Covering Spaces and Subgroups of Free Groups

Algebraic Topology I | Fall 2026 | October 19, 2026

12.1 Lifting Paths, Homotopies, and Maps

The exponential map used in Lecture 11 describes the circle by a simply connected space together with a group of translations. Covering space theory extends this description to spaces whose sufficiently small neighborhoods have no essential loops. Its classification theorem identifies connected covering spaces with subgroups of the fundamental group. The lifting results themselves require fewer hypotheses, and we prove them first.

12.1.1 Covering Maps and Unique Lifting

Write $I = [0,1]$. For composable paths α and β, the concatenation $\alpha * \beta$ traverses α first and β second. Write $\overline{\alpha}(t) = \alpha(1-t)$ for the reversed path. All homotopies of paths preserve their endpoints unless otherwise stated.

Definition 12.1.1. A *covering map* is a continuous surjection $p : Y \to X$ such that every point of X has an open neighborhood U for which $p^{-1}(U)$ is a disjoint union of open subsets of Y, each mapped homeomorphically onto U by p. Such a neighborhood is *evenly covered*, and these open subsets are its *sheets*. A *lift* of $f : T \to X$ is a continuous map $\tilde{f} : T \to Y$ satisfying $p\tilde{f} = f$.

Each fiber $p^{-1}(x)$ is discrete. Every covering map is a local homeomorphism, but the simultaneous decomposition over one neighborhood is part of the definition of a covering. For example, the restriction $\exp : (0,2) \to \mathrm{S}^1$, defined by $t \mapsto e^{2\pi i t}$, is a surjective local homeomorphism that is not a covering: the fiber over 1 has one point, whereas nearby fibers have two.

Lemma 12.1.2. Let $p : Y \to X$ be a covering map, let T be connected, and let $f : T \to X$ be continuous. If two lifts of f agree at one point, then they agree everywhere.

Proof. Let A be the set of points where the lifts agree. Near any $u \in T$, choose an evenly covered neighborhood of $f(u)$ and a neighborhood of u on which each lift lies in one sheet. If the sheets are equal, the two lifts agree throughout this neighborhood, because they are obtained by composing f with the same inverse of p. If the sheets differ, the lifts disagree throughout this neighborhood. Thus A and its complement are open. Connectedness and the assumed agreement imply $A = T$. □

Proposition 12.1.3. For every path $\alpha : I \to X$ and every $y \in p^{-1}(\alpha(0))$, there is a unique lift $\tilde{\alpha} : I \to Y$ satisfying $\tilde{\alpha}(0) = y$.

Proof. Pull back evenly covered neighborhoods along α. The Lebesgue number lemma supplies a subdivision $0 = t_0 < \cdots < t_m = 1$ such that the image of each closed interval $[t_{j-1}, t_j]$ lies in one evenly covered neighborhood U_j. On the first interval, compose α with the inverse of p on the sheet containing y. Its endpoint selects a sheet over U_2, and this determines the lift on the next interval. Continue through the finite subdivision. The resulting maps agree at their endpoints, so the pasting lemma proves continuity. Uniqueness follows from Lemma 12.1.2. □

The same local inverses lift a continuously varying family of paths. The parameter space need not be locally connected or compact.

Theorem 12.1.4. Let $p : Y \to X$ be a covering map, let T be any space, and let $F : T \times I \to X$ be a homotopy. Given a lift $g : T \to Y$ of $F|_{T\times\{0\}}$, there is a unique continuous map $\widetilde{F} : T \times I \to Y$ with $p\widetilde{F} = F$ and $\widetilde{F}(u, 0) = g(u)$.

Proof. For each u, lift the path $t \mapsto F(u, t)$ starting at $g(u)$. This gives a uniquely determined function $\widetilde{F}$. We prove that this function is continuous near $\{u_0\} \times I$ for an arbitrary $u_0 \in T$.

Choose a finite subdivision such that $F(\{u_0\} \times [t_{j-1}, t_j])$ lies in an evenly covered open set U_j. By compactness of each closed interval, there is a neighborhood N of u_0 such that

$$F(N \times [t_{j-1}, t_j]) \subseteq U_j \qquad (1 \le j \le m).$$

Indeed, the open set $F^{-1}(U_j)$ contains the compact set $\{u_0\} \times [t_{j-1}, t_j]$; finitely many product neighborhoods give the required neighborhood of u_0. Intersect these neighborhoods over the finitely many values of j.

Shrink N so that $g(N)$ lies in the sheet over U_1 containing $g(u_0)$. On $N \times [t_0, t_1]$, the desired lift is the continuous map obtained from the inverse on that sheet. Its endpoint is a continuous function of u. Shrink N again so that these endpoints lie in the sheet over U_2 containing $\widetilde{F}(u_0, t_1)$. This gives a continuous lift on the next interval. After finitely many such restrictions, the local formulas agree at every junction and define a continuous map on $N \times I$. Path-lifting uniqueness identifies this map with $\widetilde{F}$. Such neighborhoods exist for every u_0, proving continuity everywhere. Uniqueness also follows by restricting to each vertical path. □

In particular, every covering map is a Hurewicz fibration in the terminology of Lecture 10. If a homotopy fixes a point of T, its lifted track lies in a discrete fiber and is constant. Thus a homotopy of paths lifts to a homotopy with the same lifted endpoints.

12.1.2 The Fundamental Group and the Lifting Criterion

A choice of $y_0 \in Y$ determines the basepoint $x_0 = p(y_0)$ and a subgroup of $\pi_1(X, x_0)$. This subgroup records which loops return to the chosen sheet.

Proposition 12.1.5. For every covering map $p : (Y, y_0) \to (X, x_0)$, the homomorphism

$$p_* : \pi_1(Y, y_0) \longrightarrow \pi_1(X, x_0)$$

is injective. Its image consists of the classes represented by loops whose lifts starting at y_0 end at y_0.

Proof. If a loop λ in Y projects to a nullhomotopic based loop, lift a based nullhomotopy of $p\lambda$ with initial lift λ. The endpoints remain y_0. The final path is the lift of a constant path, so it is constant. This is a based nullhomotopy of λ, proving injectivity.

A loop that lifts to a loop represents an element of the image of p_*. Conversely, if α represents an element of that image, it is based homotopic to the projection of a loop in Y at y_0. Lifting this homotopy shows that its lift also ends at y_0. □

Theorem 12.1.6. Let $p : (Y, y_0) \to (X, x_0)$ be a covering map. Let T be path connected and locally path connected, with basepoint u_0, and let $f : (T, u_0) \to (X, x_0)$ be continuous. A based lift $\tilde{f} : (T, u_0) \to (Y, y_0)$ exists if and only if

$$f_*(\pi_1(T, u_0)) \subseteq p_*(\pi_1(Y, y_0)).$$

When it exists, the based lift is unique.

Proof. A based factorization $f = p\tilde{f}$ implies the subgroup inclusion by functoriality. Conversely, assume the inclusion. For $u \in T$, choose a path γ from u_0 to u and define $\tilde{f}(u)$ to be the endpoint of the lift of $f\gamma$ starting at y_0.

If γ' is another choice, the loop $f\gamma * \overline{f\gamma'}$ represents an element of $p_*(\pi_1(Y, y_0))$ and therefore lifts to a loop at y_0. Its second part, reversed, is the lift of $f\gamma'$ starting at y_0. Hence the lifts of $f\gamma$ and $f\gamma'$ have the same endpoint. The definition is independent of the path.

To prove continuity at u, choose an evenly covered neighborhood U of $f(u)$ and a path-connected open neighborhood V of u contained in $f^{-1}(U)$. For every $z \in V$, a path from u to z inside V can be appended to γ. The lifted appended path stays in the sheet W over U containing $\tilde{f}(u)$. Consequently

$$\tilde{f}|_V = (p|_W)^{-1} \circ f|_V.$$

This proves continuity. The constant path gives $\tilde{f}(u_0) = y_0$, and Lemma 12.1.2 proves uniqueness. □

The local path-connectedness hypothesis enters in choosing V. No local hypothesis on X was used in this theorem. In particular, every based map from a simply connected, locally path-connected space has a unique based lift.

12.2 Constructing and Classifying Connected Covers

The lifting criterion compares covering spaces that already exist. To realize every subgroup, we must construct a topology on suitable classes of paths. The required local condition says that small loops become trivial in the whole space, rather than necessarily in the neighborhood itself.

12.2.1 Local Hypotheses and the Path-Class Construction

The local hypotheses distinguish paths inside small neighborhoods from nullhomotopies that may leave those neighborhoods.

Definition 12.2.1. A space is *locally path connected* if its path-connected open subsets form a basis. It is *semilocally simply connected* if each point x has an open neighborhood U such that every loop in U based at x is nullhomotopic in X.

From now through Section 12.3, assume that X is connected, locally path connected, and semilocally simply connected. Connectedness and local path-connectedness imply path-connectedness: path components are open, and their complements are unions of open path components.

Call a path-connected open set $U \subseteq X$ *admissible* if every loop in U is nullhomotopic in X. Admissible sets form a basis. To see this, intersect a neighborhood from the semilocal condition with any prescribed open neighborhood and choose a path-connected open neighborhood inside the intersection. Changing the basepoint along a path in this smaller neighborhood shows that its loops at every basepoint become trivial in X.

Fix $x_0 \in X$, put $G = \pi_1(X, x_0)$, and let $H \leq G$ be a subgroup. For paths α, β starting at x_0, define

$$\alpha \sim_H \beta \quad \Longleftrightarrow \quad \alpha(1) = \beta(1) \text{ and } [\alpha * \overline{\beta}] \in H.$$

This is an equivalence relation. Reflexivity and symmetry use the identity and inverses in H; transitivity follows by cancelling the intervening path followed by its reverse. Let Y_H be the set of equivalence classes, write $[\alpha]_H$ for a class, and define $p_H : Y_H \to X$ by $p_H([\alpha]_H) = \alpha(1)$. Let y_H be the class of the constant path at x_0.

For an admissible neighborhood U of $\alpha(1)$, define

$$B([\alpha]_H, U) = \{[\alpha * \delta]_H \mid \delta \text{ is a path in } U \text{ starting at } \alpha(1)\}.$$

This set depends only on the class of α: appending the same path preserves the defining equivalence relation. We give Y_H the topology having these sets as a basis. This is a specified basis topology; no quotient topology on a function space is being assumed.

Proposition 12.2.2. The sets $B([\alpha]_H, U)$ form a basis, and the endpoint map $p_H : Y_H \to X$ is a path-connected covering space. Its associated subgroup is

$$(p_H)_*(\pi_1(Y_H, y_H)) = H.$$

Proof. First consider one admissible set U. Any two paths in U with the same endpoints differ by a loop nullhomotopic in X. After appending them to α, the resulting paths are equivalent. It follows that the restriction

$$p_H : B([\alpha]_H, U) \longrightarrow U$$

is a bijection. Surjectivity here uses path-connectedness of U.

If $[\beta]_H$ belongs to $B([\alpha]_H, U)$ and W is an admissible neighborhood of $\beta(1)$ contained in U, then

$$B([\beta]_H, W) \subseteq B([\alpha]_H, U).$$

Indeed, represent $[\beta]_H$ by $\alpha * \delta$ with δ in U, and append a path in W. Given a point in the intersection of two proposed basis sets, choose W inside the intersection of their endpoint neighborhoods. The displayed inclusion proves the basis intersection axiom.

The endpoint map is continuous: for an open set $O \subseteq X$, every class ending in O has a basic neighborhood whose endpoint neighborhood lies in O. Each displayed bijection is a homeomorphism. To check continuity of its inverse, take a basic neighborhood of a point in the sheet and shrink its endpoint neighborhood to an admissible $W \subseteq U$. The inverse then sends W into that basic neighborhood.

Fix an admissible U and a point $u \in U$. The sets $B([\alpha]_H, U)$ with $\alpha(1) = u$ cover $p_H^{-1}(U)$. In fact, for a path β ending at $x \in U$, choose a path δ in U from u to x and use the path $\beta * \overline{\delta}$ ending at u. Two of these sheets that meet are equal: append the reverse of a path in U from u to the common endpoint to see that their classes over u agree. Thus U is evenly covered. Path-connectedness of X makes p_H surjective.

For a path α starting at x_0, set $\alpha_t(s) = \alpha(ts)$. The function $t \mapsto [\alpha_t]_H$ is a continuous path from y_H to $[\alpha]_H$. Near any t_0, choose a small interval of parameters whose image under α lies in one admissible neighborhood. The classes $[\alpha_t]_H$ are then given by the local inverse on the sheet containing $[\alpha_{t_0}]_H$; reversing the short segment handles $t < t_0$. This proves continuity and path-connectedness of Y_H.

The same path is the unique lift of α from y_H. For a loop α at x_0, it closes if and only if $[\alpha]_H = y_H$, which means $[\alpha] \in H$. Proposition 12.1.5 now identifies the image subgroup with H. □

12.2.2 Based Classification, Conjugacy, and Universal Covers

An isomorphism of covers is a homeomorphism over X. For based covers it must also preserve the chosen points above x_0. The classification is more precise with those points retained.

Theorem 12.2.3. For a connected, locally path-connected, semilocally simply connected space X, the assignment

$$(p : (Y, y_0) \to (X, x_0)) \longmapsto p_*(\pi_1(Y, y_0))$$

gives a bijection from based isomorphism classes of connected covering spaces to subgroups of $G = \pi_1(X, x_0)$. If H and K are the associated subgroups, there is a based map of covers $Y_H \to Y_K$ if and only if $H \subseteq K$, and this map is unique.

Proof. A covering of a locally path-connected space is locally path connected, because its sheets are homeomorphic to open subsets of the base. Thus a connected covering space is path connected. The lifting criterion applied to the projection of one based covering space into another gives the stated existence and uniqueness of maps over X.

If two covers have the same subgroup, there are based maps over X in both directions. Each composite and the appropriate identity map are lifts of the same projection and agree at the basepoint. Unique lifting makes the composites identities, so the maps are inverse homeomorphisms. Proposition 12.2.2 realizes every subgroup. These statements prove the bijection and the assertion about maps. □

The classification also counts sheets and explains the conjugacy ambiguity. For a covering $p : Y \to X$, lift a loop representing $g \in G$ from a point $y \in p^{-1}(x_0)$ and denote its endpoint by $y \cdot g$. Homotopy lifting makes this well defined. Lifting concatenations gives a right action:

$$(y \cdot g) \cdot h = y \cdot (gh).$$

For a connected cover, the action is transitive, since a path between two points of the fiber projects to a loop. Its stabilizer at y_0 is $H = p_*(\pi_1(Y, y_0))$ by Proposition 12.1.5.

Corollary 12.2.4. For a connected based cover with subgroup $H \leq G$, the map

$$\{Hg \mid g \in G\} \longrightarrow p^{-1}(x_0), \qquad Hg \longmapsto y_0 \cdot g$$

is a bijection of right G-sets. Consequently the number of sheets is the index $[G : H]$, allowing an infinite cardinal. Unbased isomorphism classes of connected covers correspond to conjugacy classes of subgroups of G.

Proof. Transitivity proves surjectivity. The endpoints $y_0 \cdot g$ and $y_0 \cdot h$ agree if and only if the loop with class gh^{-1} lifts to a loop at y_0. This is equivalent to $gh^{-1} \in H$, and hence to $Hg = Hh$. The formula respects the right action.

If the chosen point is changed from y_0 to $y_0 \cdot g$, its stabilizer is $g^{-1}Hg$. Thus an unbased isomorphism identifies conjugate subgroups. Conversely, if the subgroup of a second based cover is $g^{-1}Hg$, change the chosen point in the first cover to $y_0 \cdot g$ and apply the based classification. This gives an unbased isomorphism. □

Here each coset Hg consists of the products hg with $h \in H$. The right action on these cosets is induced by multiplication on the right in G.

Corollary 12.2.5. The cover $p_{\{1\}} : Y_{\{1\}} \to X$ corresponding to the trivial subgroup is simply connected. It admits a unique based map over X to every connected based covering space, and this map is itself a covering map. Every simply connected connected cover is based isomorphic to $Y_{\{1\}}$. Such a cover is called a *universal cover*.

Proof. The homomorphism $(p_{\{1\}})_*$ is injective and has trivial image, so $\pi_1(Y_{\{1\}}, y_{\{1\}})$ is trivial. The classification gives the unique map to any based connected cover because the trivial subgroup is contained in every subgroup. We verify below, in Proposition 12.3.2, that every such based map between connected covers is a covering map. Conversely, a simply connected cover has trivial image subgroup, so the uniqueness statement in the classification identifies it with $Y_{\{1\}}$. □

The semilocal hypothesis is also necessary for the existence of a simply connected covering space over a locally path-connected base. Indeed, choose a path-connected evenly covered neighborhood of a base point. Every loop there lifts to a loop in one sheet; if the total space is simply connected, its projection is nullhomotopic in X.

For example, the Hawaiian earring is the union in $\mathbb{R}^2$ of the circles of radius $1/n$ centered at $(1/n, 0)$, for $n \geq 1$. Every neighborhood of their common point contains whole circles of

sufficiently small radius. Each such circle is essential: the map that fixes it and collapses every other circle to the common point is a continuous retraction, so it injects on fundamental groups. Therefore the Hawaiian earring is not semilocally simply connected and has no simply connected covering space. This distinguishes it from the CW wedge of countably many circles used below.

12.3 Deck Transformations and Intermediate Coverings

The automorphisms of a covering are constrained by path lifting: the image of one point determines the entire automorphism. The monodromy action on the fiber translates this constraint into a normalizer computation. It also identifies the quotient constructions that occur in the Galois correspondence.

12.3.1 The Deck Group and Normal Subgroups

A *deck transformation* of $p : Y \to X$ is a homeomorphism $d : Y \to Y$ satisfying $pd = p$. Write $\mathrm{Deck}(p)$ for their group, with multiplication given by composition. A connected cover is *regular* if its deck group acts transitively on the fiber over x_0. Path lifting then gives transitivity on every fiber.

Let $H \leq G$. Its *normalizer* is the subgroup

$$N_G(H) = \{g \in G \mid gHg^{-1} = H\}.$$

The subgroup H is normal in $N_G(H)$, so $N_G(H)/H$ is a group.

Theorem 12.3.1. For the connected based cover corresponding to $H \leq G$, there is a group isomorphism

$$\mathrm{Deck}(p_H) \cong N_G(H)/H.$$

This cover is regular if and only if H is normal in G. In that case its deck group is G/H. In particular, the universal cover has deck group G.

Proof. A deck transformation commutes with the monodromy action: applying it to a lifted path gives the lift of the same path starting at the transformed point. Its permutation of the fiber is therefore a G-equivariant bijection of $\{Hg \mid g \in G\}$. If it sends H to Hn, equivariance requires

$$Hg \longmapsto Hng.$$

The stabilizer of H is H, whereas the stabilizer of Hn is $n^{-1}Hn$. An equivariant bijection preserves stabilizers, so n belongs to $N_G(H)$. Unique lifting shows that the deck transformation is determined by this value at the chosen point.

Conversely, represent $n \in N_G(H)$ by a loop ν at x_0 and define

$$D_n : Y_H \longrightarrow Y_H, \qquad D_n([\alpha]_H) = [\nu * \alpha]_H.$$

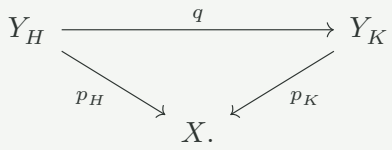
If $[\alpha * \overline{\beta}] \in H$, the difference between the new paths represents $n[\alpha * \overline{\beta}]n^{-1} \in H$. Thus the formula is well defined, and changing the representative of n does not change it. It carries every basic sheet $B([\alpha]_H, U)$ homeomorphically onto $B([\nu * \alpha]_H, U)$ and has inverse $D_{n^{-1}}$. Therefore it is a deck transformation.

The formulas give $D_n \circ D_m = D_{nm}$ and $D_n = \mathrm{id}$ if and only if $n \in H$. All deck transformations have already been accounted for, so they give the asserted isomorphism. The possible images of the chosen point are the cosets Hn with $n \in N_G(H)$. These constitute the entire fiber if and only if $N_G(H) = G$, which is the normality condition. □

No nonidentity deck transformation of a connected cover fixes a point, by unique lifting. Thus a regular finite-sheeted cover has as many deck transformations as sheets. A nonregular cover can have a much smaller deck group; an example appears in Section 12.4.2.

12.3.2 The Galois Correspondence and Circle Covers

Fix a based connected cover $p_H : Y_H \to X$. A based intermediate covering is a factorization through a connected based cover, with maps preserving all projections and basepoints:



Isomorphisms of factorizations are based isomorphisms of the middle covers commuting with both maps.

Proposition 12.3.2. The based intermediate coverings of p_H are classified by the subgroups K satisfying $H \le K \le G$. The map $q : Y_H \to Y_K$ is a covering map of degree $[K : H]$. It is regular if and only if H is normal in K, and its deck group is $N_K(H)/H$.

Proof. For $H \le K$, the unique based map is given in the path-class models by

$$q([\alpha]_H) = [\alpha]_K.$$

It is well defined because the equivalence relation for H implies that for K. Over an admissible open set $U \subseteq X$, the map carries each H-sheet homeomorphically onto the corresponding K-sheet: both sheets are identified with U by their projections. The inverse image of any K-sheet is the disjoint union of the H-sheets mapping to it. Every point of the K-sheet has a representative path and therefore a preimage. This proves that q is a covering map.

Conversely, an intermediate factorization gives

$$(p_H)_*(\pi_1(Y_H, y_H)) \subseteq (p_K)_*(\pi_1(Y_K, y_K)) \subseteq G.$$

The classification and uniqueness of based lifts identify the factorization with the one just constructed. The fiber of q over y_K consists of the cosets Hk for $k \in K$, so its degree is $[K : H]$.

The injection $(p_K)_*$ identifies $\pi_1(Y_K, y_K)$ with K and identifies the subgroup defined by q with H. The space Y_K is locally path connected by its sheet neighborhoods. A loop in a sheet over an admissible open set projects to a nullhomotopic loop in X; injectivity of $(p_K)_*$ makes the original loop nullhomotopic in Y_K. Thus Y_K is also semilocally simply connected. Applying Theorem 12.3.1 to q proves the regularity criterion and the deck-group formula. □

The direction of the correspondence is important: $K_1 \subseteq K_2$ gives a map $Y_{K_1} \to Y_{K_2}$ over X. A smaller subgroup gives a covering space lying above a larger subgroup. If the relevant indices are finite, counting fibers gives

$$[G : H] = [G : K]\,[K : H].$$

Corollary 12.3.3. Suppose H is normal in G, and identify $D = \mathrm{Deck}(p_H)$ with G/H. For every subgroup $L \le D$, let $K \le G$ be its inverse image. The orbit space Y_H/L is naturally isomorphic over X to Y_K. This gives the Galois correspondence between subgroups of D and based intermediate covers of p_H. The cover $Y_K \to X$ is regular if and only if L is normal in D, and then its deck group is D/L.

Proof. The group $L = K/H$ acts through the deck transformations D_k. Two classes $[\alpha]_H$ and $[\beta]_H$ have the same image in Y_K if and only if they have the same endpoint and $[\alpha * \overline{\beta}] \in K$. In

that case, put $k = [\alpha * \overline{\beta}]$; cancelling $\overline{\beta} * \beta$ shows that $D_k([\beta]_H) = [\alpha]_H$. Thus the fibers of $Y_H \to Y_K$ are the L-orbits.

The covering map $Y_H \to Y_K$ is open and surjective, hence is a quotient map. Its fiber description therefore identifies the quotient topology on Y_H/L with the topology of Y_K. Finally, K is normal in G if and only if $K/H = L$ is normal in $G/H = D$. The normalizer theorem and the group isomorphism $(G/H)/(K/H) \cong G/K$ prove the final assertion. □

In particular, every connected cover of X is an orbit space of the universal cover by a subgroup of its deck group. The orbit space is formed using that specified deck action, so its projection is a covering map by the preceding proof.

Example 12.3.4. For each integer $n \geq 1$, the map $p_n : \mathrm{S}^1 \to \mathrm{S}^1$ defined by $p_n(z) = z^n$ is an n-sheeted covering. Local inverse branches on short arcs give its sheets. The winding-number calculation in Lecture 11 shows that $(p_n)_*$ is multiplication by n on $\mathbb{Z}$, so the associated subgroup is $n\mathbb{Z}$. Its deck group consists of the rotations $z \mapsto \zeta z$ with $\zeta^n = 1$, and is isomorphic to $\mathbb{Z}/n\mathbb{Z}$.

The universal cover is $p : \mathbb{R} \to \mathrm{S}^1$, given by $p(t) = e^{2\pi i t}$, with deck transformations $t \mapsto t + k$ for $k \in \mathbb{Z}$. Since every subgroup of $\mathbb{Z}$ is $n\mathbb{Z}$ for a unique $n \geq 0$, these are all connected covers: $n = 0$ gives the universal cover, and $n \geq 1$ gives p_n.

For positive integers m, n, the inclusion $n\mathbb{Z} \subseteq m\mathbb{Z}$ holds if and only if m divides n. Writing $n = mk$, the corresponding intermediate map is $q : \mathrm{S}^1 \to \mathrm{S}^1$, given by $q(z) = z^k$, and $p_m q = p_n$. Its degree is k, illustrating the direction of the subgroup correspondence.

12.4 Graphs and Subgroups of Free Groups

Covering spaces turn subgroup questions into questions about spaces. For free groups, the spaces are graphs, and a maximal tree gives an explicit basis for the fundamental group. We include the infinite-graph argument because a subgroup of a finitely generated free group need not be finitely generated.

12.4.1 Maximal Trees and Infinite Graphs

A *graph* here is a one-dimensional CW complex; loops, multiple edges, and infinitely many edges are allowed. Each edge is a copy of an interval with its endpoints attached to vertices, and the topology is the CW weak topology. A *tree* is a connected graph containing no finite edge cycle. A *spanning tree* of a connected graph is a subtree containing every vertex.

Every connected graph has a spanning tree. Start with all vertices and use Zorn's lemma on the spanning subgraphs containing no finite cycle, ordered by inclusion. The union of a chain has no cycle, because a finite cycle in the union would lie in one member of the chain. A maximal such subgraph is connected: if two of its components could be joined by an edge, that edge could be added without creating a cycle. The ambient graph is connected in the edge-path sense, since its edge-path components are open and closed by the weak topology. Therefore an edge joining distinct components would exist unless the maximal subgraph were connected. Between two vertices of a tree there is a unique reduced finite edge path; two distinct such paths would give a finite cycle.

We first justify the finiteness reduction needed for loops and homotopies.

Lemma 12.4.1. Every compact subset of a graph is contained in a finite subgraph.

Proof. Suppose a compact set met infinitely many open cells, and choose points $z_1, z_2, \ldots$ in distinct such cells. Let A be their set. Every subset of A is closed in the graph: its intersection with each closed edge consists of at most the two endpoints and one point in the interior, so

its inverse image under every characteristic map is closed. The weak topology gives the assertion.

In particular, A is a closed subset of the compact set, and is therefore compact. But every subset of A is closed, so A is an infinite discrete space, which is not compact. This contradiction shows that the compact set meets only finitely many open cells. The union of their closures is a finite subgraph containing it. □

Choose a base vertex v_0 in a connected graph Γ and a spanning tree T. For each vertex v, let τ_v be the unique reduced edge path in T from v_0 to v. Choose an orientation of each edge outside T. If such an edge e runs from u to v, define the based loop

$$\lambda_e = \tau_u * e * \overline{\tau_v}.$$

Theorem 12.4.2. The fundamental group $\pi_1(\Gamma, v_0)$ is free on the classes $[\lambda_e]$, with one generator for each oriented edge outside the spanning tree T.

Proof. For a finite connected graph and a spanning tree, collapse the tree. This is a homotopy equivalence onto a finite wedge of circles, one for each remaining edge, as proved in Lecture 11. The van Kampen theorem identifies the fundamental group with the free group on these circles. Under this identification, the loops λ_e give the stated basis.

Now allow Γ to be infinite. The image of a based loop is contained in a finite subgraph by Lemma 12.4.1. Enlarge this subgraph by adjoining the finitely many tree paths from v_0 to its vertices. The result is a finite connected subgraph Γ_0 for which $T \cap \Gamma_0$ is a spanning tree. The finite case expresses the given loop class as a word in the classes $[\lambda_e]$ for edges of Γ_0 outside T. Hence these classes generate $\pi_1(\Gamma, v_0)$.

Suppose a reduced nonempty word in finitely many of these generators were nullhomotopic in Γ. The representing loops and the image of a nullhomotopy are contained in a finite subgraph, again by the compactness lemma. Adjoin the tree paths from v_0 to its vertices to obtain a finite connected subgraph Γ_1 with spanning tree $T \cap \Gamma_1$. The alleged nullhomotopy now takes place in Γ_1. In its finite-graph free basis the word is still reduced and nonempty, contradicting freeness. Thus there are no relations among the indicated generators. □

12.4.2 Covering Graphs, Free Subgroups, and Their Ranks

Graphs satisfy the local hypotheses of covering space classification. At an interior point of an edge, use a small interval. At a vertex, use an open star consisting of that vertex and a sufficiently short half-interval in each incident edge; for a loop edge, use disjoint half-intervals at its two ends. These stars form a neighborhood basis, even at a vertex of infinite valence. Radially contracting every half-interval gives a contraction of the star. Continuity at the vertex follows because a smaller star remains inside any prescribed star throughout the contraction. Thus graphs are locally path connected and semilocally simply connected.

Lemma 12.4.3. The total space of a covering of a graph has a graph structure in which vertices and edges project to vertices and edges of the base.

Proof. Take the inverse images of the vertices as the new vertices. Orient each edge of the base and lift its characteristic path $I \to \Gamma$ once from every point above its initial vertex. Each lift supplies a closed edge with endpoints among the new vertices. Over the interior of the base edge, path lifting or the lifting criterion for the interval identifies each component of its inverse image homeomorphically with the interval. Thus the lifted open edges are disjoint and cover the inverse images of all open edges.

These attachments define a CW graph with the same underlying set as the covering space. Their topologies agree. On an evenly covered interval this follows from the inverse branches

of the covering. Near a lifted vertex, choose an evenly covered open star downstairs. Its sheet is homeomorphic to that star and meets each incident lifted edge in the corresponding half-interval. The weak graph topology on such a star is the same topology, since openness can be tested separately on its incident half-intervals. These neighborhoods cover the total space, proving the claim. □

Theorem 12.4.4. Every subgroup of a free group is free.

Proof. Let F be the free group on a set S, and let $H \leq F$. Form the CW graph R_S with one vertex and one oriented loop for each element of S. Theorem 12.4.2, with the single vertex as spanning tree, identifies $\pi_1(R_S)$ with F.

Covering space classification gives a connected based cover $p : Y_H \to R_S$ whose image subgroup is H. The covering-space injection identifies $\pi_1(Y_H)$ with H. By Lemma 12.4.3, the space Y_H is a graph, and Theorem 12.4.2 makes its fundamental group free. This proves the assertion for arbitrary sets of generators and arbitrary subgroups. □

Corollary 12.4.5. Let F_r be a free group on $r \geq 1$ generators, and let $H \leq F_r$ have finite index d. Then H is free of rank

$$1 + d(r-1).$$

Proof. Use the d-sheeted connected covering of the rose with r loops. There are d vertices, one for each point over the base vertex. Each of the r oriented base edges has one lift starting at each lifted vertex, giving dr edges. A spanning tree has $d-1$ edges, so Theorem 12.4.2 gives $dr-(d-1)$ free generators. □

Example 12.4.6. Let R be the rose with loops a, b. Construct a graph with vertices $1, 2, 3$, an oriented a-edge from each vertex i to $(12)(i)$, and an oriented b-edge from each i to $(23)(i)$. Map each labeled edge to the corresponding loop of R. At every vertex there is one incoming and one outgoing edge of each label, so small vertex stars map homeomorphically to the base star. The edge interiors do as well. This is a connected three-sheeted covering.

With base vertex 1, its subgroup $H \leq F(a,b)$ consists of words whose successive edge lifts return to 1. The word b belongs to H, whereas aba^{-1} carries 1 to 3 and does not. Hence H is not normal. A deck transformation permutes the vertices and commutes with both monodromy permutations. Since (23) fixes only vertex 1, the image of vertex 1 under a deck transformation must be 1. Unique lifting then makes that transformation the identity. Thus this three-sheeted cover has trivial deck group.

There are six edges and three vertices, so H has rank four. Choose the spanning tree consisting of the a-edge from 1 to 2 and the b-edge from 2 to 3. The four edges outside this tree give the explicit free basis

$$b, \qquad a^2, \qquad abab^{-1}a^{-1}, \qquad ab^2a^{-1}.$$

The labels in these words record the projected loops $\tau_u * e * \overline{\tau_v}$ from Theorem 12.4.2.

Example 12.4.7. Consider the homomorphism $F(a,b) \to \mathbb{Z}$ sending a to 1 and b to 0. Its kernel corresponds to the graph whose vertices are the integers, with an a-edge from n to $n+1$ and a b-loop at each n. Projection to the two-loop rose is a covering, and lifting a word changes the vertex by its total exponent of a. Its associated subgroup is therefore the stated kernel.

The line of a-edges is a spanning tree. The b-loops give the free basis

$$\{a^n b a^{-n} \mid n \in \mathbb{Z}\}.$$

Thus a subgroup of a free group on two generators can be free on countably infinitely many generators. Translation of the integer vertices gives the deck group $\mathbb{Z}$, in agreement with the normal-subgroup calculation.

For covering spaces and their classification in categorical form, see J. P. May, *A Concise Course in Algebraic Topology*, Chapter 3, especially Sections 7–8, pp. 28–33. Chapter 4, pp. 35–38, treats graphs and their applications to groups. The author's text is available at https://www.math.uchicago.edu/~may/CONCISE/ConciseRevised.pdf.

For supplementary comparisons, see Allen Hatcher, *Algebraic Topology*, Section 1.3: Propositions 1.30–1.34, pp. 60–62, for lifting; the construction and classification on pp. 63–68, including Theorem 1.38; and Proposition 1.39, pp. 71–72, for deck transformations. Section 1.A, pp. 83–85, proves the graph results and the theorem that subgroups of free groups are free. The author's chapter is available at https://pi.math.cornell.edu/~hatcher/AT/ATch1.pdf.