The Fundamental Group and Its Applications

Algebraic Topology I | Fall 2026 | November 9 and 11, 2026

11.1 The Fundamental Group of the Circle

The fundamental group records the order in which a loop passes through different parts of a space. Our first calculation makes this order visible by lifting a loop in the circle to a path in the real line. We then prove a gluing theorem for fundamental groups and use it to compute groups of graphs and surfaces. The final applications concern roots of complex polynomials and antipodal points on spheres.

We retain the conventions of Lecture 9. The interval is $I = [0,1]$, and the product $\alpha * \beta$ traverses α first and β second. The reverse path is $\overline{\alpha}(t) = \alpha(1-t)$. Path homotopies fix both endpoints unless stated otherwise. For a based map f, the induced homomorphism is denoted by f_*. Write $\mathrm{S}^1 = \{z \in \mathbb{C} : |z| = 1\}$, with basepoint 1, and define

$$p : \mathbb{R} \longrightarrow \mathrm{S}^1, \qquad p(t) = e^{2\pi i t}.$$

A lift of a map $g : Y \to \mathrm{S}^1$ is a map $\tilde{g} : Y \to \mathbb{R}$ satisfying $p \circ \tilde{g} = g$. We establish the required lifting properties directly, before the general theory of covering spaces in Lecture 12.

11.1.1 Local Inverses and Lifting a Path

An angle can be chosen continuously on any sufficiently short open arc of the circle. The ambiguity between two choices is an integer. These two elementary facts provide both existence and uniqueness of lifts.

For $a \in \mathbb{R}$, let $J_a = (a - \frac{1}{4}, a + \frac{1}{4})$ and let $A_a = p(J_a)$. The restriction $p|_{J_a} : J_a \to A_a$ is a homeomorphism. Its inverse is a continuous choice of angle. Every point of $p^{-1}(A_a)$ belongs to a unique interval $J_a + k$, with $k \in \mathbb{Z}$, and restriction to each such interval is again a homeomorphism onto A_a. These assertions follow from the elementary identities

$$p(s) = p(t) \quad \Longleftrightarrow \quad s - t \in \mathbb{Z}$$

and the usual angle parametrization of an open semicircle.

Lemma 11.1.1. For every path $\gamma : I \to \mathrm{S}^1$ and every $a \in \mathbb{R}$ with $p(a) = \gamma(0)$, there is a unique lift $\tilde{\gamma} : I \to \mathbb{R}$ satisfying $\tilde{\gamma}(0) = a$.

Proof. The inverse images $\gamma^{-1}(A_b)$ form an open cover of the compact interval. The Lebesgue number lemma gives a subdivision

$$0 = t_0 < t_1 < \cdots < t_m = 1$$

such that the image of each closed subinterval lies in one arc A_b. On the first subinterval, choose the inverse of p whose value at $\gamma(0)$ is a. Once the lift has been constructed through t_{j-1}, there is a unique integer translate of an inverse on the next arc that agrees with this value. Use that inverse on $[t_{j-1}, t_j]$. The finitely many continuous pieces agree at their endpoints, so the pasting lemma gives a continuous lift on I.

For uniqueness, two lifts u, v have $u(t) - v(t) \in \mathbb{Z}$ for every t. Their difference is continuous and I is connected, so it is constant. Agreement at 0 makes that constant zero. □

The same uniqueness argument applies on any connected domain: two real-valued lifts that agree at one point agree everywhere. Existence, by contrast, needs an argument adapted to the domain.

11.1.2 Lifting a Homotopy on a Square

To distinguish loop classes, we must know that the endpoint of a lifted loop does not change during a based homotopy. Continuity in the homotopy parameter is the essential point; lifting each path separately would not yet prove it.

Lemma 11.1.2. Let $H : I \times I \to \mathrm{S}^1$ be continuous, with coordinates (s, t). Given a lift $u : I \to \mathbb{R}$ of $H(s, 0)$, there is a unique continuous lift $\widetilde{H} : I \times I \to \mathbb{R}$ satisfying $\widetilde{H}(s, 0) = u(s)$.

Proof. Apply the Lebesgue number lemma to the cover of the square by the sets $H^{-1}(A_a)$. Choose a rectangular grid fine enough that each closed rectangle Q is mapped into one such arc. Process the bottom row from left to right, then the next row from left to right, and continue upward.

On the first rectangle, choose an inverse branch over its arc and translate it by an integer so that its composition with H agrees with u at the lower-left corner. Along the bottom edge, both functions lift the same map. Their difference is integer-valued and continuous, so agreement at the corner gives agreement on the entire edge.

For each subsequent rectangle, its bottom edge is already lifted, either by u or by the preceding row. Its left edge is also already lifted unless it is the first rectangle in its row. Choose the integer translate of a local inverse that agrees at the lower-left corner. Uniqueness on a connected interval makes it agree along the entire bottom edge and, when present, the entire left edge. These edges contain its intersection with the previously processed rectangles. Thus each new formula agrees with the existing lift wherever both are defined. Finite closed pasting now gives a continuous lift on the whole square.

Two lifts of H differ by a continuous integer-valued function on the connected square. Their agreement on its bottom edge proves uniqueness. □

Corollary 11.1.3. For a based homotopy of loops at 1, the lifts that start at 0 all have the same endpoint.

Proof. Lift the homotopy by Lemma 11.1.2, starting with the lift of its bottom loop. Each vertical boundary edge has image in $p^{-1}(1) = \mathbb{Z}$. Its lift is continuous, hence constant. The left edge is therefore constantly 0, and the right edge is the common endpoint of all the lifted loops. □

11.1.3 Winding Number and Degree

A loop lifts to a path whose endpoints differ by an integer. This integer counts its signed turns around the circle. The calculation below identifies the group law as addition of these displacements.

Theorem 11.1.4. For a loop γ at 1, let $\tilde{\gamma}$ be its lift starting at 0. The assignment

$$w : \pi_1(\mathrm{S}^1, 1) \longrightarrow \mathbb{Z}, \qquad w([\gamma]) = \tilde{\gamma}(1)$$

is an isomorphism. The loop $\omega_n(t) = e^{2\pi i n t}$ represents the integer n.

Proof. The endpoint is an integer because $p(\tilde{\gamma}(1)) = 1$. Corollary 11.1.3 makes the assignment well-defined on based homotopy classes. If $\tilde{\gamma}(1) = m$ and $\tilde{\eta}(1) = n$, the lift of $\gamma * \eta$ first follows $\tilde{\gamma}$ and then follows $m + \tilde{\eta}$. Its endpoint is $m + n$, so w is a homomorphism.

The lift of ω_n starting at 0 is $t \mapsto nt$, which proves surjectivity. If the lift of γ ends at n, the straight-line homotopy

$$(t, r) \longmapsto (1 - r)\tilde{\gamma}(t) + rnt$$

fixes its endpoints 0 and n. Composing with p gives a based homotopy from γ to ω_n. Thus a loop is determined up to based homotopy by its endpoint displacement, proving injectivity. □

For a loop γ whose basepoint is any point of the circle, define its winding number by $\tilde{\gamma}(1)-\tilde{\gamma}(0)$. A different lift adds the same integer at both endpoints, so the definition is independent of the lift. If H is a homotopy through loops whose basepoints may move, square lifting still shows that

$$t \longmapsto \widetilde{H}(1,t) - \widetilde{H}(0,t)$$

is a continuous integer-valued function. Consequently winding number is invariant under these homotopies as well.

Recall that the degree of a continuous map $f : \mathrm{S}^1 \to \mathrm{S}^1$ is the integer by which it acts on $\mathrm{H}_1(\mathrm{S}^1;\mathbb{Z})$. The winding number of $t \mapsto f(e^{2\pi i t})$ computes this degree. Indeed, the degree-one Hurewicz theorem stated in Lecture 9 identifies $\pi_1(\mathrm{S}^1,1)$ with $\mathrm{H}_1(\mathrm{S}^1;\mathbb{Z})$, sending the positive loop to the positive homology generator. Naturality identifies the two multipliers. For a map not preserving 1, a rotation makes it based without changing either multiplier. The following proof obtains the properties needed here directly from winding number.

Proposition 11.1.5. Degree is invariant under homotopy. For continuous maps $f, g : \mathrm{S}^1 \to \mathrm{S}^1$, it satisfies

$$\deg(f \circ g) = \deg(f)\deg(g), \qquad \deg(z \mapsto z^n) = n \quad (n \in \mathbb{Z}).$$

If f extends to a continuous map $\mathrm{D}^2 \to \mathrm{S}^1$, then $\deg(f) = 0$.

Proof. Homotopy invariance follows from the preceding winding-number argument. The displayed power map has lift $t \mapsto nt$, giving its degree.

Lift $t \mapsto f(e^{2\pi i t})$ on $[0,1]$ and denote this lift by F. Write $d = \deg(f)$. Extend the lift to $\mathbb{R}$ by the rule

$$F(t+k) = F(t) + kd \qquad (0 \le t \le 1,\ k \in \mathbb{Z}).$$

The endpoint identity $F(1)-F(0) = d$ makes the formulas agree at integers, and hence defines a continuous function. In particular, $F(u+k) = F(u)+kd$ for every real u and integer k. Construct a lift G for g in the same manner, with displacement $e = \deg(g)$. The function $F \circ G$ lifts the composite, and its displacement is

$$F(G(1)) - F(G(0)) = F(G(0)+e) - F(G(0)) = de.$$

Finally, if $\overline{f} : \mathrm{D}^2 \to \mathrm{S}^1$ extends f, the maps $z \mapsto \overline{f}((1-r)z)$ give a homotopy from f to a constant map. Such a map has degree zero. □

11.2 Free Products and Relations

When a loop passes between two open sets, it breaks into a word of loops drawn from those sets. The relevant algebra must preserve the order of the letters. Free products provide the universal group in which words from the two factors can be multiplied without imposing extra commutativity relations.

11.2.1 Reduced Words and the Universal Property

For groups A and B, regard their nonidentity elements as distinct letters, tagged by their factor even when the underlying sets overlap. A word is reduced if consecutive letters come from different factors. The empty word is allowed.

Proposition 11.2.1. There is a group $A*B$, with homomorphisms $j_A : A \to A*B$ and $j_B : B \to A*B$, such that for every group K and every pair of homomorphisms $a : A \to K$ and $b : B \to K$, there is a unique homomorphism $h : A*B \to K$ satisfying $hj_A = a$ and $hj_B = b$. Every element of $A*B$ has a unique reduced-word expression. Both factor homomorphisms are injective.

Proof. Let $\mathcal{W}$ be the set of reduced words. For $g \in A$, define a permutation L_g of $\mathcal{W}$ by inserting g on the left and reducing at that end. More explicitly, L_1 is the identity. For $g \neq 1$, if the word is empty or starts with a B-letter, prepend g. If its first letter is $a_1 \in A$, replace that letter by ga_1 when this product is nonidentity, and delete it when the product is the identity. Define L_g for $g \in B$ in the same way.

For g, h in the same factor, these operations satisfy $L_g L_h = L_{gh}$. To verify this, a word starting in the other factor, or an empty word, acquires the single product gh, with that letter omitted if it is 1. If the first letter is in the same factor, the result is obtained by multiplying g, h and that first letter there. If h cancels the first letter, the next letter belongs to the other factor; prepending g gives the same result as first multiplying gh by the original first letter. These cases cover every word. It follows in particular that $L_g^{-1} = L_{g^{-1}}$.

Define $A * B$ to be the subgroup of the permutation group of $\mathcal{W}$ generated by these permutations. The preceding identities give the two factor homomorphisms. Every product of generators reduces by multiplying adjacent letters from the same factor and omitting identities. A reduced product $L_{g_1} \cdots L_{g_r}$ sends the empty word to $(g_1, \dots, g_r)$. Thus different reduced words represent different permutations. This proves existence and uniqueness of the normal form, including the injectivity of the factors.

Given a and b, send a reduced word to the ordered product of the images of its letters in K. Combining adjacent letters from the same factor or omitting an identity does not change that product. Hence the assignment respects multiplication and defines h. Since the two factors generate $A * B$, this homomorphism is unique. □

The construction works for any set of groups by allowing letters tagged by an arbitrary factor and requiring consecutive factors to differ. The same left-insertion proof gives unique reduced words. In particular, for a set S, the free group $F(S)$ is the free product of infinite cyclic groups with generators indexed by S. Specifying a homomorphism $F(S) \to K$ is equivalent to choosing an arbitrary element of K for each member of S.

Example 11.2.2. For two generators a, b, the word $aba^{-1}b^{-1}$ is a nonempty reduced word in $F(\{a, b\}) = \mathbb{Z} * \mathbb{Z}$. It is therefore nonidentity. Thus a free product of abelian groups need not be abelian; this is the algebraic distinction between a bouquet of two circles and a torus.

11.2.2 Normal Closures and Amalgamation

Attaching a disk makes its boundary loop nullhomotopic. A group quotient recording this fact must also make every conjugate of that loop trivial. This explains why normal closures, rather than ordinary generated subgroups, occur in topological presentations.

For a subset R of a group G, write $\langle\langle R \rangle\rangle$ for its normal closure, the smallest normal subgroup containing R. Concretely, its elements are finite products of conjugates $gr^{\varepsilon}g^{-1}$, with $g \in G$, $r \in R$, and $\varepsilon \in \{1, -1\}$. These products form a normal subgroup containing R, and every normal subgroup containing R contains them. A presentation $\langle S \mid R \rangle$ denotes $F(S)/\langle\langle R \rangle\rangle$, where the elements of R are words in the generators and their inverses.

Definition 11.2.3. For homomorphisms $u : C \to A$ and $v : C \to B$, the amalgamated free product is the quotient

$$A *_C B = (A * B) \big/ \langle\langle j_A(u(c)) j_B(v(c))^{-1} : c \in C \rangle\rangle.$$

The maps u and v need not be injective.

The universal property of the free product and the quotient says that homomorphisms $A *_C B \to K$ correspond to homomorphisms $A \to K$ and $B \to K$ whose restrictions along u and v agree. Equivalently, this quotient is the pushout of the two homomorphisms from C in the

category of groups. Even when the original maps into the free product are injective, their composites into this quotient need not be.

11.3 The Seifert–van Kampen Theorem

An open cover lets us subdivide paths and homotopies into pieces lying in individual members of the cover. Path subdivision gives generators. Subdivision of a homotopy square proves that the overlap supplies all the relations. We keep both arguments explicit because the second is what makes the theorem a computation of the group, rather than only a list of generators.

11.3.1 The Pushout Statement

Let $X = U \cup V$, where U, V are open and path-connected, and suppose that $W = U \cap V$ is nonempty and path-connected. Fix $x_0 \in W$. Denote the inclusions from W by i_U, i_V, and those into X by j_U, j_V.

Theorem 11.3.1. Under these hypotheses, the square of induced homomorphisms

$$\begin{array}{ccc} \pi_1(W, x_0) & \xrightarrow{(i_U)_*} & \pi_1(U, x_0) \\ {\scriptstyle (i_V)_*}\downarrow & & \downarrow{\scriptstyle (j_U)_*} \\ \pi_1(V, x_0) & \xrightarrow[(j_V)_*]{} & \pi_1(X, x_0) \end{array}$$

is a pushout of groups. In particular, the natural homomorphism

$$\pi_1(U, x_0) *_{\pi_1(W, x_0)} \pi_1(V, x_0) \longrightarrow \pi_1(X, x_0)$$

is an isomorphism.

Write P for the displayed amalgamated free product. Inclusion of loops gives a homomorphism $\theta : P \to \pi_1(X, x_0)$, since the two images of every overlap loop agree in X. We will construct its inverse by assigning to each loop a subdivided word in P.

11.3.2 Subdivided Loop Words and Their Choices

Take a based loop $\gamma : I \to X$. The Lebesgue number lemma gives a subdivision $0 = t_0 < \cdots < t_m = 1$ such that every restricted path γ_j lies in one chosen set $C_j \in \{U, V\}$. Here and below, restrictions to subintervals are linearly reparametrized to I.

Choose paths λ_j from x_0 to $\gamma(t_j)$ as follows. At $j = 0, m$, use the constant path. At an interior division point where the neighboring colors differ, choose λ_j in W. This is possible because the point lies in both sets and W is path-connected. Where the colors agree, choose λ_j in their common set. The loop

$$\lambda_{j-1} * \gamma_j * \overline{\lambda_j}$$

then lies in C_j. Map its class into P, and define $\Psi(\gamma)$ to be the ordered product of these m classes.

Lemma 11.3.2. The element $\Psi(\gamma) \in P$ is independent of the subdivision, the colors C_j, and the paths λ_j.

Proof. First change one interior path from λ_j to μ_j. The loop $\lambda_j * \overline{\mu_j}$ lies in the common color when the two colors agree, and in W when they differ. Write d for its image in P. The jth factor changes from a_j to $a_j d$, while the next changes from a_{j+1} to $d^{-1} a_{j+1}$. In the different-color case,

the defining overlap relation in P identifies the two images of d. Thus the complete product is unchanged. Repeating this proves independence of all the chosen paths.

Next insert a division point in a segment of color C. Choose any connecting path to it in C, keeping the two new segments of color C. Multiplication within $\pi_1(C, x_0)$ cancels this connecting path and its reverse, so the two new factors have the same product as the old one. This proves invariance under subdivision and its inverse operation.

Finally, suppose that a segment can be assigned either color. Its entire image then lies in W. Before changing its color, choose the connecting paths at its endpoints in W, using the already proved path independence. Those paths are admissible for either color and for every neighboring segment. The factor associated to the segment is now represented by a loop in W, so the overlap relation identifies its two possible images in P.

Given any two initial choices, pass to the common refinement of their subdivisions. Each refined segment lies in both of its inherited colors if those colors disagree. Change these colors one at a time by the preceding argument, then change the connecting paths. All choices therefore give the same element. □

The product of the loops defining $\Psi(\gamma)$, viewed in X, is homotopic to γ: adjacent connecting paths occur as $\overline{\lambda_j} * \lambda_j$ and contract relative to their endpoints. Consequently $\theta(\Psi(\gamma)) = [\gamma]$. In particular, θ is surjective.

11.3.3 The Relations from a Homotopy Square

It remains to show that the subdivided word depends only on the based homotopy class. For this we label all oriented edges of a sufficiently fine grid and use the relation supplied by each small rectangle.

Proof. Let $H : I^2 \to X$ be a based homotopy from γ_0 to γ_1, with path coordinate horizontal and homotopy coordinate vertical. Subdivide the square by a rectangular grid such that each closed cell has image in one of U, V, and give the cell such a color. This is possible by applying the Lebesgue number lemma to $H^{-1}(U), H^{-1}(V)$.

At each grid vertex v, choose a path λ_v from x_0 to $H(v)$. If incident cells have both colors, choose it in W; if they all have one color, choose it in that set. On the two vertical boundary edges every vertex has image x_0, and we choose the constant paths there. Each choice is possible by the assumed path-connectedness. Moreover, the chosen path at a vertex lies in the color of every cell incident to that vertex.

For an oriented grid edge $e : v \to w$, let H_e be the restriction of H along it and assign to it the element

$$a_e = [\lambda_v * H_e * \overline{\lambda_w}] \in P,$$

computed in either incident cell's group. If the two incident colors differ, the edge image and its two connecting paths lie in W, so the two answers agree in P. Reversing the edge replaces a_e by a_e^{-1}.

For one small rectangle, denote the bottom and top edge labels, oriented to the right, by a and d, and the left and right labels, oriented upward, by c and b. All four connecting paths and the whole rectangle image lie in its chosen color. Within the parameter rectangle, the bottom-then-right and left-then-top paths are homotopic relative endpoints by straight-line interpolation. Applying H gives such a homotopy in its color. Cancellation of the connecting paths at the intermediate vertices therefore gives the relation

$$ab = cd \quad \text{in } P.$$

Thus the elementary change of a grid path across this rectangle preserves its product of edge labels.

Every monotone grid path from the lower-left corner to the upper-right corner consists of rightward and upward steps. By successively exchanging an adjacent rightward-then-upward pair for an upward-then-rightward pair, any such path changes to the path that follows the left and then the top boundary. Each exchange occurs across one grid rectangle and uses the relation just proved. It follows that the bottom-then-right boundary product equals the left-then-top boundary product. The vertical boundary labels are identities, since their paths and connecting paths are constant. By Lemma 11.3.2, the remaining bottom and top products are $\Psi(\gamma_0)$ and $\Psi(\gamma_1)$. They are therefore equal.

We have proved that Ψ descends to a map $\pi_1(X, x_0) \to P$. Concatenating subdivisions for two loops shows that this map is a homomorphism: the connecting path at their shared basepoint is constant, and the resulting word is the product of their words. We already know that $\theta\Psi$ is the identity. For a loop lying entirely in U or entirely in V, choose the single-segment subdivision and constant endpoint paths. Its image under $\Psi\theta$ is its original class in P. These factor images generate P, so $\Psi\theta$ is the identity as well. Hence θ is an isomorphism, and the universal property of the amalgamated free product proves the pushout statement. □

If the overlap is simply connected, the theorem gives the ordinary free product. The overlap hypothesis has geometric content: two contractible open arcs can cover the circle, but their intersection has two components. Applying the two-set statement without path-connectedness would incorrectly give a trivial group. A groupoid version retains several basepoints and handles this situation; it is not needed for the applications below.

11.4 Graphs, Attaching Maps, and Surface Groups

The gluing theorem turns decompositions into presentations. Graphs contribute generators, and attached disks contribute relations. We begin with finite graphs so that every collapse and every open cover can be described directly.

11.4.1 Bouquets of Circles and Maximal Trees

A finite graph is a space obtained from finitely many vertices and finitely many copies of I by identifying the two endpoints of each interval with vertices. Loops and multiple edges are allowed. A tree is a connected graph with no embedded circle. A spanning tree is a subgraph containing every vertex and forming a tree.

For the bouquet $B_r = \bigvee_{j=1}^{r} \mathrm{S}^1$, choose the common vertex as basepoint. When $r \geq 2$, let U contain the first $r-1$ circles and a small open arc around the vertex in the last circle. Let V contain the last circle and small open arcs around the vertex in the others. These are open sets. They deformation retract onto B_{r-1} and S^1, while their intersection is a finite tree of arcs and contracts to the vertex. Van Kampen and induction therefore give

$$\pi_1(B_r) \cong F(\{a_1, \dots, a_r\}),$$

where a_j traverses the jth circle positively. The cases $r = 0, 1$ are a point and the circle calculation, respectively.

Theorem 11.4.1. Let Γ be a finite connected graph with v vertices and e edges. Choose a vertex x_0 and a spanning tree T. Then $\pi_1(\Gamma, x_0)$ is free of rank $e - v + 1$. A free basis consists of the loops that travel from x_0 along T to the initial vertex of an oriented edge outside T, traverse that edge, and return to x_0 along T.

Proof. A spanning tree can be constructed starting at x_0 and adjoining an edge to a new vertex whenever not all vertices have been reached. Connectedness supplies such an edge,

and adjoining a new vertex creates no cycle. The resulting tree has $v-1$ edges, one for each vertex added after x_0.

Collapse T to a point, obtaining a quotient map $q : \Gamma \to \Gamma/T$. The quotient is a bouquet with one circle for each edge outside T, hence with $e-v+1$ circles. We verify that q is a based homotopy equivalence rather than merely assuming that this collapse preserves the fundamental group.

Give each tree edge length one. Between x_0 and any point of T there is a unique simple path: two distinct simple paths would contain a cycle. Contract T by moving a point at distance d from x_0 to the point at distance $(1-t)d$ on this path. This contraction C_t is continuous, as is seen on each of the finitely many edges, and fixes x_0.

For an oriented edge a outside T, with endpoints v_a, w_a, extend this contraction over that edge at time t as follows. Traverse the tree path from $C_t(v_a)$ to v_a, then the original edge a, then the tree path from w_a to $C_t(w_a)$. Allocate parameter lengths $t/(1+2t)$, $1/(1+2t)$, and $t/(1+2t)$ to these three pieces. At $t=0$, the first and last pieces have length zero and this is the original edge map. Their images shrink to the endpoints as t tends to zero, so the family is continuous there as well. The endpoint values agree with C_t, so the edge maps and the tree contraction paste to a based homotopy on Γ.

At time one, the map is constant on T and therefore factors as sq, where $s : \Gamma/T \to \Gamma$ sends each circle to the corresponding tree-edge-tree loop. We have proved $sq \simeq \mathrm{id}_\Gamma$. On each circle of Γ/T, the map qs inserts constant initial and final segments around one traversal of that circle. A linear change of the traversal interval gives a based homotopy to the identity. These homotopies agree at the common vertex, so $qs \simeq \mathrm{id}_{\Gamma/T}$. The asserted basis and rank now follow from the bouquet computation. □

Example 11.4.2. A theta graph has two vertices joined by three edges a, b, c, all oriented from the first vertex to the second. Take the first vertex as basepoint and c as the spanning tree. The group is free on the loops $a * \bar{c}$ and $b * \bar{c}$, since $3-2+1=2$. For instance, the loop $a * \bar{b}$ represents the product $(a\bar{c})(b\bar{c})^{-1}$.

11.4.2 Attaching Disks and Higher Cells

A disk attachment identifies a boundary circle with a specified loop. The van Kampen calculation depends on an open annular neighborhood of that boundary; applying the theorem directly to the closed disk and the old space would not satisfy its hypotheses.

Proposition 11.4.3. Let X be path-connected, let $x_0 \in X$, and let $f : \mathrm{S}^1 \to X$ satisfy $f(1) = x_0$. Form the attached space by the pushout

$$\begin{array}{ccc} \mathrm{S}^1 & \xrightarrow{\ f\ } & X \\ \downarrow & & \downarrow \\ \mathrm{D}^2 & \longrightarrow & Y = X +_{\mathrm{S}^1} \mathrm{D}^2. \end{array}$$

Then inclusion induces an isomorphism

$$\pi_1(Y, x_0) \cong \pi_1(X, x_0) / \langle\langle [f] \rangle\rangle.$$

Here $[f]$ means the class of the loop $t \mapsto f(e^{2\pi i t})$.

Proof. Write $q : X + \mathrm{D}^2 \to Y$ for the quotient map. Let U be the image of X together with the disk points satisfying $\|z\| > 1/3$, and let V be the image of the disk points satisfying $\|z\| < 2/3$. Their inverse images under q are open saturated subsets, so U, V are open. They cover Y.

Radial motion toward the disk boundary gives a deformation retraction of U onto X, and V is an open disk. Their intersection is the annulus $1/3 < \|z\| < 2/3$.

Choose $b = q(1/2)$ in the intersection. Moving radially from b to $q(1) = x_0$ supplies the basepoint comparison. Under the deformation retraction of U onto X, the positive circle generating the annulus group becomes the loop f. Its image in $\pi_1(V, b)$ is trivial. Van Kampen therefore identifies $\pi_1(Y, b)$ with the quotient of $\pi_1(X, x_0)$ by the normal closure of $[f]$. Transporting the basepoint to x_0 gives the stated isomorphism, with the homomorphism induced by inclusion. □

For an attaching circle whose chosen point maps elsewhere, choose a path from x_0 to that image point. The relation is represented by the attaching loop conjugated along this path. Changing the path changes that relation by conjugacy, which leaves its normal closure unchanged. Repeated application handles any finite collection of attached disks.

Lecture 9 already established the simple connectedness of higher spheres. The following alternate proof by van Kampen also supplies the local calculation for higher-dimensional cell attachments.

Proposition 11.4.4. For every $n \geq 2$, the sphere S^n is simply connected. Attaching a cell of dimension $n \geq 3$ to a path-connected space does not change its fundamental group.

Proof. Let N and S be the two poles of S^n, and let U and V be their respective complements. Stereographic projection identifies each with $\mathbb{R}^n$, so both are contractible. Their intersection is homeomorphic to $\mathrm{S}^{n-1} \times \mathbb{R}$ and is path-connected for $n \geq 2$. Van Kampen is the pushout of two trivial groups, and hence gives $\pi_1(\mathrm{S}^n) = 1$.

For an attached n-cell, use the same two open sets as in Proposition 11.4.3, with D^n in place of D^2. Their intersection deformation retracts onto S^{n-1}. When $n \geq 3$, that intersection and the inner open ball have trivial fundamental group. The pushout therefore identifies the fundamental group of the attached space with that of the original space. □

11.4.3 Tori, Projective Planes, and Closed Surfaces

Polygonal descriptions of surfaces now become presentations: oriented edges give generators and the boundary word of the polygon gives a relation. The word must be read in order, including the orientation of each side.

Example 11.4.5. The torus is obtained from a bouquet with generators a, b by attaching a disk along $aba^{-1}b^{-1}$. Thus

$$\pi_1(\mathrm{S}^1 \times \mathrm{S}^1) \cong \langle a, b \mid aba^{-1}b^{-1} \rangle \cong \mathbb{Z}^2.$$

For the last isomorphism, the relation says that a and b commute, so every word is $a^m b^n$. Sending a to $(1, 0)$ and b to $(0, 1)$ defines a homomorphism to $\mathbb{Z}^2$, and $(m, n) \mapsto a^m b^n$ is its inverse. This also proves uniqueness of this expression.

Example 11.4.6. The real projective plane is the quotient of a disk by identifying antipodal boundary points. Its boundary quotient is $\mathbb{RP}^1 \cong \mathrm{S}^1$, and the attaching map traverses this quotient circle twice. Consequently

$$\pi_1(\mathbb{RP}^2) \cong \langle a \mid a^2 \rangle \cong \mathbb{Z}/2\mathbb{Z}.$$

Indeed, under the identification $\mathbb{RP}^1 \to \mathrm{S}^1$ given by $[e^{i\theta}] \mapsto e^{2i\theta}$, the attaching map is $z \mapsto z^2$ and has winding number 2.

For $g \geq 1$, define the orientable surface Σ_g by identifying the paired sides of a $4g$-gon according to the boundary word

$$a_1 b_1 a_1^{-1} b_1^{-1} \cdots a_g b_g a_g^{-1} b_g^{-1}.$$

Its vertices become one vertex, its sides form a bouquet of $2g$ circles, and its interior is one attached disk. If $[a,b]$ denotes $aba^{-1}b^{-1}$, the attachment calculation gives

$$\pi_1(\Sigma_g) \cong \langle a_1, b_1, \dots, a_g, b_g \mid [a_1,b_1]\cdots[a_g,b_g]\rangle.$$

For genus two, this is one relation involving four generators, rather than two separate commutativity relations. In fact, for every $g \geq 2$ there is a surjection onto the free group on two generators: send a_1, a_2 to those generators and all remaining generators to the identity. The displayed relation maps to the identity. Since a quotient of an abelian group is abelian and this free group is not, these surface groups are nonabelian.

The nonorientable surface N_h, defined by the polygon with boundary word $a_1^2 \cdots a_h^2$, similarly has

$$\pi_1(N_h) \cong \langle a_1, \dots, a_h \mid a_1^2 \cdots a_h^2 \rangle.$$

The case $h = 1$ recovers the projective plane. A usual square description of the Klein bottle instead gives $\langle a, b \mid aba^{-1}b\rangle$, or equivalently $aba^{-1} = b^{-1}$. This relation records that transport once around the a-direction reverses the b-direction.

11.5 Polynomial Roots and Antipodal Points

The winding number detects an obstruction to extending a circle map over a disk. In the fundamental theorem of algebra, a polynomial without roots would provide such an extension for a map of positive degree. In the Borsuk–Ulam theorem, an assumed separation of antipodal images would provide an extension of an odd circle map, whose degree must be odd.

11.5.1 The Fundamental Theorem of Algebra

The only estimates needed are on a sufficiently large circle, where the leading term of a polynomial dominates the sum of the lower terms. This estimate gives an explicit homotopy through nonzero complex numbers.

Theorem 11.5.1. Every nonconstant polynomial with complex coefficients has a complex root. Consequently every polynomial of degree $n \geq 1$ is a product of its leading coefficient and n linear factors, with roots repeated according to multiplicity.

Proof. Write

$$P(z) = a_n z^n + \sum_{k=0}^{n-1} a_k z^k, \qquad n \geq 1, \quad a_n \neq 0.$$

Suppose that P has no root. Choose $R \geq 1$ so large that

$$\sum_{k=0}^{n-1} |a_k| R^k < |a_n| R^n.$$

For example, any $R > \max\{1, \sum_{k<n} |a_k|/|a_n|\}$ has this property. The map $z \mapsto P(Rz)/|P(Rz)|$ is then defined on the entire closed unit disk. Its boundary restriction $f_R : \mathrm{S}^1 \to \mathrm{S}^1$ has degree zero by Proposition 11.1.5.

For $z \in \mathrm{S}^1$ and $t \in I$, set

$$P_t(z) = a_n R^n z^n + (1-t)\sum_{k=0}^{n-1} a_k R^k z^k.$$

The reverse triangle inequality gives

$$|P_t(z)| \geq |a_n| R^n - \sum_{k=0}^{n-1} |a_k| R^k > 0.$$

Thus $P_t(z)/|P_t(z)|$ is a homotopy from f_R to $z \mapsto (a_n/|a_n|)z^n$. Multiplication by the fixed unit complex number $a_n/|a_n|$ is homotopic to the identity rotation: choose an angle for that one number and rotate through that angle. Hence $\deg(f_R) = n$, contradicting $\deg(f_R) = 0$.

A root r gives a factorization $P(z) = (z - r)Q(z)$ by polynomial division. Apply the root-existence statement repeatedly to Q and the subsequent factors of decreasing degree. After n steps, the remaining factor is the leading coefficient. □

11.5.2 Odd Circle Maps Have Odd Degree

A map between spheres is odd if it commutes with the antipodal maps. On the circle, oddness determines the change in a lifted angle over a half-turn, which determines the parity of its degree.

Proposition 11.5.2. For every continuous map $g : \mathrm{S}^1 \to \mathrm{S}^1$ satisfying $g(-z) = -g(z)$, the integer $\deg(g)$ is odd.

Proof. As in the proof of Proposition 11.1.5, choose a lift $G : \mathbb{R} \to \mathbb{R}$ of $t \mapsto g(e^{2\pi i t})$ satisfying

$$G(t+1) - G(t) = \deg(g).$$

Oddness gives

$$e^{2\pi i G(t+1/2)} = -e^{2\pi i G(t)},$$

so the continuous function $G(t + 1/2) - G(t) - 1/2$ is integer-valued. Since $\mathbb{R}$ is connected, it is a constant $m \in \mathbb{Z}$. Applying this half-turn identity twice yields

$$\deg(g) = G(t+1) - G(t) = 2m + 1.$$

Thus the degree is odd. □

In particular, an odd circle map cannot extend over a disk. The obstruction concerns all continuous odd maps, with no differentiability or injectivity assumption.

11.5.3 Borsuk–Ulam in Dimensions One and Two

For a map to Euclidean space, the difference between the values at antipodal points is automatically odd. If that difference never vanishes, normalization produces a map to a smaller sphere. In dimension two, its equatorial restriction contradicts the preceding degree calculation.

Theorem 11.5.3. For $n = 1$ and $n = 2$, every continuous map $f : \mathrm{S}^n \to \mathbb{R}^n$ has a point x such that $f(x) = f(-x)$.

Proof. For $n = 1$, consider the continuous real-valued function $h(x) = f(x) - f(-x)$. If h vanishes anywhere, we are done. Otherwise choose any x and a path in the circle from x to $-x$. The values at the two ends are nonzero and opposite, because $h(-x) = -h(x)$. The intermediate value theorem gives a zero on that path, a contradiction.

For $n = 2$, suppose that $f(x) \neq f(-x)$ for every $x \in \mathrm{S}^2$. The formula

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

defines a continuous map $g : \mathrm{S}^2 \to \mathrm{S}^1$ satisfying $g(-x) = -g(x)$. Identify the equator with S^1 using $(x_1, x_2, 0) \leftrightarrow x_1 + i x_2$. The equatorial restriction is odd, so Proposition 11.5.2 gives it odd degree.

The same restriction extends over the closed northern hemisphere through g. That hemisphere is a disk: the homeomorphism from D^2 sends (u, v) to $(u, v, \sqrt{1-u^2-v^2})$ and identifies its boundary with the equator. Proposition 11.1.5 therefore gives degree zero, a contradiction.□

The proof in dimension two uses only circle lifting and an extension over a hemisphere. In higher dimensions, the analogous argument requires the higher-dimensional statement that an odd self-map of a sphere has odd degree. The low-dimensional theorem proved here is sufficient for the syllabus at this point.

Further reading. J. P. May, *A Concise Course in Algebraic Topology*, Chapter 1, Sections 5 and 7, pp. 8–11, treats circle lifting and polynomial roots. His Chapter 2, Sections 7–8, pp. 17–20, develops van Kampen as a categorical colimit, and Chapter 4, Sections 2–4, pp. 35–37, treats trees and graph groups. Further comparisons are Allen Hatcher, *Algebraic Topology*, Theorems 1.7, 1.8, and 1.10, pp. 29–33, for circle lifting and the two final applications; Theorem 1.20, pp. 43–46, for van Kampen; and Proposition 1.26, pp. 50–51, for cell attachments.