CHAPTER 1

Combinatorial Group Theory

Every group is the fundamental group of a cell complex and, not surprisingly, the topological properties of a cell complex X have algebraic significance for its fundamental group G. If X is compact, then G is finitely presented; if X is 1-dimensional, then G is free; and if X has a cut point, then G has a free product decomposition. This chapter examines each of these implications with an emphasis on the way these topological properties motivate the algebraic definitions. Other basic tools from combinatorial group theory, such as generating sets and Cayley graphs, are discussed along the way.

1.1. Manifolds and cell complexes

The fundamental groups of compact manifolds and compact cell complexes lie at heart of geometric group theory and the goal of this section is to establish various equivalent descriptions for this natural class of groups. We begin with the fundamental groups of arbitrary cell complexes before insisting on compactness. Because every group is the fundamental group of some cell complex (Exercise 1) this class of groups is completely unrestricted.

THEOREM 1.1.1 (Arbitrary groups). *The following classes of groups are equal:*

$$\begin{aligned}
G_1 &= \{\pi_1 \text{ of cell complexes}\},\\
G_2 &= \{\pi_1 \text{ of combinatorial cell complexes}\},\\
G_3 &= \{\pi_1 \text{ of simplicial complexes}\},\\
G_4 &= \{\pi_1 \text{ of } 2\text{-complexes}\},\\
G_5 &= \{\pi_1 \text{ of combinatorial } 2\text{-complexes}\}, \text{ and}\\
G_6 &= \{\pi_1 \text{ of simplicial } 2\text{-complexes}\}.
\end{aligned}$$

The definitions of the various types of complexes are reviewed in the appendix. By van Kampen's theorem the fundamental group of a cell complex is carried by its 2-skeleton (Corollary A.2.9). Thus, $G_1 = G_4$, $G_2 = G_5$ and $G_3 = G_6$. Moreover, since the descriptions are increasingly strict, $G_4 \supset G_5 \supset G_6$, so it suffices to prove $G_4 \subset G_5 \subset G_6$. The inclusion $G_5 \subset G_6$, or more generally $G_2 \subset G_3$, follows from the fact that the second barycentric subdivision of a combinatorial cell complex is a homeomorphic simplicial complex. For the final inclusion, $G_4 \subset G_5$, it suffices to show that every map from a circle to a graph is either null-homotopic or homotopic to an immersion. Since a variation on this argument is needed in Chapter 3, we include a complete proof.

PROPOSITION 1.1.2 (Simplifying loops and arcs). *Every map from a circle to a graph is homotopic to an immersion or a constant map. Similarly, every map from a closed interval to a graph is homotopic to an immersion or a constant map keeping the endpoints fixed throughout.*

PROOF. Let X be a graph and let $f : \mathbb{S}^1 \to X$ be a map. The graph X, by definition, can be given the structure of a 1-complex, and, by subdividing if necessary, we may assume that every edge in X is attached to distinct 0-cells. Next, consider the open cover $\mathcal{U}$ of X containing two types of open sets: (1) each individual (open) 1-cell and (2) a small open neighborhood around each 0-cell. To define the latter, imagine turning the graph into a metric space where each edge has unit length and then taking an ϵ-neighborhood of a vertex v with $\epsilon < \frac{1}{2}$. To minimize notation, let (v) denote this small open neighborhood of v.

If the image of f lies inside a single element of $\mathcal{U}$ then f is null-homotopic because every set in $\mathcal{U}$ is contractible. Otherwise, we can cover $\mathbb{S}^1$ by the maximal open subintervals of $\mathbb{S}^1$ whose image is contained in a single element of $\mathcal{U}$. Since $\mathbb{S}^1$ is compact, we can pass to a minimal finite subcover. Minimality of the cover implies that the intervals are not nested and thus they have a canonical cyclic ordering as we proceed around the circle. Maximality of the intervals further implies that each open interval can be labeled by the unique element of $\mathcal{U}$ that contains its image. Finally, the finite cover must strictly alternate between "edge" intervals and "vertex" intervals since the open sets of each type in $\mathcal{U}$ are pairwise disjoint. In other words, the covering of $\mathbb{S}^1$ can be summarized by a sequence $(v_0)e_1(v_1)e_2(v_2)\cdots e_n(v_n)$ where the subscripts are considered mod n and $(v_0) = (v_n)$ denotes an open vertex interval in which we start and end.

If at any point in the cyclic ordering (v_i) and (v_{i+1}) or e_i and e_{i+1} refer to the same open set in X, then f can be replaced with a homotopic map f' that is covered by strictly fewer open sets. This is because $(v) \cup e$ is contractible for any overlapping (v) and e. Continuing in this way either produces a null-homotopy or it stops at a map that can easily be locally smoothed out to an immersion. With minor modifications the same proof applies to arcs. □

Proposition 1.1.2 can be used to show that every group is the fundamental group of a 2-complex all of whose attaching maps are non-trivial immersions. We call such a 2-complex a *taut 2-complex* since its attaching maps have been pulled as tight as possible.

COROLLARY 1.1.3 (Taut 2-complexes). *Every 2-complex has a subcomplex, with the same fundamental group, that is homotopy equivalent to a taut 2-complex without altering its 1-skeleton.*

PROOF. Let X be an arbitrary 2-complex and let X' be the subcomplex of X obtained by removing all 2-cells whose attaching maps are null-homotopic in the 1-skeleton of X. Van Kampen's theorem shows their removal does not change the fundamental group. After replacing each remaining attaching map with an immersion homotopic to it (Proposition 1.1.2), the result is a taut 2-complex homotopy equivalent to X' (Theorem A.4.6). □

Since every taut 2-complex is combinatorial, Corollary 1.1.3 shows that G_4 is a subset of G_5, completing the proof of Theorem 1.1.1. There is a similar set of equivalences for fundamental groups of *compact* cell complexes. Since this class of groups contains exactly the fundamental groups of compact manifolds (including those with non-empty boundary), we call these *compact manifold groups* for now, even though they are better known as *finitely presented groups*. The equivalence will be clear once group presentations are discussed in §1.3.

Theorem 1.1.4 (Compact manifold groups). *The following classes are equal:*

$$\begin{aligned}
C_0 &= \{\pi_1 \textit{ of compact manifolds}\},\\
C_1 &= \{\pi_1 \textit{ of compact cell complexes}\},\\
C_2 &= \{\pi_1 \textit{ of compact combinatorial cell complexes}\},\\
C_3 &= \{\pi_1 \textit{ of finite simplicial complexes}\},\\
C_4 &= \{\pi_1 \textit{ of compact } 2\textit{-complexes}\},\\
C_5 &= \{\pi_1 \textit{ of compact combinatorial } 2\textit{-complexes}\}, \textit{ and}\\
C_6 &= \{\pi_1 \textit{ of finite simplicial } 2\textit{-complexes}\}.
\end{aligned}$$

Since taking barycentric subdivisions, passing to subcomplexes, and modifying attaching maps preserve compactness and finiteness, the equivalence of C_1 through C_6 follows immediately from Theorem 1.1.1. To complete the proof it suffices to show $C_0 \subset C_1$ and $C_2 \subset C_0$. The former is a consequence of the fact that every compact manifold has the homotopy type of a compact cell complex. Because the techniques would lead us too far afield, we refer the interested reader to the elegant proof in the appendix of Hatcher's book [**16**] that uses Euclidean neighborhood retracts. The final inclusion can be derived from a combinatorial version of a Whitney-type embedding theorem.

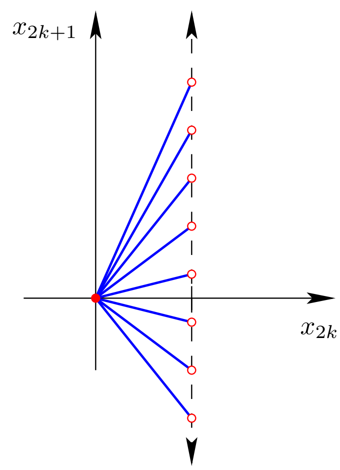

Figure 1. The projection of the embedded subdivided k-skeleton onto the coordinates x_{2k} and x_{2k+1}. The subdivided $(k-1)$-skeleton is sent to the origin and the distinct subdivided k-cells project to distinct line segments.

Theorem 1.1.5 (Embeddings). *If X is a simplicial n-complex with countably many cells, then its barycentric subdivision X' can be linearly embedded into $\mathbb{R}^{2n+1}$. As a consequence, every compact combinatorial n-complex is homotopy equivalent to a compact topological $(2n+1)$-manifold with boundary.*

Proof. Because the embedding $f : X' \to \mathbb{R}^{2n+1}$ we are constructing is supposed to be linear on each simplex of X', f is completely determined by the images of vertices. We send the vertices of X' that corresponds to the 0-cells of X to any discrete subset of points along the x_1-axis. This embeds the 0-skeleton of X into

$\mathbb{R}^1$. If an explicit map is desired, then one option would be to well-order the 0-cells of X and send the i-th 0-cell to the point on the x_1-axis with $x_1 = i$.

Next, suppose by induction that the barycentric subdivision of the $(k-1)$-skeleton of X has been embedded into the $\mathbb{R}^{2k-1}$ subspace of $\mathbb{R}^{2n+1}$ with $x_j = 0$ for all $j \geq 2k$. To extend this embedding to the subdivided k-skeleton, send the barycenters of the k-cells of X to any discrete subset of the line parallel to the x_{2k+1}-axis defined by the equations $x_{2k} = 1$ and $x_j = 0$ for all j not equal to $2k$ or $2k+1$ and then extend f linearly over the subdivided k-cells of X. By projecting onto the plane spanned by x_{2k} and x_{2k+1} (Figure 1) we see that the images of the subdivided (open) k-cells do not intersect each other or the $(k-1)$-skeleton.

The reader can verify that f is one-to-one on each subdivided k-cell and that this injection of the subdivided k-skeleton into $\mathbb{R}^{2n+1}$ is indeed a homeomorphism onto its image. The second assertion is now immediate since every compact combinatorial n-complex is homeomorphic to a finite simplicial complex and the closure of a sufficiently small ϵ-neighborhood of a finite simplicial complex linearly embedded into $\mathbb{R}^m$ is a topological m-manifold with boundary that deformation retracts back down to the original complex. □

1.2. Trees, graphs and free groups

We now shift our attention from compact complexes to those that are 1-dimensional. The main result is a classification of graphs up to homotopy and of their fundamental groups up to isomorphism. The remainder of the section is devoted to establishing the key properties that these 'free groups' possess.

1.2.1. Trees. The first step in classifying graphs up to homotopy is being able to recognize when a graph is contractible. Several equivalent conditions are recorded in Theorem 1.2.1. The graphs satisfying these conditions are called *trees.*

THEOREM 1.2.1 (Trees). *For a connected graph X, the following are equivalent:*

1. *X is contractible,*
2. *X is simply-connected,*
3. *X is minimally connected,*
4. *X does not contain an embedded circle, and*
5. *X does not contain a closed immersed path.*

Finally, for finite *connected graphs, a sixth equivalent condition is $\chi(X) = 1$.*

Most of these are self-explanatory, but condition 3 requires a definition. We call a graph X *minimally connected* if X is connected, but the removal of any 1-cell disconnects it. Theorem 1.2.1 is proved in stages. We begin by proving that the middle four conditions are equivalent.

LEMMA 1.2.2. *Let X be a connected graph. If X is not simply-connected then it contains a closed immersed path; if it contains a closed immersed path then it contains an embedded circle; if it contains an embedded circle, then it is not minimally connected; and if it is not minimally connected then it is not simply-connected. Thus conditions* 2 *through* 5 *in Theorem 1.2.1 are equivalent.*

PROOF. If X is not simply-connected and $f : [0,1] \to X$ is a closed path that represents a non-trivial element of $\pi_1(X)$, then by Proposition 1.1.2, f is homotopic to a closed immersed path. If g is a closed immersed path that is not an embedding of a circle then there is a proper subinterval whose endpoints are sent to the same vertex, and any minimal subinterval with this property can be used to construct an embedding of a circle into X. If e is an edge of an embedded circle in X, then $X \setminus \{e\}$ remains connected since any path connecting points u and v that uses e can be modified to use the remainder of the circle instead. Finally, if e is a 1-cell in X whose removal does not disconnect X, then we proceed as follows. Because, $X \setminus \{e\}$ is connected, the attaching map of e is homotopic to a constant map where both endpoints are sent to the same 0-cell v. Let X' be the complex where this altered attaching map is used for e, and note that X' is homotopy equivalent to X (Theorem A.4.6) and the union of v and e in X' is a subcomplex A homeomorphic to $\mathbb{S}^1$. Moreover, the map $r : X \to A$ that fixes A and sends every other cell to v is a continuous retraction, so the induced map r_* is surjective (Proposition A.2.4). Since $\pi_1(A) \cong \mathbb{Z}$, the group $\pi_1(X') \cong \pi_1(X)$ is non-trivial. □

For finite connected graphs conditions 3 and 6 are equivalent.

Lemma 1.2.3. *If X is a finite connected graph, then $\chi(X) \leq 1$. Moreover, $\chi(X) = 1$ if and only if X is minimally connected.*

Proof. Linearly order the 1-cells of X and attach them to the 0-skeleton one at a time. If c_i denotes the number of i-cells in X, then $c_1 \geq c_0 - 1$ because the 0-skeleton has c_0 connected components, the final result has one and attaching a 1-cell reduces the number of components by at most one. Thus, $\chi(X) \leq 1$ and $\chi(X) = 1$ if and only if each edge reduces the number of components. If $\chi(X) < 1$ then there is an edge e whose attachment does not reduce the number of components, and $X \setminus \{e\}$ remains connected since any path using e can be rewritten only using edges that occur earlier in the list. Conversely, if e is an edge whose removal does not disconnect X, then the edges can be ordered so that e occurs last. As shown above, at least $c_0 - 1$ edges were attached before e, so that $c_1 \geq c_0$ and $\chi(X) < 1$. □

The proof of Theorem 1.2.1 is nearly complete: since contractible graphs are simply-connected, it suffices to show that minimally connected graphs are contractible. For finite graphs this fact is easy to prove.[1] To prove this for arbitrary graphs, we introduce a combinatorial notion of distance.

Definition 1.2.4 (Combinatorial Distance). Let u and v be vertices in a cell complex X. The *length* of a combinatorial path from u to v is the number of 1-cells it traverses, and the *combinatorial distance* between u to v is the minimum length of a combinatorial path connecting them. Denote this value by $d_X(u, v)$ or simply $d(u, v)$ when X is implicitly understood and note that $d(u, u)$ is 0 since the constant path is considered a combinatorial path of length 0. When u and v lie in the same connected component of X, at least one such combinatorial path exists since any path from u to v is homotopic to a path in the 1-skeleton (Theorem A.1.5) that we can assume is an immersion (Proposition 1.1.2) and immersions are combinatorial. Thus, in a connected cell complex this distance d is defined for all pairs of vertices and it is easy to show that it defines a metric on the 0-skeleton of X (Exercise 3).

[1] Recall that the number of ends of edges attached to a vertex v is its *degree* and that a vertex of degree 1 is called a *free vertex*. If X contains a free vertex v, then there is a deformation retraction from X to $X \setminus \{v, e\}$ where e is the unique edge attached to v. A counting argument shows that every connected non-trivial graph with Euler characteristic 1 must have a free vertex and thus it deforms onto a proper subcomplex. Iterating this process contracts X to a point.

For later use we record the fact that minimum length paths are embedded.

Proposition 1.2.5 (Embedded Paths). *If u and v are distinct vertices in the same connected component of a cell complex, then every minimum length combinatorial path from u to v is embedded. In particular, at least one embedded combinatorial path from u to v exists.*

Proof. If a non-trivial combinatorial path is not embedded then it passes through the same vertex twice, and excising the subpath between these two occurences strictly shortens its length. □

The combinatorial distance function can be used to construct maximal contractible subgraphs of connected cell complexes better known as *spanning trees.*

Proposition 1.2.6 (Spanning trees). *Every connected cell complex contains a contractible subgraph with the same vertex set. As a consequence, every connected cell complex is homotopy equivalent to a cell complex with one vertex.*

Proof. Let v be a fixed vertex in a connected complex X. The *sphere of radius n around v* is the set of vertices u with $d(u,v) = n$ and the *ball of radius n around v* is the set of vertices u with $d(u,v) \leq n$. Denote these sets by S_n and B_n, respectively. Next, let X_n be the largest subgraph of $X^{(1)}$ with vertex set B_n. Since X is connected, the union of the graphs X_n is all of $X^{(1)}$. Inside the graphs X_n we inductively define subgraphs T_n. We start with $T_0 = X_0$ which is just v itself. The graph T_n is constructed from T_{n-1} by adding the vertices in S_n, and for each $u \in S_n$ adding a single edge connecting u to a vertex closer to v. The first edge of a path of length n connecting u to v shows that such an edge exists. Since there is an obvious deformation retraction from T_n to T_{n-1}, each T_n is contractible by induction. Finally, the subgraph $T = \bigcup_{n\in\mathbb{N}} T_n$ is a contractible subgraph (Proposition A.4.9) that contains every vertex of X. The second assertion is now immediate since the cell complex X/T has only one vertex and by Theorem A.4.5 it is homotopy equivalent to X. □

We now complete the proof of Theorem 1.2.1.

Lemma 1.2.7. *Minimally connected graphs are contractible.*

Proof. Let X be a minimally connected graph and let T be a contractible subgraph of X with the same vertex set (Proposition 1.2.6). If there is an edge e of X that is not in T, then the connected graph T is a subgraph of the disconnected graph $X \setminus \{e\}$ on the same vertex set, contradiction. Thus $X = T$ and X is contractible. □

The name *spanning tree* should now make sense. When X is a connected cell complex, a subgraph of X is contractible on the same vertex set iff it is a tree that spans the vertex set of X. A final characterization of trees is that they have unique embedded paths connecting distinct points.

Theorem 1.2.8 (Unique paths). *A connected graph is a tree iff there is a unique embedded path connecting every pair of distinct points.*

Proof. Let X be a connected graph. If X is not a tree then it contains an embedded circle and distinct points on this circle can be connected by distinct

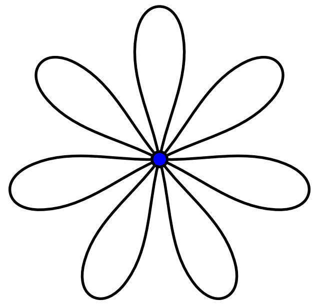

FIGURE 2. A rose with 7 edges.

embedded paths. Conversely, suppose X is tree and let x and y be distinct points of X. Since some of the equivalent conditions defining a tree are insensitive to cell structure, we may assume that x and y are 0-cells of X. At least one embedded path from x to y exists by Proposition 1.2.5. Because X is minimally connected, every edge traversed by this path would need to occur in every other path connecting x and y. Thus the only *embedded* interval containing these edges starting at x and ending at y is the one already considered, making it unique. □

1.2.2. Graphs. These chacterizations of trees quickly lead to a classification of connected graphs up to homotopy equivalence. Before establishing the classification, we note that connected graphs are classifying spaces and that the non-trivial elements of their fundamental groups are indexed by based immersed paths.

PROPOSITION 1.2.9 (Graphs as classifying spaces). *The universal cover of a connected graph is a tree. As a consequence, every connected graph is a classifying space and two connected graphs have the same homotopy type iff they have isomorphic fundamental groups.*

PROOF. The universal cover of a connected graph is both connected and simply-connected and thus a tree by Theorem 1.2.1. Since this implies it is contractible, the original graph is a classifying space for its fundamental group. The rest now follows from Theorem A.5.1. □

PROPOSITION 1.2.10 (Group elements and immersed paths). *For any connected graph X there is a natural bijection between the immersed paths in X based at x and the non-trivial elements of $G = \pi_1(X, x)$. In particular, every based immersed path in a graph represents a non-trivial element of its fundamental group.*

PROOF. Consider the function that sends each immersed path in X based at x to the element of $G = \pi_1(X, x)$ it represents. By Proposition 1.1.2 every non-trivial element of G is represented by some immersed path. On the other hand, no based immersed path represents the identity in G since it would lift to a closed immersed path in the universal cover contradicting the fact that $\widetilde{X}$ is a tree. And finally, if two distinct closed immersed paths represented the same non-trivial element $g \in G$, then they would lift to immersed paths in $\widetilde{X}$ starting at one vertex u and both ending at a different vertex v, contradicting Theorem 1.2.8. □

The simplest graphs are those with only one vertex (Figure 2). Such a graph is callled a *rose*, and its unique vertex is denoted $*$. Since the only variable in the construction of a rose is the number of edges it contains, we let R_A denote the rose whose edges are indexed by a set A and we let $\mathbb{F}_A$ denote the group $\pi_1(R_A, *)$. Since every connected graph is homotopy equivalent to a rose (Proposition 1.2.6), classifying connected graphs up to homotopy type is the same as classifying roses.

THEOREM 1.2.11 (Roses). *For sets A and B, the following are equivalent:*

1. *the sets A and B have the same cardinality,*
2. *the roses R_A and R_B are homeomorphic,*
3. *the roses R_A and R_B have the same homotopy type, and*
4. *the groups $\mathbb{F}_A$ and $\mathbb{F}_B$ are isomorphic.*

PROOF. Certainly $1 \Rightarrow 2 \Rightarrow 3$ and $3 \Leftrightarrow 4$ by Proposition 1.2.9, so we only need to prove 3 or 4 implies 1. The first observation is that compact roses and noncompact roses cannot be homotopy equivalent (Exercise 4). Thus there are two cases to consider: both R_A and R_B are compact or both R_A and R_B are noncompact. When R_A and R_B are homotopy equivalent and compact, the integers $\widetilde{\chi}(R_A) = -|A|$ and $\widetilde{\chi}(R_B) = -|B|$ must be equal by the homotopy invariance of Euler characteristics, implying $|A| = |B|$. Finally, when A is infinite, $|\mathbb{F}_A| = |A|$ (Exercise 5), so that R_A and R_B homotopy equivalent and noncompact implies $\mathbb{F}_A \cong \mathbb{F}_B$ which means $|A| = |\mathbb{F}_A| = |\mathbb{F}_B| = |B|$. □

1.2.3. Free groups. The fundamental group of a graph is called a *free group*. By Proposition 1.2.6 every free group is isomorphic to the fundamental group of a rose R_A and by Theorem 1.2.11 the cardinality of A is an invariant of the group that we call its *rank*. In fact, one way to restate Theorem 1.2.11 is that free groups are classified up to isomorphism by their rank.

COROLLARY 1.2.12 (Free groups classified). *Two free groups are isomorphic iff they have the same rank which is true iff they are fundamental groups of homotopy equivalent graphs.*

Let A be a set of cardinality κ. We use different notations for the free group of rank κ depending on the context. We continue to write $\mathbb{F}_A$ for the fundamental group of the rose R_A. On the other hand, we might write $\mathbb{F}_\kappa$ when we are only interested in the group up to isomorphism, or simply as $\mathbb{F}$ when we merely wish to indicate that the group is free. For example, if X is the 1-skeleton of a cube and x is one of its vertices, we say that $\pi_1(X, x) = \mathbb{F}_5$ since $|\widetilde{\chi}(X)| = |8 - 12 - 1| = 5$ is its rank. Several properties of free groups follow easily from their definition. The first one is known as the Nielsen-Schreier theorem.

THEOREM 1.2.13 (Free subgroups). *Subgroups of free groups are free.*

PROOF. Let G be a free group and let X be a graph with fundamental group G. Every subgroup $H \subset G$ is the fundamental group of a cover of X (Theorem A.3.10), but covers of graphs are graphs, so H is also a free group. □

THEOREM 1.2.14 (Free quotients). *Every group is a quotient of a free group. In particular, if (X, x) is a based connnected cell complex with $G = \pi_1(X, x)$ and $\mathbb{F}$ is the free group $\pi_1(X^{(1)}, x)$, then the group homomorphism $\mathbb{F} \to G$ induced by the inclusion map $X^{(1)} \hookrightarrow X$ is onto.*

PROOF. That the induced map is onto follows from the easy fact that paths starting and ending in the 1-skeleton of a cell complex are homotopic to paths entirely contained in the 1-skeleton keeping their endpoints fixed throughout. □

THEOREM 1.2.15 (Finite rank calculations). *If X is a finite connected graph, then the rank of $\pi_1(X)$ is $|\widetilde{\chi}(X)| = -\widetilde{\chi}(X)$. As a consequence, if G is a free group of rank k and H is an index d subgroup of rank l where k, d, and l are finite, then $l - 1 = d(k-1)$.*

PROOF. The first part was established during the proof of Theorem 1.2.11. Let X be a finite graph with $\pi_1(X) \cong G$ and let Y be the cover of X corresponding to H. Since Y is a d-fold cover of X, $\chi(Y) = d \cdot \chi(X)$. The fact that $\widetilde{\chi}(Y) = -l$ and $\widetilde{\chi}(X) = -k$ completes the proof. $\square$

Every free group is the fundamental group of a rose, and in these groups the elements generated by traveling along a single edge deserve special consideration.

DEFINITION 1.2.16 (Symmetric bases). Let R_A be a rose with $\mathbb{F}_A = \pi_1(R_A, *)$. Since each edge can be traversed in one of two directions, there are exactly $2|A|$ closed paths of length 1 in R_A and each path represents a distinct element in $\mathbb{F}_A$ (Proposition 1.2.10). The collection of these elements inside $\mathbb{F}_A$ is called the *symmetric basis for* $\mathbb{F}_A$ and denoted S_A.

Our use of the word 'symmetric' is one we wish to formalize. Definition 1.2.17 is not standard, but we find that including these definitions makes it easier to highlight certain aspects that would otherwise remain obscure.

DEFINITION 1.2.17 (Symmetric sets). A *symmetric set* is a set S with an implied involution $(\ )^{-1} : S \to S$, or alternatively, a set with an implied partition into blocks of size at most 2. The partition can be derived from the involution by recording the orbits of elements, and the involution can be recovered from the partition by sending each element to an element in the same block and to a distinct element whenever possible. If the involution is fixed-point free, or, equivalently, every block has size 2, we say the symmetric set is *free*. A *symmetric subset* is a subset T of a symmetric set S satisfying $T = T^{-1}$, and a *symmetry-preserving function* between symmetric sets is one that is compatible with their involutions: that is, a function $f : S \to T$ such that $t = f(s)$ implies $t^{-1} = f(s^{-1})$.

Groups are symmetric sets using the involution sending g to g^{-1} but they are never free since the identity is its own inverse. Paths in a cell complex form a symmetric set with an involution that reverses the parameterization. The symmetric basis S_A of the free group $\mathbb{F}_A$ is a free symmetric subset, and, in fact, it can be thought of as the canonical free symmetric set with blocks indexed by elements of A. Returning to Definition 1.2.16, notice that these basic paths and elements can be used to describe arbitrary paths and elements: every combinatorial path in R_A is a concatentation of these basic paths, and thus every element in $\mathbb{F}_A$ is a product of elements in S_A. The symmetric basis is easier to work with when its elements have been given explicit names. The tradition is to break symmetry by selecting a basis.

DEFINITION 1.2.18 (Bases and orientations). A subset of the symmetric basis S_A of the free group $\mathbb{F}_A$ is called a *basis* if it contains one element from each block of S_A. Topologically, selecting a basis is equivalent to *orienting* the edges of the rose R_A: the path of length 1 that crosses the edge e_a in the preferred direction represents the selected element and the path that travels in the opposite direction represents the unselected element. We call the selected element a and the other element a^{-1}. In this way the choice of a basis (or equivalently an orientation of the rose) lets us identify S_A with the set $A \cup A^{-1}$ where A denotes the selected elements and $A^{-1} = \{a^{-1} \mid a \in A\}$ collects the unselected elements. Although the free group $\mathbb{F}_A$ has many different bases inside S_A, its symmetric basis can be recovered from any one of these by simply adding in the inverses of the basis elements.

The most commonly used free groups with bases are those of finite rank and for these we introduce a simplified notation.

Remark 1.2.19 (Finite rank). When the size of A is very small we might write something like $\mathbb{F}_{\{a,b,c\}}$ to mean the free group $\mathbb{F}_A$ with basis $A = \{a, b, c\}$. More typically we use sets such as $A = \{a_i \mid i \in [n]\} = \{a_1, a_2, \ldots, a_n\}$ and we simplify the notation in this case, by writing e_i instead of e_{a_i} and $\mathbb{F}_{[n]}$ instead of $\mathbb{F}_{\{a_i \mid i \in [n]\}}$. Thus $\mathbb{F}_{[5]}$ denotes the free group of rank 5 with basis $A = \{a_1, a_2, a_3, a_4, a_5\}$.

When a free group arises as the fundamental group of a complicated graph, a basis can be selected by collapsing a spanning tree.

Proposition 1.2.20 (Selecting a basis). *Let T be a spanning tree in a based connected graph (X, x). Once the edges not in T are indexed by a set A, there is a natural isomorphism $\pi_1(X, x) \cong \mathbb{F}_A$. The elements in the symmetric basis S_A are represented by paths in X that cross over exactly one of the edges e_a concatenated, if necessary, with paths in T connecting the basepoint x with the endpoints of e_a.*

Proof. When T is collapsed to a point, the labeling identifies the quotient with the rose R_A. The homomorphism induced by the quotient map $q : X \to X/T$ is thus a map $q_* : \pi_1(X, x) \to \mathbb{F}_A$. Because trees do not contain closed immersed paths (Theorem 1.2.1), the image of a non-trivial immersed path based at x under the quotient map q is a path based at $*$ that remains non-trivial and immersed (Exercise 7). When combined with Proposition 1.2.10, this shows that q_* is injective. The surjectivity of q_* and the description of paths representing the symmetric basis elements follow from the statement and proof of the next proposition, a general result that we record for later use. □

Proposition 1.2.21 (Lifting paths). *Let $\{U_\alpha\}$ be a collection of pairwise disjoint connected subcomplexes of a connected complex U. If V is the cell complex obtained from U by collapsing each U_α to a point and $q : U \to V$ denotes the quotient map, then for every immersed path $f : I \to V^{(1)}$ there is an immersed path $g : I \to U^{(1)}$ such that $q(g)$ traces the same path as f. As a consequence, the induced map $q_* : \pi_1(U) \to \pi_1(V)$ is onto.*

Proof. Let u_α denote the point in V to which U_α collapses. Because the quotient map establishes a homeomorphism between $U \setminus \{U_\alpha\}$ and $V \setminus \{u_\alpha\}$, the portions of f that avoid the vertices u_α can be lifted to U. Moreover, we can extend these lifted portions to paths (i.e. to images of *closed* intervals) by including the vertices in the various U_α at which they start and/or end. The required path g is then patched together out of these lifted portions. See Figure 3 for an illustration. For each t in the interior of I where $f(t)$ is equal to one of the u_α, we insert, if necessary, an immersed path in $U_\alpha^{(1)}$ that connects the end point of the previous portion to the start point of next portion; if these two points are the same, we simply concatenate without inserting a path. Such connecting paths exist because each U_α is connected. Concatenating these paths produces a path $g : I \to U^{(1)}$

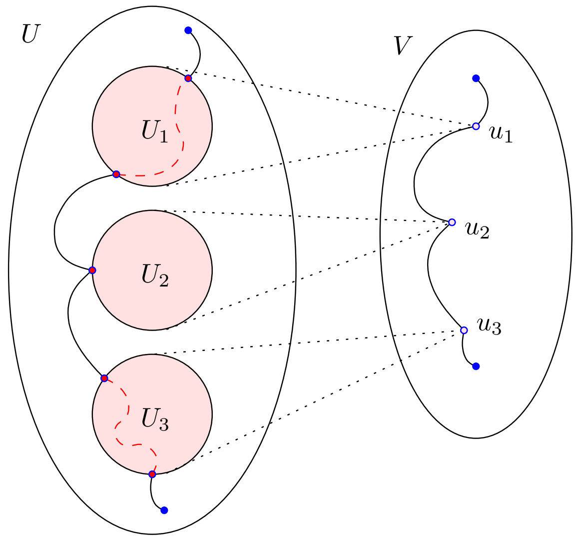

FIGURE 3. Lifting a path from V to U.

whose image under q traces out the original path in V, and the path g is immersed because (1) its image is immersed and (2) when non-trivial paths are inserted, at the transitions one edge lies in U_α and the other does not. To prove the final assertion, let u be a vertex not in any of the U_α and let v be its image in $V^{(1)}$. The induced map $q_* : \pi_1(U, u) \to \pi_1(V, v)$ is onto since every non-trivial element in $\pi_1(V, v)$ can be represented by an immersed loop $f : I \to V^{(1)}$ based at v and this loop is the image under q of a loop g based at u in U. A similar argument works when u is contained in one of the U_α, but paths inside U_α might need to be added at either end of the lifted path. $\square$

The reader should note that when one spanning tree in X is replaced with another, Proposition 1.2.20 produces a different isomorphism and a different symmetric basis is identified (Exercise 11). In fact, the situation is even more complicated. For any set A of cardinality $\kappa > 1$, the group $\mathbb{F} = \mathbb{F}_\kappa$ is isomorphic to $\mathbb{F}_A$ in an *infinite* number of distinct ways (Exercise 12) so that there are infinitely many distinct symmetric subsets of $\mathbb{F}$ that can play the role of its symmetric basis and a correspondingly infinite set of subsets that can be a basis for $\mathbb{F}$. We refer to any such subset or symmetric subset as a basis or symmetric basis for $\mathbb{F}$.

1.2.4. Alternative definitions. There are two alternative definitions of free groups that involve constructing them algebraically or defining them abstractly via their universal properties. We introduce both alternatives and prove they describe the same class of groups (Theorem 1.2.27). In order to distinguish among the different definitions, we refer to the free groups already defined as *topological* free groups. One major difference we should note is that both alternative definitions *require* the specification of a basis or symmetric basis.[2] We begin with the algebraic construction.

[2]This makes the topological version easier to apply in situations like the proof of the Nielsen-Schreier theorem. To prove Theorem 1.2.13 using one of the other definitions would have required

Definition 1.2.22 (Free groups; algebraic version). The *algebraic free group with symmetric basis* S_A is the group constructed as follows. Start with the free symmetric set S_A and consider finite sequences of elements from S_A. The collection of all such finite sequences (including the empty sequence) is denoted $(S_A)^*$. The elements of S_A are called *letters* and the finite sequences are called *words*. (More generally, for any set B we use B^* to denote the set of all 'words' built out of the 'letters' in B.) Equivalence classes are constructed based on the repeated insertion or deletion of subwords of the form aa^{-1} and the multiplication of two equivalence classes is the equivalence class of the concatenation of representatives. It is straightforward to show that this multiplication is well-defined and that the result is a group (Exercise 9). If a basis is chosen for S_A so that its elements are identified with the set $A \cup A^{-1}$, then the group we construct is the *algebraic free group with basis* A.

A non-empty word equivalent to the empty sequence is called a *Dyck word*, and one with no subwords of the form aa^{-1} is said to be *reduced*. Under the natural bijection between words in $(S_A)^*$ and combinatorial paths in the rose R_A, the reduced words correspond to the immersed paths. Thus, by Proposition 1.2.10, we can think of the reduced words in $(S_A)^*$ as parameterizing the non-trivial elements of $\mathbb{F}_A$. This bijection quickly leads to an isomorphism.

Proposition 1.2.23 (Algebraic free groups). *The algebraic free group with symmetric basis S_A is isomorphic to the fundamental group of the rose R_A. Thus, a group is free in the algebraic sense iff it is free in the topological sense.*

Proof. Let G be the algebraic free group with symmetric basis S_A and let $f : G \to \mathbb{F}_A$ be the natural homomorphism defined by identifying S_A with the symmetric basis of R_A and then interpreting the words in $(S_A)^*$ as combinatorial paths in the rose R_A that represent elements of $\mathbb{F}_A$. Since the insertion and deletion operations on words correspond to elementary homotopies on based loops, and concatenation of words corresponds to concatenation of based loops, the map f is a well-defined group homomorphism. Moreover, the canonical bijections between reduced words in $(S_A)^*$, immersed paths in R_A, and the non-trivial elements of $\mathbb{F}_A$ show that f is onto. Finally, suppose g is any non-trivial element of G. Start with any word representing g and iteratively remove subwords of the form aa^{-1}. This process must stop before it reaches the empty word since g is non-trivial. The word at which it stops is a reduced word representing g and this means that $f(g)$ is represented by a closed immersed path in R_A. By Proposition 1.2.10, $f(g)$ is a non-trivial element of $\mathbb{F}_A$, showing that f is one-to-one. $\square$

Notice that, as a consequence of our identifications, the set A^* can be viewed as a subset of the free group $\mathbb{F}_A$ with basis A since every non-empty word in A^* is automatically reduced. The non-trivial elements in A^* are called *positive words*. Our third and final definition of a free group focuses on their universal properties.

Definition 1.2.24 (Free groups; categorical version). A group G with a distinguished subset A is called a *categorical free group with basis* A if for any group H and for any function $f : A \to H$, there exists a unique extension of f to a group homomorphism $G \to H$. Similarly, a group G with a distinguished free symmetric

---

us to first find a potential basis or symmetric basis for the subgroup H and then to establish that it had the right algebraic or categorical properties.

subset S is called a *categorical free group with symmetric basis* S if for any group H and for any symmetry-preserving function $f : S \to H$, there exists a unique extension of f to a group homomorphism $G \to H$.

Categorical free groups are unique, in the appropriate sense, almost by definition, and topological free groups are used to show they exist.

PROPOSITION 1.2.25 (Uniqueness). *There is at most one categorical free group up to isomorphism for each size basis or symmetric basis. In particular, if G is a categorical free group with basis A, H is a categorical free group with basis B, and $f : A \to B$ is a bijection, then the unique homomorphism $G \to H$ extending f is an isomorphism. Similarly, if G is a categorical free group with symmetric basis S, H is a categorical free group with symmetric basis T, and $f : S \to T$ is a symmetry-preserving bijection, then the unique homomorphism $G \to H$ extending f is an isomorphism.*

PROOF. For simplicity we prove the basis version and leave the other as an exercise. Let $i : A \to G$ and $j : B \to H$ be the given inclusions. Applying the defining property of a categorical free group to the function i shows that the identity map on G is the unique homomorphism $G \to G$ fixing A pointwise. Similarly, the identity map on H is the unique homomorphism $H \to H$ fixing B pointwise. Applying the defining property to the composition $j \circ f$ shows that there is a unique homomorphism $g : G \to H$ that extends the bijection f. Similarly, using the composition $i \circ f^{-1}$ shows that there is a unique homomorphism $h : H \to G$ extending the bijection f^{-1}. Since $g \circ h$ is a homomorphism $G \to G$ fixing A pointwise, it must be the identity map on G and the composition $h \circ g$, being a homomorphism $H \to H$ fixing B pointwise, must be the identity map on H. Thus, g is injective and surjective and this unique homomorphism is an isomorphism. □

PROPOSITION 1.2.26 (Existence). *The topological free group $\mathbb{F}_A = \pi_1(R_A, *)$ with symmetric basis S_A is a categorical free group. In other words, for any group H and any symmetry-preserving function $f : S_A \to H$ there exists a unique group homomorphism $\mathbb{F}_A \to H$ extending f.*

PROOF. If X is any cell complex with $H = \pi_1(X, x)$ then we can define a map from R_A to X that sends $*$ to x and the oriented edge e_a in R_A to a based loop in X that represents the appropriate element in H. The induced homomorphism $\mathbb{F}_A \to H$ clearly extends f. It only remains to prove that this map is unique. Let $g, h : \mathbb{F}_A \to H$ be two homomorphisms that agree with f when restricted to S_A and consider the subset of $\mathbb{F}_A$ on which they agree. This set includes S_A and is closed under composition, but since every non-trivial element of $\mathbb{F}_A$ is represented by a word in $(S_A)^*$, g and h must agree on all of $\mathbb{F}_A$. □

The last three propositions taken together establish the following.

THEOREM 1.2.27 (Free groups). *A group is free in the topological sense iff it is free in the categorical sense iff it is free in the algebraic sense. Thus, the topological, algebraic, and categorical definitions define the same collection of groups.*

1.2.5. Maps and automorphisms. Bases are powerful tools that are particularly useful when describing homomorphisms from free groups to other groups.

Proposition 1.2.28 (Maps from free groups). *Let $\mathbb{F}_A = \pi_1(R_A, *)$ be a free group with basis A. For any based cell complex (X, x) with $G = \pi_1(X, x)$, the following collections are in natural bijection:*

1. *equivalence classes of based maps $(R_A, *) \to (X, x)$,*
2. *group homomorphisms $\mathbb{F}_A \to G$,*
3. *symmetry-preserving functions $S_A \to G$, and*
4. *functions $A \to G$.*

Proof. There are easy conversions among the collections that we call restrict, construct, and induce. Any group homomorphism $\mathbb{F}_A \to G$ can be restricted to its symmetric basis S_A and then further restricted to its basis A. Given any function $A \to G$, a based map $(R_A, *) \to (X, x)$ can be constructed by sending $*$ to x and each oriented edge e_a to a loop in X based at x that represents $f(a)$. Different choices for the image of e_a are equivalent up to basepoint preserving homotopy, so the map is well-defined up to equivalence. Finally, given a representative based map we can look at the induced homomorphism between their fundamental groups, which is well-defined since different representatives induce the same homomorphism. The consistency of the bijections connecting collections 2, 3, and 4 is an immediate consequence of the uniqueness part of the categorical definition of a free group, and the consistency of the bijections between collections 1 and 4 is clear: elements in G are sent to based loops in X that represent them and based loops in X are sent back to the elements in G they represent. $\square$

As a corollary, the automorphisms of a free group can be indexed by bijections between its various bases.

Corollary 1.2.29 (Automorphisms of free groups). *If $\mathbb{F}$ is a free group with basis $A \subset \mathbb{F}$ then a group endomorphism $f : \mathbb{F} \to \mathbb{F}$ is a group automorphism iff f restricted to A is a bijection between two bases for $\mathbb{F}$. Thus, for free groups of finite rank, the automorphisms in* $\text{Aut}(\mathbb{F})$ *can be indexed by the collection of ordered bases inside $\mathbb{F}$.*

Proof. If f restricted to A is a bijection between two bases for $\mathbb{F}$, then f is an isomorphism by Proposition 1.2.25. Conversely, suppose $f : \mathbb{F} \to \mathbb{F}$ is an automorphism and let $B = f(A)$. Because isomorphisms are injective, f restricted to A is a bijection and it is easy to show that B is a basis for $\mathbb{F}$ in the categorical sense (Exercise 13). For the final assertion, the orderings are an artifact used to implicitly describe the bijections between bases. First choose a standard basis A and linearly order it. A bijection $A \to B$ between bases can be used to induce a linear ordering of B and distinct bijections induce distinct orderings. Conversely, if B is any ordered basis of $\mathbb{F}$ we can reconstruct a bijection $A \to B$ by sending the first element of A to the first element of B, the second to the second, etc. $\square$

The final assertion of Corollary 1.2.29 immediately extends to free groups of infinite rank (Exercise 14) once we correct for the fact that distinct infinite ordinals (such as ω, $\omega + \omega$, and ω^ω) can have the same cardinality.

One unfortunate aspect of using a rose, or in fact any graph, as a model space for a free group $\mathbb{F}$ is that it does not treat all of its bases or symmetric bases, on an equal footing. Let $\mathbb{F} = \mathbb{F}_A$ be a non-abelian finite rank free group. By Proposition 1.2.28 every automorphism $\mathbb{F} \to \mathbb{F}$ can be represented by an equivalence class of based maps $(R_A, *) \to (R_A, *)$ but only finitely many of these classes contain a representative map that is a homeomorphism of the rose (Exercise 18). The images of A under these maps are the bases of $\mathbb{F}_A$ inside its symmetric basis S_A and this illustrates how they are qualitatively different from the others bases of $\mathbb{F}$. There are better models in higher dimensions as we briefly indicate. For further information on this topic see our discussion of free group automorphisms in the Epilogue.

Remark 1.2.30 (Other model spaces for free groups). Let A be a finite set with more than one element and pick a relatively nice embedding of R_A into $\mathbb{R}^2$ or $\mathbb{R}^3$. If we replace R_A with a regular neighborhood of its image, then the resulting manifold with boundary is homotopy equivalent to the original rose. The result in $\mathbb{R}^2$ is called a *planar surface* and the manifold in $\mathbb{R}^3$ is known as a *handlebody*. If we call the resulting planar surface P and the handlebody H then by Proposition 1.2.28 every automorphism $\mathbb{F}_A \to \mathbb{F}_A$ corresponds to an equivalence class of based maps $(R_A, *) \to (P, p)$ or $(R_A, *) \to (H, h)$, respectively. In the planar surface case infinitely many but not all automorphisms can be represented by maps that embed R_A into P (Exercise 19), and for handlebodies, every automorphism can be represented by a map that embeds R_A into H (Exercise 20). In fact, each of these embeddings can be chosen so that there is a deformation retraction from P or H onto the image of R_A. As a result, there is a precise sense in which the planar surface model treats infinitely many but not all bases on an equal footing and the handlebody model treats all bases equally.

1.3. Complexes and presentations

The geometric group theorist Martin Bridson began his address to the 2006 International Congress of Mathematicians as follows: "When viewed through the eyes of a topologist, a finite group-presentation $\Gamma = \langle \mathcal{A} | \mathcal{R} \rangle$ is a concise description of a compact, connected 2-dimensional CW-complex K with one vertex: the generators $a \in \mathcal{A}$ index the (oriented) 1-cells and the defining relations $r \in \mathcal{R}$ describe the loops along which the boundaries of the 2-cells are attached. Γ emerges as the group of deck transformations of the universal cover $\widetilde{K}$ and the Cayley graph $\mathcal{C}_{\mathcal{A}}(\Gamma)$ is the 1-skeleton of $\widetilde{K}$." This succinctly summarizes the ideas discussed in this section. The traditional algebraic machinery of presentations with generators and relations is introduced, but with an emphasis on the topological structures to which these concepts correspond.

1.3.1. Generating sets. A generating set for a group G is, essentially, a surjection onto G from a free group with a specified basis. They arise whenever G is viewed as the fundamental group of cell complex and there is a certain rough equivalence between generating sets for G, cell complexes with fundamental group G, and G-actions on graphs (Theorem 1.3.9). The first claim is easy to illustrate. Let G be the fundamental group of a connected cell complex X. If T is a spanning tree for $X^{(1)}$ and the edges not in T are indexed by A, then by Theorem 1.2.14 and Proposition 1.2.21 the inclusion map $X^{(1)} \hookrightarrow X$ induces a surjection $f : \mathbb{F}_A \twoheadrightarrow G$.

Definition 1.3.1 (Generating sets). Let $\mathbb{F}_A$ be a free group with basis A and recall that A^* can be viewed as a subset of $\mathbb{F}_A$. If $f : \mathbb{F}_A \to G$ is onto, then the function f (or, equivalently, its restriction $A \to G$) is said to *generate* G since we can generate every element of G from the image of $S_A = A \cup A^{-1}$. With the typical abuse of notation, the map to G often goes unmentioned. We say instead that A *generates* G and is a *generating set*. Similarly, S_A *symmetrically generates* G and is a *symmetric generating set*. And when f restricted to A^* remains onto, A *generates* G *as a monoid* and is a *monoid generating set*.

These conditions have several easily established reformulations (Exercise 23).

PROPOSITION 1.3.2 (Detecting generating sets). *If $\mathbb{F}_A$ is a free group with basis A and $f : \mathbb{F}_A \to G$ is a map with $B = f(A)$ and $T = B \cup B^{-1} = f(S_A)$, then the map f is onto iff no proper subgroup contains B iff no proper submonoid contains T which is true iff every element of G is represented by some word in $T^* = (B \cup B^{-1})^*$. Similarly, f restricted to A^* remains onto iff no proper submonoid contains B which is true iff every element of G is represented by some word in B^*.*

REMARK 1.3.3 (Generating sets as subsets). If $A \to G$ generates G and B is the image of A, then the inclusion $B \hookrightarrow G$ also generates G: the corresponding homomorphism $\mathbb{F}_B \to G$ must be onto since no proper subgroup of G contains B. Thus, in principle at least, generating sets for G can be replaced with generating subsets of G and it is tempting to make this assumption part of the definition. We do not do so precisely because closed paths in complicated spaces can be unexpectedly null-homotopic, making it difficult to determine whether the map $\mathbb{F} \to G$ derived from the inclusion $X^{(1)} \hookrightarrow X$ is injective on a basis of the free group $\mathbb{F}$.

The different types of generating sets can be illustrated using the integers.

EXAMPLE 1.3.4 (Generating sets for $\mathbb{Z}$). Consider the subsets $A = \{1\}$, $B = \{-1\}$, $C = \{2, 3\}$, and $D = \{-2, -3\}$ in $\mathbb{Z}$. Each of the four is a generating set for $\mathbb{Z}$. None of the four is a symmetric generating set or a monoid generating set. The combinations $A \cup B$ and $C \cup D$ symmetrically generate $\mathbb{Z}$, and a set such as $A \cup D$ is a monoid generating set that is not symmetric.

A group with a finite generating set is said to be *finitely generated*, and it should be clear from our earlier construction that the fundamental group of a cell complex with a finite 1-skeleton is an example of such a group. We have highlighted how cell complexes create generating sets for their fundamental groups; this process can also be reversed.

LEMMA 1.3.5 (Complexes from generating sets). *For each map $f : \mathbb{F}_A \twoheadrightarrow G$ there is connected cell complex X with R_A as its 1-skeleton, G as its fundamental group, and with f as the homomorphism induced by the inclusion $X^{(1)} \hookrightarrow X$. As a consequence, every finitely generated group is the fundamental group of a cell complex with a finite 1-skeleton.*

PROOF. Let K be the kernel of f and construct X as follows. Start with the rose R_A and for each non-trivial $k \in K$ attach a 2-cell to R_A. The element k, being a non-trivial element of $\mathbb{F}_A$, corresponds to a reduced word in $(S_A)^*$, and thus to a closed immersed path in R_A. This is the loop we use as the attaching map for the 2-cell indexed by k. Let $g : \mathbb{F}_A \twoheadrightarrow \pi_1(X, *)$ be the group homomorphism induced by the inclusion $(R_A, *) \hookrightarrow (X, *)$. By van Kampen's theorem (Theorem A.2.8) the kernel of g is the normal subgroup generated by K, but since K is already normal, the kernel of g is the kernel of f, and consequently $\pi_1(X, *) \cong G$. □

Generating sets also arise from free actions on cell complexes and we prove two different versions of this result. The first is more general and produces a generating set simply from a fundamental domain for the action. The second gives an alternative proof using the theory of covering spaces in the special case where the complex is a graph and the quotient by the group action is a rose.

LEMMA 1.3.6 (Generating sets from actions, I). *If a group G acts on a connected cell complex X and $\mathcal{F}$ is both a subcomplex and a fundamental domain for the action, then the set $\{g \in G \mid \mathcal{F} \cap (g \cdot \mathcal{F}) \neq \emptyset\}$ is a symmetric monoid generating set for G. More specifically, when the action of G on X is proper and cocompact, G is finitely generated.*

PROOF. Let $S = \{g \in G \mid \mathcal{F} \cap (g \cdot \mathcal{F}) \neq \emptyset\}$ and let x be a vertex in the subcomplex $\mathcal{F}$. Because X is connected, for each non-trivial $g \in G$ there is an immersed path in the 1-skeleton of X connecting x and $g \cdot x$. Using this path we can find a finite sequence $\{(g_0 \cdot \mathcal{F}), (g_1 \cdot \mathcal{F}), (g_2 \cdot \mathcal{F}), \ldots, (g_n \cdot \mathcal{F})\}$ of translates of $\mathcal{F}$, such that g_0 is the identity of G, g_n is g, and for each $i \in [n]$, $(g_{i-1} \cdot \mathcal{F}) \cap (g_i \cdot \mathcal{F})$ is non-empty. Because of the G-action, for each $i \in [n]$, $\mathcal{F} \cap (a_i \cdot \mathcal{F})$ is non-empty where $a_i = (g_{i-1})^{-1} g_i$. Thus each a_i is in S and the factorization $g = a_1 a_2 \cdots a_n$ shows that g is represented by a word in S^*. Since g was arbitrary, S is a monoid generating set for G. The G-action also shows that S is symmetric since $\mathcal{F} \cap (a \cdot \mathcal{F})$ is non-empty iff $(a^{-1} \cdot \mathcal{F}) \cap \mathcal{F}$ is non-empty. For the second assertion, note that the implicit assumption that the action of G on X is cellular means that S is precisely the subset of G that sends some vertex of $\mathcal{F}$ to another vertex in $\mathcal{F}$. Because the action is cocompact, our fundamental domain is a subcomplex with a finite set of vertices, and because the action is proper, for any pair of vertices u and v there are only a finite number of group elements that send u to v. The set S is thus a finite union of finite sets. □

LEMMA 1.3.7 (Generating sets from actions, II). *When a group G acts freely on a connected graph Γ with the rose R_A as its quotient, this induces a map $\mathbb{F}_A \twoheadrightarrow G$.*

PROOF. Because the action of G on Γ is free, the quotient map $p : \Gamma \to R_A$ is a covering map and the group G can be identified as the group of deck transformations of p. Moreover, the transitivity of the action on vertices means that Γ is a regular cover of R_A. If we pick a vertex $v \in \Gamma$ and define $K = \pi_1(\Gamma, v)$, then, by the theory of covering spaces, the induced homomorphism p_* embeds K as a normal subgroup of $\mathbb{F}_A = \pi_1(R_A, *)$. The quotient of $\mathbb{F}_A$ by $p_*(K)$ is isomorphic to G (Proposition A.3.9) and the required map is $\mathbb{F}_A \twoheadrightarrow \mathbb{F}_A / p_*(K) \cong G$. □

Using these lemmas, it is easy to establish that several natural topological conditions are equivalent to being finitely generated.

THEOREM 1.3.8 (Finitely generated). *For each G the following are equivalent:*

1. *G is finitely generated;*
2. *G is the fundamental group of a cell complex with a finite 1-skeleton.*
3. *G acts freely and cocompactly on a connected graph;*
4. *G acts properly and cocompactly on a connected cell complex;*

PROOF. Lemmas 1.3.6 and 1.3.5 show $4 \Rightarrow 1 \Rightarrow 2$. If $G \cong \pi_1(X, x)$ with $X^{(1)}$ finite, then the action of G on the 1-skeleton of $\widetilde{X}$ is both free and cocompact, so $2 \Rightarrow 3$. Finally, free actions are proper and graphs are cell complexes, so $3 \Rightarrow 4$. □

We can also clarify when a set A can generate a group G.

Theorem 1.3.9 (Generating sets). *For each set A and group G consider the following three collections.*

1. *homomorphisms $\mathbb{F}_A \twoheadrightarrow G$;*
2. *cell complexes X with $\pi_1(X, *) = G$ and $X^{(1)} = R_A$;*
3. *free G-actions on connected graphs Γ with quotient R_A.*

A homomorphism in collection 1 converts to a cell complex in collection 2, which converts to a action on a graph in collection 3 which converts to a homomorphism in collection 1. Thus, one collection is non-empty iff they are all non-empty. In fact, up to the appropriate notions of equivalence, these three collections are in natural bijection.

Proof. The conversions from 3 to 1 and from 1 to 2 use Lemma 1.3.7 and Lemma 1.3.5, respectively, and to convert from 2 to 3 let G act on the 1-skeleton of the universal cover of X. We leave the proof that they are in natural bijection up to equivalence as an exercise. For the record, the appropriate notions of equivalence are as follows:

1. Two functions $f_A : \mathbb{F}_A \twoheadrightarrow G$ and $f_B : \mathbb{F}_B \twoheadrightarrow G$ that generate G are considered equivalent if there is a symmetry-preserving bijection $S_A \to S_B$ between their symmetric bases that extends to an isomorphism $i : \mathbb{F}_A \to \mathbb{F}_B$ with $f_B \circ i = f_A$.

2. Let X_A and X_B be cell complexes with 1-skeletons R_A and R_B, $\pi_1(X_A, *) = G = \pi_1(X_B, *)$, and induced homomorphisms $f_A : \mathbb{F}_A \to G$ and $f_B : \mathbb{F}_B \to G$, respectively. These complexes are considered (very roughly) equivalent if there is a homeomorphism $R_A \to R_B$ between their 1-skeletons that induces an isomorphism $i : \mathbb{F}_A \to \mathbb{F}_B$ with $f_B \circ i = f_A$.

3. Let G act freely on two graphs Γ and Γ'. These are considered equivalent if there is an isomorphism of the underlying graphs that is compatible with the G-action. In other words, there is an isomorphism $f : \Gamma \to \Gamma'$ such that for all $g \in G$ and for every cell $\sigma \in \Gamma$, $f(g \cdot \sigma) = g \cdot f(\sigma)$. □

1.3.2. Cayley graphs. The graphs with group actions that occur in the statement of Theorem 1.3.9 are called Cayley graphs.

Definition 1.3.10 (Cayley graphs). Let G be a group. A connected graph Γ with a free and vertex-transitive G-action is called a *Cayley graph for G*. The quotient of Γ by the action of G is a rose. If we index its edges by a set A, then by Theorem 1.3.9 there is a corresponding map $f : \mathbb{F}_A \twoheadrightarrow G$ that generates G, and Γ is called the *Cayley graph for G with respect to f*. Finally, if an orientation is added to the edges of R_A, then this defines a basis A for the free group $\mathbb{F}_A$ and we call Γ the *Cayley graph for G with respect to A*.

According to Theorem 1.3.9 a Cayley graph for G can be constructed from a map $\mathbb{F}_A \twoheadrightarrow G$ or from a one vertex cell complex with G as its fundamental group.

Example 1.3.11 (Free groups). Since the rose R_A is a one vertex complex with fundamental group $\mathbb{F}_A$, its universal cover is a Cayley graph for the free group $\mathbb{F}_A$. The universal cover is, of course, a tree and each vertex has valence $2|A|$. A portion of a Cayley graph for $\mathbb{F}_3 \cong \mathbb{F}_{\{a,b,c\}}$ is sketched on the lefthand side of Figure 4. The drawing conventions are as follows: the edges with negative slope represent a (when moving up and to the left) or a^{-1} (when moving down and to the right). Those with positive slope similarly represent b or b^{-1} and the vertical edges represent c

or c^{-1}. The length of each edge has been scaled according to its distance from the base vertex, to enable the graph to be drawn without self-intersections.

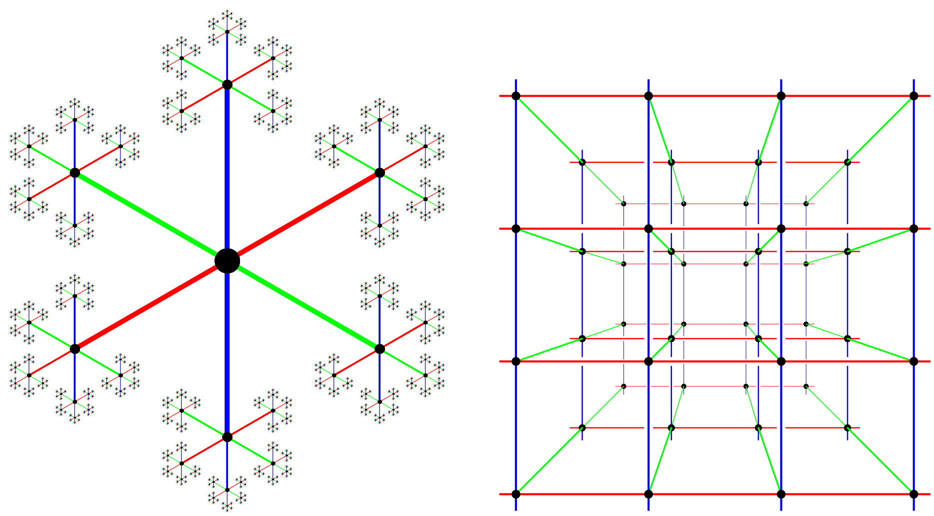

FIGURE 4. Portions of Cayley graphs for $\mathbb{F}_3$ (left) and $\mathbb{Z}^3$ (right).

EXAMPLE 1.3.12 (Free abelian groups). The vector space $\mathbb{R}^n$ can be given a cell structure where it looks like unit n-dimensional cubes stacked up in all directions with a vertex at every point with integer coordinates and edges in the coordinate directions connecting points distance 1 apart. Moreover, there is a free cellular $\mathbb{Z}^n$-action on $\mathbb{R}^n$ where the action is by rigid translation. The action is transitive on vertices so the quotient of $\mathbb{R}^n$ by action of $\mathbb{Z}^n$ is a one vertex complex called an *n-torus* and denoted T^n. The cell complex T^n has $\binom{n}{i}$ i-cells for each $i \in \{0, 1, 2 \ldots, n\}$. Since the space $\mathbb{R}^n$ is simply-connected and the $\mathbb{Z}^n$-action is free, (1) $\pi_1(T^n, *) = \mathbb{Z}^n$, (2) the quotient map is a covering projection, (3) $\mathbb{R}^n$ is the universal cover of T^n and (4) its 1-skeleton is a Cayley graph for $\mathbb{Z}^n$. A portion of this Cayley graph for $\mathbb{Z}^3$ is sketched on the righthand side of Figure 4.

Given a map $f : \mathbb{F}_A \twoheadrightarrow G$ there is a direct construction of the corresponding graph Γ with a free G-action that eliminates the need for the intermediate complex X, but to make the construction precise we need the notion of a symmetric edge labeling of a graph.

DEFINITION 1.3.13 (Edge labelings). In any graph Γ, the edges of Γ form a set and the oriented edges of Γ (i.e. the combinatorial paths of length 1) form a symmetric set using the involution that reverses orientation. An *edge labeling of* Γ *by* A is a bijection between the set A and the edges of Γ. Similarly, a *symmetric edge labeling of* Γ *by* S is a symmetry-preserving bijection between the symmetric set S and the oriented edges of Γ.

DEFINITION 1.3.14 (Labeling the 1-skeleton of $\widetilde{X}$). Let X be a cell complex with $X^{(1)} = R_A$ and $\pi_1(X, *) = G$, let $\widetilde{X}$ be its universal cover, and let $\tilde{x}$ be one of its vertices. The group G acts on $\widetilde{X}$ as the group of deck transformations of the projection map $p : \widetilde{X} \to X$ and this action can be used to induce a vertex labeling and a symmetric edge labeling on the 1-skeleton of $\widetilde{X}$. Since the action of G on $\widetilde{X}$ is free and vertex-transitive, for each vertex v in $\widetilde{X}$ there exists a unique element $g \in G$ with $v = g \cdot \tilde{x}$. We call this vertex v_g and the action of G on the vertex set is described by $g \cdot v_h = v_{gh}$. Similarly, every path of length 1 in $\widetilde{X}$ is uniquely determined by its starting point and the path of length 1 in R_A to which it projects. Since the vertices of $\widetilde{X}$ are indexed by G and the paths of length 1 in R_A are indexed by the symmetric basis S_A, the paths of length 1 in $\widetilde{X}$ are indexed by $G \times S_A$. If we use $[a]$ to denote the image of a under the induced map $f : \mathbb{F}_A \to G$, then the oriented edge $e_{(1,a)}$, almost by definition, starts at $\tilde{x} = v_1$ and ends at $v_{[a]}$. More generally, using the G-action, the oriented edge $e_{(g,a)}$ starts at v_g and ends at $v_{g\cdot[a]}$ and the G-action on oriented edges is described by $g \cdot e_{(h,a)} = e_{(gh,a)}$. From this we can see that the appropriate involution to define on $G \times S_A$ to turn it into a symmetric set is the one sending (g, a) to $(g \cdot [a], a^{-1})$.

The key observation is that the vertex labeling and the symmetric edge labeling describe the structure of the graph $\widetilde{X}^{(1)}$ in a way that only depends on the map $f : \mathbb{F}_A \to G$ and the multiplication in G. In particular, it and its G-action can be completely reconstructed with no mention of cell complex X.

Definition 1.3.15 (Cayley graphs from generating sets). The previous discussion shows that if $A \to G$ generates G and $[a]$ denotes the image of $a \in A$ under this map, then the corresponding Cayley graph of G with respect to A can be constructed as follows. Start with a vertex v_g for each $g \in G$, then add an edge connecting v_g to $v_{g\cdot[a]}$ for each $(g, a) \in G \times A$, and call the resulting graph Γ. To recover the symmetric labeling of the oriented edges of Γ, we let $e_{(g,a)}$ label the path of length 1 that travels along the edge indexed by (g, a) from v_g to $v_{g\cdot[a]}$ and label the same edge with the opposite orientation by $e_{(g\cdot[a],a^{-1})}$. This gives a symmetric edge labeling of Γ by $G \times S_A$, where the latter is a symmetric set under the involution with $(g, a)^{-1} = (g \cdot [a], a^{-1})$. Finally, there is a natural (left) action of G on this graph Γ that is defined on vertices and edges by the equations $g \cdot v_h = v_{gh}$ and $g \cdot e_{(h,a)} = e_{(gh,a)}$.

Remark 1.3.16 (Left and right). The Cayley graph constructed above is sometimes called the *right* Cayley graph of G with respect to $f : \mathbb{F}_A \to G$ since the edges record what happens when you right multply by $[a] \in G$. The switch between a right multiplication $(\cdot[a])$ that defines the edges and a left multiplication $(g\cdot)$ that defines the G-action is crucial for their compatibility. One could define a *left* Cayley graph for G generated by A, but it would only have a natural *right* G-action.

Remark 1.3.17 (Covers, Cayley graphs and groups). Let X be a one vertex cell complex with $\pi_1(X, *) = G$. Because X is a cell complex, we know that a universal cover $\widetilde{X}$ exists, but that does not mean that we know how to construct it. The main difficulty is being able to construct the 1-skeleton of $\widetilde{X}$, i.e. the Cayley graph of G with respect to the generating set that arises from the 1-skeleton of X. In a way that we make precise in Chapter 3, arbitrarily large portions of $\widetilde{X}$ can be constructed iff arbitrarily large portions of its Cayley graph can be constructed, which can be done iff we truly understand how to multiply elements inside its fundamental group. In particular, whenever we know something about the structure of $\widetilde{X}$, it can usually be translated into algebraic information about the group G.

REMARK 1.3.18 (Encoding the group action). If $A \to G$ generates G and Γ is the Cayley graph for G with respect to A as constucted above, then a few simple decorations can be added to Γ that encode the group action. For example, it is sufficient to indicate which vertex v corresponds to the identity in G and to label each oriented edge $e_{(g,a)}$ by its second coordinate a. For each $a \in A \cup A^{-1}$, the action of $[a]$ on Γ is then the unique label-preserving motion that sends v to the other end of the unique oriented edge starting at v and labeled by a. The fact that A generates G means that the motions corresponding to the other group elements are compositions of these basic motions.

Before leaving the subject of group actions on graphs, we briefly indicate how close an arbitrary free action on a graph is to being a true Cayley graph.

DEFINITION 1.3.19 (Partial and non-standard Cayley graphs). If Γ is a graph with a free G-action, then Γ can be viewed as a Cayley graph that is *partial* and *non-standard*. The word 'partial' indicates that Γ need not be connected and 'non-standard' indicates that the G-action need not be vertex transitive. When a distinction needs to be drawn, ordinary Cayley graphs are said to be *full* and *standard*.

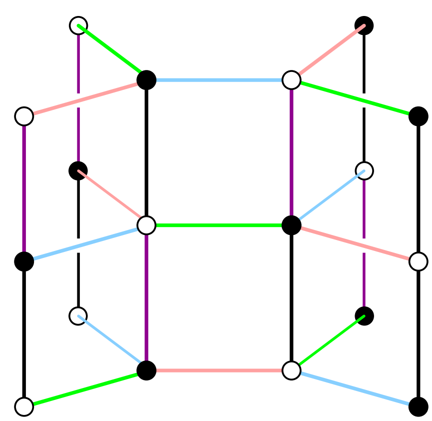

FIGURE 5. A portion of a non-standard Cayley graph for the fundamental group of the complement of the trefoil knot.

Many of the earlier results on Cayley graphs immediately extend to the partial and non-standard ones, and the places where they do not extend only serve to highlight how the additional assumptions were used. For example, every full non-standard Cayley graph corresponds to the 1-skeleton of the universal cover of a connected cell complex with more than one vertex. Thus, the 1-skeleton of the complex $\widetilde{\mathcal{D}}$ that we examined in the prologue is a non-standard Cayley graph for the fundamental group of the complement of the trefoil knot. We can convert a full non-standard Cayley graph into standard one by either contracting a spanning tree in the complex X before we construct its universal cover and restrict to the 1-skeleton, or, more directly, we can simply contract all preimages of this tree inside $\widetilde{X}^{(1)}$. This extra flexibility is particularly useful when the non-standard Cayley graph is easier to visualize, as in the case of the trefoil knot. Next, partial standard Cayley graphs arise when we consider arbitrary maps $\mathbb{F}_A \to G$ that need not be onto, and they can be converted to full Cayley graphs by enlarging the set A to a generating set. Finally, in complete generality, let X be a connected complex with $\pi_1(X,x) = G$ and covering projection $p : \widetilde{X} \to X$. The action of G on the preimage under p of a *portion* of the 1-skeleton of X is a partial non-standard Cayley graph, and, up to disjoint union, every partial non-standard Cayley graph (i.e. every free G-action on a graph) arises in this way.

1.3.3. Presentations. Group presentations, like generating sets and Cayley graphs, have a topological definition, a corresponding algebraic formalism, and a group action interpretation.

Definition 1.3.20 (Topological presentations). Let G be a group and let (X,x) be a based connected combinatorial 2-complex with $\pi_1(X,x) = G$. When X has only one vertex, it is called a (topological) *presentation of* G, and when X has only a finite number of cells (i.e. when X is compact), we say X is a *finite topological presentation of* G and G is *finitely presented by* X.

Note that Theorem 1.1.4 and Proposition 1.2.6 prove that the class of groups with finite topological presentations is the same as the class of compact manifold groups. Echoing the distinctions for Cayley graphs, we add the adjective *non-standard* when X has more than vertex. The Dehn complex of a knot diagram, for example, is a non-standard presentation for the fundamental group of the knot complement since it has, by construction, 2 distinct vertices.

Example 1.3.21 (Finite groups). If G is a finite group then G has a finite presentation. In particular, the complex described in Exercise 1 is a finite non-standard three vertex 2-complex with fundamental group G.

Before turning to the algebraic version, we note that the correspondence between the definitions is much closer when relators can be listed more than once. The notion of a multiset in introduced to make this precise.

Definition 1.3.22 (Multisets). Let S be a set. A *multiset selected from* S is a function $m : S \to \mathbb{N}$, where the value $m(s)$ indicates the number of times that $s \in S$ is selected. Intuitively, a multiset is a cross between a list and set: repetition is allowed but the ordering is irrelevant. Subsets of S corresponds to multisets with range in $\{0,1\} \subset \mathbb{N}$. In the other direction, every multiset $m : S \to \mathbb{N}$ has an associated subset formed by collecting together all elements of S selected at least once. This is equivalent to removing any redundancies. For more on the combinatorics of multisets, see [**28**].

Definition 1.3.23 (Algebraic presentations). Let A be a set and let $\mathcal{R}$ be a multiset selected from $(A \cup A^{-1})^*$. Since each r in $\mathcal{R}$ represents an element of $\mathbb{F}_A$, $\mathcal{R}$ implicitly describes a subset of $\mathbb{F}_A$. Let N be the smallest normal subgroup of $\mathbb{F}_A$ containing this subset and let G be the quotient group $\mathbb{F}_A/N$. The pair $\mathcal{P} = \langle A|\mathcal{R}\rangle$ is called an *algebraic presentation of* G, the elements of $\mathcal{R}$ are called *relators*, and $\mathcal{R}$ itself is a *set of defining relators*. The quotient map $\mathbb{F}_A \twoheadrightarrow G$ shows that A generates G. When both A and $\mathcal{R}$ are finite, we say that $\mathcal{P}$ is a *finite algebraic presentation of* G.

Converting from an algebraic to a topological presentation is straightforward.

Definition 1.3.24 (Relators to 2-cells). If $\mathcal{P} = \langle A|\mathcal{R}\rangle$ is an algebraic presentation of a group G, then we construct a 2-complex X starting with the oriented rose R_A and attaching one 2-cell to R_A for each $r \in \mathcal{R}$, attaching it along the closed combinatorial path in R_A to which the word r corresponds. As in the proof of Lemma 1.3.5, by van Kampen's theorem (Theorem A.2.8) the kernel of the map $g : \mathbb{F}_A \to \pi_1(X, *)$ is the normal subgroup N generated by the subset of $\mathbb{F}_A$ that $\mathcal{R}$ implicitly represents. Because $\langle A \mid \mathcal{R} \rangle$ is an algebraic presentation of G, $\pi_1(X, *) = \mathbb{F}_A/N = G$ and X is a topological presentation of G.

Remark 1.3.25 (Redundant 2-cells). When the multiset $\mathcal{R}$ is not a subset of $(A \cup A^{-1})^*$, the corresponding 2-complex has *redundant* 2-*cells* (distinct 2-cells attached along the same closed path). Redundant 2-cells can be removed without changing the fundamental group, but they are not easy to avoid completely since covers of complexes with no redundant 2-cells can have redundant 2-cells. The classic example is the complex for $\langle a | a^n \rangle$. Its single 2-cell is not redundant, but its universal cover has n distinct 2-cells attached to the same closed path.

Definition 1.3.26 (Useful conventions). When giving explicit examples it is convenient to use uppercase roman letters, such as 'A', 'B', 'C', to denote the inverse of their lowercase equivalents, 'a', 'b', 'c'. We write, for example, $abcABC$ instead of $abca^{-1}b^{-1}c^{-1}$ because the first form is significantly easier to parse and absorb. The fact that we use 'A' to denote both the alphabet of symbols and the inverse of $a \in A$ should not cause any problem since the context makes clear which is meant. A second convenient convention is to allow relators such as $abAB$ to be given implicitly via relations such as $ab = ba$. A *relation* is an equation of the form $r = s$ where r and s are words in $(A \cup A^{-1})^*$, and the implicit relator is the word rs^{-1}. The extra flexibility can be used to highlight aspects that would otherwise be opaque. In our example, the relation $ab = ba$ makes clear that (the group elements represented by) a and b commute. This is less clear from the relator $abAB$.

The conversion in the other direction is similarly straightforward.

Definition 1.3.27 (2-cells to relators). If X is a standard topological presentation of G, we can index and oriented its edges to identify $X^{(1)}$ with an oriented rose R_A. Next, for each 2-cell, its attaching map is a combinatorial map from a subdivided circle to R_A. If we pick an orientation of the circle and a preimage of the vertex $*$ as our basepoint, then the attaching map can described using the closed combinatorial path that starts at the lifted basepoint and travels around the circle in the chosen direction. This combinatorial path is associated with a word $r \in (A \cup A^{-1})^*$. If $\mathcal{R}$ collects the multiset of such words, one for each 2-cell of X, then $\mathcal{P} = \langle A | \mathcal{R} \rangle$ is an algebraic presentation of G.

It should be clear that these conversions are compatible in the following sense.

Proposition 1.3.28 (Presentations). *Every standard topological presentation of a group G can be can be converted into a algebraic presentation of G, from which the topological presentation can be recovered. Under these conversions, the number of* 1*-cells and* 2*-cells in the topological presentation correspond to $|A|$ and $|\mathcal{R}|$, respectively, in the algebraic presentation. In particular, G has a finite topological presentation iff it has a finite algebraic presentation.*

Algebraic presentations produce standard topological presentations with only one vertex. At the end of the chapter we introduce an alternative procedure that efficiently constructs a large and important class of non-standard complexes. We conclude our discussion of presentations with the observation that finitely presented groups can be characterized via their actions on cell complexes.

Theorem 1.3.29 (Presentations as actions). *For each group G, there is a natural bijection between connected cell complexes with fundamental group G and 1-connected cell complexes with free G-actions. Moreover, the complexes with one vertex correspond to the G-actions that are vertex-transitive and the complexes that are compact correspond to the actions that are cocompact. As a consequence, a group G has a finite presentation iff it acts freely and cocompactly on a 1-connected cell complex.*

Proof. If G is the fundamental group of a cell complex X then G acts freely on its 1-connected universal cover $Y = \widetilde{X}$. Conversely, if G acts freely on a 1-connected cell complex Y then by Proposition A.3.9 the quotient of Y by its G-action is a cell complex X with G as its fundamental group and Y as its universal cover. The remaining assertions are immediate. □

Remark 1.3.30 (Proper actions). There is a more expansive characterization of finitely presented groups as those groups capable of acting *properly* and cocompactly on a 1-connected cell complex. The easy direction is clear from Theorem 1.3.29 and in Chapter 7 we establish the more difficult implication.

The reformulation of a presentation as an action naturally leads to the notion of a Cayley complex. The name highlights the fact that Cayley complexes are to presentations as Cayley graphs are to generating sets.

Definition 1.3.31 (Cayley complexes). A *Cayley complex* for a group G is a 1-connected 2-complex Y with a free and vertex-transitive G-action. The 1-skeleton of a Cayley complex is a Cayley graph and they are created in similar ways. In particular, when X is a topological presentation of G, the universal cover of X with its natural free G-action is a Cayley complex for G.

Indexed to here

1.4. Cut points and free products

In this section we focus on a third topological feature of a space that impacts the structure of its fundamental group: the existence of a cut point. A cell complex with a cut point can be viewed as a collection of simpler pieces that have been wedged together and the goal of this section is to show that the fundamental group of such a wedge product is built out of the fundamental groups of its pieces in an understandable way.

1.4.1. Wedge products. The *wedge product* of a collection of based spaces is usually defined as the quotient of their disjoint union in which their base points are identified (§A.2), but the key results are easier to prove and easier to visualize when this standard construction is replaced with a non-standard variation.

Definition 1.4.1 (Non-Standard wedge products). The *non-standard wedge product* of based connected spaces (X_α, x_α) is a based space (Y, y) created by adding a new vertex y to the disjoint union of the X_α and adding a new edge e_α for each α that connects y to x_α. The non-standard wedge product of three copies of $\mathbb{R}P^2$ is schematically shown in Figure 6. The subcomplex formed by the x_α, e_α and the vertex y is a subtree of Y that we call its *backbone*.

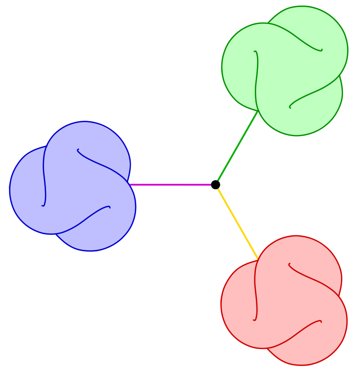

FIGURE 6. A non-standard wedge product of three projective planes

Collapsing the backbone of a non-standard wedge product to a point produces the standard wedge product. So long as each X_α is a cell complex, the two wedge products are homotopy equivalent (Theorem A.4.5), but this need not be the case in general (Exercise 24). Assume from here on that each X_α is a cell complex and let G denote the group $\pi_1(X,x) \cong \pi_1(Y,y)$. To better understand the group G, we build the universal cover of Y. A portion of the universal cover of the non-standard wedge product of three projective planes is shown in Figure 7. In this example it should be clear that when each 2-sphere is collapsed to a point, the result is a tree. In fact, the universal cover of a non-standard wedge product always collapses to a tree in exactly this fashion. The first step is to inductively construct the universal cover of Y.

Let Y_1 be a copy of the backbone of Y and let $p_1 : Y_1 \hookrightarrow Y$ be the natural inclusion map. The map p_1 is an immersion and it is a local homeomorphism except at the vertices x_α in Y_1. To remedy this we attach a copy of the universal cover $\widetilde{X}_\alpha$ (which exists since X_α is a cell complex) to each deficient vertex x_α in Y_1 and we extend p_1 using the covering maps $\widetilde{X}_\alpha \to X_\alpha$. Call the resulting space Y_2 and the map $p_2 : Y_2 \to Y$. The new map remains an immersion and it is a local homeomorphism except at preimages of x_α in the newly attached copies of the $\widetilde{X}_\alpha$ not already attached to a copy of the backbone. To remedy this we attach copies of the backbone to each of these vertices and extend the projection to Y in the obvious way. The result is a local immersion $p_3 : Y_3 \to Y$ that is now a local homeomorphism except at the preimages of x_α in the recently attached copies of the backbone that are not already attached to a copy of the appropriate $\widetilde{X}_\alpha$. At each such vertex we attach a copy of the appropriate $\widetilde{X}_\alpha$ and extend the projection to Y as before. Continuing in this way forever (alternately attaching copies of the backbone and copies of the universal covers of the cell complexes used to produce the wedge product) eventually constructs a space Y' and a map $p : Y' \to Y$ that is a local homeomorphism everywhere and thus a cover (Proposition A.3.6). Figure 7 shows the intermediate space Y_5 in this example; backbones have been attached to the spheres that were attached to the backbones attached to the spheres attached to the initial backbone. If final result Y' is simply-connected then it is the universal cover of Y (Proposition A.3.11). The next lemma shows that the intermediate stages, at least, are simply-connected.

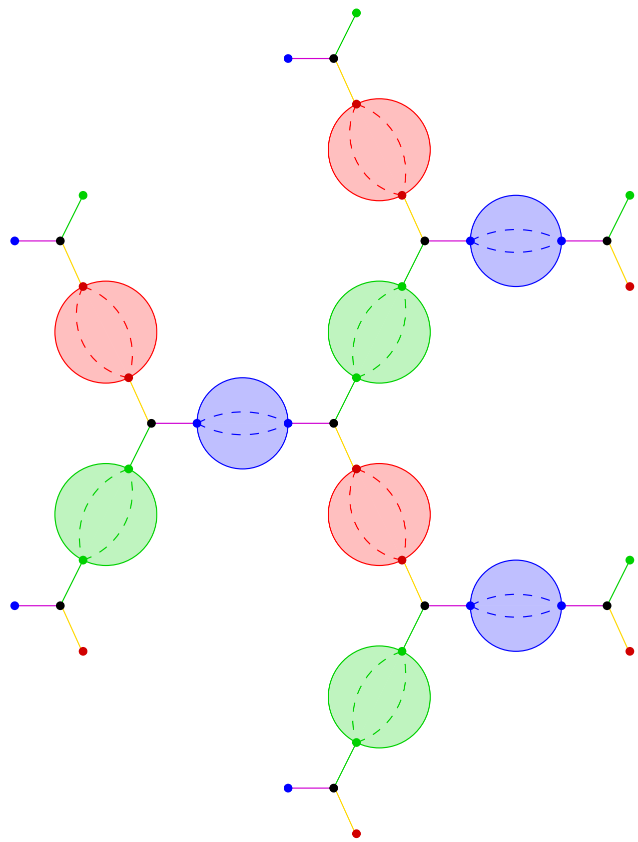

FIGURE 7. A portion of the universal cover of the non-standard wedge product of three projective planes.

PROPOSITION 1.4.2 (Attaching simply-connected spaces). *If V is a simply-connected cell complex, $\{(U_\alpha, u_\alpha)\}_{\alpha \in A}$ is a collection of based simply-connected cell complexes, and $\{v_\alpha\}_{\alpha \in A}$ is a collection of distinct points in V, then the space U, formed by attaching each U_α to V by identifying u_α with v_α, is simply-connected.*

PROOF. Pick a vertex $u \in V \subset U$ and consider an element $g \in \pi_1(U, u)$. It can be represented by an immersed loop $f : I \to U^{(1)}$ based at u. See Figure 8. Since $U_\alpha \cap V = u_\alpha$ and the various subcomplexes U_α are pairwise disjoint, the maximal subpaths of f in U_α start and end at u_α. But each $\pi_1(U_\alpha, u_\alpha)$ is trivial, so these subpaths are null-homotopic and can be excised without changing the fact that the loop represents g. After excising all of these subpaths, the result is a loop that remains in V. Because V is simply-connected, g is trivial, and U is simply-connected. $\square$

By Proposition 1.4.2 the intermediate stages in the construction are simply-connected, and thus by Proposition A.4.9 the end result is simply-connected. This means that Y' is indeed the universal cover of Y and, as a consequence, we now know the local structure of $\widetilde{Y}$.

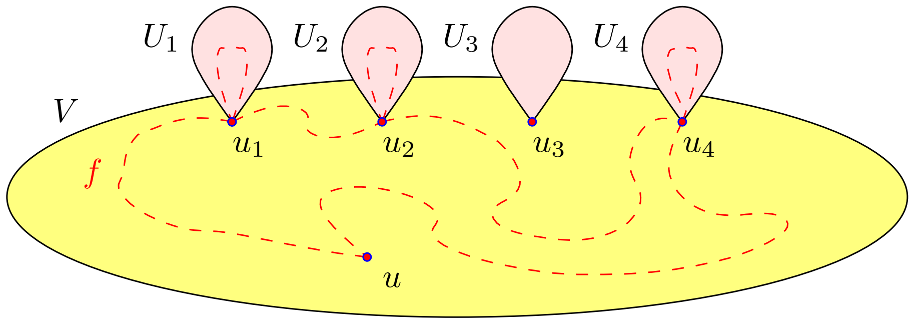

FIGURE 8. Attaching simply-connected spaces.

LEMMA 1.4.3 (Structure of $\widetilde{Y}$). *Let (Y, y) be the non-standard wedge product of a collection of based connected cell complexes (X_α, x_α) and let S be its backbone. If $p : \widetilde{Y} \to Y$ is its universal cover, then each component of $p^{-1}(X_\alpha)$ is a subcomplex homeomorphic to $\widetilde{X}_\alpha$ and each component of $p^{-1}(S)$ is homeomorphic to S.*

The universal cover of the standard wedge product X is, of course, closely related to $\widetilde{Y}$. In fact, $\widetilde{X}$ can be obtained from $\widetilde{Y}$ by collapsing each component of the preimage of the backbone to a point, just as X can be obtained from Y by collapsing the backbone to a point. If we let X' denote the space obtained by quotienting $\widetilde{Y}$ in this way, it is easy to see that the composition $\widetilde{Y} \to Y \to X$ factors through X' to produce a map $X' \to X$ that is a local homeomorphism and thus a cover. Finally, by Proposition 1.2.21 the trivial fundamental group of $\widetilde{Y}$ maps onto the fundamental group of X' making X' simply-connected and the universal cover of X. This establishes the local structure of $\widetilde{X}$.

COROLLARY 1.4.4 (Structure of $\widetilde{X}$). *Let (X, x) be the wedge product of based connected cell complexes (X_α, x_α), let (Y, y) be the non-standard wedge product of this collection with backbone S, and let $p : \widetilde{Y} \to Y$ be its universal cover. If each component of $p^{-1}(S)$ is collapsed to a point then the resulting complex is $\widetilde{X}$. As a consequence, the inclusion map $X_\alpha \hookrightarrow X$ lifts to an inclusion $\widetilde{X}_\alpha \hookrightarrow \widetilde{X}$.*

The fact that we can construct the universal cover of X from the universal covers of the spaces X_α means that the fundamental group of a wedge product can be understood once we understand the fundamental groups of the individual spaces (Remark 1.3.17). In particular, the fact that $\widetilde{X}_\alpha$ embeds in $\widetilde{X}$ immediately proves the following basic result.

THEOREM 1.4.5 (Fundamental groups inject). *If (X, x) is a wedge product of based connected cell complexes (X_α, x_α), then the homomorphism $i_\alpha : \pi_1(X_\alpha, x_\alpha) \to \pi_1(X, x)$ induced by the inclusion map is injective for each α.*

PROOF. Each non-trivial $g \in \pi_1(X_\alpha, x_\alpha)$ is represented by a loop in X that lifts to an open path in a copy of $\widetilde{X}_\alpha$ inside $\widetilde{X}$ proving that $i_\alpha(g)$ is nontrivial. $\square$

In addition to collapsing onto $\widetilde{X}$, the space $\widetilde{Y}$ also collapses onto a tree.

COROLLARY 1.4.6 (Tree-like). *Let (Y, y) be the non-standard wedge product of a collection of based connected cell complexes (X_α, x_α) and let $p : \widetilde{Y} \to Y$ be its* *universal cover. If for each α each component of $p^{-1}(X_\alpha)$ is collapsed to a point, then the resulting graph is a tree.*

PROOF. The result is a graph since the high dimensional cells have disappeared and it is a simply-connected since its fundamental group is a quotient of the trivial group $\pi_1(\widetilde{Y})$ (Proposition 1.2.21). By Theorem 1.2.1 the quotient is a tree. □

Let T denote the tree obtained from $\widetilde{Y}$ in this way. When similar collapsing are carried out in Y itself, the result looks like its backbone S. The relations among the spaces and maps constructed so far are best illuminated by a diagram.

$$\begin{array}{ccccc} \widetilde{X} & \leftarrow & \widetilde{Y} & \rightarrow & T \\ \downarrow & & \downarrow & & \downarrow \\ X & \leftarrow & Y & \rightarrow & S \end{array}$$

Each of these arrows represents a quotient map that has already been described with the exception of the vertical arrow from T to S. The group $G \cong \pi_1(X, x) \cong \pi_1(Y, y)$ acts freely on $\widetilde{Y}$ by the fundamental theorem of covering spaces and it acts on $\widetilde{X}$ and T because the horizontal quotient maps commute with the G-action on $\widetilde{Y}$. In addition, each space in the bottom row can be viewed as the quotient of the space directly above it by this G-action. Because the action of G on $\widetilde{X}$ is free, the map $\widetilde{X} \to X$ is a cover; the map from $T \to S$ is not since the G-action on T has non-trivial stabilizers. We return to this picture in Chapter 7 since T is a simple example of a Bass-Serre tree and S, with the addition of the stabilizer information, is a simple example of a graph of groups.

1.4.2. Normal forms. Now that the tree-like nature of $\widetilde{Y}$ has been firmly established, we use this structure to create a canonical factorization of each nontrivial element in $G = \pi_1(Y, y)$. To facilitate the proof, we introduce additional notation.

DEFINITION 1.4.7 (Backbone vertices and their labels). Every vertex in the nonstandard wedge product Y belongs to exactly one of the cell complexes X_α except for the vertex y at the center of the backbone. We call y the *backbone vertex*, and the others we call *cell complex vertices*. Using the covering map $p : \widetilde{Y} \to Y$ and the quotient map $q : \widetilde{Y} \to T$ we can extend this partitioning of vertices to $\widetilde{Y}$ and then to T: preimages and images of backbone / cell complex vertices are backbone / cell complex vertices. These distinctions are particularly striking in T where every edge connects a backbone vertex to a cell complex vertex. Finally, we pick a backbone vertex $\tilde{y} \in \widetilde{Y}$ as our base point, and, as in Definition 1.3.14, we then use the G-action to label each backbone vertex of $\widetilde{Y}$ as y_g where g is the unique element in G with $y_g = g \cdot \tilde{y}$. The image of y_g under the quotient map q is called t_g.

LEMMA 1.4.8 (Paths to factors). *Each immersed path of length* 2 *connecting backbone vertices in T corresponds in a canonical way to a non-trivial element $g \in i_\alpha(\pi_1(X_\alpha, x_\alpha))$ for some particular α.*

PROOF. Let t_g and $t_{g'}$ be the backbone vertices in T at either end of the path and let v_α be the cell complex vertex it passes through. As in Proposition 1.2.21 we lift this path to $\widetilde{Y}$ by inserting a path in $q^{-1}(v_\alpha)$ connecting the appropriate endpoints. The lifted path projects to a loop in Y and then to a loop in $X_\alpha \subset X$ using the quotient maps already described and thus corresponds to an element

$g \in i_\alpha(\pi_1(X_\alpha, x_\alpha))$. See Figure 9. Although the lifting process involves a choice, any two such insertion paths are homotopic relative to their endpoints precisely because $q^{-1}(v_\alpha)$ is a copy of $\widetilde{X}_\alpha$ (Lemma 1.4.3) and thus simply-connected. $\square$

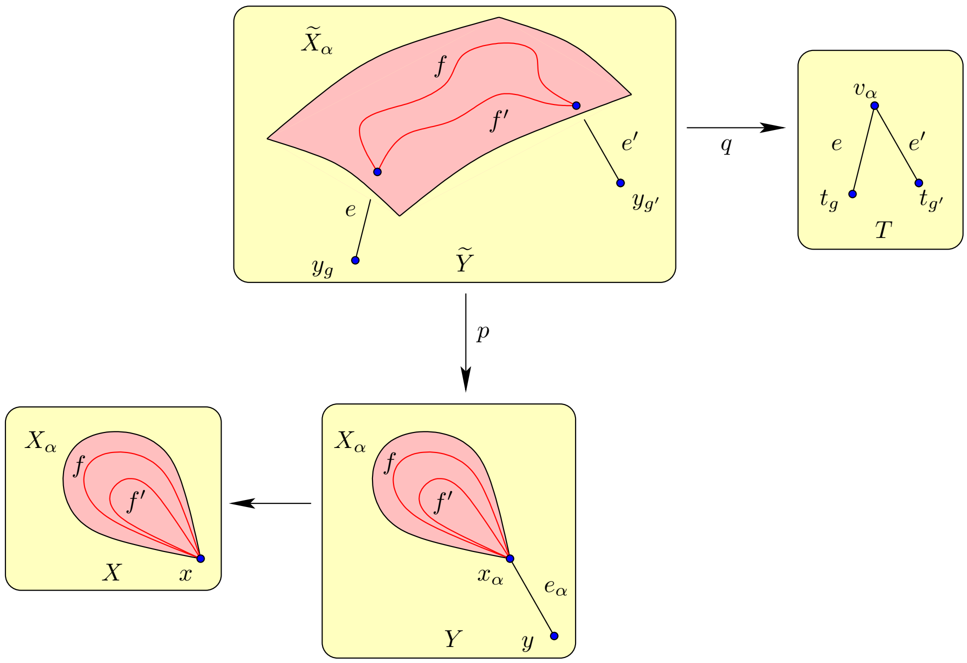

FIGURE 9. The correspondence betweeen paths in T and loops in X.

COROLLARY 1.4.9 (Existence). *If (X, x) is a wedge product of based complexes (X_α, x_α) and i_α is the group homomorphism induced by the inclusion $X_\alpha \hookrightarrow X$, then for each non-trivial element $g \in \pi_1(X, x)$ there is a canonical factorization of g as $g_1 g_2 \cdots g_k$ where each g_i is a non-trivial element of $i_\alpha(\pi_1(X_\alpha, x_\alpha))$ for some α and consecutive g_i's belong to distinct subgroups of this form.*

PROOF. Start with the unique immersed path in T from t_1 to t_g. Because it is immersed as it passes through cell complex vertices, Lemma 1.4.8 can be used to convert it step by step into a factorization; because it is immersed as it passes through backbone vertices, the α's involved in consecutive factors are distinct; and because the lifted path connects y_1 to y_g, the factorization produced is a factorization of g. $\square$

Uniqueness follows from the reversibility of this process.

LEMMA 1.4.10 (Factors to paths). *Every non-trivial element $g \in i_\alpha(\pi_1(X_\alpha, x_\alpha))$ corresponds in a canonical way to an immersed path of length 2 in T from the base point t_1 to the backbone vertex t_g.*

PROOF. For the appropriate α, pick an immersed path $f : I \to X_\alpha^{(1)} \subset X$ based at x representing g. Lift f to a loop based at y in Y by adding the edge e_α both before and after the loop f in $X_\alpha \subset Y$. There is unique lift of this new loop to $\widetilde{Y}$ starting at $\tilde{y} = y_1$ (Theorem A.3.7) and that lift projects to a path in T starting at t_1. Since the loop in $\widetilde{Y}$ only crosses two edges in the backbone, the projected path to T has length 2. Because g is non-trivial, the lift to $\widetilde{Y}$ is not closed and the two edges in the projection to T are distinct. Finally, the end result is independent of the path chosen to represent g since the possible lifts to $\widetilde{Y}$ have the same endpoints and project to the same path in T. $\square$

We now complete the proof of the main result.

THEOREM 1.4.11 (Wedge product normal form). *If (X, x) is a wedge product of based connected cell complexes (X_α, x_α), then for every non-trivial element $g \in \pi_1(X, x)$ there is one and only one way to write it in the form $g = g_1 g_2 \cdots g_k$ where each g_i is a non-trivial element in $i_\alpha(\pi_1(X_\alpha, x_\alpha))$ for some α and consecutive g_i's belong to distinct subgroups of this type.*

PROOF. Corollary 1.4.9 proves the existence of such factorizations, so we only need to show uniqueness. Let $g = g_1 g_2 \cdots g_k$ be a factorization where each g_i is a non-trivial element in $i_\alpha(\pi_1(X_\alpha, x_\alpha))$ for some α and consecutive g_i's belong to distinct subgroups of this type. Because each g_i is non-trivial, Lemma 1.4.10 can be used to produce an immersed length 2 path in T starting at any particular backbone vertex. If we rechoose the base point in $\widetilde{Y}$ and T at each step to be the endpoint of the previous lift, then these lifted paths can be concatenated and the result is immersed as it passes through each cell complex vertex. The fact that consecutive g_i's belong to distinct subgroups ensures that the concatenated path is also immersed as it passes through each backbone vertex. The result is an immersed path from in T from t_1 to t_g. Since the conversion process is deterministic in both directions, there must be a one-to-one correspondence between immersed paths from t_1 to t_g and factorizations of g of the desired type. But there is only one such path in a tree, so there is only one such factorization. $\square$

Since the rose R_A can be viewed as a wedge product of A circles, Theorem 1.4.11 immediately implies the following normal form for elements of free groups.

COROLLARY 1.4.12 (Free group normal form). *If R_A is an oriented rose with $\mathbb{F}_A = \pi_1(R_A, *)$ and $A \subset \mathbb{F}_A$, then every non-trivial element of $\mathbb{F}_A$ can be uniquely written in the form $a_1^{n_1} a_2^{n_2} \cdots a_k^{n_k}$ where each a_i is in A, each n_i is a nonzero integer, and adjacent a_i's are distinct.*

This corollary should not be surprising since the free group normal form given above is just a way of writing reduced words so that the transitions between letters are highlighted. The subwords $a_i^{n_i}$ highlighted in the corollary are known as *syllables*. Wedge products now can be used to define a product operation on groups.

DEFINITION 1.4.13 (Free products). Given an arbitrary collection $\{G_\alpha\}$ of groups, we can select a collection of based, connected cell complexes $\{(X_\alpha, x_\alpha)\}$ with $\pi_1(X_\alpha, x_\alpha) \cong G_\alpha$ for each α, and then define the *free product* of the collection, denoted $G = *_\alpha G_\alpha$, as $\pi_1(X, x)$ where $(X, x) = \vee_\alpha (X_\alpha, x_\alpha)$. By Theorem 1.4.11 the group that results is independent of the cell complexes chosen to represent each G_α, so the group G is well defined. In this notation the free group $\mathbb{F}_A$ is a free product of the form $*_{\alpha \in A}\mathbb{Z}$, and the fundamental group of the space shown in Figure 6 is the group $(\mathbb{Z}/2\mathbb{Z}) * (\mathbb{Z}/2\mathbb{Z}) * (\mathbb{Z}/2\mathbb{Z})$.

Reversing this construction leads to the notion of a free decomposition.

DEFINITION 1.4.14 (Free decompositions). A group G is *freely decomposable* if it can be written as a free product of non-trivial groups and any particular way of writing G as $*_\alpha G_\alpha$ is called a *free decomposition* of G. The groups G_α are called *free factors of* G. Topologically, a group is freely decomposable iff it can be represented as the fundamental group of a wedge product of two connected but not simply-connected spaces. A group that cannot be freely decomposed is said to be *freely indecomposable*.

1.4.3. Vertex links. We have seen that cell complexes with cut points are wedge products and that their fundamental groups are free products of the fundamental groups of the pieces. A similar result holds when a cell complex has a local cut point (Theorem 1.4.19).

DEFINITION 1.4.15 (Local cut points). A point x in a topological space X is called a *cut point* if X is connected but $X \setminus \{x\}$ is disconnected, and x is called a *local cut point* if there is a neighborhood U of x such that U is connected but $U \setminus \{x\}$ is disconnected.

When X is a combinatorial cell complex, it can be subdivided so that x is a vertex, and the structure of X near x is encoded in a lower dimensional complex called its *link*. Although the link of a vertex is slightly tricky to define, the idea is easy to explain, at least in the presence of a reasonable metric.

DEFINITION 1.4.16 (Vertex links; metric intuition). Let X be a combinatorial cell complex with a metric compatible with its topology and let v be a vertex in X. The *link of* v is the set of points in X at distance exactly ϵ from v (ϵ being a small positive number) with the induced topology and cell structure. It should be intuitively clear that so long as the metric on X remains reasonably nice and ϵ sufficiently small, the link, denoted $\text{Link}(v, X)$, is a cell complex whose structure is independent of ϵ and independent of the metric on X.

The idea behind Definition 1.4.16 can be made rigorous when X is a simplicial complex. An extended technical definition that applies to arbitary combinatorial cell complexes is also sketched.

DEFINITION 1.4.17 (Vertex links; technical version). If σ is a single simplex with the regular Euclidean metric and v is one of its vertices, then $\text{Link}(v, \sigma)$ is a simplex of one lower dimension cannonically homeomorphic (via the projection using straight lines through v) to the simplex spanned by the remaining vertices of σ. As a consequence, the link of a vertex v in a simplicial complex X can be idenitified with (or defined as) the set of simplices not containing v that are nonetheless contained in simplices that do contain v. The set of simplices containing v is called the *star of* v and it is homeomorphic to the ball of radius ϵ around v.

Vertex links in arbitrary combinatorial cell complexes can be defined using subdivision. If X is a combinatorial cell complex, then its second barycentric subdivision is a simplicial complex, and the link of v in the second barycentric subdivision of X is the second barycentric subdivision of the combinatorial cell complex one would want to call the link of v. The details of this procedure are left as an exercise.

Figure 10 illustrates the correspondence between the two definitions.

EXAMPLE 1.4.18 (Vertex links in 2-complexes). In combinatorial 2-complexes, vertex links have a very simple description: $\text{Link}(v, X)$ is a graph with a vertex

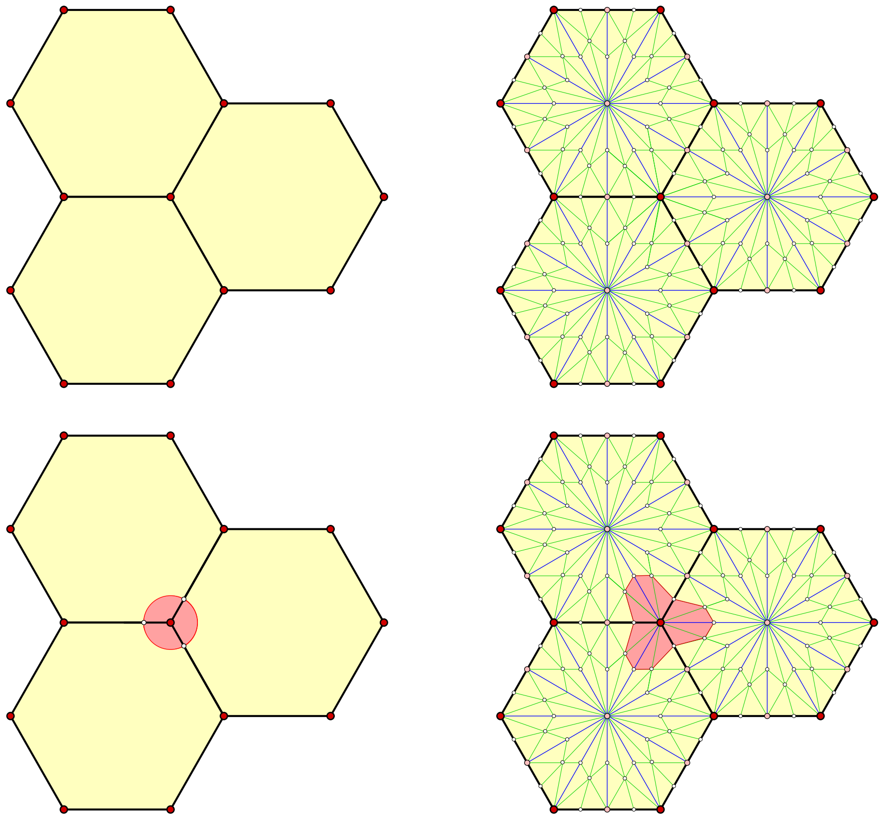

Figure 10. The top row shows a simple cell complex and its second barycentric subdivision; the second row shades the ball of radius ϵ around its central vertex and its star in the second barycentric subdivision, respectively.

for each end of a 1-cell attached to v and an edge for each occurrence of v in the boundary cycle of a 2-cell of X.

Since the link of v can be thought of as the sphere of radius ϵ around v and its structure is independent of ϵ as ϵ shrinks to 0, the ball of radius ϵ around v can be identified with the topological cone over its link. As a result, v is a local cut point iff the link of v is disconnected. We call a combinatorial cell complex X *link-connected* when all of its vertex links are connected cell complexes, and we note that this is true iff X has no local cut points. Using this characterization Theorem 1.4.19 converts local cut points into wedge products. A concrete illustration of the proof is given in Example 1.4.20 and shown in Figure 11.

Theorem 1.4.19 (Splitting 2-complexes). *Every group is the fundamental group of a wedge product of circles and link-connected 2-complexes.*

Proof. For every group G there is a taut, connected, one vertex 2-complex X with $\pi_1(X, *) = G$ (Proposition 1.2.6 and Corollary 1.1.3). Let $L = \text{Link}(*, X)$. If L is connected then we are done. Otherwise, let A and B be sets that index the connected components of L and $X \setminus \{*\}$, respectively, and note that since L

can be viewed as the boundary of an ϵ-neighborhood of $*$ in X, there is a well-defined map $f : A \twoheadrightarrow B$. We construct a new 2-complex Y by pulling the connected components of L in different directions. More specifically, start with a tree T that has 0-cells indexed by $A \sqcup \{*\}$ and an edge e_α from v_* to v_α for each $\alpha \in A$. The rest of Y is built by adding a 1-cell or 2-cell to T for each 1-cell and 2-cell in X in such a way that the complex obtained by contracting T to a point is equal to X. Concretely, for each 1-cell of X we add a 1-cell to T with each end attached to the vertex v_α in T where $\alpha \in A$ indexes the component of L through which this end approaches $*$ in X. This completes the 1-skeleton of Y. For each 2-cell of X we attach a 2-cell to $Y^{(1)}$ along the lift of its attaching map in X as constructed by Proposition 1.2.21. Because paths of length 2 in the boundary cycles of 2-cells create edges in L, the ends of these adjacent edges belong to the same component α, their lifts are attached to the same vertex v_α, and thus these lifts are concatenated without inserting additional edges. The quotient map from Y to X is a homotopy equivalence by Theorem A.4.5.

The remaining steps are straightforward. Since $Y \setminus T$ is homeomorphic to $X \setminus \{*\}$ under the quotient map, its connected components are also indexed by B. For each $\beta \in B$ select an edge e_α with $f(\alpha) = \beta$ and then reattach all unselected edges in T so that both of their endpoints are at v_*. See the lower righthand corner of Figure 11. The result is homotopy equivalent to Y by Theorem A.4.6 since there is a path from the other endpoint to v_* that travels through a component of $Y \setminus T$ and then back to v_* along a selected edge. The last step is to contract the tree formed by the selected edges to a point and to note that the result is a wedge product of circles and complexes indexed by B. Every vertex link in a complex indexed by B is connected since, by construction, it can be identified with a connected component of the original link L. Finally, if desired, simply-connected complexes can be removed from the wedge product without changing its fundamental group. $\square$

Example 1.4.20 (Splitting 2-complexes). Let X be the quotient of $\mathbb{S}^2 \sqcup \mathbb{S}^2$ that identifies two distinct points in the first 2-sphere and three distinct points in the second 2-sphere to a single point. The quotient X can be given a cell structure so that it is a taut connected one vertex 2-complex, but the exact cell structure is irrelevant. The link of the unique vertex $*$ in X has 5 connected components and $X \setminus \{*\}$ has 2. In other words $|A| = 5$ and $|B| = 2$. Figure 11 illustrates the sequence of steps used to show that X is homotopy equivalent to $\mathbb{S}^2 \vee \mathbb{S}^2 \vee \mathbb{S}^1 \vee \mathbb{S}^1 \vee \mathbb{S}^1$ and that $\pi_1(X, *) = \pi_1(\mathbb{S}^1 \vee \mathbb{S}^1 \vee \mathbb{S}^1, *) = \mathbb{F}_3$.

One corollary of Theorem 1.4.19 is that every freely indecomposable group is either infinite cyclic or the fundamental group of a link-connected 2-complex. The converse, however, is false (Exercise 27).

1.5. Constructions and examples

This final section contains a way to easily describe many 2-complexes with multiple vertices, and discusses examples of groups that arise from simple topological constructions.

1.5.1. Presentations revisited. Some cell complexes are easy to describe: a rose corresponds to a set A, and a standard topological presentation can be

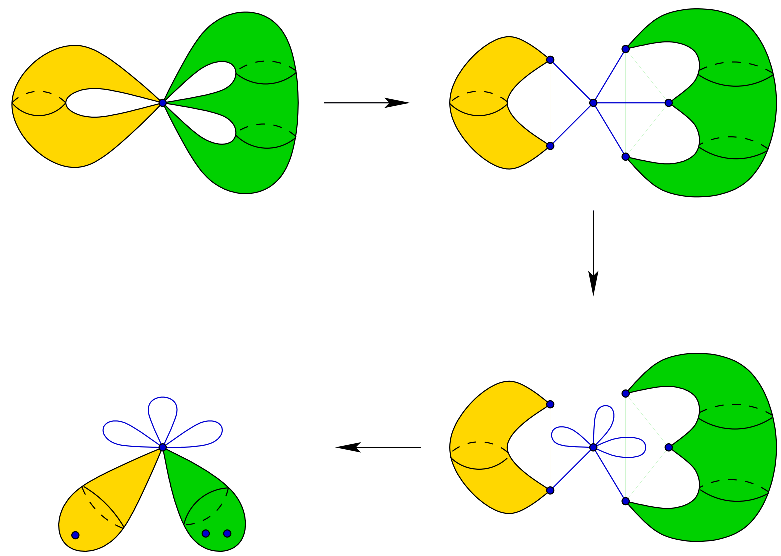

FIGURE 11. An illustration of the homotopy equivalences used to convert an arbitrary 2-complex into a wedge product of circles and link-connected 2-complexes.

constructed from an algebraic presentation $\langle A|\mathcal{R}\rangle$. When we try to describe 2-complexes with multiple vertices using similar techniques, there are two issues that arise. First, there is no standard way to quickly describe a complicated 1-complex with edges oriented and labeled by a set A, and second, even once such a 1-skeleton is given, not all words in $(A \cup A^{-1})^*$ can be used to describe closed paths, making it easy to list collections of words that are incompatible with the given graph. In the absence of local cut points, however, there is a simple procedure that avoids both of these difficulties. It constructs a multi-vertex link-connected combinatorial 2-complex from any multiset of words, and every such complex can be constructed in this way. Such a process is sufficient for most purposes since by Theorem 1.4.19 the only 2-complexes excluded are those that are homotopy equivalent to a non-trivial wedge product in an obvious way. The construction begins with polygons.

DEFINITION 1.5.1 (Polygons). A *polygon* is a 2-disc whose boundary cycle has been given the structure of a graph. When its boundary cycle has combinatorial length n it is called an *n-gon*, and traditional names, such as monogon, bigon, triangle, square, pentagon and hexagon, are used when n is small.

Polygons arise naturally in the construction of combinatorial 2-complexes.

REMARK 1.5.2 (Polygons and 2-complexes). Let X be an arbitrary 2-complex and recall that X is defined as $X^{(1)} \sqcup_F E$ where $X^{(1)}$ is a 1-complex, $E = \sqcup\, \mathbb{D}^2$ is a disjoint union of 2-discs, one for each 2-cell of X, and E is attached to $X^{(1)}$ along the induced map $F : \partial E \to X^{(1)}$ that collects together all of the individual attaching maps (Definition A.1.3). When X is combinatorial the boundaries of the 2-discs in E can be subdivided into vertices and edges so that $F : \partial E \to X$ is a cellular map. Under this subdivision, E is a disjoint union of polygons and a combinatorial 2-complex in its own right. Moreover, the induced map $E \to X$ is cellular and a quotient map. Note also that there is a natural bijection between vertices in E and the edges in the vertex links of X.

There is an intermediate complex constructed from the edges identifications.

Definition 1.5.3 (Edge identifications). Let X be a combinatorial 2-complex and let E be the disjoint union of polygons used to construct X. There is a third cell complex Y, between E and X, defined as follows. Identify pairs of 1-cells in E iff they are sent to the same 1-cell in X, and identify them in the same fashion. For Y to be a cell complex certain vertex identifications must also be made, but only make those that are forced by the edge identifications. The quotient map $E \twoheadrightarrow X$ factors into quotient maps $E \twoheadrightarrow Y \twoheadrightarrow X$, and we say that Y is *constructed from* X *by edge identifications*. Notice that since $E \to Y$ is a factor of $E \to X$, the only vertices in E that can be identified in Y are those with the same image in X.

The key observation is the following.

Lemma 1.5.4 (Vertex identifications). *If X, E, F and Y are defined as above and v and v' are vertices in E with $F(v) = F(v') = u$ in X, then v and v' are identified in Y iff the edges of* $\text{Link}(u, X)$ *corresponding to v and v' belong to the same connected component.*

Proof. Both directions are straightforward. If the corresponding edges belong to the same connected component then there is a finite length path connecting them in the link. This path encodes a finite sequence of individual edge identifications that force v and v' to be identified in Y. Conversely, identifying vertices iff the corresponding edges belong to the same connected component of the link produces a cell complex in which all the edge identifications can be performed with no further vertex identifications. Thus, no additional vertex identifications are forced. □

The following properties follow immediately from Lemma 1.5.4.

Proposition 1.5.5 (Edge identifications). *If X is a combinatorial* 2*-complex and Y is constructed from X by edge identifications, then Y is always link-connected and the quotient map $Y \to X$ is a homeomorphism iff X is link-connected.*

We are now ready for the general construction.

Definition 1.5.6 (Combinatorial descriptions). Let A be a set and let $\mathcal{R}$ be a multiset selected from $(A \cup A^{-1})^*$. First, let E be a disjoint union of polygons indexed by the words in $\mathcal{R}$ so that the polygon corresponding to a word of length n is an n-gon. Next, choose a vertex and a direction for each polygon in E and then use the corresponding word to orient and label the edges of this polygon so that starting at the chosen vertex and proceeding in the chosen direction, the labels and orientations encountered represent the associated word. Finally, define Y as the quotient of E that identifies edges according to label and orientation, and identifies vertices only when this is needed to make the quotient a cell complex. We call Y the *complex constructed from* $[\mathcal{R}]$ and $[\mathcal{R}]$ is a *combinatorial description of* Y. Square brackets are used in place of angled ones to highlight the distinction between combinatorial descriptions and algebraic presentations, and it is “combinatorial” rather than “algebraic” since the letters used do not correspond to the generators of a group.

Theorem 1.5.7 (Words and 2-complexes). *Every combinatorial description $[\mathcal{R}]$ constructs a link-connected combinatorial 2-complex X and every such X can be converted into a multiset of words from which X can be recovered. Under these conversions, the number of 2-cells in X corresponds to $|\mathcal{R}|$. In particular, X is compact iff $\mathcal{R}$ is a finite list.*

Proof. The main assertions have similar proofs. Either let X be the complex constructed from a given combinatorial description $[\mathcal{R}]$, or let $[\mathcal{R}]$ be the combinatorial description derived from a given link-connected combinatorial 2-complex X as follows. Let E be the disjoint union of polygons used to construct X. Orient and index the 1-cells of X by a set A, and then induce an orientation and labeling of the 1-cells in E by pulling these features back through the quotient map $E \twoheadrightarrow X$. Next, for each polygon in E, select a vertex and a direction and then reduce the oriented labeling of its boundary cycle to a word in $(A \cup A^{-1})^*$. Let $\mathcal{R}$ denote the multiset of words produced in this way. Under either scenario, we claim that the complex described by $[\mathcal{R}]$ is identical to the complex constructed from X by edge identifications since they both make the same identifications. Moreover, this common complex is link-connected and equal to X by Proposition 1.5.5. $\square$

The useful conventions for algebraic presentations listed in Definition 1.3.26 also apply to combinatorial descriptions. The main distinction between combinatorial descriptions and algebraic presentations is highlighted by the following example.

Example 1.5.8 (Combinatorial descriptions vs. algebraic presentations). The complex constructed by $[abcABC]$ is a non-standard torus with two vertices and its fundamental group is $\mathbb{Z}^2$. The algebraic presentation $\langle a, b, c \mid abcABC \rangle$, on the other hand, constructs the quotient of this torus with its vertices identified. Using Theorem 1.4.19, the latter complex is homotopy equivalent to a wedge product of a torus and a circle and its fundamental group is $\mathbb{Z}^2 * \mathbb{Z}$. See Exercise 33 for a generalization.

1.5.2. Simple examples. The only groups that are fundamental groups of 1-complexes are, by definition, the free groups, and their algebraic structure is reasonably well understood. On the other hand, every group is the fundamental group of a 2-complex, and, in a very precise sense, many of them are difficult or impossible to understand. See Chapter 3. As a first step into the world of 2-complexes, we consider 2-complexes that have combinatorial descriptions consisting of a single word and their corresponding *one-relator* fundamental groups. We begin with surfaces.

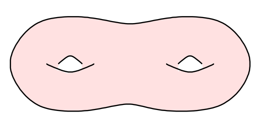

Figure 12. A surface of genus 2.

say more about surfaces topologically, universal covers, classification, etc

Example 1.5.9 (Compact surfaces). Classification, genus, orientation, distinctions between universal covers

The ones with boundary deformation retract to graphs and thus their fundamental groups are free.

Let $\mathcal{R}$ be a finite list of words. If every letter that occurs, occurs at most twice (in either orientation), then $\mathcal{R}$ describes a compact surface, possibly disconnected, with a nonempty boundary iff there is a letter that only occurs once (Exercise 35). Simple examples such as $[a]$, $[aa]$, $[aA]$, $[aab]$, $[abAB]$, $[abaB]$, $[abcABC]$, and $[abcB]$ produce a disc, a real projective plane, a 2-sphere, a Möbius strip, a torus, a Klein bottle, another torus, and another Möbius strip, respectively.

Our next examples take a surface with boundary and wrap each boundary cycle multiple times around a circle. These attaching maps are determined, up to homotopy, by an integer called its degree.

Definition 1.5.10 (Maps between circles). If we view $\mathbb{S}^1$ as the set of unit complex numbers, then for each $n \in \mathbb{Z}$ we can define a map $f_n : \mathbb{S}^1 \to \mathbb{S}^1$ that sends $z \mapsto z^n$. Topologically this is just a map that wraps one circle $|n|$ times around the other with no backtracking where the sign of n indicates which way to proceed. The number n is called its *degree* and if $f : \mathbb{S}^1 \to \mathbb{S}^1$ is any map homotopic to f_n, then f is called a *degree n map*. It is easy to show that every map $\mathbb{S}^1 \to \mathbb{S}^1$ is homotopic to exactly one such f_n (Exercise XXX), so every map between circles has a unique degree.

The simplest surface with boundary is a disc.

Example 1.5.11 (Discs and finite cyclic groups). Let X be the space that results when a disc is attached to a circle by an attaching map of degree n (Figure 13). A combinatorial description of X is $[a^n]$ and an algebraic presentation is $\langle a|a^n\rangle$. These are derived by giving the circle the simplest possible graph structure with one vertex and one edge. The fundamental group of X is $\mathbb{Z}/n\mathbb{Z}$, the finite cyclic group of order n, and its universal cover looks like n distinct n-gons with their boundary cycles identified.

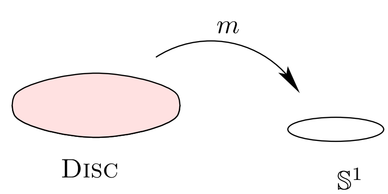

Figure 13. A disc attached to a circle.

Example 1.5.12 (Annuli and torus knot groups). The one vertex version $[a^m = b^n]$ is hard to understand (but not /that/ hard really) but the two-vertex presentation $[a^m t = tb^n]$ is trivial since its universal cover is a tree cross the reals.

Example 1.5.13 (Möbius bands and one-relator Artin groups). If a Möbius strip is attached to a circle along its boundary cycle then the fundamental group of the resulting space is called a *one-relator Artin group*. The name, of course, is derived from the theory of a larger class of groups. Topologically these groups are very simple, and there should probably be a better name and notation for them.

(explain the presentation) $\langle a, b \mid ababa... = babab...\rangle$

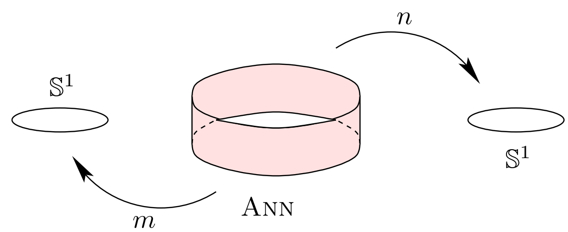

Figure 14. An annulus attached to two different circles.

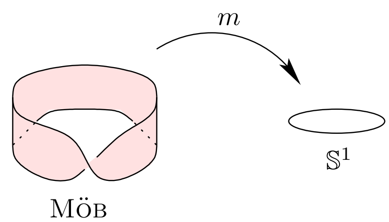

Figure 15. A Möbius band attached to a circle.

Example 1.5.14 (Annuli and Baumslag-Solitar groups). The Baumslag-Solitar group, $BS(m, n)$ is the group defined by the single relation $[a^m b = ba^n]$. Topologically they are the fundament groups of the spaces constructed by attaching both ends of an annulus to the same circle, one attaching map with degree m and the other with degree n. Despite their elementary definition, these groups have a number of surprising properties.

These groups were first systematically studied by Gilbert Baumslag and Donald Solitar in 1969 (check this and add some history / refs, quote John's book).

They have a number of quite interesting properties, and their analysis is not nearly so elementary as one might think.

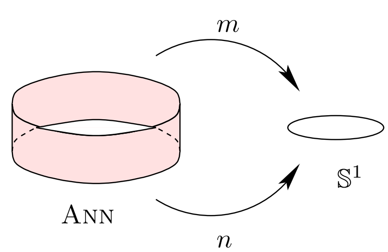

Figure 16. An annulus attached to a single circle.

They have proved an interesting object of study, even after 30 years. For example, it was only recently that it was completely determined which pairs of Baumslag-Solitar groups were quasi-isometric to one another (and the answer was slightly surprising). [Amenable ones by Benson Farb and Lee Mosher in 1998 [**11**] and the non-amenable ones by Kevin Whyte in 2001 [**30**]]

Remark 1.5.15 (3-manifolds groups). Every group is the fundamental group of a combinatorial 2-complex, but not every group is the fundamental group of a 1-complex or of a manifold with dimension at most 3.

Add remarks about 3-manifolds being special somewhere

Notes

Historical notes and other comments will eventually go here.

Exercise 1 is a baby version of Milnor's construction of Eilenberg-Maclane spaces for groups.

Exercises

Cell complexes

1. (Groups as fundamental groups) Let G be a group and let Y be the simplicial complex with vertices indexed by $G \times \{1,2,3\}$ and a simplex spanning every subset of vertices with pairwise distinct second coordinates. Let X be the quotient of Y by the natural G-action defined by $g \cdot v_{(h,a)} = v_{(gh,a)}$.
   a. Show that when $|G| = 2$, Y is the boundary of an octahedron, the G-action is antipodal, and X is homeomorphic to $\mathbb{R}P^2$.
   b. Prove that for every group G, Y is connected and simply-connected, the G-action is free and cellular, and thus X is a cell complex with fundamental group G.
2. (Whitney embeddings) Find explicit embeddings of (subdivisions of) the graphs K_5 and $K_{3,3}$ into $\mathbb{R}^3$ using the proof of Theorem 1.1.5. Similarly, choose a cell structure for $\mathbb{R}P^2$ and linearly embed a subdivision into $\mathbb{R}^5$.

Graphs and trees

3. (Metrics on graphs) Prove that the combinatorial distance function $d_X(u,v)$ defines a metric on the 0-skeleton of any connected graph X. Next, show that there is a natural extension of the combinatorial distance function that defines a metric on all of X.
4. (Finite versus infinite rank) Prove that when A is finite and B is infinite, R_A and R_B are not homotopy equivalent, and conclude by the theory of Eilenberg-Maclane spaces that $\mathbb{F}_A$ and $\mathbb{F}_B$ are not isomorphic groups.
5. (Free group cardinality) Recall from cardinal arithmetic that if at least one of κ and λ is an infinite cardinal, then $\kappa \cdot \lambda = \max\{\kappa, \lambda\}$. In particular, if $\aleph_0$ denotes the cardinality of the natural numbers, n denotes a finite cardinal (any cardinal $n < \aleph_0$) and κ denotes an infinite cardinal (any cardinal $\kappa \geq \aleph_0$), then $n \cdot \kappa = \aleph_0 \cdot \kappa = \kappa \cdot \kappa = \kappa$.
   a. Prove that in a uniformly κ-branching tree there are exactly $\kappa(\kappa-1)^n$ vertices distance $n+1$ from a given vertex v, κ arbitrary.
   b. Prove that $|\mathbb{F}_A| = \aleph_0 \cdot |A| = \max\{\aleph_0, |A|\}$, and conclude that $|\mathbb{F}_A| = \aleph_0$ when A is finite and $|\mathbb{F}_A| = |A|$ when A is infinite.
6. (Maps between roses) A map $f : X \to Y$ is called a π_1-*injection*, a π_1-*surjection*, or a π_1-*isomorphism* when the induced map f_* between fundamental groups is injective, surjective or an isomorphism, respectively. By Theorem 1.2.11, $\exists$ a π_1-isomorphism $f : R_A \to R_B$ iff $|A| = |B|$.
   a. Prove $\exists$ a π_1-surjection $f : R_A \to R_B$ iff $|A| \geq |B|$.
   b. Prove $\exists$ a π_1-injection $f : R_A \to R_B$ iff $|\mathbb{F}_A| \leq |\mathbb{F}_B|$.
7. (Tree removal) Let T be a tree in a graph X and let $q : X \to X/T$ be the corresponding quotient map. Use Theorem 1.2.1 to show that every non-trivial immersed closed path based at $x \in X$ is sent by q to a non-trivial immersed closed path based at $q(x)$.

Free groups

8. (Basic properties) Prove that for every cardinal $\kappa > 1$, $\mathbb{F}_\kappa$ is infinite, non-abelian and has trivial center.
9. (Algebraic definition) Prove that Definition 1.2.22 produces a group.
10. (Symmetric bases) State and prove versions of Proposition 1.2.25 and Proposition 1.2.26 that hold for categorical free groups with symmetric bases. In

particular, prove that a group G is a categorical free group with symmetric basis S iff there is a topological free group $\mathbb{F}_A = \pi_1(R_A, *)$ with symmetric basis S_A and an isomorphism $f : G \to \mathbb{F}_A$ with $f(S) = S_A$.

11. (Comparing bases) Let (X, x) be a based connected graph, let T and T' be two different spanning trees in X, and let A and B index the edges not in T and T', respectively. Determine which elements of $\pi_1(X, x)$ are contains in both symmetric bases S_A and S_B under the isomorphisms $\mathbb{F}_A \cong \pi_1(X, x) \cong \mathbb{F}_B$. In particular, prove that distinct spanning trees identify distinct bases for $\pi_1(X, x)$.
12. (Infinitely many bases) Prove that $\mathbb{F}_\kappa \cong \textsc{Inn}(\mathbb{F}_\kappa) \subset \textsc{Aut}(\mathbb{F}_\kappa)$ for any cardinal $\kappa > 1$. Conclude that every non-abelian free group has an infinite number of bases. What happens for $\kappa \leq 1$?

Free group automorphisms

13. (Finite rank automorphisms) Complete the proof of Corollary 1.2.29.
14. (Infinite rank automorphisms) Let α be any ordinal of cardinality κ. Prove that the automorphisms of $\mathbb{F}_\kappa$ are in one-to-one correspondence with the well-orderings of the bases of $\mathbb{F}_\kappa$ that have order type α.
15. (Abelianization) Let $\mathbb{Z}^A$ denote the direct sum of A copies of the integers (whose elements are functions $A \to \mathbb{Z}$ with only finitely many non-zero values). Show that the abelianization of $\mathbb{F}_A$ is $\mathbb{Z}^A$ and that the abelianization map $\mathbb{F}_A \to \mathbb{Z}^A$ sends a basis of $\mathbb{F}_A$ to a basis of $\mathbb{Z}^A$ viewed as a free $\mathbb{Z}$-module. Conclude that there is a group homomorphism from $\textsc{Aut}(\mathbb{F}_A)$ to $\textsc{Aut}(\mathbb{Z}^A)$ and note that the latter is the group $GL_\kappa(\mathbb{Z})$ when $\kappa = |A|$ is finite.
16. (Primitive elements) An element in a free group is *primitive* if it belongs to some free basis. Find an element in $\mathbb{F}_2$ that is not primitive (and prove that it is not primitive).
17. (Bases and graphs) Let X be a connected graph and let T be a spanning tree in X. Show that the edges of X not in T form a basis in the following sense. [Fundamental groups of connected graphs are free groups but they do not have obvious bases when there is more than one vertex. For example, if X is the 1-skeleton of a cube and x is one of its vertices, then $\pi_1(X, x)$ is isomorphic to $\mathbb{F}_5$ (since its rank is $|\widetilde{\chi}(X)| = |8 - 12 - 1| = 5$), but there is no obvious choice for a five element basis or ten element symmetric basis. One possibility is to contract a spanning tree in X to create a rose with 5 leaves, but doing so involves several asymmetrical choices.]
18. (Rose homeomorphisms)
19. (Planar surface model)
20. (Handlebody model)
21. (Develop some elementary automorphism of free group stuff in the exercises. Include exercises on the various model spaces for free groups)
22. (Develop some elementary Stallings foldng exercises as well)

Generating sets and Cayley graphs

23. (Detecting Generating Sets) Complete the proof of Proposition by showing that the three collections are in natural bijection up the listed notions of equivalence.

Wedge products and free products

Add a figure so its clear that x is *not* the cone point

24. (Standard and non-standard wedge products) Let (X, x) be a cone on the Hawiian earring where x is the so-called 'bad point' in the base of the cone. Let (Y, y) be another copy of the same based space. Show that the standard and non-standard wedge products of (X, x) and (Y, y) are not homotopy equivalent by showing that the non-standard wedge product is simply-connected but that the standard wedge product has a non-trivial fundamental group.
25. (Non-abelian) Use the normal form theorem to prove that every non-trivial free product is non-abelian, and conclude that abelian groups are freely indecomposable.
26. (Local Cut Points) Let x be a point in a topological space X. Show that if U is a connected neighborhood of x such that $U \setminus \{x\}$ is disconnected and $V \subset U$ is another connected neighborhood of x, then $V \setminus \{x\}$ is also disconnected. Thus being a local cut point only depends on arbitrarily small neighborhoods of x.
27. (Decomposable and link-connected) Give an example of a combinatorial 2-complex that is link-connected but whose fundamental group can be decomposed as a free-product of non-trivial groups.
28. (Normal form algorithm) Describe an algorithm that inputs an aribitrary product of elements in a free product, outputs its unique normal form, and only basic knowledge about elements in the factor groups. In particular, your algorithm may assume (and in fact it must assume) that the algebraic structure of the factor groups is well understood. How is your algorithm related to the process for simplifying paths in trees?

    What effect do reparsing and 1-elimination have on the corresponding path in T.

Presentations

29. (The quick brown fox) Prove that the complex $[The, quick, brown, fox, jumped, over, a, lazy, dog]$ is connected and that its fundamental group is free. Find its rank.
30. (English) Let X be the 2-complex defined by the list of the 50,000 or so words in the English language (picking some official list of words in order to make this precise). Prove that X is connected, simply-connected, and has only one vertex.
31. (Your name here) Let X be the 2-complex constructed from your full name. Find $\chi(X)$. What do you know about $\pi_1(X)$? Is it free? If so, what is its rank? Warning: for some names these later questions might be hard to answer.
32. (Retracts and finite presentations) Prove that the retract of a finitely presented group is finitely presented (and that a presentation can be found via the retraction map). The idea for this exercise is from the Groves-Wilton paper/presentation.
33. (Standard versus non-standard) Let $[\mathcal{R}]$ be a combinatorial description, let A be the set of letters that occur in $\mathcal{R}$ and let X be the complex described by $[\mathcal{R}]$. Show that if X is connected, then the group presented by $\langle A|\mathcal{R}\rangle$ is the free product of $\pi_1(X)$ and a free group $\mathbb{F}$ of rank $|X^{(0)}| - 1$.

Simple examples

34. (Trefoil knot Dehn complex) Prove that the complex $[acbd, adbe, aebc]$ is the Dehn complex of the trefoil knot described in the prologue. How do the edges a through e listed above correspond to the edges E_0 through E_4 used in the prologue?
35. (Letters and surfaces) Prove that a list of words in which each letter is used at most twice describes a surface. When is the surface closed? When is it connected? If it is a single word, how can you detect orientability?
36. Describe in detail the universal covers of each of the following basic complexes (move the simple surface-like examples here).
37. (Classification) We should outline the classification of surfaces as an exercise.

    Turn some of this into exercises.
    $\langle a_1, a_2, \ldots, a_n \mid a_1 a_2 \cdots a_n = a_n \cdots a_2 a_1 \rangle$
    $\langle a_1, a_2, \ldots, a_n \mid a_1^2 a_2^2 \cdots a_n^2 = 1 \rangle$

    In the first case, it is easy to check that these surfaces are orientable (because the two occurences of each letter have opposite orientations in the boundary). In the second case, these surfaces are non-orientable since the presence of the subword a_1^2 already implies that there is a Möbius strip inside the surface. (comment about what happens when n is odd /even in the first type: 1-vertex versus 2).

    By the classification of compact surfaces, the complexes for these presentations include at least one representative of each compact surface, and the only ones which are homeomorphic are Type I with $2n$ and $2n + 1$.