COVERING SPACES, GRAPHS, AND GROUPS

CARSON COLLINS

Abstract. We introduce the theory of covering spaces, with emphasis on explaining the Galois correspondence of covering spaces and the deck transformation group. We focus especially on the topological properties of Cayley graphs and the information these can give us about their corresponding groups. At the end of the paper, we apply our results in topology to prove a difficult theorem on free groups.

Contents

1. Covering Spaces

The aim of this paper is to introduce the theory of covering spaces in algebraic topology and demonstrate a few of its applications to group theory using graphs. The exposition assumes that the reader is already familiar with basic topological terms, and roughly follows Chapter 1 of Allen Hatcher's *Algebraic Topology*.

We begin by introducing the covering space, which will be the main focus of this paper.

Definition 1.1. A **covering space** of X is a space $\widetilde{X}$ (also called the covering space) equipped with a continuous, surjective map $p : \widetilde{X} \to X$ (called the covering map) which is a local homeomorphism. Specifically, for every point $x \in X$ there is some open neighborhood U of x such that $p^{-1}(U)$ is the union of disjoint open subsets V_λ of $\widetilde{X}$, such that the restriction $p|_{V_\lambda}$ for each V_λ is a homeomorphism onto U.

Informally, every neighborhood of a point in a covering space must "look like" a neighborhood in the space it covers, and every point in the space being covered

Date: August 26, 2018.

must have a neighborhood that "looks like" some neighborhood of the covering space.

We also have terminology that allows us to more easily refer to the properties of the covering space.

Definition 1.2. An open set $U \subset X$ is said to be **evenly covered** by the covering space $p : \widetilde{X} \to X$ if $p^{-1}(U)$ is the union of disjoint open subsets of $\widetilde{X}$ mapped homeomorphically onto U by p. These disjoint open subsets of $\widetilde{X}$ are called the **sheets** of $\widetilde{X}$ over U.

By the definition of a covering space, every point $x \in X$ has some evenly covered neighborhood U. We will often refer to $p : \widetilde{U} \to U$ and $p^{-1} : U \to \widetilde{U}$, the homeomorphisms we obtain by restricting the domain of p to a single sheet $\widetilde{U}$ of $\widetilde{X}$ over U.

After digesting these definitions, one might wonder what covering spaces can possibly tell us about the spaces they cover. We will begin to explore this relationship by considering a few examples of covering spaces which will prove useful throughout this paper.

First, consider the unit circle S^1. Locally, the space S^1 near a point x must look the same as its covering space $\widetilde{X}$ near a point $y \in p^{-1}(\{x\})$. Essentially, the space $\widetilde{X}$ should look like a line near y, since an open neighborhood containing x is homeomorphic to the open unit interval. Actually, S^1 can be a covering space of itself using covering maps other than the identity. For example, if we treat S^1 as the unit circle in the complex plane, then the map $p : S^1 \to S^1$ given by $p(z) = z^6$ satisfies all the conditions of a covering space. The sheets of $U = S^1 \setminus \{1\}$ are six disjoint open arcs of angle $\frac{\pi}{3}$ in the covering space S^1, and $p^{-1}(\{1\})$ contains the six endpoints of these arcs. Another covering space of S^1 is $\mathbb{R}$ using the covering map $p : \mathbb{R} \to S^1$ given by $p(t) = \cos(2\pi t) + i \sin(2\pi t)$. This covering space can be viewed as an infinite helix of radius 1, being projected onto the unit circle.

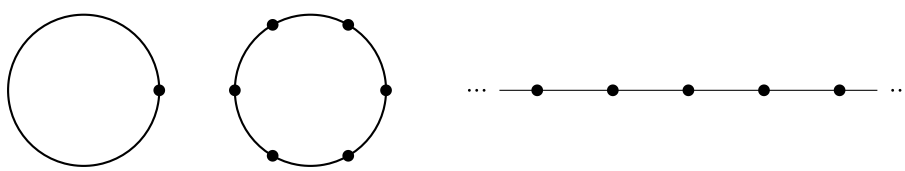

FIGURE 1. Covering spaces of S^1, with covering maps $p(z) = z$, $p(z) = z^6$, and $p(t) = \cos(2\pi t) + i \sin(2\pi t)$. The points in bold are elements of the fiber of 1.

While we used geometric language to describe the previous spaces, they can also be thought of as graphs. Specifically, the topological definition of a graph which we will use is:

Definition 1.3. A **graph** is a topological space formed from a discrete set X_0 and copies of the closed unit interval I_λ as follows: we take the disjoint union of X_0 and the intervals, and then for every I_λ in this space, we identify its endpoints with points in X_0. The resulting space is the graph X, with vertex set X_0 and edges corresponding to the intervals I_λ. A subset U of a graph X with edges $\{e_\lambda\}$ is open if and only if for each e_λ, $U \cap e_\lambda$ is open in e_λ.

Note that this definition gives a space consistent with our general interpretation of a graph, and we may frequently use terminology from graph theory to refer to graphs in topology. One important note is that a graph is connected in the graph theory sense if and only if it is connected in the topological sense. Also, a connected graph is path-connected.

Under this definition, the unit circle can be seen as the graph of one vertex and one edge. The covering map given by $p(z) = z^6$ no longer maps a space to itself, but instead maps the cycle on six vertices to the cycle on one vertex. The infinite helix of $\mathbb{R}$ becomes the infinite linear graph. The fact that these spaces can all be described as graphs will allow future results to characterize them much more strongly than a generic space. Topologically, graphs are important in two regards: they are simple spaces whose structure is easily understood, and they can be constructed to give geometric or topological representation of the algebraic structure of a group. Along these latter lines:

Definition 1.4. The **Cayley graph** Γ of a group G generated a set S is the graph with vertices corresponding to the elements of G, and a directed edge from $g_1 \in G$ to $g_2 \in G$ if and only if $g_1 s = g_2$ for some $s \in S$.

Directed edges are not topologically different from ordinary edges, but it is helpful to be able to refer to this concept in constructions. In particular, note that this definition forbids a graph from having a bidirectional edge; instead, it would have two oppositely directed edges joining the same two vertices.

We have already seen two Cayley graphs: the cycle on six vertices is the Cayley graph of $\mathbb{Z}/6\mathbb{Z}$ generated by 1, and the infinite linear graph is the Cayley graph of $\mathbb{Z}$ generated by 1.

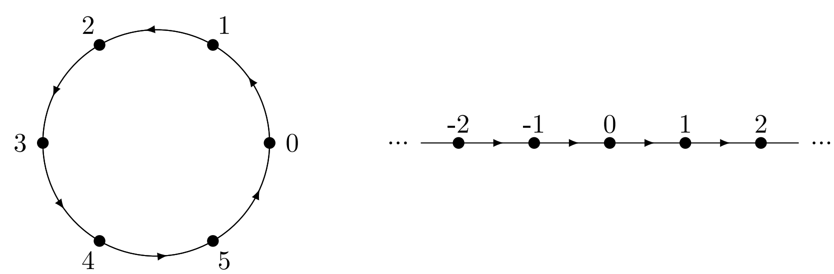

FIGURE 2. The Cayley graphs of $\mathbb{Z}/6\mathbb{Z}$ and $\mathbb{Z}$, both for generating set $\{1\}$.

It is worth noting that a group does not necessarily have a unique Cayley graph: the Cayley graph of a group depends on the choice of generators. In the cases we consider, our choice of generating set will be clear from context. Every group G has at least one Cayley graph, given by letting every element of the group be a generator; this is the complete graph on $|G|$ vertices.

We will discuss a few more covering spaces in terms of Cayley graphs, but first, we review a special kind of group that will appear frequently in this paper.

Definition 1.5. The **free group** F_S is the group generated by the elements of S (or by $|S|$ elements), such that two products of generators are equal if and only if the group axioms require them to be equal.

That is, a free group is generated by some set of elements such that no product of generators is the identity, except the trivial case where each generator in the product meets its own inverse. We often describe the elements of the free group as words; for example, the free group on two generators F_2 has elements like ab, $b^3a^{-2}b$, and $ababab$. In this model, the group law of the free group is concatenation of words. We can assume that the elements of a free group are fully reduced; e.g. $aba^2a^{-1}b$ should be written instead as $abab$.

The next examples will be covers of $S^1 \vee S^1$, where $\vee$, called the wedge sum, is the union of spaces joined together at a point. We note that $S^1 \vee S^1$ can also be regarded as a graph with one vertex and two edges.

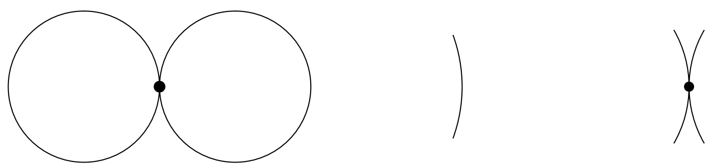

FIGURE 3. The wedge sum of two circles, along with the two basic kinds of open neighborhoods of a point in $S^1 \vee S^1$.

Near any point, $S^1 \vee S^1$ resembles either a line or two intersecting lines, and its covering space must be similar. A space which meets this requirement is the lattice graph Γ on $\mathbb{Z} \times \mathbb{Z}$. If we let p be the map which takes each vertex of Γ to the basepoint of $S^1 \vee S^1$, each vertical edge to one circle, and each horizontal edge to the other circle, then it is not hard to show that p is a covering map. Of course, Γ can also be seen as the Cayley graph of the group $\mathbb{Z} \times \mathbb{Z}$ generated by $(1, 0)$ and $(0, 1)$.

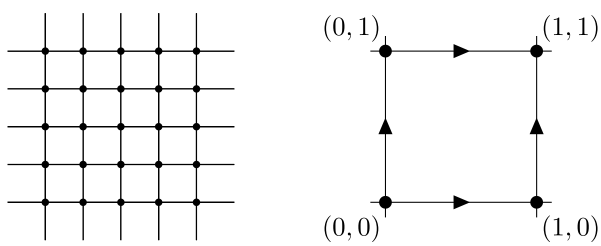

FIGURE 4. The Cayley graph of $\mathbb{Z} \times \mathbb{Z}$ along with an example labeling for generators $(1, 0)$ and $(0, 1)$.

A similar covering space is constructed from a base vertex by constructing edges emanating in four directions, placing vertices at $\frac{2}{3}$ the length of each edge from the base vertex, and then repeating the process for each of these vertices. The infinite graph Γ resulting from this process, shown in Figure 5, is also a covering space of $S^1 \vee S^1$ with very similar covering map: we define $p : \Gamma \to S^1 \vee S^1$ to be the map taking each vertex of Γ to the basepoint of $S^1 \vee S^1$, each vertical edge to the one circle of $S^1 \vee S^1$, and each horizontal edge to the other circle.

This space is the Cayley graph of F_2 corresponding to generators a, b; associating right edges with a and upward edges with b gives a bijection between vertices and elements of F_2, and the construction of the graph guarantees that it contains no cycles, so no two elements of the free group are identified and it has an empty relations set. The property that a graph contains no cycles, and more generally that a covering space contains no nontrivial loops, marks a very special kind of covering space. Compare the Cayley graph of $\mathbb{Z}$ as another example of a space with this property, called simply-connectedness, which we will explore later.

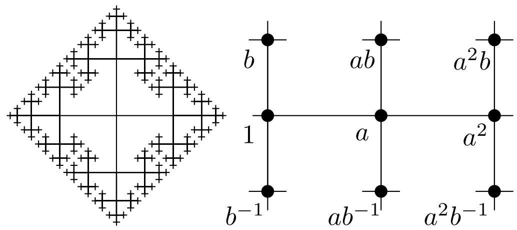

FIGURE 5. The Cayley graph of F_2, along with an example labeling for generators a and b.

A common theme in these examples has been understanding how the sheets of a covering space are connected. One result we will eventually prove is that the nature of this property uniquely defines a covering space up to a certain kind of isomorphism. However, first we need to establish a formal way to explore this property.

2. The Fundamental Group

Recall that a path is a continuous function mapping the unit interval I into a space X, and a loop is a path f with $f(0) = f(1)$. Our key insight will be that if f is a path connecting two points in the covering space $\widetilde{X}$ such that $p(f(0)) = p(f(1))$, then $p \circ f$ must be a loop in X. Therefore, the loops of a space are related to how the sheets of its covering spaces are connected.

In particular, using path homotopy, we can partition the loops with basepoint x_0 into equivalence classes. We can also define the composition of loops γ, η given by $\gamma \cdot \eta$ to be the loop which traverses γ first, then η. A basic theorem of topology gives that the equivalence classes of loops with multiplication rule $[\gamma] \cdot [\eta] = [\gamma \cdot \eta]$ forms a group.

Definition 2.1. The **fundamental group** of a space X with basepoint x_0, denoted $\pi_1(X, x_0)$, consists of the equivalence classes of loops in X with basepoint x_0, with the group law given by composition of loops, as described above.

The subscript 1 indicates that this is the first of many homotopy groups that can be associated to the space X; however, we will have no need to consider any higher subscripts in this paper. Also, from now on we may use the notation (X, x_0) to refer to a space equipped with a given basepoint. Note that we will always require that $p(\tilde{x}_0) = x_0$ for a covering space $p : (\widetilde{X}, \tilde{x}_0) \to (X, x_0)$.

For an example, consider the plane $\mathbb{R}^2$. For any two loops f, g in this space sharing a basepoint x_0, we can use the homotopy $h(t, s) = (1-s)f(t) + sg(t)$ to send each point $f(t)$ along the segment joining it to $g(t)$, continuously deforming f into g. This construction tells us that all loops are homotopic in $\mathbb{R}^2$, so $\pi_1(\mathbb{R}^2, x_0) = \{0\}$.

On the other hand, we may fail to have homotopy between loops when our space has certain kinds of holes. For example, now consider the plane with a removed point, $\mathbb{R}^2 \setminus \{p\}$. If f is a loop which goes around p and g is a loop which does not go around p, then no homotopy takes f to g, because either f will not be a loop when the homotopy passes it over p, or the homotopy will fail to be continuous. One way to imagine this is that the hole in the plane is a peg, and the loop f is a circle of string lying in the plane and wrapped around the peg. We cannot possibly stretch or shift the string so that it is no longer wrapped around the peg, all while keeping the string in the plane.

In short, the fundamental group tells us about the holes in a space, since these are what prevent two paths or loops from being homotopic. To close this section, we list a few results about the fundamental group which are not too difficult and will prove useful in subsequent discussions.

Proposition 2.2. *A continuous map* $f : (Y, y_0) \to (X, x_0)$ *induces a homomorphism* $f_* : \pi_1(Y, y_0) \to \pi_1(X, x_0)$.

Specifically, for any loop in Y based at y_0, its composition with f is a loop in X based at x_0, and composition with f maps a product of loops to the composition of their images. Thus, f_* maps $\pi_1(Y, y_0)$ into a subgroup of $\pi_1(X, x_0)$. The most important induced homomorphism for our purposes will be $p_* : \pi_1(\widetilde{X}, \tilde{x}_0) \to \pi_1(X, x_0)$, the homomorphism induced by a covering map $p : \widetilde{X} \to \tilde{x}_0$.

Proposition 2.3. *Let* X *be path-connected, and* $x_1, x_2 \in X$. *Then* $\pi_1(X, x_1) \simeq \pi_1(X, x_2)$.

As such, we will occasionally refer to the fundamental group as $\pi_1(X)$ when the choice of basepoint is not relevant to the discussion at hand.

Proposition 2.4. *If two spaces are homotopy equivalent, then they have isomorphic fundamental groups.*

Exact proofs of the previous remarks can be found in [1].

3. Lifts

Now that we have the covering space and the fundamental group, we can begin to prove the key relation between the two which makes covering spaces so useful to study. Our first task along these lines will be to introduce lifts, which give us another powerful way to relate the base space to its covering space using the covering map.

Definition 3.1. A lift of a map $f : Y \to X$ to the covering space $\widetilde{X}$ with covering map $p : \widetilde{X} \to X$ is a continuous function $\tilde{f} : Y \to \widetilde{X}$ such that $p \circ \tilde{f} = f$.

As the figure shows, the lift of a loop is not necessarily a loop. Exactly the loops in X contained by classes in $p_*(\pi_1(\widetilde{X}, \tilde{x}_0))$ lift to loops in $\widetilde{X}$ beginning at $\tilde{x}_0$. Lifts have a number of useful properties, including that a lift always exists given a condition on the fundamental groups of Y and $\widetilde{X}$ and that a lift is uniquely defined by the value it takes at a single point. In the above example, we could have lifted our loop to a path beginning at any of the six points in the fiber of $(1, 0)$.

The following result is a construction to show that any homotopy f_t can be lifted given a lift of f_0. Having this construction will allow us to quickly show the existence of other kinds of lifts. The key idea of the proof is that we can lift the homotopy on an open neighborhood around each point in its domain with finitely

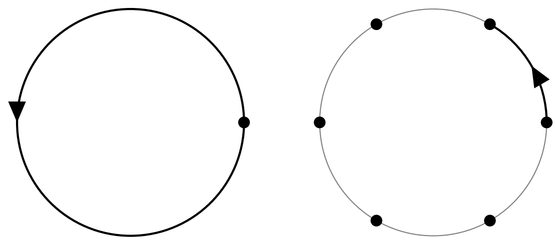

FIGURE 6. A loop going once around S^1, counterclockwise, and its lift beginning at $(1,0)$ under the covering map $p : S^1 \to S^1$ with $p(z) = z^6$.

many steps. Then, by showing that two lifts of the homotopy at a point must agree, we can paste these lifts together to obtain a lift of the entire homotopy. The complete details of this proof will not prove important, so the reader need only understand the results to safely proceed to the next section.

Theorem 3.2. *Given a covering space* $p : \widetilde{X} \to X$*, a homotopy* $F : Y \times I \to X$*, and a map* $\widetilde{F} : Y \times \{0\} \to \widetilde{X}$ *lifting* $F|Y \times \{0\}$*, there exists a unique lift* $\widetilde{F} : Y \times I \to \widetilde{X}$ *of* F *which restricts to the given* $\widetilde{F}$ *on its domain.*

Proof. For any $y_0 \in Y$ and $t \in I$, the point $F(y_0, t)$ has an evenly covered open neighborhood V_t in X, and the continuity of F guarantees that $F^{-1}(V_t)$ is an open neighborhood. Fixing y_0 and varying t, we can write each $F^{-1}(V_t)$ as $N_t \times (a_t, b_t)$, where N_t is an open neighborhood of y_0 in Y and (a_t, b_t) is an open neighborhood of t in I. Since $\{y_0\} \times I$ is compact, finitely many of the $N_t \times (a_t, b_t)$ cover it. Let N be the open neighborhood which is the intersection of the finitely many N_t, and let $\{t_0, t_1 \ldots, t_m\}$ be a partition of I taken from the finitely many a_t and b_t. Then for each i, $N \times (t_i, t_{i+1}) \subset N_t \times (a_t, b_t)$ for some t, so $F(N \times (t_i, t_{i+1})) \subset V_t$. Therefore, $F(N \times (t_i, t_{i+1}))$ is contained in an evenly covered open neighborhood, which for convenience we will label U_i.

Suppose inductively that $\widetilde{F}$ has been defined on $N \times [0, t_i]$. Since U_i is evenly covered, there is some open $\widetilde{U}_i \subset \widetilde{X}$ containing $\widetilde{F}(y_0, t_i)$ which is mapped homeomorphically by p onto U_i. Let $N_1 \times \{t_i\} = (\widetilde{F}|N \times \{t_i\})^{-1}(\widetilde{U}_i) \subset N \times \{t_i\}$, and note that N_1 is also an open neighborhood of y_0, with the property that $F(N_1 \times \{t_i\}) \subset \widetilde{U}_i$. Then, define $\widetilde{F}$ on $N_1 \times [t_i, t_{i+1}]$ to be $p^{-1} \circ F$ where $p^{-1} : U_i \to \widetilde{U}_i$ is the inverse of the covering map restricted to $\widetilde{U}_i$, where it is bijective. Since our partition is finite, the induction obtains some open neighborhood $N_f \subset \cdots \subset N_1 \subset N$ of y_0 so that $\widetilde{F}$ is defined on $N_f \times I$. Furthermore, $\widetilde{F}$ is a lift of F, since $F = p \circ \widetilde{F}$ on every $N_f \times [t_i, t_{i+1}]$ and $\widetilde{F}$ is pasted continuously at each t_i.

Now, we wish to show the uniqueness of this lift at any point in Y. To do so, consider the case where the set Y from the statement is a point, and where we can consider F to be a function on I. Suppose two lifts of F satisfy $\widetilde{F}(0) = \widetilde{F}'(0)$, and let $\{t_0, \ldots, t_m\}$ be a partition of I such that each $F([t_i, t_{i+1}])$ is contained in an evenly covered neighborhood U_i. We proceed by induction, assuming that $\widetilde{F} = \widetilde{F}'$ on $[0, t_i]$. Then the continuity of both lifts means that they preserve the connectedness of $[t_i, t_{i+1}]$. In particular, each lift can only map into one of the $\widetilde{U}_i$, since these are disjoint open sets, and their union is disconnected. But since

$\widetilde{F}(t_i) = \widetilde{F}'(t_i)$, they map into the same $\widetilde{U}_i$, and since p is injective when restricted to $\widetilde{U}_i$, we have that $p\widetilde{F} = p\widetilde{F}' = F$ implies $\widetilde{F} = \widetilde{F}'$ on $[t_i, t_{i+1}]$. By induction, the lift of F restricted to a point in Y is unique.

Thus, we have shown that we can lift F on some open neighborhood N around each $y_0 \in Y$, and now we have that if two such neighborhoods intersect, their respective lifts must agree when restricted to any point in the intersection. It follows that we obtain a lift $\widetilde{F} : Y \times I \to \widetilde{X}$ from pasting together these lifts, which is continuous since it is continuous on each $N \times I$ and unique since it is unique when restricted to any point in Y. □

Now given a function, its lift, and a function homotopic to the first, we can find a lift of this last function homotopic to the lift of the first. One very useful corollary of this result is that if the loop γ in X lifts to a loop $\tilde{\gamma}$ with basepoint $\tilde{x}_0$ in $\widetilde{X}$, then every loop in the equivalence class $[\gamma]$ lifts to a loop with basepoint $\tilde{x}_0$ in $\widetilde{X}$. This allows us to prove a stronger result regarding relationship between the fundamental groups of a space and its covering space.

Theorem 3.3. *Given a space (X, x_0) and a covering space $p : (\widetilde{X}, \tilde{x}_0) \to (X, x_0)$, the induced homomorphism $p_* : \pi_1(\widetilde{X}, \tilde{x}_0) \to \pi_1(X, x_0)$ is injective. The elements of $p_*(\pi_1(\widetilde{X}, \tilde{x}_0))$ are exactly the equivalence classes of loops in X with basepoint x_0 which lift to loops in $\widetilde{X}$ with basepoint $\tilde{x}_0$.*

Proof. The kernel of p_* consists of loops $\tilde{f}_0$ such that $f_0 = p\tilde{f}_0$ is homotopic to the trivial loop. We can lift such a homotopy to $\widetilde{X}$, and find that $\tilde{f}_0$ is homotopic to a lift of the trivial loop, which is clearly trivial in $\widetilde{X}$. Thus, only the trivial loop maps to the trivial loop, and p_* is injective.

In one direction, a class containing a loop γ in X which lifts to a loop $\tilde{\gamma}$ in $\widetilde{X}$ is clearly a member of $p_*(\pi_1(\widetilde{X}, \tilde{x}_0))$, since $p \circ \tilde{\gamma}$ explicitly gives the corresponding element of the group. On the other hand, any element of $p_*(\pi_1(\widetilde{X}, \tilde{x}_0))$ is a class of loops in X which contains at least one loop which is the image of a loop in $\widetilde{X}$ under p, and thus has a lift to a loop. Therefore, for any other loop in the same class there is a homotopy which takes that loop to the loop with a lift, and the lift of the homotopy shows that this loop must also have a lift to a loop in $\widetilde{X}$. □

We can also use our construction of the homotopy lift to prove the existence of general lifts, given certain conditions on the domain of the function we are lifting. The key idea of these conditions is that the easiest way to define a lift of a function is by lifting paths in X with a designated basepoint. Our theorem on homotopy lifts guarantee that the lift of a path exists and is unique. However, we need to know that two paths with the same endpoint in Y, mapped by f to two paths with the same endpoint in X, will lift to two paths with the same endpoint in $\widetilde{X}$, and this is exactly equivalent to asking that a loop in X which is the image of a loop in Y also be the image of a loop in $\widetilde{X}$.

Theorem 3.4 (The Lifting Criterion)**.** *Let Y be path-connected and locally path-connected, with basepoint y_0. Let $p : (\widetilde{X}, \tilde{x}_0) \to (X, x_0)$ be a covering space. A lift $\tilde{f} : (Y, y_0) \to (\widetilde{X}, \tilde{x}_0)$ of a continuous map $f : (Y, y_0) \to (X, x_0)$ exists if and only if $f_*(\pi_1(Y, y_0)) \subset p_*(\pi_1(\widetilde{X}, \tilde{x}_0))$ where $f_* : \pi_1(Y, y_0) \to \pi_1(X, x_0)$ is the fundamental group homomorphism induced by f.*

Proof. Suppose $p(\tilde{f}) = f$. If $[\gamma] \in \pi(Y, y_0)$ satisfies $[f \circ \gamma] \in \pi_1(X, x_0)$, then $[p \circ \tilde{f} \circ \gamma] \in \pi_1(X, x_0)$. Recalling how f_* and p_* are defined, we obtain $f_*(\pi_1(Y, y_0)) \subset p_*(\pi_1(\widetilde{X}, \tilde{x}_0))$.

If $f_*(\pi_1(Y, y_0)) \subset p_*(\pi_1(\widetilde{X}, \tilde{x}_0))$, then we construct $\tilde{f}$ in the following way: let $y \in Y$ and let γ be a path from y_0 to y; then $f \circ \gamma$ is the image of this path in X, and there exists a unique lift $\widetilde{f\gamma}$ of this path starting at $\tilde{x}_0$. We then define $\tilde{f}(y) = \widetilde{f\gamma}(1)$.

We can show that the function is well-defined by considering two paths γ and γ' in Y from y_0 to y. The two paths form a loop in Y which f maps to a loop in X, and since $f_*(\pi_1(Y, y_0)) \subset p_*(\pi_1(\widetilde{X}, \tilde{x}_0))$, this loop is the image of a loop in $\widetilde{X}$ under p. Therefore, the lifts of our two paths have the same endpoint in $\widetilde{X}$, since the lifts form a loop.

Finally, we prove that $\tilde{f}$ is continuous. Let $U \subset X$ be an evenly covered open neighborhood of $f(y)$ for some $y \in Y$, such that $\widetilde{U}$ is a sheet of $\widetilde{X}$ mapped homeomorphically by p onto U. Using the continuity of f and the local path-connectedness of Y, there exists a path-connected open neighborhood V of y with $f(V) \subset U$. Given a fixed path γ from y_0 to y and any path η in V from y to a point in V, we obtain the path $f\gamma \cdot f\eta$ in X, which lifts to a path with endpoint in $\widetilde{U}$. Since we let $\eta(1)$ be arbitrary in V, it follows that $\tilde{f}(V) \subset \widetilde{U}$, and $\tilde{f}$ is continuous. □

The other major property of lifts we want to show is that a lift is unique up to choice of basepoint. To do so, we show that two lifts agree on an open set and disagree on an open set, so connectedness gives that two lifts agree everywhere or nowhere.

Theorem 3.5 (The Unique Lifting Property)**.** *If two lifts $\tilde{f}_1, \tilde{f}_2 : Y \to \widetilde{X}$ of $f : Y \to X$ agree at one point of Y and Y is connected, then $\tilde{f}_1 = \tilde{f}_2$.*

Proof. Let $y \in Y$, and let $U \subset X$ be an evenly covered open neighborhood of $f(y)$. Then $p^{-1}(U)$ is the union of disjoint sheets of $\widetilde{X}$, so we must have $\tilde{f}_1(y) \in \widetilde{U}_1$ and $\tilde{f}_2(y) \in \widetilde{U}_2$ for two such sheets $\widetilde{U}_1$ and $\widetilde{U}_2$. Since both lifts are continuous, the preimages of $\widetilde{U}_1$ and $\widetilde{U}_2$ respectively are open in Y and share the point y, so we can find an open neighborhood N mapped by $\tilde{f}_1$ into $\widetilde{U}_1$ and mapped by $\tilde{f}_2$ into $\widetilde{U}_2$. If $\tilde{f}_1(y) \neq \tilde{f}_2(y)$, then $\widetilde{U}_1 \neq \widetilde{U}_2$, and hence they are disjoint. It follows that $\tilde{f}_1 \neq \tilde{f}_2$ throughout N, which implies that the set of points where $\tilde{f}_1 = \tilde{f}_2$ is the complement of a union of open sets; i.e. a closed set. On the other hand, if $\tilde{f}_1(y) = \tilde{f}_2(y)$, then $\widetilde{U}_1 = \widetilde{U}_2$. Since $p \circ \tilde{f}_1 = p \circ \tilde{f}_2$, it follows that $\tilde{f}_1 = \tilde{f}_2$ on N since p is injective on $\widetilde{U}_1 = \widetilde{U}_2$. This implies that the set of points where $\tilde{f}_1 = \tilde{f}_2$ is a union of open sets, and thus open. Since Y is connected, the only open and closed sets in Y are $\emptyset$ and Y, but by hypothesis $\tilde{f}_1$ and $\tilde{f}_2$ agree at one point, so they agree on all of Y. □

4. The Universal Covering Space

At this point, we have proven a number of useful technical results regarding lifts, and now we can finally return and use these to develop covering spaces. The first step will be to construct a covering space with a number of special properties, called the universal covering space. The first step will be a pair of new definitions.

Definition 4.1. A space X is **simply-connected** if it is path-connected and has a trivial fundamental group. A space X is **semi-locally simply-connected** if every point $x \in X$ has an open neighborhood U such that the homomorphism $\pi_1(U, x) \to \pi_1(X, x)$ induced by the inclusion U into X is trivial.

We will show in this section that every path-connected, locally path-connected, and semi-locally simply-connected space has a simply-connected covering space, which is unique up to isomorphism. This covering space is known as the universal covering space. The universal covering space will allow us to easily construct and classify all covering spaces of its base space.

Informally, simply-connectedness is a relatively rare property which more or less entails that a space have no holes in its interior. Semi-local simply-connectedness is a common property which only requires that a space have no arbitrarily small holes. Semi-local simply-connectedness does not necessarily imply that $\pi_1(U, x)$ is trivial, which would be equivalent to X being locally simply-connected. Rather, whatever equivalence classes of loops may make up the fundamental group of U, these are all homotopic to the trivial loop when considered in the larger space X.

It is not hard to show that semi-local simply-connectedness is necessary for the existence of the universal cover. Suppose $\widetilde{X}$ is a simply-connected cover of X, and $U \subset X$ is open. Then there exists an evenly covered $V \subset U$, and we can choose a sheet $\widetilde{V} \subset \widetilde{X}$ of V. Any loop in V lifts to a loop in $\widetilde{V}$, which the covering map takes to a loop in X. However, $\widetilde{X}$ is simply-connected and contains no nontrivial loops, so the lift of the loop must be trivial, and the loop will be trivial in X.

Theorem 4.2. *A path-connected, locally path-connected, semi-locally simply-connected space X with basepoint x_0 has a simply-connected covering space given by:*
$\widetilde{X} = \{[\gamma] \mid \gamma \text{ is a path in } X \text{ starting at } x_0\}$
$p : \widetilde{X} \to X \text{ where } p([\gamma]) = \gamma(1)$

Proof. Let $\widetilde{X}$ and p be as defined in the statement. If we take $\mathcal{U}$ to be the collection of path-connected open sets U in X with $\pi_1(U) \to \pi_1(X)$ trivial, then any open path-connected subset V of such a U has $\pi_1(V) \to \pi_1(X)$ also trivial, giving $V \in \mathcal{U}$. For any open subset S of X, we then can describe S as the union of sets in $\mathcal{U}$, since X is locally path-connected and semi-locally simply-connected. Thus $\mathcal{U}$ is a basis of the topology on X. For each $U \in \mathcal{U}$ and each path γ in X with $\gamma(0) = x_0$ and $\gamma(1) \in U$, we define a corresponding subset of $\widetilde{X}$:

$$U_{[\gamma]} = \{[\gamma \cdot \eta] \mid \eta \text{ is a path in } U \text{ with } \eta(0) = \gamma(1)\}$$

It is somewhat lengthy but not difficult to show that the collection of $U_{[\gamma]}$ form the basis of a topology on $\widetilde{X}$.

We now consider the restriction of the covering map to $p : U_{[\gamma]} \to U$. We can see that the set $U_{[\gamma]}$ is mapped surjectively by p onto U, since U is path-connected. We can also see that $p : U_{[\gamma]} \to U$ is injective since U is semi-locally simply-connected, so that all paths between two points in U are homotopic in X. It follows that p is a bijection, which takes an open set $V_{[\gamma']} \subset U_{[\gamma]}$ to the open set $V \subset U$, so it follows that p restricted to $U_{[\gamma]}$ is a homeomorphism. For any path γ', if $U_{[\gamma]} \cap U_{[\gamma']} \neq \emptyset$, then it follows from the path-connectedness of X that $U_{[\gamma]} = U_{[\gamma']}$. Therefore, $p^{-1}(U)$ is a union of disjoint open sets mapped homeomorphically by p onto U, and p is a covering map.

Let $[\gamma] \in \widetilde{X}$, and define γ_t to be the path γ on $[0, t]$ and the constant path $\gamma(t)$ on $[t, 1]$. The continuous map given by $f(t) = [\gamma_t]$ is thus a path in $\widetilde{X}$ from $[x_0]$ to

$[\gamma]$, and a lift of γ since $p(f(t)) = \gamma_t(1) = \gamma(t)$. It follows that $\widetilde{X}$ is path connected, since this procedure can be performed for any path. For any loop γ in X which lifts by this method to a loop in $\widetilde{X}$, we must have the endpoints of the lift be equal, or $[x_0] = [\gamma]$. But then only loops equivalent to the trivial loop lift to loops, so $p_*(\pi_1(\widetilde{X})) = 0$, and $\widetilde{X}$ is simply-connected. □

The most important takeaway from the construction of the universal cover is that given two paths with common endpoints which are not homotopic, we can lift these paths to paths beginning at the same point in the universal cover and ending in different places. Consider the example of the Cayley graph X on $\mathbb{Z}\times\mathbb{Z}$, for which the Cayley graph $\widetilde{X}$ of F_2 is a simply-connected covering space.

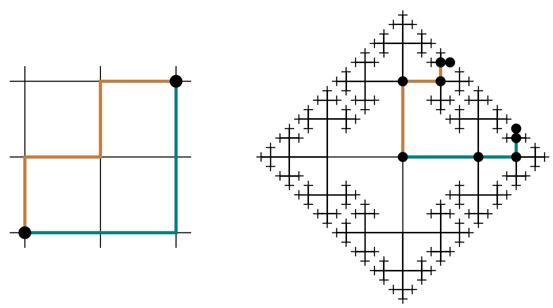

FIGURE 7. Two paths in the Cayley graph of $\mathbb{Z} \times \mathbb{Z}$ which are not homotopic lift to paths with different endpoints in the Cayley graph of F_2. Those endpoints can be associated with the homotopy class of the respective path.

Any path in X from $(0, 0)$ to another vertex is homotopic to a unique path consisting of a sequence of edges, where no edge is followed immediately by itself. This gives a straightforward bijection between such paths and the elements of F_2, written as reduced words (let a be right, b be up, and the inverses similarly). Thus, each path corresponds to a point in the Cayley graph of F_2 and also lifts to a path from the origin of the graph to that point.

With access to the universal cover, it becomes possible to construct and classify all other covering spaces in terms of the universal cover.

Theorem 4.3 (The Galois Correspondence for Covering Spaces)**.** *Let X be a space with a simply-connected covering space $\widetilde{X}$. Let H be a subgroup of $\pi_1(X, x_0)$. There exists a covering space $p : X_H \to X$ with $p_*(\pi_1(X_H, x_H)) = H$ that is unique up to isomorphism. Thus, there is a bijection between subgroups of $\pi_1(X, x_0)$ and covering spaces of X.*

Proof. We define $[\gamma] \sim [\gamma']$ for paths γ, γ' in X if $\gamma(1) = \gamma'(1)$ and $[\gamma \cdot \overline{\gamma'}] \in H$. The group properties of H guarantee that $\sim$ is an equivalence relation. Now, we consider the quotient space $X_H = \widetilde{X}/\sim$. For any γ and γ' starting at the basepoint of X_H and any other path η with $\eta(0) = \gamma(1) = \gamma'(1)$, we have $[\gamma] \sim [\gamma']$ if and only if $[\gamma \cdot \overline{\gamma'}] = [\gamma \cdot \eta \cdot \overline{\eta} \cdot \overline{\gamma'}] \in H$, which in turn is equivalent to $[\gamma \cdot \eta] \sim [\gamma' \cdot \eta]$. Thus, any point in a path connected neighborhood of $\gamma(1)$ in X has preimages in path connected neighborhoods of $[\gamma]$ and $[\gamma']$ in $\widetilde{X}$ which are identified in X_H as a single point in the neighborhood of $[\gamma] \sim [\gamma']$. This property means that the quotient map is a local homeomorphism, since it is continuous, takes open neighborhoods to open neighborhoods, and is locally invertible. Thus, the map $p : X_H \to X$ given by $p([\gamma]) = \gamma(1)$ inherited from $\widetilde{X}$ is also a local homeomorphism, and thus a covering map.

An element of the fundamental group of X_H is the image of a path in $\widetilde{X}$ beginning at its basepoint $[c]$ with endpoint $[\gamma] \sim [c]$, where c is the constant path in X at x_0. But $[\gamma] \sim [c]$ implies that $\gamma(0) = \gamma(1) = x_0$ and $[\gamma \cdot \bar{c}] = [\gamma] \in H$. Conversely, every loop γ in X with $[\gamma] \in H$ can be lifted to a path $\tilde{\gamma}$ ending at $[\gamma]$ in X_H, which will invariably be a loop since $[\gamma \cdot c] = [\gamma] \in H$. Therefore, $p_*(\pi_1(\widetilde{X}, \tilde{x}_0)) = H$.

Suppose covering spaces (X_1, x_1) and (X_2, x_2) with covering maps p_1, p_2 both have fundamental group H. Then the lifting criterion implies the existence of a lift $f : X_1 \to X_2$ and a lift $g : X_2 \to X_1$ which each maps basepoint to basepoint. Thus the composition $g \circ f$ maps x_1 to x_1 and is a lift, since $p_1 \circ g \circ f = p_2 \circ f = p_1$, so the unique lift property gives that $g \circ f = id_{X_1}$. Symmetric reasoning gives that $f \circ g = id_{X_2}$. Thus, f is a homeomorphism between X_1 and X_2 which preserves the covering map since $p_1 = p_2 \circ f$, so f is an isomorphism and the two covering spaces are isomorphic. □

This bijection between covering spaces and subgroups of the fundamental group of X is called the Galois Correspondence. We can strengthen this bijection by placing partial orderings on both sets. The convenient choice for groups is to use subgroup inclusion. For covering spaces, we say one covering space is less than the other if it is a covering space of the other; that is, for $p_1 : \widetilde{X}_1 \to X$ and $p_2 : \widetilde{X}_2 \to X$, if there exists a covering map $p : \widetilde{X}_1 \to \widetilde{X}_2$ such that $p_1 = p_2 p$, then $p_1 < p_2$. The remarkable property of this correspondence is that it preserves the partial ordering: if $H_1 \subset H_2 \subset \pi_1(X, x_0)$ are subgroups, then X_{H_1} covers X_{H_2}, which covers X. An immediate corollary is that the universal covering space covers every other covering space, using exactly the quotient maps from the proof of the theorem.

5. The Deck Transformation Group

There is one final tool to understand covering spaces in terms of the fundamental group of the covered space. We previously discussed the notion of an isomorphism between covering spaces $p_1 : \widetilde{X}_1 \to X$ and $p_2 : \widetilde{X}_2 \to X$ as a homeomorphism $f : \widetilde{X}_1 \to \widetilde{X}_2$ such that $p_1 = p_2 f$. This naturally leads to a definition of automorphism for a covering space, and thus a group of automorphisms of a given covering space. We call these automorphisms deck transformations.

Definition 5.1. A **deck transformation** is an isomorphism from a covering space to itself. We will use the notation $Aut(p)$ to refer to the group of deck transformations of the covering space $p : \widetilde{X} \to X$ with group law given by composition of maps.

An important note here is that a deck transformation can also be seen as a lift of the covering map p. As a result of the unique lifting property, it follows that any deck transformation is given by its value at a single point. We will often characterize deck transformations by their action on the fiber of x_0; a deck transformation bijectively maps this set to itself, so we obtain a homomorphism from $Aut(p)$ to the permutation group of the fiber.

A particularly interesting question then is whether the deck transformations of a covering space can map any member of the fiber of x_0 to any other member.

Definition 5.2. A **regular** or **normal** covering space satisfies that for any $x \in X$ and any $\tilde{x}_1, \tilde{x}_2 \in p^{-1}(x)$, there exists a deck transformation taking $\tilde{x}_1$ to $\tilde{x}_2$.

As one might expect from the name, a covering space is regular or normal if and only if it corresponds to a normal subgroup of $\pi_1(X, x_0)$. The proof of this fact will use the idea that a loop at x_0 also determines a permutation of the fiber of x_0, since we can lift it to each point in the fiber and take the endpoints of the lifts as the image of the permutation. We can also use this idea to construct an isomorphism between the group of deck transformations of $\widetilde{X}$ and the quotient group $N(p_*(\pi_1(\widetilde{X}, \tilde{x}_0)))/p_*(\pi_1(\widetilde{X}, \tilde{x}_0))$.

Theorem 5.3. *Let $H = p_*(\pi_1(\widetilde{X}, \tilde{x}_0))$ for a covering space $p : (\widetilde{X}, \tilde{x}_0) \to (X, x_0)$. $Aut(p)$ is isomorphic to the quotient $N(H)/H$ where $N(H)$ is the normalizer of H in $\pi_1(X, x_0)$. The covering space $p : (\widetilde{X}, \tilde{x}_0) \to (X, x_0)$ is regular if and only if H is a normal subgroup of $\pi_1(X, x_0)$.*

Proof. First, we consider the effect of a change of basepoint from $\tilde{x}_0 \in p^{-1}(x_0)$ to $\tilde{x}_1 \in p^{-1}(x_0)$ on the fundamental group of $\widetilde{X}$. Let $\tilde{\gamma}$ be a path from $\tilde{x}_0$ to $\tilde{x}_1$, and note that the image of $\tilde{\gamma}$ under p is a loop γ with basepoint x_0. Then for any loop $\tilde{\eta}$ in $\widetilde{X}$ about $\tilde{x}_0$, we have that $\tilde{\gamma}\tilde{\eta}\overline{\tilde{\gamma}}$ is a loop about $\tilde{x}_1$, and thus $[\gamma\eta\overline{\gamma}] \in p_*(\pi_1(\widetilde{X}, \tilde{x}_1))$. More generally, if $H_0 = p_*(\pi_1(\widetilde{X}, \tilde{x}_0))$ and $H_1 = p_*(\pi_1(\widetilde{X}, \tilde{x}_1))$, then this gives that $\gamma H_0 \gamma^{-1} \subset H_1$, where we now use notation more traditional of conjugation. Symmetric reasoning gives that $\gamma^{-1} H_1 \gamma \subset H_0$, and conjugating both sides by γ yields that $H_1 \subset \gamma H_0 \gamma^{-1}$, so $\gamma H_0 \gamma^{-1} = H_1$.

We can view a deck transformation as a lift of p, and since an isomorphism has an inverse that is also an isomorphism, the lifting criterion gives that a deck transformation $(\widetilde{X}, \tilde{x}_0) \to (\widetilde{X}, \tilde{x}_1)$ exists if and only if $p_*(\pi_1(\widetilde{X}, \tilde{x}_0)) = p_*(\pi_1(\widetilde{X}, \tilde{x}_1))$. By the previous result, however, these two groups are conjugate, so they are equal if and only if the loop γ which lifts to a path from $\tilde{x}_0$ to $\tilde{x}_1$ is an element of the normalizer of $H = p_*(\pi_1(\widetilde{X}, \tilde{x}_0))$.

Let $\phi : N(H) \to Aut(p)$ take the loop $[\gamma]$ which lifts to a path $\tilde{\gamma}$ from $\tilde{x}_0$ to $\tilde{x}_1$ to the deck transformation τ which takes $\tilde{x}_0$ to $\tilde{x}_1$. The previous paragraph makes clear that τ exists given such a γ, and that this mapping is surjective since every deck transformation maps $\tilde{x}_0$ to some other point in :$p^{-1}(x_0)$, giving a path in $\widetilde{X}$ which maps to some loop in X which is in the normalizer of H. Furthermore, since a deck transformation is itself a lift of p and we have fixed its value at $\tilde{x}_0$, the unique lift property guarantees the uniqueness of τ, and thus shows that ϕ is well-defined.

Let γ, γ' be loops in $N(H)$ which lift to paths $\tilde{\gamma}$ and $\tilde{\gamma}'$. If τ is the deck transformation which takes $\tilde{x}_0$ to $\tilde{\gamma}(1)$ and τ' is the deck transformation which takes $\tilde{x}_0$ to $\tilde{\gamma}'(1)$, then $\tilde{\gamma} \cdot \tau(\tilde{\gamma}')$ is a lift of the path $\gamma \cdot \gamma'$, ending at $\tau(\tilde{\gamma}'(1)) = \tau(\tau'(\tilde{x}_0))$. Therefore, by the unique lifting property, we have that $\phi([\gamma])\phi([\gamma']) = \phi([\gamma\gamma'])$, so ϕ is a group homomorphism. Also, $\phi([\gamma]) = id$ if and only if γ lifts to a loop at $\tilde{x}_0$, which is equivalent to $[\gamma] \in H$, so ϕ has kernel H. Since we have shown that ϕ is a surjective group homomorhpism with kernel H, it follows that the quotient $N(H)/H$ is mapped isomorphically by ϕ onto $Aut(p)$.

If H is a normal subgroup of $\pi_1(X, x_0)$, then it is invariant under conjugation by any loop, so each deck transformation taking $\tilde{x}_0$ to another point in $p^{-1}(x_0)$ exists. Composition of these deck transformations and their inverses leads to a deck transformation from $\tilde{x}_1$ to $\tilde{x}_2$ for any $\tilde{x}_1, \tilde{x}_2 \in p^{-1}(x)$, so $\widetilde{X}$ is a regular covering space. On the other hand, if $\widetilde{X}$ is regular, then the same reasoning with the lifting criterion implies that $\widetilde{X}$ has the same fundamental group under p_* at each basepoint in $p^{-1}(x_0)$. But this means that conjugation by any loop in X fixes H, so it follows that H is normal. □

Note that the preceding proof identifies a loop which lifts to a path from $\tilde{x}_0$ to $\tilde{x}_1$ with the deck transformation that takes $\tilde{x}_0$ to $\tilde{x}_1$. In general, however, the loop and the deck transformation need not have the same action on the fiber of x_0.

As a counterexample, consider the Cayley graph $\widetilde{X}$ of F_2 generated by a, b as a covering space of $S^1 \vee S^1$. Since $\widetilde{X}$ contains no nontrivial loops, it is simply-connected, and hence the universal cover of $S^1 \vee S^1$. Recall that the covering map of this graph sent every vertex to the basepoint of $S^1 \vee S^1$, every vertical edge (every edge labeled b) to one circle, and every horizontal edge (every edge labeled a) to the other. Thus, we can define a loop γ in $S^1 \vee S^1$ which goes once around the circle corresponding to a. The action of γ on any vertex g will take it to ga, since γ lifts to the segment connecting g to its adjacent right vertex. But if we consider the vertex gb, which is located directly above g, we notice that the deck transformation associated with the action of γ on g must take gb to gab, which is directly above ga, in order to be continuous, rather than gba, which is not joined by an edge to ga. Thus, the deck transformation corresponding to γ depends on our choice of basepoint in $\widetilde{X}$ and may have a different action on the fiber of x_0 from γ.

There is also a way to generalize some of these ideas in defining a way for any group G to act on the space X. Essentially, let $Homeo(X)$ be the group of homeomorphisms from X to itself (not to be confused with isomorphisms which preserve the covering map), and let $\phi : G \to Homeo(X)$ be a group homomorphism. We then define the action of $g \in G$ on $x \in X$ to equal $\phi(g)(x)$, the homeomorphism assigned to g evaluated at x, which we will hence refer to as $g(x)$. Thus, we have a way to consider G as giving a set of maps from X to itself, which we can further constrain to obtain valuable results.

Definition 5.4. A **covering space action** is a group action of the group G on the space X such that for every $x \in X$ there is some open neighborhood U of x such that if $g_1, g_2 \in G$ and $g_1 \neq g_2$, then $g_1(U) \cap g_2(U) = \emptyset$.

Note that many sources instead use the phrase *properly discontinuous action* to refer to this type of action. "Covering space action" is not entirely standard, but is the choice of [1], which we will follow for this exposition.

The deck transformation group of a covering space has a natural action where each element already is the homeomorphism we assign to its action. We can then take U to be a sheet on which the covering map is injective, at which point any distinct deck transformations map U to disjoint sheets.

We also have the ability to define a quotient map in terms of a group action on a space. We define the orbit of x to be the set $\{g(x) \mid g \in G\} = Gx$. One can verify that a point being in the orbit of another point is an equivalence relation using the group axioms, and it follows that the orbit space X/G consisting of the orbits of the points in X is a well-defined quotient space of X. We can prove several results about the orbit space of a covering space action.

Theorem 5.5. *The quotient map $p : X \to X/G$ of a covering space action is a normal covering space. If X is path-connected, then this covering space satisfies $Aut(p) = G$. If X is path-connected and locally path-connected, then $G \simeq \pi_1(X/G)/p_*(\pi_1(X))$.*

Proof. For a particular open neighborhood U of $x \in X$, $g(U)$ is open for each $g \in G$ by homeomorphism and disjoint between any two elements of G, and it follows by quotient space topology that $p(U)$ is open and that p is a homeomorphism when restricted to $g(U)$. Furthermore, because the sheets of U are exactly the sets $g(U)$ for $g \in G$, it follows that each $g \in G$ is a deck transformation taking one sheet to another and preserving the covering map. Then we can take $g_1(U)$ to $g_2(U)$ by the deck transformation $g_2 g_1^{-1}$, so p is normal.

From this last discussion, G is isomorphic to a subgroup of $Aut(p)$. If X is also path-connected, then the unique lifting property guarantees that a deck transformation taking x to $x_1 \in p^{-1}(p(x)) = G(x)$ is uniquely determined, and this bijection gives that G is isomorphic to $Aut(p)$ as a whole.

By Theorem 5.3, $Aut(p)$ is isomorphic to the quotient of the normalizer of $p_*(\pi_1(X))$ by $p_*(\pi_1(X))$. But since X is a normal covering space, it follows by 5.3 that this is simply $\pi_1(X/G)/p_*(\pi_1(X))$. ☐

Note that if X is also simply-connected, then the 3rd statement reduces to $G = \pi_1(X/G)$. This gives us a powerful way to compute the fundamental group of a space, if we can identify a covering space action on its universal cover. For example:

Theorem 5.6. *Let X be the Cayley graph of a group G with respect to generating set S. Then G has a covering space action on X such that X/G is the wedge sum of $|S|$ circles.*

Proof. Let $g \in G$. We define $g : X \to X$ to take the vertex $v \in G$ to gv, and the edge between v_1 and v_2 homeomorphically to the edge between gv_1 and gv_2. An edge exists between v_1 and v_2 if and only if $v_1 = v_2 s$ for some generator s, in which case $gv_1 = gv_2 s$, so this map is well-defined. Furthermore, it is surjective, since for every $g_1 \in G$, $g^{-1}g_1$ is a vertex of X. It is injective, since $gv_1 = gv_2$ if and only if $v_1 = v_2$, and thus edges map one-to-one. It follows that g bijective and continuous, with continuous inverse g^{-1}, and thus it is a homeomorphism.

Let a be a vertex of X, and U be an open neighborhood of a containing no other vertices. For any distinct $g_1, g_2 \in G$, $g_1 a \neq g_2 a$, so $g_1 U \cap g_2 U = \emptyset$. Let e be the edge excluding vertices joining $a, b \in X$ for a generator s with $as = b$. Then $g_1(e) \cap g_2(e) \neq \emptyset$ only if we have equal endpoints; that is, $g_1 a = g_2 a$ and $g_1 b = g_2 b$, or $g_1 a = g_2 b$ and $g_1 b = g_2 a$. We have shown that the first case is not possible if $g_1 \neq g_2$. The second case also fails, because g_1 maps e to a directed edge from $g_1 a$ to $g_1 b = (g_1 a)s$, while g_2 maps e to a directed edge from $g_2 a$ to $g_2 b = (g_2 a)s$, or equivalently from $g_1 b$ to $g_1 a = (g_1 b)s$. From our Cayley graph definition, oppositely directed edges between two vertices are distinguished, so these two edges are disjoint. Therefore, the images of e under the actions of g_1 and g_2 are disjoint, and the action we have defined is a covering space action.

For any $a, b \in G$, the element ba^{-1} takes the vertex a to the vertex b, so it follows that all vertices in X are identified in X/G. The image of an edge labeled by generator s is also labeled by generator s, since if $b = as$, then $gb = (ga)s$. Since ga can be any vertex, all edges corresponding to s are identified, and no edges corresponding to one generator map to those for another. Therefore, X/G is a graph on one vertex, with an edge starting and ending at that vertex for each generator in S. A homeomorphism takes this space to the wedge sum of $|S|$ circles. □

Corollary 5.7. *The fundamental group of a wedge sum of circles is a free group generated by elements corresponding to the circles.*

Let X be a wedge sum of circles, F be the free group on generators corresponding to the circles, and $\widetilde{X}$ be the Cayley graph of F. Then F has a covering space action on $\widetilde{X}$, so the quotient map $p : \widetilde{X} \to \widetilde{X}/F$ is a regular covering space, and $F \simeq \pi_1(\widetilde{X}/F)/p_*(\widetilde{X})$. However, the Cayley graph of a free group is simply-connected, since every element has a unique representation as a reduced product of generators and thus there is only one homotopy class of paths joining two points in the graph. It follows that $p_*(\widetilde{X})$ is trivial, and $F \simeq \pi_1(\widetilde{X}/F)$. From the previous result, $\widetilde{X}/F$ will be a wedge sum of circles corresponding to the generators of F, and thus in bijection with the circles of X. Therefore, $X = \widetilde{X}/F$, and $\pi_1(X) = F$.

6. The Fundamental Group of a Cayley Graph

We now have many resources for computing fundamental groups, using covering spaces, deck transformation groups, and covering space actions. In this section, we will prove a general method for computing the fundamental group of a graph, and use this method to explore the relationship between a group and the fundamental group of its Cayley graph.

We will need another definition before our first result.

Definition 6.1. A **tree** is a graph in which every two vertices are joined by exactly one path. A **spanning tree** T of X is a subgraph of X which is a tree and which contains all vertices of X.

It follows that a tree T is connected, that $\pi_1(T, t_0) = \{0\}$, and that T is contractible.

We next show the existence of a spanning tree for any connected graph. For finite graphs, we can directly construct the tree by induction: Starting with a vertex, draw edges to all vertices which can be reached from the current subgraph and which are not yet elements of the subgraph. After each step, we have a tree with more vertices, and since the graph is finite and connected, we eventually no longer can reach more vertices, at which point we must have a spanning tree. For infinite graphs, the proof requires the axiom of choice, as follows:

Theorem 6.2. *A connected graph X contains a spanning tree.*

Proof. Let X be a connected graph, and let $\{T_\lambda \mid \lambda \in I\}$ be a collection of trees in X that is totally-ordered by proper subgraph inclusion.

Then any edge of $\bigcup_{\lambda \in I} T_\lambda$ is contained in some T_{λ_0}. In particular, subgraph inclusion means that for any loop in the union, there is a tree T_{λ_0} which contains every edge of the loop, and thus the loop itself, so T_{λ_0} is not a tree. Contradiction, and the union contains no loops.

Similarly, if the union is disconnected, then there are two vertices in the union which are not joined by a path, and there is some tree T_{λ_1} which contains both vertices, so T_{λ_1} is not connected and not a tree. Contradiction, and the union is connected.

Thus, $\bigcup_{\lambda \in I} T_\lambda$ is connected with no loops, so it is a tree, and it contains every tree in the collection $\{T_\lambda \mid \lambda \in I\}$. Since every collection of trees totally ordered by subgraph inclusion has an upper bound, it follows by Zorn's Lemma that the set of all trees of X has a maximal element T. Since X is connected, if T does not contain every vertex of X, we can find an edge from a vertex in T to a vertex in $X \setminus T$, which yields a larger tree which contains T. Therefore, it follows that T contains every vertex of X, and so T is a spanning tree. □

We now have the existence of a contractible subset T of any graph X, such that the quotient X/T identifies all vertices of X. The key idea of the next theorem will be to use the corresponding quotient map to study the fundamental group of X. This study, however, requires the fact that X and X/T are homotopy equivalent.

Proposition 6.3. *Given a graph X and any spanning tree $T \subset X$, the spaces X and X/T are homotopy equivalent.*

The proof of this statement follows from a more general theorem in [1]. To avoid unnecessary digression, we provide a partial summary here.

Since T is a tree and thus contractible, there is a homotopy $g_t : T \to T$ such that g_0 is the identity on T and g_1 maps T to a point. We can prove that it is possible to extend this homotopy to a homotopy $f_t : X \to X$ such that f_t restricts to g_t on T and f_0 is the identity of X. Intuitively, as this homotopy contracts the edges of T, it also stretches the edges of $X \setminus T$, until T vanishes to a point x_0 and $f_1(X \setminus T) = X \setminus \{x_0\}$. Let $q : X \to X/T$ be the quotient map, which like f_1 maps T to a point, and we can show along those intuitive lines that f_1 induces a map $p : X/T \to X$ such that $p \circ q = f_1$. Here, f_1 is homotopic to f_0, the identity on X, giving one direction of homotopy equivalence. For any t, the map $q \circ f_t$ sends T to a point; from this, we can prove that there exists an alternative factoring of this map given by $h_t \circ q$ for some homotopy $h_t : X/T \to X/T$. Similar to before, we can show that $q \circ p = h_1$ which is homotopic to h_0, the identity on X/T. Therefore, X and X/T are homotopy equivalent.

We will assume this result to prove the next theorem.

Theorem 6.4. *The fundamental group of a graph is a free group. Specifically, if the graph X has spanning tree T, then this free group is generated by elements corresponding to the edges of $X \setminus T$.*

Proof. Since X and X/T are homotopy equivalent by the proposition, they have isomorphic fundamental groups. Since T spans X, the space X/T consists of a single vertex and edges which begin and end at that vertex corresponding to the edges of $X \setminus T$. This space is homeomorphic to the wedge sum of circles S^1 corresponding to the edges of $X \setminus T$, which has a free fundamental group generated by elements corresponding to each circle. □

This result is very powerful, allowing us to compute the fundamental group of a graph just by knowing how many vertices and edges it has. Another application is that we can finally explore the fundamental group of a Cayley graph, and see how it relates to its respective group. The proof of the next theorem will be very illustrative in this regard.

Theorem 6.5. *Every group G is isomorphic to the quotient of a free group by a normal subgroup.*

Proof. Let X be a Cayley graph of G. We have already shown that G has a covering space action on X, such that X/G is a wedge sum of circles with free fundamental group. Then by Theorem 5.5, the quotient map $p : X \to X/G$ is a normal covering space of X/G and $G \simeq \pi_1(X/G)/p_*(\pi_1(X))$. By Theorem 5.3, we have that $p_*(\pi_1(X))$ is a normal subgroup of $\pi_1(X/G)$. □

This theorem is not particularly strong; it is equivalent to the claim that every group has a group presentation. However, the proof demonstrates the key topological relations between a group G on generating set S and its Cayley graph X. From Lemma 5.6, the space X/G consists of a single vertex and edges corresponding to each element of S, so $\pi_1(X/G) = F_S$. On the other hand, a loop in X based at the identity deforms to a sequence of edges, such that the product of their generator labels is the identity. We also know from how edges in X map to edges in X/G that $p_* : \pi_1(X) \to \pi_1(X/G)$ will take a sequence of edges to the word formed by the edges' generator labels. In other words, $p_*(\pi_1(X))$ is a subgroup of F_S containing exactly the words that correspond to products of elements of S which evaluate to 1 in G. Therefore, $F_S/p_*(\pi_1(X)) = G$, and G has a presentation given by $< S \mid p_*(\pi_1(X)) >$. We can even obtain a smaller, possibly finite relations set by simply listing the images of the generators of $\pi_1(X)$, computed using Theorem 6.4.

For an example of this idea in action, consider the dihedral group of order 8, D_8. With geometry, we can show that this group is generated by a rotation r and a reflection s in such a way as to give the following Cayley graph X:

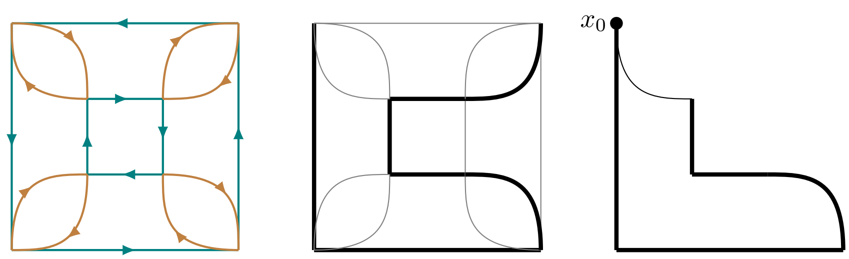

Figure 8. The Cayley graph X of D_8 with edges colored blue for r and brown for s. Also shown: a spanning tree of X and one of the generating loops of $\pi_1(X)$.

From the spanning tree in the diagram, Theorem 6.4 gives that $\pi_1(X) = F_9$, the free group on generators corresponding to each edge not in the tree. However, we also have that $\pi_1(X/D_8) = F_2$, so the covering map gives us a homomorphism $p_* : F_9 \to F_2$ which maps loops in X to words of r and s. For example, the generator loop in the diagram above based at x_0 would be mapped to $r^2s^{-1}r^2s$. If R is the set of generator loops of $\pi_1(X)$, those loops homotopic to the union of the spanning tree and a single edge not on the tree, then $p_*(\pi_1(X))$ is a normal subgroup of F_2 generated by $p_*(R)$, such that $D_8 \simeq F_2/p_*(\pi_1(X))$. Therefore, the

Cayley graph X on $S = \{r, s\}$ gives us the group presentation $< S \mid p_*(R) >$ of D_8.

7. The Nielsen-Schreier Theorem

The final result of this paper will use graphs to understand the subgroup structure of free groups. To accomplish this, we require one final lemma to characterize the covering spaces of a graph.

Theorem 7.1. *Every covering space of a graph is a graph whose edges and vertices are respective lifts of the edges and vertices of the base graph.*

Proof. Let X be a graph and $p : \widetilde{X} \to X$ be a covering space of X. By the definition of a graph, we can write X as $X^0 \bigsqcup_\lambda I_\lambda$, where each I_λ is a copy of the unit interval joining vertices taken from the discrete set X^0. If we take $p^{-1}(X^0)$ to be a set of vertices, then we can treat the edges I_λ as paths $I \to X$ and lift each edge to unique paths $I_\lambda \to \widetilde{X}$ corresponding to each point in the preimage of that path's starting point. Every point in $\widetilde{X}$ must map to a point on an edge or vertex of X, so our lifts of these objects surject onto $\widetilde{X}$. Since p is a local homeomorphism, $\widetilde{X}$ must also locally have the same topology as X, so it follows that $\widetilde{X}$ is a graph, with edges and vertices given by the lifts of those of X. □

The next theorem was originally proven by Jakob Nielsen in 1921 using strictly group theoretic techniques. With the previous two results as well as our work in covering spaces, the proof becomes greatly simplified.

Theorem 7.2 (The Nielsen-Schreier Theorem)**.** *Every subgroup of a free group is free.*

Proof. Let G be a free group, and H be a subgroup of G. Then we can construct a wedge sum of circles corresponding to the generators of G to obtain a graph Γ with $\pi_1(\Gamma) = G$. Using the Galois correspondence, for our subgroup H there exists a covering space $p : \widetilde{\Gamma} \to \Gamma$ such that $p_*(\pi_1(\widetilde{\Gamma})) = H$, and since p_* is an injective homomorphism it follows that $\pi_1(\widetilde{\Gamma})$ is isomorphic to H. Since the covering space of a graph is a graph and the fundamental group of a graph is free, it follows that the fundamental group of $\widetilde{\Gamma}$ is free. Therefore, H is free. □

Thus, applying topology to graphs reveals a fairly surprising result about free groups. The limitations on how much a covering space can diverge from its base space are closely related to the limitations on how much a subgroup can diverge from a free group.

After seeing this theorem, the obvious corollary is to apply it to directly compute the subgroups of a free group. This proves to be fairly easy with the tools we have developed, leading to the result below.

Corollary 7.3. *Every free group of rank at least 2 contains a subgroup of countable rank.*

Proof. Consider the Cayley graph X of $\mathbb{Z} \times \mathbb{Z}$ generated by $(1, 0)$ and $(0, 1)$. As in our previous examples, X is a covering space of $S^1 \vee S^1$, with a covering map $p : X \to S^1 \vee S^1$. We can construct a spanning tree of X from the union of a vertical line and every horizontal line, leaving countably many edges not included in the tree, so that $\pi_1(X)$ is isomorphic to the free group on countable generators. But we know that p_* is an injective homomorphism, so the free generators of $\pi_1(X)$ are mapped to the free generators of $p_*(\pi_1(X))$, and F_2 contains a subgroup of countable rank. Any free group of greater rank contains a subgroup isomorphic to F_2 by restricting to two generators, and the claim follows. □

Since the free group of countable rank contains every free group of finite rank as a subgroup via restriction to finitely many generators, it follows that a free group of rank at least 2 contains every free group of at most countable rank as a subgroup. This exactly characterizes the subgroups of such a free group, since the set of words from countable generators is itself countable.

For readers curious about the construction in Corollary 7.3, we can characterize the countable rank subgroup of F_2 in terms of the generating loops of $\pi_1(X)$. With the spanning tree of one vertical line and every horizontal line, our generating loops consist of some integer number of steps along the vertical line, some nonzero number of steps along the horizontal line, a step up along the vertical edge not included in the tree, and then the return to the origin along the horizontal and vertical lines in the tree. The result gives us a subgroup of F_2 generated by $\{b^n a^m b a^{-m} b^{-n-1} \mid m, n \in \mathbb{Z} \text{ and } m \neq 0\}$.

Acknowledgments

I would like to thank my mentor, Shiva Chidambaram, for helping me learn topology and advising me through the paper-writing process. Through his help, I was able to ask the right questions each step along the way and achieve much more than I could have on my own. I would also like to thank Peter May for organizing the 2018 REU and allowing me to be a part of this unique mathematical experience.

References

[1] Allen Hatcher. Algebraic Topology. 2001. https://pi.math.cornell.edu/ hatcher/AT/AT.pdf
[2] James Munkres. Topology: A First Course. Prentice Hall. 1975.