and $\varphi(g.y) = g.\varphi(y)$ for all $y \in Y, g \in G$, then we say that φ is an *isomorphism* of G-coverings and that the G-coverings Y and Y' are *isomorphic*.

4. The Fundamental Group

Our next definition introduces one of the most important concepts in algebraic topology.

Definition 4.1. A *loop* is a path $f : [0, 1] \to X$ such that $f(0) = f(1) = x_0$, $x_0 \in X$. We call x_0 the *basepoint* of the loop. We denote the set of homotopy classes of loops at the basepoint x_0 by $\pi_1(X, x_0)$. We call $\pi_1(X, x_0)$ the *fundamental group* of X.

The previous definition seemingly makes two assumptions. The first is that $\pi_1(X, x_0)$ is a group, and the second is that $\pi_1(X, x_0)$ does not depend, as a group, on the choice of x_0. As it turns out, $\pi_1(X, x_0)$ is always a group, and if X is path-connected, $\pi_1(X, x_0)$ will not depend on the choice of x_0 (up to group isomorphism). This is the content of the next two propositions.

Proposition 4.2. [2, Proposition 1.3]. *$\pi_1(X, x_0)$ forms a group with respect to the product $[f][g] = [f \cdot g]$.*

Proof. Fix $x_0 \in X$. We already know that the product of two loops f and g based at x_0 exists and that the product $[f][g] = [f \cdot g]$ is well-defined. We must check the three conditions for a set to be a group under a binary operation.

Existence of identity element: Let c be the *constant loop* at x_0, i.e., $c(x) = 0$ for all $x \in [0, 1]$. Then $f \cdot c \simeq f$ by the reparametrization s_1 given by $s_1(x) = 2x$ for $x \in [0, 1/2]$ and $s_2(x) = 1$ for $x \in [1/2, 1]$. Similarly, $c \cdot f \simeq f$ by the reparametrization $s_2(x) = 0$ for $x \in [0, 1/2]$ and $s(x) = 2x - 1$ for $x \in [1/2, 1]$. Then $[c]$ is the identity in $\pi_1(X, x_0)$.

Associativity of the group operation: Given loops f, g, h based at x_0, we have $(f \cdot g) \cdot h \simeq f \cdot (g \cdot h)$ by the reparametrization s given by $s(x) = (1/2)x$ for $x \in [0, 1/2]$, $s(x) = x - 1/4$ for $x \in [1/2, 3/4]$, and $s(x) = 2x - 1$ for $x \in [3/4, 1]$.

Existence of inverses: Let $\overline{f}(x) = f(1 - x)$. Then $f \cdot \overline{f} \simeq c$, where c is the constant loop defined earlier by the homotopy $g_t h_t$, where g_t equals f on $[0, 1 - t]$ and stays at $f(1 - t)$ on $[1 - t, 1]$ and where h_t is, for each value of t, the reverse of g_t. For example, $g_0(x) = f(x)$ for all x, $g_1(x) = f(0) = x_0$ for all x, and $g_{1/2}$ is given by $g_{1/2}(x) = f(x)$ for $x \in [0, 1/2]$ and $g_{1/2}(x) = f(1/2)$ for $x \in [1/2, 1]$. Thus it is clear that $g_t h_t$ is a homotopy from $f \cdot \overline{f}$ to $c \cdot \overline{c} = c$, so that we indeed have $f \cdot \overline{f} \simeq c$. Switching f and $\overline{f}$ in the above gives the result that $\overline{f} \cdot f \simeq c$ as well. We have shown that $[f]^{-1} = [f^{-1}]$. □

Notice that if $x_0, x_1 \in X$ lie in the same path component of X, i.e., have a path h between them, then, given a loop f based at x_1, $h \cdot f \cdot \overline{h}$ is a loop based at x_0. (Notice that the distinction between the two orders in which this product may be formed, $h \cdot (f \cdot \overline{h})$ and $(h \cdot f) \cdot \overline{h}$, is moot here as we are only working up to homotopy now, and the two resulting paths are homotopic.) We will now show that this relationship between loops at x_0 and x_1 gives an isomorphism between $\pi_1(X, x_0)$ and $\pi_1(X, x_1)$, so that the fundamental group does not depend on choice of basepoint for path-connected spaces.

Proposition 4.3. [2, Proposition 1.5]. *If X is path-connected, the map $\beta_h : \pi_1(X, x_1) \to \pi_1(X, x_0)$ given by $\beta_h([f]) = [h \cdot f \cdot \overline{h}]$ is an isomorphism.*

Proof. The assumption that X is path-connected is needed so that the required path h between x_0 and x_1 exists. Now, given two loops based at x_1 and a homotopy f_t between them, $h \cdot f_t \cdot \overline{h}$ is a homotopy of loops based at x_0. Thus β_h is well-defined. Moreover, we calculate

$$\begin{aligned}\beta_h[f \cdot g] &= [h \cdot f \cdot g \cdot \overline{h}] \\ &= [h \cdot f \cdot \overline{h} \cdot h \cdot g \cdot \overline{h}] \\ &= [h \cdot f \cdot \overline{h}][h \cdot g \cdot \overline{h}] \\ &= \beta_h[f]\beta_h[g],\end{aligned}$$

so β_h is a homomorphism. Finally,

$$\begin{aligned}\beta_h\beta_{\overline{h}}[f] &= \beta_h[\overline{h} \cdot f \cdot h] \\ &= [h \cdot \overline{h} \cdot f \cdot h \cdot \overline{h}] \\ &= [f],\end{aligned}$$

and similarly $\beta_{\overline{h}}\beta_h[f] = [f]$. Because $\beta_{\overline{h}}$ is also a homomorphism, we conclude that β_h is an isomorphism with inverse $\beta_{\overline{h}}$, and therefore that $\pi_1(X, x_1)$ and $\pi_1(X, x_0)$ are isomorphic groups. □

We now calculate the fundamental group of S^1 from first principles to demonstrate the inconvenience of doing so.

Theorem 4.4. [2, Theorem 1.7]. *The map $\phi : \mathbb{Z} \to \pi_1(S^1)$ given by $\phi(n) = [\omega_n(s)]$, where*

$$\omega_n(s) = (\cos 2\pi ns, \sin 2\pi ns)$$

is a loop based at $(1, 0)$, is an isomorphism. That is, $\pi_1(S^1) \simeq \mathbb{Z}$.

Proof. Notice that we have a projection $p : \mathbb{R} \to S^1$ given by $p(s) = (\cos 2\pi s, \sin 2\pi s)$. We can think of this as a helix being projected down onto a circle:

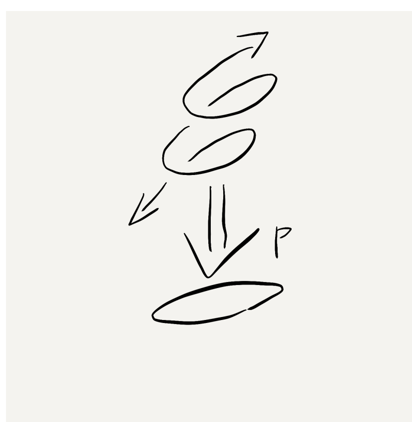

Notice that ω_n may be obtained as the composition $p\tilde{\omega}_n$, where $\tilde{\omega}_n : [0, 1] \to \mathbb{R}$ is the path $\tilde{\omega}_n(s) = ns$, $s \in [0, 1]$.

We may think of ϕ slightly differently if we set $\phi(n)$ to be the homotopy class of $p\tilde{f}$ where $\tilde{f}$ is any path in $\mathbb{R}$ from 0 to n, as such an $\tilde{f}$ is homotopic to $\tilde{\omega}_n$ by the homotopy $(1-t)\tilde{f}+t\tilde{\omega}_n$, $t \in [0, 1]$, so $p\tilde{f} \simeq p\tilde{\omega}_n = \omega_n$, which agrees with our first definition of ϕ. This reformulation simplifies our subsequent calculations slightly.

Now suppose we have a loop $f : [0,1] \to S^1$ with $f(0) = f(1) = x_0$. By Proposition 2.7 we can lift this to a path $\tilde{f}$ starting at 0, as $0 \in p^{-1}(x_0)$. Moreover, $\tilde{f}(1)$ must be some integer n because $p\tilde{f}(1) = f(1) = x_0$ and $p^{-1}(x_0) = \mathbb{Z}$. But the path $\tilde{\omega}_n$ that we previously defined also goes from 0 to n, and $\tilde{f} \simeq \tilde{\omega}_n$ by the homotopy $(1-t)\tilde{f} + t\tilde{\omega}_n$, so we may compose this homotopy with p to get $[f] = [\omega_n]$.

It remains to prove that the value of n is uniquely determined by $[f]$. Suppose $f \simeq \omega_n$ and $f \simeq \omega_m$ for $m, n \in \mathbb{Z}$. We know $\omega_m \simeq \omega_n$, so choose a homotopy f_t achieving this. By Proposition 2.10, we can lift f_t to a homotopy $\tilde{f}_t$ in $\mathbb{R}$. By the uniqueness part of Proposition 2.7, we have $\tilde{f}_0 = \tilde{\omega}_m$, $\tilde{f}_1 = \tilde{\omega}_n$. But $\tilde{f}_t(1)$ is independent of $t \in [0,1]$, so $\tilde{f}_0(1) = m = \tilde{f}_1(1) = n$, which was what we wanted. □

Clearly the method used above is cumbersome to generalize to other more complicated topological spaces. We should therefore seek a more general way to calculate fundamental groups. This is the purpose of the Seifert–van Kampen Theorem, which we will encounter in §8.

5. Universal Coverings

The following definitions establish the topic of this section.

Definition 5.1. A topological space X is *simply connected* if it is path-connected and for each $x \in X$, $\pi_1(X, x)$ is trivial.

Definition 5.2. A cover $\tilde{X}$ of a topological space X is *universal* if it is simply connected.

A universal cover exists for a topological space X under some fairly weak conditions.[1] However, we must build up some more covering space theory before we can prove this.

Given a covering map $p : (\tilde{X}, \tilde{x}_0) \to (X, x_0)$ (i.e., a covering map $p : \tilde{X} \to X$ such that $p(\tilde{x}_0) = x_0$), we may speak of the induced homomorphism $p_* : \pi_1(\tilde{X}, \tilde{x}_0) \to \pi_1(X, x_0)$. This homomorphism is defined by sending the equivalence class of a loop based at $\tilde{x}_0$ to the equivalence class of its image under p, which will be a loop based at x_0. We then have the following result.

Proposition 5.3. [2, Proposition 1.31]. *The induced homomorphism $p_* : \pi_1(\tilde{X}, \tilde{x}_0) \to \pi_1(X, x_0)$ is injective, and the subgroup $p_*(\pi_1(\tilde{X}, \tilde{x}_0))$ of $\pi_1(X, x_0)$ consists of homotopy classes of loops based at x_0 that remain loops when lifted to $\tilde{X}$.*

Proof. What does an arbitrary element in $\ker p_*$ look like? It can be represented by a loop $\tilde{f}$ whose image in X is homotopic to the trivial loop.Thus there is a lifted homotopy of loops $\tilde{f}_t$ that starts with $\tilde{f}$ and ends with the constant loop in $\tilde{X}$. So $[\tilde{f}] = [c]$ and p_* is injective.

Clearly loops at x_0 that lift to loops at $\tilde{x}_0$ are in $p_*(\pi_1(\tilde{X}, \tilde{x}_0))$. On the other hand, if a loop represents an element of $p_*(\pi_1(\tilde{X}, \tilde{x}_0))$, then it must be homotopic to a loop that lifts to a loop, so by Proposition 2.10 it must have such a lift itself. □

[1]Indeed, under these conditions, the universal cover is unique. However, we will not need this fact, so we will not prove it. For a reference, see the paragraph after [2, Theorem 1.38].

The following definition will be important in describing the conditions under which a topological space has a universal cover.

Definition 5.4. A topological space X is said to be *locally path-connected* if for every $x \in X$ and each neighbourhood U of x there exists an open path-connected neighbourhood $V \subset U$ of x. X is said to be *semilocally simply connected* if for each $x \in X$ there is a neighbourhood U of x such that every loop in U is homotopic in X to the constant loop, or equivalently, the map $\pi_1(U, x) \to \pi_1(X, x)$ induced by the inclusion $U \subseteq X$ is trivial.

The next example gives a space with no universal cover, as will follow from Theorem 5.7.

Example 5.5. Consider the *Hawaiian earring*, a topological space X given by the following construction. For each $n \in \mathbb{N}$, let C_n be the circle of radius $1/2^n$ in $\mathbb{R}^2$ centered at $(1/2^n, 0)$, and let $X = \bigcup_{n\in\mathbb{N}} C_n$. Here is a picture:

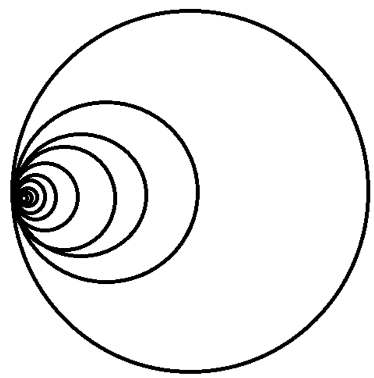

Then X is clearly connected and locally path-connected. However, it is not semilocally simply connected. Indeed, if we take an ϵ-neighbourhood of the origin for some $\epsilon > 0$, we can choose n large enough that $\frac{1}{2^n} < \epsilon$. Thus the neighbourhood will contain C_k for all $k \geq n$, and we may choose a loop based at the origin that winds around one of these C_k once counterclockwise. It is clear that this loop is not homotopic to the constant loop in X.

We will need the following lemma.

Lemma 5.6. [5]. *A topological space that is connected and locally path-connected is path-connected.*

Proof. Let X be such a space. Let $p \in X$. Let $C \subseteq X$ be the set of points that can be joined by a path to p. C is non-empty because X is locally path-connected; indeed, $p \in C$. We will show that C is both closed and open, which will prove $C = X$.

Let $c \in C$, and let U be an open path-connected neighbourhood of c. If $u \in U$, we can give a path from u to c, then join that path to a path from p to c. Therefore, $u \in C$, so C is open.

Now let $c \in \overline{C}$, the closure of C, and choose an open path-connected neighbourhood U of c again. The set $C \cap U$ is non-empty, so choose $q \in C \cap U$. Join c to q by a path, and join q to p by a path. Therefore, p can be joined to c by a path, so $c \in C$. This proves C is closed. Therefore, $C = X$, which gives the result. $\square$

We are almost ready to prove the existence of universal coverings under certain conditions. First, because in algebraic topology we are often concerned with having paths in our spaces, we will assume our topological space is locally path-connected. Next, we will assume it is connected because otherwise the results we prove will reduce to results about the connected components of our space.

We have proved that the induced map p_* is injective and it remains to prove it is surjective. In particular, we might ask whether the trivial subgroup of $\pi_1(X, x_0)$ can be obtained as $p_*(\pi_1(\tilde{X}, \tilde{x}_0))$ for some covering space $\tilde{X}$. By the aforementioned injectivity of p_*, it suffices to prove the existence of a simply connected covering space of X, that is to say, a universal cover.

For X to have a universal cover, it must be semilocally simply connected. Indeed, suppose $p : \tilde{X} \to X$ is a universal cover. For each $x \in X$, there exists an open neighbourhood U containing x such that $p^{-1}(U)$ is a union of open sets homeomorphic to U. Given a loop in U, we can lift it to a loop in $p^{-1}(U)$, and because $\pi_1(\tilde{X}) = 0$, the lifted loop must be homotopic to the constant map. But composing this homotopy with p, we get that the original loop is also homotopic to the constant map in X. We conclude that each point $x \in X$ has a neighbourhood U such that the map $\pi_1(U, x) \to \pi_1(X, x)$ induced by inclusion is trivial. This is precisely the statement that X is semilocally simply connected.

We finally have the following theorem.

Theorem 5.7. [1, Theorem 13.20]. *If X is connected and locally path-connected, it has a universal cover if and only if it is semilocally simply connected.*

Proof. In the paragraph preceding the theorem, we showed that if X is connected, locally path-connected, and has a universal cover, then it must be semilocally simply connected. It remains to show the converse.

Suppose X is connected, locally path-connected, and semilocally simply connected. We give an explicit construction of its universal cover. Choose $x \in X$ and define $\tilde{X}$ to be the space of homotopy classes of paths $\gamma : [0, 1] \to X$ such that $\gamma(0) = x$. In a few paragraphs, we will define a topology on $\tilde{X}$ (given by sets of the form $U_{[\gamma]}$, also soon to be defined), but first we require some more exposition. Define a map $p : \tilde{X} \to X$ by $p([\gamma]) = \gamma(1)$. Recall that we require homotopies to fix endpoints, so p is well-defined. Since X is connected and locally path-connected, it is path-connected by Lemma 5.6, so p is surjective.

Define S to be the collection of path-connected open subsets $U \subseteq X$ such that the inclusion-induced map $\pi_1(U) \to \pi_1(X)$ is trivial. Note that whether this map is trivial or not does not depend on choice of basepoint because X is path-connected. Moreover, given a path-connected open subset $V \subseteq U$, we have that the map $\pi_1(V) \to \pi_1(U) \to \pi_1(X)$ is also trivial, so $V \in S$ as well. Because X is semilocally simply connected, every point has a neighbourhood in S, and because X is locally path-connected, arbitrarily small open neighbourhoods of $x \in X$ will also be in S; more precisely, given any neighbourhood W of x, we can find an open neighbourhood $W' \subseteq W$ containing x such that $W' \in S$. It follows that S is a basis for the topology on X.

Now suppose we are given a set $U \in S$ and a path $\gamma : [0,1] \to X$ such that $\gamma(0) = x_0$ and $\gamma(1) \in U$. Define

$$U_{[\gamma]} := \{[\gamma \cdot \sigma] | \sigma : [0,1] \to U \text{ is a path and } \sigma(0) = \gamma(1)\} \subset \tilde{X}.$$

Notice that $p : U_{[\gamma]} \to U$ is surjective because U is path-connected and injective because that X is semilocally simply connected implies that any two different choices of σ from $\gamma(1)$ to a point in U are homotopic in X. Therefore, p is bijective on sets of the form $U_{[\gamma]}, U \in S$.

Moreover, if $[\gamma'] \in U_{[\gamma]}$, then $U_{[\gamma]} = U_{[\gamma']}$ because if $\gamma' = \gamma \cdot \sigma$, then elements of $U_{[\gamma']}$ look like $[\gamma \cdot \sigma \cdot \eta] \in U_{[\gamma]}$ and elements of $U_{[\gamma]}$ look like

$$\begin{aligned}[\gamma \cdot \eta] &= [\gamma \cdot \sigma \cdot \overline{\sigma} \cdot \eta] \\ &= [\gamma' \cdot \overline{\sigma} \cdot \eta] \in U_{[\gamma']}.\end{aligned}$$

Suppose we are given open sets $U, V \in S$ and paths $\gamma, \gamma' : [0,1] \to X$ such that $\gamma(0) = \gamma'(0) = x_0$ and $\gamma(1) \in U, \gamma'(1) \in V$. Then we can consider $U_{[\gamma]}, V_{[\gamma']}$. Given $[\gamma''] \in U_{[\gamma]} \cap V_{[\gamma']}$, we have $U_{[\gamma]} = U_{[\gamma'']}$, $V_{[\gamma']} = V_{[\gamma'']}$ by the previous paragraph. Thus if $W \in S$, $W \subseteq U \cap V$, and $\gamma''(1) \in W$, then $W_{[\gamma'']} \subseteq U_{[\gamma'']} \cap V_{[\gamma'']}$ and $[\gamma''] \in W_{[\gamma'']}$. Therefore, sets of the form $U_{[\gamma]}$, $U \in S$, form a basis for $\tilde{X}$ because sets of the form $U \in S$ form a basis for X.

Suppose $[\gamma'] \in U_{[\gamma]}$, γ' has its other endpoint in V, and $V_{[\gamma']} \subseteq U_{[\gamma']} = U_{[\gamma]}$. Then $p(V_{[\gamma']}) = V$ and $p^{-1}(V) \cap U_{[\gamma]} = V_{[\gamma']}$, so $p : U_{[\gamma]} \to U$ gives a bijection between subsets $V_{[\gamma']} \subseteq U_{[\gamma]}$ and sets $V \in S$ such that $V \subseteq U$. Thus $p : U_{[\gamma]} \to U$ is a homeomorphism, and since sets of the form $U_{[\gamma]}$ form a basis for $\tilde{X}$, $p : \tilde{X} \to X$ is continuous.

Moreover, if we fix $U \in S$, we have that sets of the form $U_{[\gamma]}$ partition $p^{-1}(U)$ because $[\gamma''] \in U_{[\gamma]} \cap U_{[\gamma']}$ implies $U_{[\gamma]} = U_{[\gamma'']} = U_{[\gamma']}$. Thus $\tilde{X}$ is indeed a covering space of X.

It is only left to show that $\tilde{X}$ is simply connected. Suppose $[\gamma] \in \tilde{X}$. We want to show that $\tilde{X}$ is path-connected. We also want to show that $[\gamma] = [x_0]$, as this will imply that $p_*(\pi_1(\tilde{X}, \tilde{x}_0))$ is trivial, and thus, by injectivity of p_*, that $\pi_1(\tilde{X}, \tilde{x}_0)$ is trivial.

Let γ_t be the path in X that equals γ on $[0,t]$ and equals $\gamma(t)$ on $[t,1]$. Then $t \mapsto [\gamma_t]$ is a lift of γ in $\tilde{X}$ that starts at $[x_0]$ and ends at $[\gamma]$. Therefore, $\tilde{X}$ is path-connected.

Now, given an element in the image of p_*, we can represent it by a loop γ at x_0 that lifts to a loop in $\tilde{X}$ at $[x_0]$. We already gave a path $t \mapsto [\gamma_t]$ which lifts γ starting at $[x_0]$. If this path is a loop, then $[\gamma_1] = [x_0]$. But $\gamma_1 = \gamma$, so $[\gamma] = [x_0]$, which was what we wanted. Therefore, $\tilde{X}$ is a universal cover of X. □

6. Deck Transformations

Suppose X is a topological space and $p : Y \to X$ is a covering where Y is path-connected. In this section, we will define an action of the fundamental group $\pi_1(X, x)$ on Y. The fundamental group will be said to act by *deck transformations* on Y. This will allow us to construct G-coverings from the universal covering of a topological space in the next section, as the (somewhat more complicated) action defined there will rely in the first coordinate on the deck transformation action of this section.

Example 6.1. We have seen that if $x \in S^1$, then $\pi_1(S^1, x) \simeq \mathbb{Z}$. We have also seen that $\mathbb{R}$ is a covering space of S^1 and that we can visualize of it as a spiral lying above S^1. Imagine rotating this spiral counterclockwise or clockwise by 360 degrees; the fibre $p^{-1}(y)$ for any $y \in S^1$ will be sent to itself by this rotation. Indeed, repeating such a rotation n times will still preserve the fibres of the covering. If we think of rotating clockwise n times as rotating counterclockwise $-n$ times, this defines an action of $\mathbb{Z}$ on the covering space $\mathbb{R}$. This is an example of a deck transformation.

Definition 6.2. Recall that we were considering a covering $p : Y \to X$, Y path-connected. Suppose $[\sigma] \in \pi_1(X, x)$ and $z \in Y$. We want to define $[\sigma].z$. We do it as follows.

Let $y \in p^{-1}(x)$. Lift $[\sigma]$ to a path in Y going from y to y'. Let γ be a path from y to z in Y. (This is where we use path-connectedness of Y.) Now the path $p \circ \gamma$ starts at x (since $p(y) = x$) and ends at $p(z)$, but by definition of y', we have $p(y') = p(y) = x$. We can therefore lift $p \circ \gamma$ to a map starting at y' instead, and we can define $[\sigma].z := w(z, \sigma, \gamma)$, where $w(z, \sigma, \gamma)$ is the endpoint of this lift of $p \circ \gamma$.

Lemma 6.3. [1, §13b]. *If we have another path γ' that goes from y to z in Y and if $p_*(\pi_1(Y, y))$ is a normal subgroup of $\pi_1(X, x)$, then w is well-defined, i.e., $w(z, \sigma, \gamma) = w(z, \sigma, \gamma')$.*

Proof. We want the lifts of $\sigma \circ (p \circ \gamma')$ and $\sigma \circ (p \circ \gamma)$ that start at y to end at the same point. By Proposition 5.3, this is the case precisely when the class $[(\sigma \circ (p \circ \gamma'))(\sigma \circ (p \circ \gamma))^{-1}]$ is an element of $p_*(\pi_1(Y, y))$. Now by definition of γ', the composition $\gamma' \circ \gamma^{-1}$ is a loop based at y, so we have

$$\begin{aligned}[(\sigma \circ (p \circ \gamma'))(\sigma \circ (p \circ \gamma))^{-1}] &= [(\sigma \circ (p \circ \gamma') \circ (p \circ \gamma^{-1}) \circ \sigma^{-1}] \\ &= [\sigma][p \circ \gamma' \circ \gamma^{-1}][\sigma]^{-1} \\ &= [\sigma]p_*([\gamma' \circ \gamma^{-1}])[\sigma]^{-1} \in [\sigma]p_*(\pi_1(Y, y))[\sigma]^{-1}.\end{aligned}$$

If $p_*(\pi_1(Y, y))$ is a normal subgroup of $\pi_1(X, x)$, then $[(\sigma \circ (p \circ \gamma'))(\sigma \circ (p \circ \gamma))^{-1}] \in p_*(\pi_1(Y, y))$ by the above calculation. We therefore assume in what follows that $p_*(\pi_1(Y, y))$ is normal. In that case our mapping $\pi_1(X, x) \times Y \to Y$ is well-defined. □

Lemma 6.4. [1, §13b]. *The mapping $\pi_1(X, x) \times Y \to Y$ defined above is a (continuous) group action.*

Proof. Suppose $[\sigma], [\tau] \in \pi_1(X, x)$. We want to show that $([\sigma][\tau]).z = [\sigma].([\tau].z)$. Suppose γ is a path from y to z. Lift $\tau \circ (p \circ \gamma)$ to a path starting at y; this path then ends at $[\tau].z$ by definition. Thus we see that $[\sigma]([\tau].z)$ is the endpoint of the lift, starting at y, of $\sigma \circ (\tau \circ (p \circ \gamma))$. But $(\sigma \circ \tau) \circ (p \circ \gamma)$ is homotopic to $\sigma \circ (\tau \circ (p \circ \gamma))$, so this endpoint is $([\sigma][\tau]).z$. The identity with respect to our action is just the deck transformation induced by the trivial element of $\pi_1(X, x)$, which fixes every point of Y.

It remains to prove that if $[\sigma]$ is fixed, then the map $z \mapsto [\sigma].z$ is continuous. To do this, we will need to assume that X is locally path-connected. (If you are beginning to lose track of our assumptions at this point, do not worry, because we will summarize everything with a nice theorem a few paragraphs from now.) Given $z \in Y$, take a path-connected and evenly-covered neighbourhood N of $p(z)$. (To see that such a neighbourhood exists, produce a path-connected neighbourhood of $p(z)$, which exists by local path-connectedness, and produce an evenly-covered neighbourhood of $p(z)$, which exists by definition of a covering projection, then take their intersection, which will have both properties.) Let V and V' be the components of $p^{-1}(N)$ containing z and $z' := [\sigma].z$, respectively. It suffices to show that $[\sigma].V \subseteq V'$. (To see this works, recall the following definition of continuity: f is continuous at x if for any neighbourhood $N'(f(x))$ there is a neighbourhood $N(x)$ such that $f(y) \in N'(f(x))$ whenever $y \in N(x)$.) Now suppose $v \in V$ and α is a path in V from z to v. Let γ be a path from y to z in Y. Then $\gamma \circ \alpha$ is a path from y to v. Moreover, $[\sigma].v$ is the endpoint of the lift of $p \circ \alpha$ starting at z', and this lift is in V' by definition. This shows $z \mapsto [\sigma].z$ is continuous. □

Let $\mathrm{Aut}(Y/X)$ denote the group of covering transformations, which is the automorphism in the category of coverings. (We know by Remark 3.2 that this is indeed a group.) We have thus far produced a homomorphism $\pi_1(X, x) \to \mathrm{Aut}(Y/X)$.

Lemma 6.5. [1, §13b]. *The homomorphism $\pi_1(X, x) \to Aut(Y/X)$ is surjective.*

Proof. Suppose $\varphi : Y \to Y$ is a covering transformation with $\varphi(y) = y'$. Since Y is assumed connected, a covering transformation is determined by the image of one point, so it suffices to produce $[\sigma] \in \pi_1(X, x)$ such that $[\sigma].y = y'$. But this is simple: let γ be a path from y to y' in Y, and let $\sigma = p \circ \gamma$. Then $[\sigma].y = y'$ by definition. □

Finally, to measure the failure of injectivity, we compute $\ker(\pi_1(X, x) \to \mathrm{Aut}(Y/X))$. Since Y is connected, it suffices to characterize the $[\sigma] \in \pi_1(X, x)$ such that $[\sigma].y = y$. But this happens precisely when the lift of σ to Y is a path from y to y, i.e., when $[\sigma] \in p_*(\pi_1(Y, y))$ by Proposition 5.3. We have proved the following theorem, which characterizes deck transformations.

Theorem 6.6. [1, Theorem 13.11]. *Suppose Y is connected, X is locally path-connected, and $p : Y \to X$ is a covering projection. Let $y \in Y$ with $p(y) = x$. Suppose $p_*(\pi_1(Y, y))$ is a normal subgroup of $\pi_1(X, x)$. Then there is an isomorphism*

$$\pi_1(X, x) \backslash p_*(\pi_1(Y, y)) \simeq Aut(Y/X).$$

7. Constructing G-Coverings from the Universal Covering

Let X be a connected, locally path-connected, semilocally simply connected space, so that X has a universal covering $\tilde{X}$. Let $u : \tilde{X} \to X$ be the corresponding covering map. We take all our spaces to have basepoints and all our maps to take basepoints to basepoints. We denote the basepoint of X by x and the basepoint of $\tilde{X}$ over x by $\tilde{x}$.

Proposition 7.1. [1, §14a]. *Given a homomorphism $\rho : \pi_1(X, x) \to G$, where G is an arbitrary group, we can construct a G-covering $p_\rho : Y_\rho \to X$ with basepoint $y_\rho \in Y_\rho$ over x, where Y_ρ is a topological space equipped with the quotient topology under an action of $\pi_1(X, x)$.*

Proof. First, we endow G with the discrete topology so that $\tilde{X} \times G$ is a product of copies of $\tilde{X}$ with one copy per element of G. Given $[\sigma] \in \pi_1(X, x)$, $[\sigma]$ acts on $\tilde{X} \times G$ by

$$[\sigma].(z, g) := ([\sigma].z, g\rho([\sigma])^{-1}), \quad (*)$$

where $(z, g) \in \tilde{X} \times G$ and $[\sigma].z$ denotes the action of $\pi_1(X, x)$ on $\tilde{X}$ described in the previous section. This is a common way of extending an action on the first coordinate to an action on the Cartesian product; the inverse in $\rho([\sigma])^{-1}$ is just there so that the action is associative.

Next, define

$$Y_\rho := \pi_1(X, x)\backslash(\tilde{X} \times G),$$

where the quotient is by the action, and let y_ρ be the image of $(\tilde{x}, e)$ in Y_ρ, where e denotes the group identity in G. Let $\langle z, g\rangle$ denote the image of (z, g) in Y_ρ. Note that, retaining the notation of Equation (*), we have

$$\langle [\sigma].z, g\rangle = \langle z, g\rho([\sigma])\rangle.$$

Indeed, $[\sigma]^{-1}\langle [\sigma].z, g\rangle = \langle z, \rho([\sigma])\rangle$, but since we quotient out by the action of $\pi_1(X, x)$ in Y_ρ, these two pairs are identical.

Define $p_\rho : Y_\rho \to X$ by $p_\rho(\langle z, g\rangle) := u(z)$. (We recall that u is the universal covering map.)

Now G acts on Y_ρ by

$$h.\langle z, g\rangle := \langle z, hg\rangle,$$

where $h, g \in G, z \in \tilde{X}$. This is an action because $h = e$ gives the identity, because the map $z \mapsto hz$ is continuous on G, and because

$$(h_1h_2).\langle z, g\rangle = \langle z, h_1h_2g\rangle$$

whereas

$$\begin{aligned} h_1.(h_2.\langle z, g\rangle) &= h_1.\langle z, h_2g\rangle \\ &= \langle z, h_1h_2g\rangle. \end{aligned}$$

We claim that this action is even, so that p_ρ is a G-covering. Indeed, let N be an evenly-covered open set in X (with respect to the covering u). Then $u^{-1}(N)$ is homeomorphic to the product covering $N \times \pi_1(X, x)$ (i.e., the covering map is $P : N \times p_1(X, x) \to N$ given by taking the first coordinate). Thus we have homeomorphisms

$$p_\rho^{-1}(N) \simeq (N \times \pi_1(X, x)) \times (\pi_1(X, x)\backslash G) \simeq N \times G,$$

where the last homeomorphism is obtained from the map $(N \times \pi_1(X, x)) \times (\pi_1(X, x)\backslash G) \to N \times G$ given by $\langle (u, [\sigma]), g\rangle \mapsto (u, g\rho([\sigma]))$ and the map $N \times G \to (N \times \pi_1(X, x)) \times (\pi_1(X, x)\backslash G)$ given by $(u, g) \mapsto \langle (u, e), g\rangle$, where e is the identity, as usual. It follows from the homeomorphism $p_\rho^{-1}(N) \simeq N \times G$ that G acts evenly on $p_\rho^{-1}(N)$. Since such open sets N cover X, p_ρ is a G-covering of X. □

Proposition 7.2. [1, §14a]. *Given a G-covering $p : Y \to X$ with $p(y) = x$, there exists a homomorphism $\rho : \pi_1(X, x) \to G$.*

Proof. Let $[\sigma] \in \pi_1(X, x)$ and define $\rho([\sigma]).y$ to be the endpoint of the lift of σ to y. (This is independent of the choice of representative σ of $[\sigma]$, so it is well-defined. Moreover, because the action is even, defining $\rho([\sigma]).y$ is enough to uniquely specify an image $\rho([\sigma])$ of $[\sigma]$ in G.) Indeed, let us adopt the convenient notation

$$y * \sigma$$

for the endpoint of the lift of σ starting at y. We then observe that if $z \in p^{-1}(x)$, σ is a loop based at x, and τ is a path starting at x, then

$$(z * \sigma) * \tau = z * (\sigma \cdot \tau), \quad (i)$$

which follows immediately from the definition of $*$. It is less obvious that if $g \in G$, $z \in p^{-1}(x)$, and γ is a path starting at x, then

$$g.(z * \gamma) = (g.z) * \gamma. \quad (ii)$$

We show this as follows. Suppose $\tilde{\gamma}$ is a lift of γ starting at z. Consider the path $t \mapsto g.\tilde{\gamma}(t)$, $0 \leq t \leq 1$. This is a lift of γ starting at $g.z$.[2] By definition of $*$, this path ends at $(g.z) * \gamma$, but by definition of $g.\tilde{\gamma}(t)$, it ends at $g.\tilde{\gamma}(1)$, so $g.\tilde{\gamma}(1) = (g.z) * \gamma$. But we have $\tilde{\gamma}(1) = z * \gamma$, so it follows that $g.(z * \gamma) = (g.z) * \gamma$.

It remains to prove that the map ρ is a homomorphism. We show this using Equations (i) and (ii). Indeed,

$$\begin{aligned}\rho([\sigma]).(\rho([\tau]).y) &= \rho([\sigma]).(y * \tau)\\ &= (\rho([\sigma]).y) * \tau\\ &= (y * \sigma) * \tau\\ &= y * (\sigma \circ \tau)\\ &= \rho([\sigma][\tau]).y.\end{aligned}$$

□

We are now ready to prove the correspondence between homomorphisms $\pi_1(X, x) \to G$ and G-coverings up to isomorphism.

Proposition 7.3. [1, Proposition 14.1]. *Suppose X is connected, locally path-connected, and semilocally simply connected.*[3] *Then there is a one-to-one correspondence between the set of homomorphisms from $\pi_1(X, x)$ to a group G and the set of based G-coverings, up to G-covering isomorphism.*

Proof. We want to show that if we are given a G-covering $p : Y \to X$ with basepoints $p(y) = x$, from which we can produce a homomorphism $\rho : \pi_1(X, x) \to G$ as above, then the given covering is isomorphic (as a G-covering) to the covering $p_\rho : Y_\rho \to X$, constructed from ρ as above. To do this, we must produce a map $Y_\rho \to Y$. It suffices to map $\tilde{X} \times G \to Y$ and show that elements in the same orbit of $\pi_1(X, x)$ have the same image. Recall from Theorem 5.7 that we can think of the universal cover $\tilde{X}$ as the space of homotopy classes of paths in X starting at x. Given such a class of paths $[\gamma] \in \tilde{X}$ and $g \in G$, we have a map $\tilde{X} \times G \to Y$ given by $([\gamma], g) \mapsto g.(y * \gamma) = (g.y) * \gamma$, where we recall that $y * \gamma$ means the endpoint of the lift of γ starting at y.

We must map Y_ρ to Y. Thus we must map $\tilde{X} \times G$ to Y and show that orbits by $\pi_1(X, x)$ have the same image. We know from Theorem 5.7 that the universal cover $\tilde{X}$ can be thought of as the space of homotopy classes of paths starting at $x \in X$. We define a map $\tilde{X} \times G \to Y$

[2] We are slightly abusing notation here, as G acts on X, not Y. By $g\tilde{\gamma}(t)$, we really mean the lift of $g\gamma(t)$, $0 \leq t \leq 1$, starting at $g.z$.

[3] We will need these assumptions to produce a universal cover of X.

by $([\gamma], g) \mapsto g.(y * \gamma) = (g.y) * \gamma$, where we recall that $y * \gamma$ is the endpoint of the lift of γ starting at y. The map $X \times G \to Y$ factors into two maps $X \times G \to Y \times G \to Y$ given by $([\gamma], g) \mapsto (y * \gamma, g) \mapsto g.(y * \gamma)$. Since each of these maps is continuous, the map $\tilde{X} \times G \to Y$ is continuous as well.

We now check that the point $[\sigma].([\gamma], g) = ([\sigma][\gamma], g\rho([\sigma]^{-1})$, which is in the same orbit as $([\gamma], g)$ under the action of $\pi_1(X, x)$, maps to the same point in Y. We apply Equations (i) and (ii) to calculate

$$\begin{aligned}(g.\rho([\sigma])^{-1})(y * (\sigma \circ \gamma)) &= (g.\rho([\sigma])^{-1})((y * \sigma) * \gamma) \\ &= ((g.\rho([\sigma])^{-1})(y * \sigma)) * \gamma \\ &= (g.((y * \sigma) * \sigma^{-1})) * [\gamma] \\ &= (g.(y * (\sigma \circ \sigma^{-1}))) * [\gamma] \\ &= (g.y) * [\gamma],\end{aligned}$$

which was what we wanted. This shows that our map gives the same value in each $\pi_1(X, x)$ orbit, so it is well-defined as a map of G-coverings $Y_\rho \to Y$.

On the other hand, given a homomorphism $\rho : \pi_1(X, x) \to G$, we showed in the previous discussion how to construct a G-covering $Y_\rho \to X$. Suppose we then construct another homomorphism $\tilde{\rho}$ from this covering using the construction described earlier. We must check that $\tilde{\rho} = \rho$. Let $[\sigma] \in \pi_1(X, x)$ and $\langle \tilde{x}, g \rangle \in Y_\rho$. Then, by our previous definitions, we calculate

$$\begin{aligned}\tilde{\rho}([\sigma]).\langle \tilde{x}, g \rangle &= \langle \tilde{x}, g \rangle * \sigma \\ &= \langle \tilde{x} * \sigma, g \rangle \\ &= \langle [\sigma].\tilde{x}, g \rangle \\ &= \langle \tilde{x}, g.\rho([\sigma]) \rangle \\ &= \rho([\sigma]).\langle \tilde{x}, g \rangle.\end{aligned}$$

This shows that $\tilde{\rho} = \rho$, as desired. □

Remark 7.4. (Gluing Coverings). Suppose $X = U \cup V$ where U, V are open, and suppose $p : \tilde{X} \to X$ is a covering of X. Then $p|_U : p^{-1}(U) \to U$ and $p|_V : p^{-1}(V) \to V$ are coverings of U and V, respectively, and are clearly isomorphic over $U \cap V$, i.e., the restrictions over $U \cap V$ are isomorphic. This shows that restriction of coverings works as we would expect it to; indeed, in the language of algebraic geometry, coverings of X form a *presheaf*.

Conversely, suppose we are given coverings $p_1 : \tilde{U} \to U, p_2 : \tilde{V} \to V$ and have an isomorphism $f : p_1^{-1}(U \cap V) \to p_2^{-1}(U \cap V)$. We claim we can construct a covering $p : \tilde{X} \to X$ and isomorphisms of coverings

$$f_1 : \tilde{U} \to p^{-1}(U), f_2 : \tilde{V} \to p^{-1}(V)$$

such that $f = f_1^{-1} \circ f_2$ over $U \cap V$.

To this end, we define the equivalence relation $\sim$ such that a point $y_1 \in p_1^{-1}(U \cap V)$ is identified with $f(y_1) \in p_2^{-1}(U \cap V)$. (Although we used the notation $\sim$ for previous equivalence relations, it will be clear from context which one we mean.) Then we set $\tilde{X} := \tilde{U} \coprod \tilde{V} / \sim$ and equip $\tilde{X}$ with the quotient topology (i.e., the open sets in $\tilde{X}$ are precisely those sets whose inverse image in $\tilde{U} \coprod \tilde{V}$ is open). Define $p : \tilde{X} \to X$ as follows: if $x \in \tilde{U}$, let $p(x) := p_1(x)$, and if $x \in \tilde{V}$, let $p(x) := p_2(x)$.

The map p is well-defined because if two points x, y in $\tilde{U}$ and $\tilde{V}$, respectively, are identified in $\tilde{X}$, then $y = f(x)$ and $p_2(y) = p_2(f(x)) = p_1(x)$. Also, since the inclusion $i_1 : \tilde{U} \to \tilde{X}$ is a homeomorphism onto its image $p^{-1}(U)$ and likewise for the inclusion of $\tilde{V}$ in $\tilde{X}$, we see that $p|_{p^{-1}(U)} \simeq p_1$ and that $p|_{p^{-1}(U)} \simeq p_2$, where $\simeq$ now denotes isomorphism of coverings. Therefore, p is a covering. Similarly, if p_1, p_2 are G-coverings for some group G and f is an isomorphism of G-coverings, then p is a G-covering.

8. The Seifert–van Kampen Theorem

We have the following theorem.

Theorem 8.1. (Seifert–van Kampen) [1, Theorem 14.4]. *Let $X = U \cup V$, U, V open. Suppose X, U, V, and $U \cap V$ are path-connected, locally path-connected, and semilocally simply connected, and let $x \in U \cap V$. Let $i_1 : \pi_1(U \cap V, x) \to \pi_1(U, x)$ and $i_2 : \pi_1(U \cap V, x) \to \pi_1(V, x)$ be the homomorphisms induced by inclusion, and define $j_1 : \pi_1(U, x) \to \pi_1(X, x)$ and $j_2 : \pi_1(V, x) \to \pi_1(X, x)$ similarly. Given a group G and homomorphisms $h_1 : \pi_1(U, x) \to G$, $h_2 : \pi_1(V, x) \to G$ such that $h_1 \circ i_1 = h_2 \circ i_2$, there is a unique homomorphism $h : \pi_1(X, x) \to G$ such that $h \circ j_1 = h_1$ and $h \circ j_2 = h_2$.*

Proof. First, it is helpful to visualize the situation through the following commutative diagram depicting the homomorphisms in question:

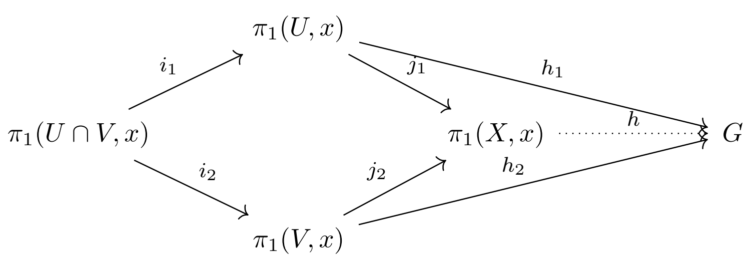

Our assumptions guarantee that X, U, V, and $U \cap V$ each have universal covering spaces by Theorem 5.7. The proof of Proposition 7.3 used the existence of a universal covering, so this guarantees that h_1 and h_2 determine G-coverings $Y_1 \to U$ and $Y_2 \to V$ as well as basepoints y_1 and y_2 over x. Since $h_1 \circ i_1 = h_2 \circ i_2$, the restrictions of these coverings to $U \cap V$ are isomorphic G-coverings. (Note that the inclusion-induced maps i_1, i_2, j_1, j_2 are injective by Proposition 5.3.) Also, since $U \cap V$ is connected, there is an isomorphism between these G-coverings that maps y_1 to y_2. By Remark 7.4, these two G-coverings can be patched together, using the aforementioned isomorphism over $U \cap V$. Thus we obtain a G-covering $Y \to X$ that restricts to the two given G-coverings and has the same basepoint. By Proposition 7.3, we obtain a homomorphism $h : \pi_1(X, x) \to G$, and the fact that the restricted coverings agree means $h \circ j_1 = h_1$ and $h \circ j_2 = h_2$ as desired. Uniqueness of h follows from uniqueness of the G-covering. □

We now prove a consequence of the Seifert–van Kampen Theorem that will subsequently be of more use to us than the fully general theorem.

Corollary 8.2. *If the conditions of Theorem 8.1 hold, then for any group G there is a bijection*

$$Hom(\pi_1(X,x),G) \simeq \{(h_1,h_2) \in Hom(\pi_1(U,x),G) \times Hom(\pi_1(V,x),G) | h_1 i_1 = h_2 i_2\}.$$

In category-theoretic language, the diagram

$$\begin{array}{ccc} \pi_1(U\cap V,x) & \longrightarrow & \pi_1(U,x) \\ \downarrow & & \downarrow \\ \pi_1(V,x) & \longrightarrow & \pi_1(X,x) \end{array}$$

is a pushout diagram in the category of groups.

Proof. We define a map

$$F : \mathrm{Hom}(\pi_1(U,x),G) \times \mathrm{Hom}(\pi_1(V,x),G) \to \mathrm{Hom}(\pi_1(X,x),G)$$

and show that F is one-to-one and onto. Given $(h_1,h_2) \in \mathrm{Hom}(\pi_1(U,x),G) \times \mathrm{Hom}(\pi_1(V,x),G)$ such that $h_1 i_1 = h_2 i_2$, let

$$F(h_1,h_2) := h$$

where $h \in \mathrm{Hom}(\pi_1(X,x),G)$ is the unique homomorphism from the statement of the Seifert–van Kampen Theorem satisfying $h \circ j_1 = h_1$ and $h \circ j_2 = h_2$.

Suppose $F(h_1,h_2) = F(h_1',h_2') = h$, where $(h_1',h_2') \in \{(h_1,h_2) \in \mathrm{Hom}(\pi_1(U,x),G) \times \mathrm{Hom}(\pi_1(V,x),G) | h_1 i_1 = h_2 i_2\}$. Then $h_1' = h \circ j_1 = h_1$ and $h_2' = h \circ j_2 = h_2$, so $(h_1,h_2) = (h_1',h_2')$. Thus F is one-to-one.

Now suppose $h \in \mathrm{Hom}(\pi_1(X,x),G)$. Then take $(h_1,h_2) := (h\circ j_1, h\circ j_2)$. Clearly $F(h_1,h_2) = h$ in this case, and clearly $(h_1,h_2) \in \mathrm{Hom}(\pi_1(U,x),G) \times \mathrm{Hom}(\pi_1(V,x),G) | h_1 i_1 = h_2 i_2\}$ as well. Thus F is surjective. ☐

9. Free Groups and Free Products

In this section, we will demonstrate how to use Corollary 8.2 to calculate fundamental groups. Suppose X is a topological space and $X_1, X_2 \subseteq X$ are path-connected open subsets. Let $X_0 := X_1 \cap X_2$, and suppose X_0 has exactly two distinct path-components, say Y and Z. Define $\tilde{Z}$ by gluing X_1 and X_2 along Z. Define $\tilde{X}$ by gluing $Y \times [0,1]$ to $\tilde{Z}$ as follows: identify $Y \times \{i\}$ with the copy Y_{i+1} of Y in X_{i+1} for $i = 0,1$. Choose $y \in Y$, and let y_i be its image in Y_i. Let $\ell : [0,1] \to \tilde{Z}$ be a path with $\ell(0) = y_1, \ell(1) = y_2$. We then have homomorphisms

$$\alpha_1 : \pi_1(Y,y) \simeq \pi_1(Y_1,y_1) \to \pi_1(\tilde{Z},y_1),$$

$$\alpha_2 : \pi_1(Y,y) \simeq \pi_1(Y_2,y_2) \to \pi_1(\tilde{Z},y_2) \simeq \pi_1(\tilde{Z},y_1),$$

where the last isomorphism $\pi_1(\tilde{Z},y_2) \simeq \pi_1(\tilde{Z},y_1)$ of the second line is induced by the path ℓ. (Given a loop representative σ based at y_1, $\ell^{-1}\sigma\ell$ is based at y_2, and similarly given a loop at y_2 we can produce a loop at y_1. This induces the isomorphism of fundamental groups.) We claim that with this notation the following holds.

Proposition 9.1. [6, Proposition 1.2]. *If G is an arbitrary group, we have a bijection*

$$Hom(\pi_1(\tilde{X}, y_1), G) \simeq \{(f_1, t) \in Hom(\pi_1(\tilde{Z}, y_1), G) \times G | f_1\alpha_2(p) = t^{-1} f_1\alpha_1(p)t \ \ \forall p \in \pi_1(Y, y)\}.$$

Proof. We consider $W := \tilde{Z} \cup (y \times [0,1])$. Notice that $W = \tilde{Z} \cup (y \times [0,1])$ and $\tilde{Z} \cap (\ell([0,1]) \cup (y \times [0,1])) = \ell$. By Corollary 8.2 and the fact that the circle has fundamental group $\mathbb{Z}$, $\pi_1(W, y_1)$ is the pushout of $\mathbb{Z} \leftarrow 1 \rightarrow \pi_1(\tilde{Z}, y_1)$. We *define* this pushout to be the *free product* $\mathbb{Z} * \pi_1(\tilde{Z}, y_1)$, which we will discuss in more detail later.

We now observe that $\tilde{X} = W \cup (Y \times [0,1])$ and $W \cap (Y \times [0,1]) = (Y \times \{0\})) \cup (y \times [0,1]) \cup (Y \times \{1\})$. Therefore, applying Corollary 8.2 again, we have a pushout diagram

$$\begin{array}{ccc} \pi_1(Y,y) * \pi_1(Y,y) & \xrightarrow{\phi} & \mathbb{Z} * \pi_1(\tilde{Z}, y_1) \\ \downarrow & & \downarrow \\ \pi_1(Y,y) & \longrightarrow & \pi_1(\tilde{X}, y_1) \end{array}$$

Here the map ϕ acts on the first coordinate as α_1 and on the second coordinate as $c \mapsto t\alpha_2(c)t^{-1}$. (This follows from the $h_1 i_1 = h_2 i_2$ condition in the statement of Corollary 8.2.) This gives the result. □

In the proof above, we think of f_1 as the composition $\pi_1(\tilde{Z}, y_1) \rightarrow \pi_1(\tilde{X}, y_1) \rightarrow G$ and t as the image of the class of the loop $\ell \cup y \times [0,1]$ in $\pi_1(\tilde{X}, y_1)$

Example 9.2. Take two copies of S^1, each with a distinguished basepoint, and identify the two basepoints. The resulting space is $S^1 \vee S^1$, the *wedge product* of S^1 and S^1. Taking X_1 and X_2 to be the two circles and applying Proposition 9.1, we obtain that $\mathrm{Hom}(\pi_1(X,x), G) \simeq G \times G$ for any group G. $\pi_1(X,x)$ is called the *free group* on generators t_1 and t_2, where t_i is the class in $\pi_1(X,x)$ of X_i. This group is denoted F_2. In general, the free group on n generators is denoted F_n and is the fundamental group of the wedge product of n based circles with all of the basepoints identified.

We will now give a concrete description of F_2. We consider the symbols $t, u, \overline{t}, \overline{u}$. We say a *word* is a finite (possibly empty) sequence of these symbols, and we say a word is *reduced* if it does not contain the strings $t\overline{t}$, $\overline{t}t$, $u\overline{u}$, $\overline{u}u$. We now state and prove the result.

Theorem 9.3. [6, Theorem 1.4]. *There is a bijection between the set W of reduced words and the free group F on generators t, u. In particular, the symbols $\overline{t}, \overline{u}$ correspond to t^{-1},u^{-1}, respectively, under this bijection.*

Proof. First, F contains t, u, so because it is a group, it contains any product of t, u, t^{-1}, and u^{-1}. Let H be the subgroup of such products. The embedding $F \rightarrow H \subseteq F$ coincides with the identity on t and u, but t, u generate F, so $F = H$. Each element of H can be represented by a reduced word (where we are identifying $\overline{x}$ with x^{-1} for $x \in \{t, u\}$) because we can cancel products tt^{-1}, uu^{-1}, etc. So we have a surjection $\alpha : W \rightarrow F$. We must show that α is injective.

Let S be the symmetric group of the set W. We define permutations τ, σ as follows. Suppose $w \in W$ ends in $\overline{t}$ (resp. $\overline{u}$). Then let $w\tau$ (resp. $w\sigma$) be w with the last letter deleted. Otherwise, let $w\tau$ (resp. $w\sigma$) be w followed by t (resp. u). It is clear that the inverses τ^{-1} (resp. ω^{-1}) are defined the same way as τ (resp. ω) but with t replacing $\overline{t}$ in the definition (resp. u replacing $\overline{u}$).

By definition of F as the free group on 2 generators (see Example 9.2), we have $\mathrm{Hom}(F,G) \simeq G \times G$ for any group G, so there is a unique homomorphism $\phi : F \to S$ with $\phi(t) = \tau$, $\phi(u) = \sigma$. We define $\beta : F \to W$ by $\beta(g) := \phi(g)(1)$. We claim that $\beta(\alpha(w)) = w$ for all reduced words w. If w is the empty word, this holds. If w is a reduced word of length 1, then $w \in \{t, u, t^{-1}, u^{-1}\}$, and it is easy to check that $\beta(\alpha(w)) = w$. Suppose this holds for all reduced words of length $\leq n$. Let w be a word of length $n+1$, so that it is a word of length n with an element of $\{t, u, t^{-1}, u^{-1}\}$ as its $(n+1)$'st letter. Then by our inductive assumption it is easy to check that $\beta(\alpha(w)) = w$. For example, if $w = xt$ where x is a reduced word of length n (and in particular x doesn't end in t^{-1}, or w wouldn't be reduced), then (slightly abusing notation) $\alpha(w) = xt$ as well, and

$$\beta(\alpha(w)) = \phi(xt)(1) = \phi(x)\phi(t)(1) = \phi(x)\tau(1) = \phi(x)t = w,$$

where $\phi(x)$ doesn't end in $\bar{t}$ since w doesn't end in t^{-1}. The other cases are similar. We thus obtain that $\beta(\alpha(w)) = w$ for all reduced words w. This proves α is injective, which completes the proof. $\square$

The proof just given readily generalizes by induction to the case of F_n for any positive integer n. In the notation of group presentations familiar from an introductory group theory course, we write $F_n = \langle a_1, ..., a_n \rangle$. If X is an infinite generating set, we define

$$F(X) := \bigcup \{F(Y) | Y \subset X, Y \text{ finite}\}.$$

Suppose $y_i \in F(Y_i) \subset F(X)$ for $i = 1, 2, 3$ where the Y_i are finite. We can then define $y_1 y_2$ to be the product in $F(Y_1 \cup Y_2)$. Associativity of the product follows from $F((Y_1 \cup Y_2) \cup Y_3) = F(Y_1 \cup (Y_2 \cup Y_3))$. This definition makes it clear that if G is an arbitrary group, then a homomorphism $F(X) \to G$ simply corresponds to a map of sets $X \to G$, as $F(X)$ is generated by X, so $F(X)$ is indeed the free group on X.

Free groups were particularly nice cases of fundamental groups calculated by applying Corollary 8.2 to certain topological spaces. We now give two more general constructions which are ever-present in geometric group theory.

In the context of Corollary 8.2 and Proposition 9.1, we concerned ourselves with the following two diagrams:

$$\begin{array}{ccc} C & \xrightarrow{\alpha_1} & A \\ \downarrow{\scriptstyle \alpha_2} & & \vdots{\scriptstyle \beta_1} \\ B & \xrightarrow{\beta_2} & G \end{array}$$

and

$$C \underset{\alpha_2}{\overset{\alpha_1}{\rightrightarrows}} A \xrightarrow{\beta} G$$

We introduce the following definitions.

Definition 9.4. If α_1 and α_2 are injective maps[4], the universal group G (i.e., the group G such that if there is a group H satisfying the same commutative diagram but with possibly

[4]In both of the commutative diagrams, we want α_1 and α_2 to be injective.

different maps $A \to H$, $B \to H$, then there is a unique map $G \to H$) is called the *free product* of A and B *amalgamated along* C in the case of the first diagram, and is called the *HNN extension*[5] of A *along* C in the case of the second diagram. For the amalgamated free product, we write $A *_C B$, and for the HNN extension of A along C, we write $A*_C$. In particular, the free product $A * B$ is the special case where C is the trivial group.

Although there are concrete constructions of the amalgamated free product and of HNN extensions just as there are for free groups, we will not use them in this paper. We refer the interested reader to p. 145–147 of [6] for the details.

10. Graphs of Groups

In this section, we will need to introduce some definitions that are fairly elementary but have not been needed until now. We will suppose all the topological spaces we deal with in proofs in this section and the next have the three properties required for having a universal cover from the statement of Theorem 5.7.

Definition 10.1. Let X and Y be topological spaces. We say they are *homotopy equivalent* if there exist continuous maps $f : X \to Y$ and $g : Y \to X$ such that $f \circ g$ is homotopic to the identity on Y and $g \circ f$ is homotopic to the identity on X.

Definition 10.2. Let X be a topological space with subspace $A \subseteq X$. A continuous map $r : X \to A$ is a *retraction* if $r|_A = \mathrm{id}_A$. A *deformation retraction* $F : X \times [0,1] \to X$ is a homotopy between a retraction and id_A, i.e., it is a continuous map such that for every $x \in X$ and every $a \in A$, we have $F(x,0) = x$, $F(x,1) \in A$, and $F(a,1) = a$.

Definition 10.3. A *CW complex* is a topological space constructed as follows.

(i) Start with a set of points, also known as 0-cells, X^0.

(ii) Inductively, construct the *n-skeleton* X^n from X^{n-1} by attaching spaces homeomorphic to open n-disks, or *n-cells*, e^n_α via continuous maps $\varphi_\alpha : S^{n-1} \to X^{n-1}$.

(iii) Either stop at some finite n, or inductively set $X := \bigcup_{n\geq 0} X^n$. In the latter case, equip X with the *weak topology*, i.e., $U \subseteq X$ is open iff $U \cap X^n$ is open for each $n \geq 0$.

Example 10.4. The sphere S^n is a CW complex with a 0-cell e^0 and an n-cell e^n, where the n-cell is attached by the constant map $S^{n-1} \to e^0$. Indeed, most topological spaces are homotopy equivalent to some CW complex, although the Hawaiian earring isn't. (We won't prove this here.)

Proposition 10.5. [6, p. 142]. *Given a finitely-generated group G, there exists a CW complex X such that $\pi_1(X,x) \simeq G$.*

Proof. Let $F(X)$ be the free group generated by a set X. Produce one circle for each element of X, each circle having an identified basepoint. Take the wedge product of the circles along their basepoints. By our earlier discussion of free groups, this will have fundamental group $F(X)$. Let K^1 be this complex (the 1 indicates it is a 1-skeleton). We wish to see what

[5]This stands for *Higman–Neumann–Neumann*, as HNN extensions were introduced by Graham Higman, B.H. Neumann, and Hanna Neumann in [3].

happens to the fundamental group when we attach a 2-cell e^2.

The cell e^2 is contractible, and a neighbourhood of K^1 meets e^2 in a copy of $S^1 \times [0,1]$ where I is an interval. Denote this neighbourhood by $N(K^1)$. The map $\alpha : S^1 \to N(K^1) \to K^1$ is then homotopic to the attaching map of e^2, and there is an induced map $\alpha_* : \pi_1(S^1) = \mathbb{Z} \to \pi_1(K^1)$. Suppose $\alpha_*(1) = r$. Then $\pi_1(K^1 \cup e^2)$ is the pushout of $\{1\} \leftarrow \mathbb{Z} \to \pi_1(K^1)$, where the right arrow is given by α_*.

Giving the same argument when we attach other 2-cells e_j, we see that if their attaching maps give classes $r_j \in \pi_1(K^1)$, then the resulting 2-skeleton K^2 has the property that if G is an arbitrary group, then

$$\mathrm{Hom}(\pi_1(K^2), G) = \{f \in \mathrm{Hom}(\pi_1(K^1), G) | f(r_j) = 1 \quad \forall j\}.$$

We know that $\pi_1(K^1)$ is free on some generators $\{g_i\}$, so f is determined by the values $f(g_i)$. This gives a group presentation $\langle g_i | r_j \rangle$. The g_i are arbitrary, and the r_j are arbitrary as well; to see that the latter are arbitrary, note that given an arbitrary product of generators, we can attach along the loop defined by that product to produce the corresponding relator in the presentation of the fundamental group of the resulting CW complex. Since, from elementary group theory, every group has a presentation in generators and relators, this concludes the proof. □

We now define graphs. For our purposes, graphs will have twice as many edges as they typically do (e.g., in the context of graph theory). This convention is typical in Bass–Serre theory, the theory of group actions on graphs.

Definition 10.6. A *graph* consists of two sets $E(\Gamma)$ and $V(\Gamma)$, the set of edges and of vertices, respectively, as well as an involution $E(\Gamma) \to E(\Gamma)$ given by $e \mapsto \overline{e}$, where $\overline{e} \neq e$. (The picture to have in mind is that $\overline{e}$ is the oriented edge with opposite orientation to e.) We also have a map $\partial_0 : E(\Gamma) \to V(\Gamma)$ which sends each edge to an adjacent vertex (i.e., for each edge we can make an arbitrary choice between the two vertices to send it to), and we define $\partial_1 e := \partial_0 \overline{e}$, saying that e joins $\partial_0 e$ to $\partial_1 e$. An *orientation* of a graph is a choice of one edge from each pair $(e, \overline{e})$.

Remark 10.7. One might wonder why we want to have twice as many edges as in the usual definition of a graph (i.e., why the map $e \mapsto \overline{e}$ is useful). To the author's best knowledge, this is intended to allow us to easily prohibit actions that contain "inversions". For this notion, see the first paragraph of §11.

We can now define a graph of groups.

Definition 10.8. A *graph of groups* is a graph Γ, always assumed for our purposes to be connected, and a function $\mathcal{G}$ which assigns a group G_v to each vertex v of Γ and a group G_e to each edge e. We also suppose that $G_e = G_{\overline{e}}$ and that there is an injective homomorphism $f_e : G_e \to G_{\partial_0 e}$ for each e.

Similarly, a *graph of topological spaces* is a connected graph such that to each vertex we assign a topological space X_v and to each edge we assign a topological space X_e. We require that $X_e = X_{\overline{e}}$ for all edges e, but we do not impose a requirement corresponding to the injective homomorphism in the definition of a graph of groups.

Given a graph $\mathcal{X}$ of topological spaces, we define the *total space* X_Γ as the quotient of

$$\bigcup\{X_v | v \in V(\Gamma)\} \cup \bigcup\{X_e \times [0,1] | e \in E(\Gamma)\}$$

by identifications

$$X_e \times [0,1] \to X_{\overline{e}} \times [0,1] \text{ by } (x,t) \mapsto (x, 1-t),$$
$$X_e \times \{0\} \to X_{\partial_0 e} \text{ by } (x,0) \mapsto f_e(x).$$

If we have a graph of connected topological spaces, we can take fundamental groups to obtain a graph of groups $\mathcal{G}$. The total space X_Γ has a fundamental group, which we will denote by G_Γ. We *define* G_Γ to be the fundamental group of $\mathcal{G}$. In particular, suppose we have vertex groups A and B connected by an edge group C. Then it follows from Corollary 8.2 that the fundamental group of the graph of groups is just the amalgamated free product $A *_C B$. Similarly, if we have a single vertex group A connected to itself by an edge group C, it follows from Corollary 8.2 that the fundamental group of the graph of groups is the HNN extension $A*_C$.

We will now demonstrate that G_Γ does not depend on our choice of $\mathcal{X}$. Given any $\mathcal{G}$, we can, by Proposition 10.5, choose connected 2-dimensional CW complexes such that $\pi_1(X_v) \simeq G_v$ and $\pi_1(X_e) \simeq G_e$. (They can be chosen to be 2-dimensional because the construction of Proposition 10.5 only used 2-dimensional complexes. Because we rely on this construction, we will assume in what follows that all fundamental groups we work with are finitely-generated.) This gives a graph $\mathcal{X}$ of topological spaces that gives rise to $\mathcal{G}$. Given such an $\mathcal{X}$, we can attach cells of dimension ≥ 3 to the $\{X_v\}$ and $\{X_e\}$ to produce spaces $\{K_e\}$ and $\{K_v\}$ which are *aspherical*, i.e., have trivial homotopy groups π_n for $n > 1$.[6] (This is the higher-dimensional analogue of adding 2-cells to make the fundamental group of a CW complex trivial, and the added higher-dimensional cells are not detected by the fundamental group, which depends only on the 1- and 2-skeletons of the CW complex.) The maps $f_e : X_e \to X_{\partial_0 e}$ then extend to maps $k_e : K_e \to K_{\partial_0 e}$, so we have a new graph $\mathcal{K}$ of aspherical topological spaces, which still induces $\mathcal{G}$. What is more, the total space K_Γ has the same fundamental group as X_Γ, and the homotopy groups of the vertex and edge spaces are now determined entirely by the collections $\{G_v\}$ and $\{G_e\}$, respectively. It follows that G_Γ does not depend on our choice of $\mathcal{X}$.

To prove the next proposition, we will need Whitehead's theorem from algebraic topology. The proof of this theorem would take us too far off course as it involves higher homotopy groups, so we simply cite a proof.

Theorem 10.9. (Whitehead). *If X and Y are CW complexes, and if a continuous map $f : X \to Y$ induces isomorphisms on the n'th homotopy groups of X and Y for each $n \geq 1$, then f is a homotopy equivalence.*

Proof. Refer to [2, Theorem 4.5]. □

Proposition 10.10. [6, Proposition 3.6]. *(i) If $\mathcal{G}$ is a graph of groups as defined in the previous discussion, then $G_v \to G_\Gamma$ is injective.*

[6]Higher homotopy groups are defined in many places, e.g., in [2, §4.1]. The first homotopy group of a topological space is the fundamental group.

(ii) If $\mathcal{K}$ is a graph of aspherical spaces as discussed above, then the total space K_Γ is aspherical.

Proof. Suppose we are given the graph $\mathcal{K}$ of aspherical spaces. Let v be a vertex of Γ. Then the space $L_v := K_v \cup \bigcup_{\partial_0 e = v}(K_e \times [0,1])$ deformation retracts to K_v. In particular, this is a homotopy equivalence. Since homotopy groups are only defined up to homotopy (as one might expect from the name), this induces isomorphisms of n'th fundamental groups for all n, so we may apply Theorem 10.9 to obtain that the universal cover $\tilde{L}_v$ is contractible (since L_v deformation retracts to K_v, which is aspherical and thus has contractible universal cover). Since the maps $G_e \to G_v$ are injective, the universal cover $\tilde{L}_v$ can be obtained from the universal cover $\tilde{K}_v$ by attaching copies of $\tilde{K}_e \times [0,1]$, the product of the universal cover of K_e with the unit interval.

We now inductively construct a space $Y = \bigcup Y_n$; we will equip it with the weak topology. Choose a vertex v_0 of Γ and set $Y_0 := \tilde{L}_{v_0}$. For any $n \geq 1$, we will have attached copies of $\tilde{K}_e \times [0,1]$ to form Y_{n-1}, and we define Y_n to be the union of Y_{n-1} with a copy of $\tilde{L}_{\partial_1 e}$ for every such copy of $\tilde{K}_e \times [0,1]$, where we are attaching it along $\tilde{K}_e \times [0,1]$. Each Y_n is contractible as we have only attached contractible sets along contractible sets. We claim that $Y = \bigcup Y_n$ is also contractible. Indeed, a well-known result[7] states that a CW complex that is the union of an increasing sequence of subcomplexes such that the inclusion of each subcomplex in the sequence into the next is homotopic to the constant map is contractible. Since the inclusions $Y_i \to Y_{i+1}$ are homotopic to the constant map for $i \geq 0$, it follows that Y is contractible.

We have a canonical projection $Y \to K_\Gamma$. This projection evenly covers K_Γ by construction. This shows that K_Γ is aspherical. To get (i), observe that for each $K_v \subseteq K_\Gamma$, the induced covering of K_v contains the universal covering, so taking fundamental groups, the map $G_v \to G_\Gamma$ is injective. ☐

11. Bass–Serre Theory

We now come to our final topic: group actions on trees. Suppose a group G acts on a graph Γ. We write $G \curvearrowright \Gamma$ and take our action to be continuous and *without inversions*, this second condition meaning that whenever g fixes an edge $e \in E(\Gamma)$, it fixes every point of e; in particular, this rules out the "inversion" $g.e = \overline{e}$. The following example is another classic geometric group theory construction.

Definition 11.1. Let G be a group, $S \subseteq G$ a subset of G, and let $\Gamma = \Gamma(G; S)$ be the graph with vertex set G and edge set defined as follows: for each $(g, s) \in G \times S$, include a single edge $e(g, s)$ from g to gs. We say Γ is an *S-graph* for G. When S is a generating set for G, we say Γ is a *Cayley graph* for G with respect to S.

Example 11.2. Let S be a generating set for a group G. We have an action $G \curvearrowright \Gamma(G; S)$ given as follows: if $h \in G$, then h maps the vertex corresponding to $g \in G$ to the vertex hg and the edge $e(g, s)$ to the edge $e(hg, hs)$. Then G acts freely without inversions on Γ. Note that we do not identify the edges $e(g, s)$ and $e(gs, s)$ even if $s^2 = 1$.

Now we will continue our discussion of graphs of groups.

[7] It is given, for example, in [2, §4.1].

Lemma 11.3. [6, §4]. *Let $\mathcal{G}$ be a graph of groups. (Recall that we then have vertex groups $\{G_v\}$ and edge groups $\{G_e\}$ such that $G_e = G_{\bar{e}}$ for all e and such that we have injections $\alpha(e) : G_e \to G_{\partial_0 e}$.) As discussed earlier, we can produce a graph of connected spaces X_Γ with vertex spaces $\{X_v\}$, edge spaces $\{X_e\}$, and fundamental group G_Γ. Let $\tilde{X}_\Gamma$ be the universal cover of X_Γ, and define $\tilde{X}_v$, $\tilde{X}_e$ similarly. (These exist because a graph of groups is a graph and therefore satisfies all the conditions required to have a universal cover.) Then $\tilde{X}_\Gamma$ can be written as a union of copies of $\tilde{X}_v$ and $\tilde{X}_e \times [0,1]$.*

Proof. By Proposition 10.10, the maps $G_v \to G_\Gamma$ and $G_e \to G_\Gamma$ are injective, so $\tilde{X}_\Gamma$ can be written as a union of copies of $\tilde{X}_v$ and $\tilde{X}_e \times [0,1]$ by the proof of Proposition 10.10. □

Definition 11.4. Retaining the notation of Lemma 11.3, identify each copy of $\tilde{X}_v$ with a point and each copy of $\tilde{X}_e \times [0,1]$ with the interval $[0,1]$. Call the resulting space Z. Clearly Z is a graph, and we have given a projection $\pi : \tilde{X}_\Gamma \to Z$. We now define a map $j : Z \to \tilde{X}_\Gamma$. For each vertex (resp. edge) of Z, choose a point v (resp. a *point* e) in the copy of $\tilde{X}_v$ (resp. $\tilde{X}_e$) corresponding to it. Next, divide each edge of Z into three parts in an arbitrary way. Define j to map the middle third of the edge e to $e \times [0,1]$ and to map the end thirds to paths in the space $\tilde{X}_v$ (which is of course connected) that join the corresponding points e, v. (This is fine since e was mapped to a point, not an edge.) Then $\pi \circ j$ is homotopic to the identity. It follows that Z is connected and simply connected and is therefore a tree.

The projection π commutes with the action of G_Γ on $\tilde{X}_\Gamma$, so this action descends to an action on Z, which has no inversions. By construction, if a vertex in Z is obtained by collapsing $\tilde{X}_v$, then its stabilizer is a conjugate of G_v, and similarly the stabilizer of an edge obtained by collapsing $\tilde{X}_e \times [0,1]$ is a conjugate of G_e. The space $G_\Gamma \backslash Z$ is then also a graph because the action is without inversions.

We now make the following claim.

Lemma 11.5. [6, §4]. *The space $G_\Gamma \backslash Z$ coincides with the geometric realization $|\Gamma|$ of our original graph Γ, i.e., the graph Γ in the graph-theoretic sense, where we do not consider a multigraph with inverse edges $\bar{e}$.*

Proof. We have a canonical surjection $G_\Gamma \backslash Z \to |\Gamma|$, and for each vertex (resp. edge) v (resp. e) of $|\Gamma|$, we have X_v (resp. $X_e \times [0,1]$) in X_Γ, so a collection of copies of $\tilde{X}_v$ (resp. $\tilde{X}_e \times [0,1]$) in $\tilde{X}_\Gamma$, and G_Γ acts evenly on this collection. We thus obtain a single vertex (resp. edge) of $G_\Gamma \backslash Z$. We can therefore recover our graph of groups $\mathcal{G}$ from G_Γ acting on Z. □

Lemma 11.6. [6, §4]. *Suppose we have a group G acting on a (topological) tree Y with no inversions.Then there exists a graph of groups whose fundamental group is G.*

Proof. By Proposition 10.5, we can produce a connected CW complex U with fundamental group G. By the construction from Proposition 10.5, this complex will have a universal cover $\tilde{U}$, and G will act freely on it. Therefore, G will act on $\tilde{U} \times Y$ by acting on each coordinate. Let

$$X := G \backslash (\tilde{U} \times Y).$$

We then have a projection $X \to G \backslash Y =: \Gamma$. Because G acts without inversions on Y, Γ is a graph. Moreover, if v is a vertex (resp. e is an edge) of Y with stabilizer G_v (resp. G_e), then $G \backslash (\tilde{U} \times v)$ (resp. $G \backslash (\tilde{U} \times e)$) has fundamental group G_v (resp. G_e) by our construction.

This shows that X is a graph of connected spaces which is the topological realization of a graph of groups $\mathcal{G}$. G acts freely on $\tilde{U} \times Y$, which is simply connected, so $G = \pi_1(X)$, which is the fundamental group of $\mathcal{G}$. This was what we wanted. □

Given a CW complex K, we can consider its 1-skeleton $K^{(1)}$, which is a graph and therefore contains a maximal tree T. The edges of K that are not in T then correspond to generators of $\pi_1(K)$. This situation motivates the following result.

Proposition 11.7. [6, Proposition 4.2]. *Suppose G acts without inversions on a connected graph Y. Then $X := G\backslash Y$ is a graph, so we can choose a subtree T of X containing a vertex v, and we let $\tilde{v} \in Y$ lie above v. Then there exists a lift $j : T \to Y$ of the inclusion of T in X such that $j(v) = \tilde{v}$.*

Proof. Consider pairs (S, f) where S is a subtree of T containing v and $f : S \to Y$ is a lift of the inclusion of T in X such that $f(v) = \tilde{v}$. (The set of such pairs is non-empty because we have the trivial inclusion, which lifts the tree consisting of v alone to $\tilde{v}$.) Define a partial order on such pairs by saying that $S_1 \leq S_2$ if $S_1 \subseteq S_2$ and $f_1 \leq f_2$ if $f_2|_{S_1} = f_1$. Then every totally ordered subset of such pairs contains a maximal element given by taking unions of trees and of maps, where $f_1 \cup f_2 := f_2$ if $f_1 \leq f_2$. By Zorn's lemma, we now get a maximal such pair (T', j').

Suppose $T' \neq T$. Choose a vertex w in $T - T'$. T is connected, so we can join w to v by a path, and at least one edge in the path, say e, has a vertex v_0 which is in T' and a vertex which is not in T'. Let $\tilde{e}$ be a lift of e in Y. Then v_0 has a corresponding lift $\tilde{v}_0$. By definition of j', both $j'(v_0)$ also lies over v_0, so it is in the same G-orbit as $\tilde{v}_0$; we write $g.\tilde{v}_0 = j'(v_0)$. But we can now define an extension j of j' to e by setting $j(e) = g.\tilde{e}$, contradicting the maximality of j'. This proves the result. □

Now, if G acts on the tree Y as before, we choose a maximal tree T in $\Gamma := G\backslash Y$ and a lift $j : T \to Y$ with $j(T) =: \tilde{T}$. Over each vertex v of Γ, we now have just one lift $\tilde{v}$ of $\tilde{T}$. We *define* G_v to be the stabilizer of $\tilde{v}$. Similarly, for an edge e of T, we have the edge $\tilde{e} := j(e)$ and the stabilizer G_e of $\tilde{e}$. For the other edges e of Γ, choose a lift $\tilde{e}$ in Y with $\widetilde{\partial_0 e} = \partial_0\tilde{e}$ and choose $g_e \in G$ with $\partial_1\tilde{e} = g_e.\widetilde{\partial_1 e}$, then define G_e to be the stabilizer of $\tilde{e}$. Recall from earlier in this section that in a graph of groups we have injections $\alpha(e) : G_e \to G_{\partial_0 e}$. Let us write $\alpha_i(e)$ for the injection $G_e \to G_{\partial_i e}$, where $i \in \{0, 1\}$. Then we have that $\alpha_0(e)$ is the natural inclusion map and $\alpha_1(e)$ is induced by conjugating by g_e.

Conversely, we have shown that if G acts on a tree Y without inversions, then we can produce a graph of groups over $\Gamma = G\backslash Y$. The theory culminates in the following fundamental result.

Theorem 11.8. [6, Theorem 4.3]. *Given a graph of groups $\mathcal{G}$ with fundamental group G_Γ and underlying space X_Γ, we have an action of G_Γ on the tree Z defined in Definition 11.4.*

Conversely, given a group G acting on a tree Y without inversions, we can produce a graph of groups over $\Gamma = G\backslash Y$. The vertex groups (resp. edge groups) of this graph of groups are vertex-stabilizers (resp. edge-stabilizers) of images of vertices (resp. edges) of the maximal subtree T of Γ under the map j from Proposition 11.7, where they are stabilizers with respect to the induced action of G on $j(T)$.

The two constructions described above are inverses of each other up to isomorphism and up to replacing the maps $\alpha_i(e)$ *(*$i \in \{0,1\}$*) in the graph of groups by conjugate homomorphisms.*

Proof. Many details of the proof were already given in the above discussion, but we summarize it here. Given the action of G on Y, we have seen that the graph of groups over $\Gamma = G\backslash Y$ corresponds to a graph of spaces with total space $G\backslash(\tilde{U} \times Y)$. Here U is a connected CW complex with fundamental group G. We can then collapse the $\tilde{U} \times Y$ to obtain the tree Y and a G-action on it.

On the other hand, suppose we are given a graph of groups $\mathcal{G}$ and construct, as described earlier, an action of the corresponding fundamental group G_Γ on a tree Z. We noticed earlier that the vertex- and edge-stabilizers agree up to conjugation by some element with the images in G_Γ of the G_v's and G_e's.

We need to be more careful in dealing with the injections $\alpha_i(e)$ described above because it is not *a priori* clear when $\alpha_1(e)$ will give an inclusion. Take $\mathcal{X}$ to be a graph of spaces with choices of basepoints, and identify Γ as a subset of $\mathcal{X}$ with a choice of maximal tree T of Γ. T satisfies the properties for existence of a universal cover, so we have such a cover $\tilde{T} \subseteq \tilde{X}_\Gamma$. Given an edge e of $\Gamma - T$, we can lift it uniquely to $\tilde{e} \subset \tilde{X}_\Gamma$ with $\partial_0\tilde{e} \in \tilde{T}$, and we can also produce a unique $g_e \in G_\Gamma$ with $g_e^{-1}.\partial_0\tilde{e} \in \tilde{T}$. Then $\mathcal{G}$ is isomorphic to a graph of groups for which every $\alpha_0(e)$ and those $\alpha_1(e)$ where $e \in T$ are inclusions. This means we can consider the G_v's and G_e's as subgroups of G_Γ.

The image of the tree $\tilde{T}$ under the projection $\pi : \tilde{X}_\Gamma \to Z$ will be an isomorphic tree, which we will also write as $\tilde{T}$, over T. We have seen that G_Γ acts on Z and that the resulting graph of groups will have the same subgroups G_v and G_e. We have also seen in the preceding paragraph that the $\alpha_0(e)$ and each $\alpha_1(e)$ with $e \in T$ are inclusions. If $e \notin T$, then $\alpha_1(e)$ is induced by conjugating by g_e, and in the given graph of groups, $\alpha_1(e)$ was induced by a map $f_e^1 : (X_e, *) \to (X_{\partial_1 e}, *)$, where $*$ represents the arbitrary basepoint. Since T is a maximal tree, we can find a unique path p in T that joins $\partial_1 e$ to $\partial_0 e$. We have identified G_e with a subgroup of $G_{\partial_0 e}$, so f_e^1 does not preserve basepoints, as these are translated along p. It follows that $\alpha_1(e)$ is induced by conjugation by $g' \in G_\Gamma$ corresponding to the path $p.e$ (i.e., the path starting at the identity and then doing p from there). In $\tilde{T}$, p lifts to a path joining $\widetilde{\partial_1 e}$ to $\widetilde{\partial_0 e} = \partial_0\tilde{e}$; the lift of $p.e$ therefore joins it to $\partial_1\tilde{e}$. It follows that $g'.\widetilde{\partial_1 e} = \widetilde{\partial_1 e}$, so g' is identified with the element g_e from earlier, which shows the correspondence holds. □

We conclude this article with two special cases of the above result that are particularly commonly-encountered. These are Corollaries 4.4 and 4.5 in [6].

Corollary 11.9. *If* $G\backslash Y$ *is a single edge* e *between distinct vertices* v *and* v'*, then* $G \simeq G_v *_{G_e} G_{v'}$*, where* G_i *now indicates the stabilizer of* i *(and* i *is either a vertex or an edge). If* $G\backslash Y$ *is a single loop, then* $G \simeq G_v *_{G_e}$*.*

Example 11.10. I borrow this example from [8], although J.P. Serre gives similar examples in [7].

The group $\mathrm{SL}_2(\mathbb{Z})$ of 2×2 integer-valued matrices with determinant 1 acts on the complex upper half plane $\mathbb{H}^2 = \{z \in \mathbb{C} \mid \mathrm{Im}(z) > 0\}$ by fractional linear transformations, i.e., by

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix} .z := \frac{az+b}{cz+d}$$

for all $z \in \mathbb{C}$. Quotienting out by the scalar transformations with unit determinant (i.e., the linear transformations I and $-I$) gives the projective special linear group $\mathrm{PSL}_2(\mathbb{Z})$, which also acts on $\mathbb{H}^2$. This action is generated by the translation $z \mapsto z+1$ and the inversion $z \mapsto -\frac{1}{z}$, so a fundamental domain (i.e., a region containing a unique representative for each orbit) is given by

$$\{z \in \mathbb{C} \mid |z| \geq 1, |\mathrm{Re}(z)| \leq \frac{1}{2}\}.$$

This domain contains the segment $[i, e^{2\pi i/3}]$, whose images under the action form a tree on which $\mathrm{PSL}_2(\mathbb{Z})$ acts without inversions. The stabilizer of this edge is trivial, while the stabilizer of the vertex i is generated by $z \mapsto -\frac{1}{z}$, which is of order 2, and the stabilizer of the vertex $e^{2\pi i/3}$ is generated by $z \mapsto 1 - \frac{1}{z}$, which is of order 3. It follows from Corollary 11.9 that

$$\mathrm{PSL}_2(\mathbb{Z}) \simeq (\mathbb{Z}/2\mathbb{Z}) * (\mathbb{Z}/3\mathbb{Z})$$

and therefore that

$$\mathrm{SL}_2(\mathbb{Z}) \simeq (\mathbb{Z}/4\mathbb{Z}) *_{\mathbb{Z}/2\mathbb{Z}} (\mathbb{Z}/6\mathbb{Z}).$$

12. Acknowledgements

The author thanks his advisor, Inna Bumagin, for suggesting the topic of this paper, for providing many useful comments and suggestions, and for running a reading seminar in which the author was required to present about [6] as well as several other group-theoretic articles. The author also thanks his second reader, Charles Starling, for the many suggested changes that have been incorporated into this work. Finally, the author thanks Thomas Dykes for running the aforementioned reading seminar with him during Professor Bumagin's sabbatical and for his recurrent expatiation upon the elegance of Cayley graphs and Bass–Serre trees, which helped spark the author's interest in geometric group theory.

References

[1] Fulton, William. *Algebraic Topology: A First Course*. Springer Science+Media, Inc. (1995).
[2] Hatcher, Allen. *Algebraic Topology*. Cambridge University Press (2002).
[3] Higman, Graham; Neumann, B.H.; and Neumann, Hanna. *Embedding Theorems for Groups*. Journal of the London Mathematical Society. s1-24 (4): 247–254. (1949).
[4] Munkres, James. *Topology: Second Edition*. Prentice Hall, Inc. (2000).
[5] *Planet Math*. Available at `http://planetmath.org/aconnectedandlocallypathconnectedspaceispathconnected`. (2013).
[6] Scott, Peter and Wall, Terry. *Topological methods in group theory*. Published in London Mathematical Society Lecture Notes Series, vol. 36. Cambridge University Press. (1979).
[7] Serre, Jean-Pierre. *Trees*. Springer-Verlag. (1977).
[8] Wilton, H.J.R. *Bass–Serre Theory*. Talk given on June 2, 2004. Notes available at `https://www.dpmms.cam.ac.uk/~hjrw2/Talks/bassserre.pdf`.