


Geometric Group Theory

Jon McCammond



[Image by Bill Casselman: <http://www.math.ubc.ca/~cass/coxeter/rank3/cd.gif>]

Contents

Introduction	iii
Prologue: Trefoil Knot Group	vii
Exercises	xi
Chapter 1. Combinatorial Group Theory	1
1.1. Manifolds and cell complexes	1
1.2. Trees, graphs and free groups	4
1.3. Complexes and presentations	15
1.4. Cut points and free products	24
1.5. Constructions and examples	33
Notes	39
Exercises	40
Chapter 2. Metrics on Groups	45
2.1. Metrics and quasi-isometries	45
2.2. Geometries and geometric actions	46
2.3. Geometric properties of groups	52
Exercises	59
Chapter 3. Dehn's Fundamental Problems	61
3.1. Disc diagrams and van Kampen's lemma	61
3.2. Algorithms and decidability	62
3.3. Dehn's algorithm for surface groups	63
3.4. Combinatorial Curvature	64
Exercises	67
Chapter 4. Hyperbolic Geometry	69
4.1. Circle-preserving maps	69
4.2. Models of hyperbolic space	73
4.3. Isometries of hyperbolic space	74
4.4. Triangles in hyperbolic space	74
Chapter 5. Gromov's Hyperbolic Groups	77
5.1. δ -hyperbolic spaces	78
5.2. Quasi-geodesics	78
5.3. Equivalent definitions	79
5.4. Rips' complex	80
5.5. Finite subgroups	81
Chapter 6. Ends and Boundaries	83

6.1. The Space of Ends	83
6.2. The Boundary at Infinity	84
Exercises	84
Chapter 7. Splittings and Quasiconvexity	85
7.1. Actions on trees	85
7.2. Scott and Wall's approach	85
7.3. Bass-Serre Theory	85
7.4. Amalgamations and Quasiconvexity	85
Epilogue: Where to go from here	87
Appendix A. Algebraic Topology	89
A.1. Cell complexes and Euler characteristics	89
A.2. Fundamental groups and van Kampen's theorem	91
A.3. Group actions and covering spaces	93
A.4. Homotopy invariants and Whitehead's theorem	95
A.5. Classifying spaces and Hurewicz's theorem	97
Exercises	100
Appendix B. Hints	103
Chapter 1	103
Appendix A	103
Bibliography	105

Introduction

Geometric group theory classifies groups by the nature of the spaces on which the groups act geometrically.

James W. Cannon [7]

Geometric group theory is a relatively young field, but it has deep roots in the study of groups from combinatorial and topological perspectives. For almost one hundred years combinatorial group theorists have viewed groups as essentially topological objects and they have used the topological invariants of combinatorial cell complexes to study their associated fundamental groups. Since the mid-1980s, spurred on by the seminal ideas of Jim Cannon and Misha Gromov, group theorists have paid increasing attention to the geometric structures these cell complexes can carry. Finitely generated groups are now also viewed as inherently metric objects.

The addition of a geometric perspective has been tremendously successful at solidifying previously disparate results, generating new questions for researchers to investigate, and enabling rapid progress on many fronts. An unfortunate corollary of this rapid expansion has been a separation between the background acquired by graduate students in their standard courses and the conceptual tools used by current researchers in the field. This book is my attempt to partially fill this gap.

Groups as Actions

One way to appreciate the naturalness of the geometric group theory approach is to take a step back and consider the way in which groups arise in mathematics more generally. Group theory comes from the study of symmetry, where a *symmetry* of an object P (or an equation, or a geometric configuration, or any other mathematical structure) is a non-trivial invertible map f from P to P that preserves the properties we wish to consider. The collection of all such maps, trivial or not, is clearly closed under function composition (automatically an associative operation), it includes the identity map, and it includes the inverses of these maps by definition. These *symmetry groups* are where the subject began. To this day, groups are often first introduced through a careful examination of the symmetry groups of specific geometric objects, such as regular n -gons, regular n -simplices, or the unit n -sphere. The resulting groups are the dihedral, symmetric, and orthogonal groups in our examples, or, if we restrict our attention to only those symmetries realizable as continuous motions inside \mathbb{R}^2 , \mathbb{R}^n , or \mathbb{R}^{n+1} , respectively, we get the cyclic, alternating, and special orthogonal groups. Geometric objects, of course, are not the only mathematical structures that have symmetry groups. The symmetries of a vector space V form the *general linear group* $GL(V)$ and, more generally, the symmetries of any mathematical structure is called its *automorphism group*.

But if the abstractions pursued by twentieth century mathematicians have taught us anything, it is that mathematical structures should always be considered in conjunction with their structure-preserving homomorphisms and these maps also have symmetry groups! If $f : P \rightarrow Q$ is any structure-preserving homomorphism, then the collection of all invertible structure-preserving maps $g : P \rightarrow P$ such that $f \circ g = f$ form a group, as do the collection of all invertible structure-preserving maps $h : Q \rightarrow Q$ such that $h \circ f = f$. We can think of these groups as the *right* and *left stabilizer groups* of the map f , respectively. These types of groups also occur throughout mathematics. If $f : k \rightarrow K$ is a (necessarily injective) field homomorphism, for example, then its left stabilizer is better known as the *Galois group* of K over k . A second example, and one that is particularly important in our context, is when X is a path-connected topological space that has a universal cover \tilde{X} and $p : \tilde{X} \rightarrow X$ is the natural covering projection. The right stabilizer of p is the *group of deck transformations* of p , and it happens to be isomorphic to the *fundamental group* of X .

In each of the situations described above, the group under consideration is acting on some mathematical object via structure-preserving maps. The structure of the object upon which the group is acting can then be used to extract detailed information about the group itself. In some sense, this is the main way that groups occur “in nature”, as mathematicians like to say, and it is primarily through such actions, or *representations*, that groups are studied.

Finitely Presented Groups

Groups are investigated via representations as actions, but the type of representation varies with the type of group under consideration. For finite groups, group actions on finite sets (called *permutation representations*) or on vector spaces (known as *linear representations*) are highly effective and extensively used.¹ Geometric group theorists, on the other hand, focus their attention on groups that can be analyzed using actions on topological spaces—particularly cell complexes and metric spaces—and these often have infinitely many elements.

Infinite groups remain mysterious to many mathematics majors, since the groups encountered in a typical abstract algebra course are mostly finite. This is partly out of necessity: the main tools used to study infinite groups require more topology and geometry than can be presumed at that point. Moreover, when studying infinite groups, the algebraic structure often recedes into the background as topological, geometrical and logical considerations play a greater role.

Once infinite groups are under consideration, logical and informational issues immediately arise. Which infinite groups should be studied? If we are too inclusive in our scope, set theoretic issues could easily play a dominating role. On the other hand, the scope should be broad enough to include interesting examples, such as the fundamental groups of compact manifolds with or without boundary. One approach would be to limit our attention to precisely these groups. The obvious follow-up question is which groups are these? It turns out that this particular class of groups has several equivalent characterizations. Algebraically, they are the groups G that

¹These types of representations have been particularly important in the classification of the finite simple groups. See Michael Aschbacher’s book on finite group theory [1] for an excellent illustration of this approach and its benefits.

can be *finitely presented* in the following sense: (1) there exists some finite set of elements that generate all of G and (2) the relations that hold among the words in these *generators* can be derived from a finite list of basic rules or *relations*. Two other descriptions that describe the same class of groups are the fundamental groups of compact cell complexes and the fundamental groups of finite simplicial complexes. In other words, the following four collections of groups are identical.

$$\begin{aligned} \{ \text{finitely presented groups} \} &= \{ \pi_1 \text{ of compact manifolds} \} \\ &= \{ \pi_1 \text{ of compact cell complexes} \} \\ &= \{ \pi_1 \text{ of finite simplicial complexes} \} \end{aligned}$$

This natural class of groups will be our primary focus, although it is sometimes convenient to consider groups that are finitely generated but not finitely presented, or even groups where no finite subset generates the whole group.

While it is certainly possible to develop the theory of finitely presented groups using the algebraic description with only a passing mention of topology and geometry, doing so makes many of the fundamental properties of infinite groups unnecessarily difficult to express and even harder to establish. As a geometric group theorist, I have tried instead to highlight the geometric and topological aspects as much as possible.

Scope and Prerequisites

As it has grown over the past twenty years, geometric group theory has developed strong connections with geometry, topology, analysis and logic and each of these facets is currently undergoing rapid development. It would be nearly impossible at this point to give a truly comprehensive introduction to geometric group theory in a single volume and the text you have before you is not intended as one.² I have tried instead to produce a book that thoroughly covers a cohesive subset of fundamental ideas, focusing on a selection of elementary and intermediate topics that I feel are absolutely essential. Such a selection is, of course, highly subjective. While I am confident about the centrality of the included topics, the reader should not infer that excluded ones are less important.

The foundational ideas in geometric group theory are fairly accessible and the required prerequisites are correspondingly minimal: the algebraic topology covered in Hatcher's book [16] is more than sufficient. In fact, if the reader is willing to take a few of the theorems listed in Appendix A on faith, the entire book can be understood after completing a course on fundamental groups and covering spaces.

²In fact, extensive volumes already exist or are nearing completion on several topics that are mentioned here only in passing. See the Epilogue for an extended discussion of these additional resources.

Structure of the Text

The structure of the text is relatively straightforward. After a prologue designed to whet the reader's appetite, there is one introductory chapter, two chapters that present the core philosophy behind geometric group theory, two chapters that examine the special role played by hyperbolic metrics, and two chapters that cover more advanced topics. Finally, there is an epilogue that tries to ease the transition into the research literature, and an appendix that reviews those aspects of basic algebraic topology that serve as a foundation for the subject.

Acknowledgements

(These acknowledgements listed below are as preliminary and incomplete as the book itself.) Many people have had a hand in shaping this text, but foremost among them has been John Meier. Early on John and I had many long and fruitful conversations about structure, content, level and tone, and many of his ideas have been incorporated into the final text. A second source of inspiration has been (and will be) the UC Santa Barbara graduate students in the courses based on this material in Spring 2005 and Fall 2009. And finally, the dedication (of course) will be to my partner Mary Bucholtz.

Jon McCammond

Santa Barbara, CA
September 2009

Prologue: Trefoil Knot Group

The simplest non-trivial knot is the trefoil knot shown in Figure 1. As a way to introduce the flavor of geometric group theory we ask: *What can we say about the fundamental group of its complement?* Our primary goal is to illustrate how geometric arguments can be used to prove purely algebraic results. The arguments are merely sketched, but the reader should be able to go back and fill in the details as they work their way through the text.




FIGURE 1. The trefoil knot.

To establish notation, let K denote the knot shown and let G be the group $\pi_1(\mathbb{S}^3 \setminus K)$. (For technical reasons it is cleaner and more symmetric to work in \mathbb{S}^3 , the 1-point compactification of \mathbb{R}^3 , than in \mathbb{R}^3 itself.) The first thing to notice is that G is the fundamental group of a compact two-dimensional complex. To show this we construct a 2-complex \mathcal{D} inside $\mathbb{S}^3 \setminus K$ and then deform $\mathbb{S}^3 \setminus K$ down to \mathcal{D} . Since deformation retractions do not alter fundamental groups, $G = \pi_1(\mathbb{S}^3 \setminus K) \approx \pi_1(\mathcal{D})$. There are two common constructions for retracting arbitrary knot complements onto two-dimensional subspaces, usually attributed to Dehn and Wirtinger. Since we are using Dehn's procedure, the final result is known as a *Dehn complex*.

The Dehn complex for the trefoil knot shown in Figure 1 has two vertices, five edges and three 2-cells. To construct it we think of K as living in a small neighborhood of the xy -plane (or rather in a small neighborhood of its 1-point compactification, an equatorial 2-sphere inside \mathbb{S}^3), and we place a vertex v_+ above K and a vertex v_- below K . Next, we add an edge for each of the five regions of the xy -plane determined by the projection of K . More concretely, the regions in Figure 1 have been numbered and we add an edge E_i that connects v_- and v_+ , oriented from v_- to v_+ , passing through region i . Finally, we add a 2-cell for each of the three crossings. Each 2-cell is a square folded to look like Figure 2. The boundary of the square is then identified with the four edges corresponding to the four adjacent regions. The crossing at the top left of Figure 1, for example, creates




FIGURE 2. The 2-cell at a crossing

a square whose boundary follows the path $E_0^{-1}E_3E_1^{-1}E_2$. This can be interpreted as claiming that the loop $E_0^{-1}E_3$ based at v_+ is homotopic to the loop $E_2^{-1}E_1$. The other two 2-cells are attached along $E_0^{-1}E_4E_1^{-1}E_3$ and $E_0^{-1}E_2E_1^{-1}E_4$, respectively. See Figure 3.

The deformation retraction from $\mathbb{S}^3 \setminus K$ to \mathcal{D} alluded to above expands away from K like adding air into a long balloon. Parts of this retraction are easy to visualize. In Figure 2, for example, the complement of K inside this tent clearly retracts onto the folded square. Piecing together these local pictures, we find that G is the fundamental group of \mathcal{D} and, as a consequence, that G acts freely on its universal cover $\tilde{\mathcal{D}}$.




FIGURE 3. The three 2-cells in \mathcal{D} . The open circles represent v_+ and the closed circles represent v_- . To make these 2-cells look like the one shown in Figure 2, fold up along the dashed lines and down along the dotted ones.

The next key idea is that *if we understand the geometry of $\tilde{\mathcal{D}}$ and the way G acts on it, then we gain insight into the algebraic structure of G as a group.* The geometry of $\tilde{\mathcal{D}}$ is quite elegant. Since \mathcal{D} contains only three 2-cells, $\tilde{\mathcal{D}}$ has only three equivalence classes of 2-cells under the action of G . For convenience we refer to these as the green, yellow and blue 2-cells, reading left to right in Figure 3. The edges E_0 and E_1 are both contained in all three 2-cells, while the other three edges only occur in two of the three 2-cells. In fact, if you fix a particular lift of the green 2-cell in $\tilde{\mathcal{D}}$, oriented as shown, then there is a unique yellow 2-cell below it,

followed by a unique blue 2-cell, followed by a unique green 2-cell, and so on. The two sides of this infinite strip consist of lifts of the edges E_0 and E_1 , alternating on both sides. With a bit more work one sees that the local structure of $\tilde{\mathcal{D}}$ looks like Figure 4 and that as a topological space $\tilde{\mathcal{D}}$ is homeomorphic to the direct product of the real line and an infinite, trivalent tree.

Actually, even more is true. We can add a metric to $\tilde{\mathcal{D}}$ by making each 2-cell isometric to a unit Euclidean square. The metric space $\tilde{\mathcal{D}}$ still splits as a direct product, this time of the real line with the standard metric and a metric trivalent tree where each edge has length 1. In other words, if we let \mathcal{T}_3 denote the infinite, trivalent tree with edges of unit length, then $\tilde{\mathcal{D}}$ is isometric to $\mathcal{T}_3 \times \mathbb{R}$. The action of G preserves the metric as well as the product structure on $\mathcal{T}_3 \times \mathbb{R}$ so that by projecting onto the first or the second factor, the group G acts by isometries on \mathcal{T}_3 and it acts by isometries on \mathbb{R} .




FIGURE 4. The local structure in $\tilde{\mathcal{D}}$.

The last bit of preparation we need is to find a presentation of the group G . Any $g \in G$, acting on $\tilde{\mathcal{D}}$, takes lifts of v_+ to lifts of v_+ . In order to get a generating set for G it suffices to pick enough elements of G so that any lift of v_+ can be moved to any other using some composition of the actions of these elements and their inverses. Let v be a fixed lift of v_+ in $\tilde{\mathcal{D}}$, and let a , b and c represent the unique elements of G that move v diagonally up and across the unique green, yellow and blue squares, respectively, that have v as a bottom corner. To see that a , b and c generate G , let V be the orbit of v under the action of the subgroup generated by a , b and c . Suppose that u is in V , g is the element of G that sends v to u , and u' is a lift of v_+ connected to u along the diagonal of a single square. Since g sends the vertices connected to v by a diagonal to the vertices connected to u by a diagonal, we can find an element h in the set $\{a, b, c, a^{-1}, b^{-1}, c^{-1}\}$ so that gh sends v to u' . Notice that we are *precomposing* g with h which involves *right* multiplication by h . This is because we always assume that our groups act on the left. See Appendix A. In any case, this shows that the vertices in V are closed under adjacency. Geometrically, it is now clear that every lift of v_+ lies in V , and thus a , b and c generate all of G .

Finally, suppose that v is the open circle on the bottom of Figure 4 slightly to the right of center. The reader can check that the words ab , bc and ca all move v to

the open circle directly above it, so these words all represent the same element in G . When projected to \mathcal{D} , the products ab , bc and ca are homotopic, and represent the loop that passes through the central region of the trefoil knot and then returns to v_+ via the exterior region. With only a bit more work, one can show that the presentation $\langle a, b, c \mid ab = bc = ca \rangle$ is a presentation for G .

The group G acts freely and cocompactly on a contractible complex $\tilde{\mathcal{D}}$. We understand the structure of $\tilde{\mathcal{D}}$ and the action of G , and we have “words” we can use to describe the elements of G . We can now establish the following:

THEOREM. *If G is the fundamental group of the trefoil knot complement then*

1. *We can efficiently determine whether a word in the generators represents the identity;*
2. *The group G contains no nontrivial element of finite order;*
3. *The kernel of the map $f : G \rightarrow \mathbb{Z}$ sending a , b and c to $1 \in \mathbb{Z}$ is a free group of rank two;*
4. *The element $z = (ab)^3 = (bc)^3 = (ca)^3$ is central in G ;*
5. *The group G contains a finite index subgroup isomorphic to $\mathbb{F}_2 \times \mathbb{Z}$;*
6. *The group G is residually finite, meaning that the intersection of all finite index subgroups of G is the trivial subgroup $\{1\}$.*
7. *The element z generates the center, so that $Z(G) = \langle z \rangle$;*
8. *The quotient $G/Z(G)$ is isomorphic to $PSL_2(\mathbb{Z})$;*

How are such claims established? We outline one approach and completely ignore the technical details.

SKETCH OF PROOF: To tell whether an element in the generators represents the identity, simply trace out its effect on a lift of v_+ inside $\tilde{\mathcal{D}}$. If this lift ends where it started then this word represents the identity; otherwise, it does not. The reason this works is because the process of constructing the universal cover $\tilde{\mathcal{D}}$ secretly encodes a solution to the “word problem” for G . See Chapter 3 for details.

To prove 2 we combine the action of G on the factors of $\tilde{\mathcal{D}} \approx \mathcal{T}_3 \times \mathbb{R}$ with the fact that any finite order isometry of a metric tree must fix a point. (This fact is proved in Chapter 5.) Thus, any $g \in G$ of finite order fixes a point in \mathcal{T}_3 and it fixes a point in \mathbb{R} , so it fixes a point in $\tilde{\mathcal{D}}$. But the action of G on $\tilde{\mathcal{D}}$ is free so g is the identity.

Let H be the kernel of the map $f : G \rightarrow \mathbb{Z}$ described in item 3. Since the action of H on $\tilde{\mathcal{D}}$ projects to a free action on the tree \mathcal{T}_3 , the fundamental group of the quotient of \mathcal{T}_3 by this action is isomorphic to H . This quotient has two vertices, three edges and its fundamental group is \mathbb{F}_2 .

The action of the element $z = (ab)^3 = (bc)^3 = (ca)^3$ on $\tilde{\mathcal{D}}$ is a rotationless vertical translation. It can then be checked that pre- and post-composing any $g \in G$ with z results in the same action on $\tilde{\mathcal{D}}$, hence both expressions describe the same element of G . (Actually, it is sufficient to check that this is true for a , b , and c since they generate G .) This proves 4. Item 5 is now immediate since the subgroup generated by H and z is isomorphic to $\mathbb{F}_2 \times \mathbb{Z}$ and index 6 in G .

Next, the easiest way to prove 6 is to combine item 5 with two easily proved facts: free groups are residually finite, and the class of residually finite groups is closed under direct product and finite extension.

To prove 7, let v be a particular lift of v_+ inside $\tilde{\mathcal{D}}$ and consider the orbit of v under the action of $Z(G)$. Because of the symmetry of the situation with respect to a , b and c , the orbit of $Z(G)$ must be invariant under a $2\pi/3$ rotation around the vertical line through v . On the other hand, since the free group of rank 2 has trivial center, $Z(G) \cap H$ only contains the identity element and thus the orbit of v can have at most one element at each height. Combining these two ideas shows that 1) the orbit of v under the action of $Z(G)$ is contained in the vertical line through v , and 2) $Z(G)$ must be generated by the element that produces the smallest possible positive vertical change when applied to v . By 4, z is central and it moves v up six steps. There are exactly two elements of G , namely ab and $(ab)^2$, that move v to a lift of v_+ that is both on the vertical line through v and between v and its image under z . After checking that ab and $(ab)^2$ are not central, we conclude that $Z(G) = \langle z \rangle$.

Finally, to prove 8 we note that the action of $Z(G)$ on the \mathcal{T}_3 factor is trivial. Thus, we get a well-defined action of the quotient group $G/Z(G)$ on the trivalent tree \mathcal{T}_3 . There is a well-known action of $\text{PSL}_2(\mathbb{Z})$ on \mathcal{T}_3 , and, by comparing the two actions, we can see that the groups are identical. \square

Exercises

1. (Details) Fill in as many of the details of the proof of the theorem as you can. Alternatively, make a list of the arguments that seem unclear to you or imprecise at this point.
2. (Figure 8 knot) Let K be the knot shown in Figure 5.
 - a. Construct the Dehn complex \mathcal{D} for K .
 - b. Draw a small portion of $\tilde{\mathcal{D}}$ and try to understand its structure. Be forewarned that this is more difficult than it was for the trefoil knot.




FIGURE 5. The figure 8 knot

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:*

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

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 $f : \mathbb{S}^1 \rightarrow X$ 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. \square

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). \square

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:*

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

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 $f : X' \rightarrow \mathbb{R}^{2n+1}$ 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. \square

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 $f : [0, 1] \rightarrow X$ 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 $r : X \rightarrow A$ 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. \square

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 $\chi(X) < 1$. \square

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).

¹Recall 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 occurrences strictly shortens its length. \square

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 . \square

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. \square

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. \square

1.2.2. Graphs. These characterizations 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. \square

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 \tilde{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 \tilde{X} starting at one vertex u and both ending at a different vertex v , contradicting Theorem 1.2.8. \square

The simplest graphs are those with only one vertex (Figure 2). Such a graph is called 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 $\tilde{\chi}(R_A) = -|A|$ and $\tilde{\chi}(R_B) = -|B|$ 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 $|A| = |\mathbb{F}_A| = |\mathbb{F}_B| = |B|$. \square

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}_κ 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 $|\tilde{\chi}(X)| = |8 - 12 - 1| = 5$ 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. \square

THEOREM 1.2.14 (Free quotients). *Every group is a quotient of a free group. In particular, if (X, x) is a based connected cell complex with $G = \pi_1(X, x)$ and \mathbb{F} is the free group $\pi_1(X^{(1)}, x)$, then the group homomorphism $\mathbb{F} \rightarrow G$ 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. \square

THEOREM 1.2.15 (Finite rank calculations). *If X is a finite connected graph, then the rank of $\pi_1(X)$ is $|\tilde{\chi}(X)| = -\tilde{\chi}(X)$. 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 $\tilde{\chi}(Y) = -l$ and $\tilde{\chi}(X) = -k$ 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 \rightarrow 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 $f : S \rightarrow T$ 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 concatenation 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 $A = \{a_i \mid i \in [n]\} = \{a_1, a_2, \dots, a_n\}$ 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 $q : X \rightarrow X/T$ is thus a map $q_* : \pi_1(X, x) \rightarrow \mathbb{F}_A$. 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. \square

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 $q : U \rightarrow V$ denotes the quotient map, then for every immersed path $f : I \rightarrow V^{(1)}$ there is an immersed path $g : I \rightarrow U^{(1)}$ such that $q(g)$ traces the same path as f . As a consequence, the induced map $q_* : \pi_1(U) \rightarrow \pi_1(V)$ 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 $g : I \rightarrow U^{(1)}$

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 $q_* : \pi_1(U, u) \rightarrow \pi_1(V, v)$ is onto since every non-trivial element in $\pi_1(V, v)$ can be represented by an immersed loop $f : I \rightarrow V^{(1)}$ 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.

²This 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 $f : G \rightarrow \mathbb{F}_A$ 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 $f : A \rightarrow H$, there exists a unique extension of f to a group homomorphism $G \rightarrow H$. 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 $f : S \rightarrow H$, there exists a unique extension of f to a group homomorphism $G \rightarrow H$.

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 $f : A \rightarrow B$ is a bijection, then the unique homomorphism $G \rightarrow H$ 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 $f : S \rightarrow T$ is a symmetry-preserving bijection, then the unique homomorphism $G \rightarrow H$ extending f is an isomorphism.*

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

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 $f : S_A \rightarrow H$ there exists a unique group homomorphism $\mathbb{F}_A \rightarrow H$ 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 $\mathbb{F}_A \rightarrow H$ clearly extends f . It only remains to prove that this map is unique. Let $g, h : \mathbb{F}_A \rightarrow H$ 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 . \square

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 $(R_A, *) \rightarrow (X, x)$,*
2. *group homomorphisms $\mathbb{F}_A \rightarrow G$,*
3. *symmetry-preserving functions $S_A \rightarrow G$, and*
4. *functions $A \rightarrow G$.*

PROOF. There are easy conversions among the collections that we call restrict, construct, and induce. Any group homomorphism $\mathbb{F}_A \rightarrow G$ can be restricted to its symmetric basis S_A and then further restricted to its basis A . Given any function $A \rightarrow G$, a based map $(R_A, *) \rightarrow (X, x)$ 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 $f : \mathbb{F} \rightarrow \mathbb{F}$ 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 $f : \mathbb{F} \rightarrow \mathbb{F}$ 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 $A \rightarrow B$ 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 $A \rightarrow B$ 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 $\mathbb{F} \rightarrow \mathbb{F}$ can be represented by an equivalence class of based maps $(R_A, *) \rightarrow (R_A, *)$ 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 $\mathbb{F}_A \rightarrow \mathbb{F}_A$ corresponds to an equivalence class of based maps $(R_A, *) \rightarrow (P, p)$ or $(R_A, *) \rightarrow (H, h)$, 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 \tilde{K} and the Cayley graph $\mathcal{C}_{\mathcal{A}}(\Gamma)$ is the 1-skeleton of \tilde{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 $f : \mathbb{F}_A \rightarrow G$ is onto, then the function f (or, equivalently, its restriction $A \rightarrow G$) 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 $f : \mathbb{F}_A \rightarrow G$ 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 $A \rightarrow G$ generates G and B is the image of A , then the inclusion $B \hookrightarrow G$ also generates G : the corresponding homomorphism $\mathbb{F}_B \rightarrow G$ 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 $\mathbb{F} \rightarrow G$ 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 $\pi_1(X, *) \cong G$. \square

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 $\{(g_0 \cdot \mathcal{F}), (g_1 \cdot \mathcal{F}), (g_2 \cdot \mathcal{F}), \dots, (g_n \cdot \mathcal{F})\}$ 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. \square

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 $\mathbb{F}_A \rightarrow G$.*

PROOF. Because the action of G on Γ is free, the quotient map $p : \Gamma \rightarrow R_A$ 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 $\mathbb{F}_A \twoheadrightarrow \mathbb{F}_A/p_*(K) \cong G$. \square

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 \tilde{X} is both free and cocompact, so $2 \Rightarrow 3$. Finally, free actions are proper and graphs are cell complexes, so $3 \Rightarrow 4$. \square

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 an 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 $S_A \rightarrow S_B$ between their symmetric bases that extends to an isomorphism $i : \mathbb{F}_A \rightarrow \mathbb{F}_B$ 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 $f_A : \mathbb{F}_A \rightarrow G$ and $f_B : \mathbb{F}_B \rightarrow G$, respectively. These complexes are considered (very roughly) equivalent if there is a homeomorphism $R_A \rightarrow R_B$ between their 1-skeletons that induces an isomorphism $i : \mathbb{F}_A \rightarrow \mathbb{F}_B$ 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 $f : \Gamma \rightarrow \Gamma'$ such that for all $g \in G$ and for every cell $\sigma \in \Gamma$, $f(g \cdot \sigma) = g \cdot f(\sigma)$. \square

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.

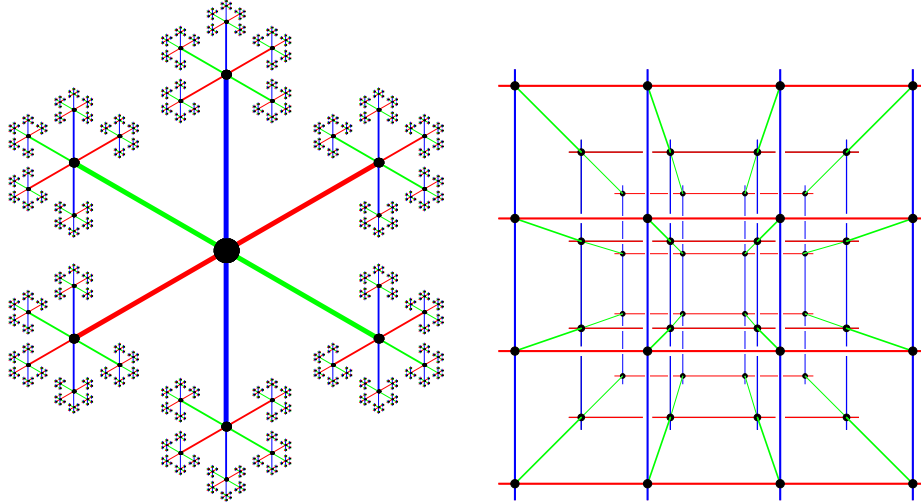


FIGURE 4. Portions of Cayley graphs for \mathbb{F}_3 (left) and \mathbb{Z}^3 (right).

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 $i \in \{0, 1, 2, \dots, n\}$. 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 $f : \mathbb{F}_A \rightarrow G$ 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 \tilde{X}). Let X be a cell complex with $X^{(1)} = R_A$ and $\pi_1(X, *) = G$, let \tilde{X} be its universal cover, and let \tilde{x} be one of its vertices. The group G acts on \tilde{X} as the group of deck transformations of

the projection map $p : \tilde{X} \rightarrow X$ and this action can be used to induce a vertex labeling and a symmetric edge labeling on the 1-skeleton of \tilde{X} . Since the action of G on \tilde{X} is free and vertex-transitive, for each vertex v in \tilde{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 \tilde{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 \tilde{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 \tilde{X} are indexed by $G \times S_A$. If we use $[a]$ to denote the image of a under the induced map $f : \mathbb{F}_A \rightarrow G$, 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 $\tilde{X}^{(1)}$ in a way that only depends on the map $f : \mathbb{F}_A \rightarrow G$ 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 $A \rightarrow G$ 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 $f : \mathbb{F}_A \rightarrow G$ since the edges record what happens when you right multiply 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 \tilde{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 \tilde{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 \tilde{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 \tilde{X} , it can usually be translated into algebraic information about the group G .

REMARK 1.3.18 (Encoding the group action). If $A \rightarrow G$ generates G and Γ is the Cayley graph for G with respect to A as constructed 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 $\tilde{\mathcal{D}}$ 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 $\tilde{X}^{(1)}$. 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 $\mathbb{F}_A \rightarrow G$ 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 $p : \tilde{X} \rightarrow X$. 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 is introduced to make this precise.

DEFINITION 1.3.22 (Multisets). Let S be a set. A *multiset selected from S* is a function $m : S \rightarrow \mathbb{N}$, 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 $m : S \rightarrow \mathbb{N}$ 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 $g : \mathbb{F}_A \rightarrow \pi_1(X, *)$ 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 $\langle a \mid a^n \rangle$. 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 orient 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 be 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 $\mathcal{P} = \langle A \mid \mathcal{R} \rangle$ 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 converted into an 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 $Y = \tilde{X}$. 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. \square

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 \tilde{X}_α (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 $\tilde{X}_\alpha \rightarrow X_\alpha$. Call the resulting space Y_2 and the map $p_2 : Y_2 \rightarrow Y$. 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 \tilde{X}_α 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 $p_3 : Y_3 \rightarrow Y$ 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 \tilde{X}_α . At each such vertex we attach a copy of the appropriate \tilde{X}_α 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 $p : Y' \rightarrow Y$ 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 $f : I \rightarrow U^{(1)}$ 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 \tilde{Y} .




FIGURE 8. Attaching simply-connected spaces.

LEMMA 1.4.3 (Structure of \tilde{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 $p: \tilde{Y} \rightarrow Y$ is its universal cover, then each component of $p^{-1}(X_\alpha)$ is a subcomplex homeomorphic to \tilde{X}_α 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 \tilde{Y} . In fact, \tilde{X} can be obtained from \tilde{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 \tilde{Y} in this way, it is easy to see that the composition $\tilde{Y} \rightarrow Y \rightarrow X$ factors through X' to produce a map $X' \rightarrow X$ that is a local homeomorphism and thus a cover. Finally, by Proposition 1.2.21 the trivial fundamental group of \tilde{Y} maps onto the fundamental group of X' making X' simply-connected and the universal cover of X . This establishes the local structure of \tilde{X} .

COROLLARY 1.4.4 (Structure of \tilde{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 $p: \tilde{Y} \rightarrow Y$ be its universal cover. If each component of $p^{-1}(S)$ is collapsed to a point then the resulting complex is \tilde{X} . As a consequence, the inclusion map $X_\alpha \hookrightarrow X$ lifts to an inclusion $\tilde{X}_\alpha \hookrightarrow \tilde{X}$.*

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 \tilde{X}_α embeds in \tilde{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 $i_\alpha: \pi_1(X_\alpha, x_\alpha) \rightarrow \pi_1(X, x)$ 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 \tilde{X}_α inside \tilde{X} proving that $i_\alpha(g)$ is nontrivial. \square

In addition to collapsing onto \tilde{X} , the space \tilde{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 $p: \tilde{Y} \rightarrow Y$ 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 $\pi_1(\tilde{Y})$ (Proposition 1.2.21). By Theorem 1.2.1 the quotient is a tree. \square

Let T denote the tree obtained from \tilde{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.

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

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 \tilde{Y} by the fundamental theorem of covering spaces and it acts on \tilde{X} and T because the horizontal quotient maps commute with the G -action on \tilde{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 \tilde{X} is free, the map $\tilde{X} \rightarrow X$ is a cover; the map from $T \rightarrow S$ 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 \tilde{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 $p : \tilde{Y} \rightarrow Y$ and the quotient map $q : \tilde{Y} \rightarrow T$ we can extend this partitioning of vertices to \tilde{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 $\tilde{y} \in \tilde{Y}$ as our base point, and, as in Definition 1.3.14, we then use the G -action to label each backbone vertex of \tilde{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 \tilde{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 \tilde{X}_α (Lemma 1.4.3) and thus simply-connected. \square




FIGURE 9. The correspondence between 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 $f : I \rightarrow X_\alpha^{(1)} \subset X$ 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 \tilde{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 \tilde{Y} only crosses two edges in the backbone, the projected path to T has length 2. Because g is non-trivial, the lift to \tilde{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 \tilde{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 \tilde{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 arbitrary 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 canonically 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 identified 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 $F : \partial E \rightarrow X^{(1)}$ 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 $F : \partial E \rightarrow X$ 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 $E \rightarrow X$ 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 $E \rightarrow Y$ is a factor of $E \rightarrow X$, 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. \square

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 $Y \rightarrow X$ 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 $E \rightarrow X$. 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 $f_n : \mathbb{S}^1 \rightarrow \mathbb{S}^1$ 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 $f : \mathbb{S}^1 \rightarrow \mathbb{S}^1$ is any map homotopic to f_n , then f is called a *degree n map*. It is easy to show that every map $\mathbb{S}^1 \rightarrow \mathbb{S}^1$ 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 $[a^m t = t b^n]$ 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) $\langle a, b \mid ababa\dots = babab\dots \rangle$




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 $[a^m b = b a^n]$. Topologically they are the fundamental 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.

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 $f : X \rightarrow Y$ 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 $f : R_A \rightarrow R_B$ iff $|A| = |B|$.
 - a. Prove \exists a π_1 -surjection $f : R_A \rightarrow R_B$ iff $|A| \geq |B|$.
 - b. Prove \exists a π_1 -injection $f : R_A \rightarrow R_B$ iff $|\mathbb{F}_A| \leq |\mathbb{F}_B|$.
7. (Tree removal) Let T be a tree in a graph X and let $q : X \rightarrow X/T$ 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}_κ 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 $f : G \rightarrow \mathbb{F}_A$ 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 contained 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 $\mathbb{F}_\kappa \cong \text{INN}(\mathbb{F}_\kappa) \subset \text{AUT}(\mathbb{F}_\kappa)$ 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}_κ are in one-to-one correspondence with the well-orderings of the bases of \mathbb{F}_κ 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 $A \rightarrow \mathbb{Z}$ with only finitely many non-zero values). Show that the abelianization of \mathbb{F}_A is \mathbb{Z}^A and that the abelianization map $\mathbb{F}_A \rightarrow \mathbb{Z}^A$ 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 $\text{AUT}(\mathbb{F}_A)$ to $\text{AUT}(\mathbb{Z}^A)$ 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 $|\tilde{\chi}(X)| = |8 - 12 - 1| = 5$), 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 folding 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 Hawaiian 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 arbitrary 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.

$$\langle a_1, a_2, \dots, a_n \mid a_1 a_2 \cdots a_n = a_n \cdots a_2 a_1 \rangle$$

$$\langle a_1, a_2, \dots, a_n \mid a_1^2 a_2^2 \cdots a_n^2 = 1 \rangle$$

In the first case, it is easy to check that these surfaces are orientable (because the two occurrences 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 $2n + 1$.

CHAPTER 2

Metrics on Groups

There is a long mathematical history—dating back at least to the late 1800s and Felix Klein’s Erlanger Programme—of exploring the connections between geometric spaces and their groups of isometries. Here are two striking results that indicate the influence that geometry can have on group theory:

1. If G acts properly discontinuously and cocompactly on \mathbb{R}^n by isometries, then G contains a finite index subgroup isomorphic to \mathbb{Z}^n . (Bieberbach)
2. If G acts properly discontinuously and cocompactly on hyperbolic space \mathbb{H}^n by isometries, then G does not contain a subgroup isomorphic to \mathbb{Z}^2 . (Gromoll and Wolf?)

Even if you have only a passing familiarity with hyperbolic spaces (we describe them in Chapter 4) the message should be clear: If a group admits a properly discontinuous and cocompact action on a particularly nice geometry, then there are algebraic consequences. Although such results may seem to be special to particular geometries, if one adopts a fairly flexible notion of what constitutes a “geometry” then the techniques and results hinted at in these two examples apply broadly to infinite groups. In this chapter we introduce such a broad notion of a geometry (due to Jim Cannon) along with the appropriate types of maps between such geometries: quasi-isometries. The goal is to prove two fundamental results: first, that every finitely generated group acts on a geometry, and second that this geometry is unique up to quasi-isometry.

2.1. Metrics and quasi-isometries

DEFINITION 2.1.1 (Isometries). An *isometric embedding* of one metric space into another is an injective map $f : X \rightarrow Y$ such that $d_X(x, y) = d_Y(f(x), f(y))$ for all $x, y \in X$. If f is also onto, then f is an *isometry* and X and Y are said to be *isometric*.

DEFINITION 2.1.2 (bi-Lipschitz equivalence). Two metric spaces X and Y are *bi-Lipschitz equivalent* if there is a function $f : X \rightarrow Y$ and a fixed constant $K \geq 1$ such that

1. The function is a *bi-Lipschitz embedding* meaning that for any two points $x, x' \in X$, one has

$$\frac{1}{K}d_X(x, x') \leq d_Y(f(x), f(x')) \leq K \cdot d_X(x, x').$$

2. The function f is *quasi-onto* meaning that for any $y \in Y$, there is an $x \in X$ such that $d_Y(f(x), y) \leq K$.

When it is interesting to note the constant, one refers to *K -bi-Lipschitz* maps.

We leave the proof of the following lemma and a related result (Proposition 2.1.5) to the exercises.

LEMMA 2.1.3. *The notion of bi-Lipschitz equivalence is an equivalence relation.*

nets

The group \mathbb{Z}^n has a natural geometry upon which it acts geometrically, namely, \mathbb{R}^n . The also acts geometrically on its Cayley graphs. While any two Cayley graphs are bi-Lipschitz equivalent, none are bi-Lipschitz equivalent to \mathbb{R}^n ($n \geq 2$). In order to pass between discrete subgroups and extensions, we need a slightly broader notion of equivalence.

DEFINITION 2.1.4 (Quasi-isometry). A function $f : X \rightarrow Y$ between metric spaces is a K -quasi-isometric embedding if for all $x, x' \in X$,

$$\frac{1}{K} \cdot d_X(x, x') - K \leq d_Y(f(x), f(x')) \leq K \cdot d_X(x, x') + K$$

If f is also K -quasi-onto, then f is a quasi-isometry.

The way that geometric group theorists often talk about these constants is that multiplying and dividing by K means a bounded amount of “stretching” and the addition/subtraction of K means a bounded amount of “tearing”.

PROPOSITION 2.1.5 (Quasi-isometry is an equivalence relation). *The relation of being quasi-isometric is an equivalence relation on (qualified) metric spaces. In particular, the identity map is a quasi-isometry, every quasi-isometry has a quasi-inverse, and the composition of two quasi-isometries is a quasi-isometry.*

One of the more difficult steps in establishing this proposition is the construction of the quasi-inverse. Let $f : X \rightarrow Y$ be a K -quasi-isometry. Since f is K -quasi-onto, every $y \in Y$ is contained in a metric ball $B_K(f(x))$ for at least one x in X . Define $f^{-1}(y)$ to be any such x . The reader who objects that such a map is almost guaranteed to be discontinuous has not yet grasped the rough nature of quasi-isometries! We leave the verification that $f^{-1} : Y \rightarrow X$ is a quasi-isometry, as well as the proof of the rest of this proposition, as an exercise.

2.2. Geometries and geometric actions

Of particular importance is the case where the metric space being embedded is an interval of the reals.

DEFINITION 2.2.1 (Embeddings and geodesics). An isometric embedding of an interval $[a, b] \subset \mathbb{R}$ into X is a *geodesic segment*¹. If the embedding is $f : [a, b] \rightarrow X$ then this is a *geodesic segment of length $b - a$ from $f(a)$ to $f(b)$* . Let (X, d) be a fixed metric space, let x and y be fixed distinct points of X , and consider the infimum of the lengths of paths from x to y . When there is at least one such path, it is clear that this infimum exists. The space (X, d) is a *path metric space* if the distance between two points equals the infimum of path lengths, and this infimum is realized by at least one geodesic.

¹If you studied geodesics in a course on differential geometry, you probably recall that they were defined to be locally length minimizing curves or a curve that satisfies a certain differential equation. The global length minimizing property used in this definition is often presented as a consequence of certain curvature conditions.

lem: bilip-equiv

Explain Cannon's alternative definition or put it in an exercise.

DEFINITION 2.2.2 (Geometries and geometric actions). Let (X, d) be a metric space. The *closed ball of radius $r \in \mathbb{R}^+$ centered at $x \in X$* is the set

$$B_{\leq r}(x) = \{y \in X \mid d(x, y) \leq r\}.$$

The space X is *proper* if $B_{\leq r}(x)$ is compact for every choice of r and x . A proper metric space whose metric coincides with the path metric, is a *geometry*. An action $G \curvearrowright X$ of a group on a geometry is *geometric* if it is cocompact and properly discontinuous.

This definition includes the classic geometries—Euclidean, spherical and hyperbolic spaces—but it is a good deal more broad than that. It includes each of the eight geometries that Thurston championed for the study of 3-manifolds as well as many cell-complexes endowed with metrics, as we discuss below.

DEFINITION 2.2.3 (Graph metric). Any connected graph (i.e. any connected 1-dimensional cell complex) can be turned into a metric space by declaring that every 1-cell has the same metric as the unit interval $[0, 1]$ in the reals. The distance between two points x and y in Γ is defined to be the minimum length of a path from x to y . Notice that when x and y are distinct vertices this minimal distance is a positive integer that counts the smallest number of edges traversed in a path from x to y .

It is an easy exercise to check that this is a metric: Because paths are reversible, the function is symmetric; because all edges are isometric, the infimum of path lengths between points is bounded away from zero; because paths can be concatenated, the function satisfies the triangle inequality.

EXAMPLE 2.2.4. Let G be a group with finite generating set S , and let $\Gamma = \text{CAYLEY}(G, S)$ be the associated Cayley graph. The graph metric turns Γ into a geometry and the action $G \curvearrowright \Gamma$ is geometric. Thus every finitely generated group acts geometrically on some geometry.

This metric on the Cayley graph is often conflated with the *word metric* on the group G . The word metric on G is defined as

$$d(g, h) = \|gh^{-1}\|$$

where $\|g\|$ denotes the minimal length of a word in S and S^{-1} that produces the element g . Restricting the graph metric to the vertices of Γ , and identifying the vertices with the elements of G , shows that these two metrics are identical. Note, however, that while Γ is a geometry, the word metric does not make G a geometry, as this metric space is not path connected

The graph metric is a basic example of taking a cell complex and forming a metric space by endowing each cell with a metric structure. In higher dimensions it is important to see some of the details of this construction.

DEFINITION 2.2.5. Let S be a collection of compact polyhedral subsets of any of the standard spaces of constant curvature—spheres, \mathbb{R}^n or \mathbb{H}^n —where we presume the set S is closed under the operation of taking faces. The set S is the *set of shapes*.

A *piecewise polyhedral complex* X is a disjoint union of a set of shapes, $\coprod S$, modulo an equivalence relation \sim , which satisfies certain conditions. Let $p : \coprod S \rightarrow$

$X = \coprod S / \sim$ denote the quotient map and p_σ the restriction to any particular polyhedron $\sigma \in S$. To insure that X is a regular cell complex and that the polyhedral metrics have been assigned consistently, we assume:

1. Each p_σ is locally injective; and
2. If $p_\sigma(\sigma) \cap p_{\sigma'}(\sigma') \neq \emptyset$ then there is an isometry ϕ from the support of σ to the support of σ' such that $p_\sigma(x) = p_{\sigma'}(x')$ iff $\phi(x) = x'$.

If all the shapes come from a single space of constant curvature, then X would be referred to as a *piecewise spherical, Euclidean, or hyperbolic complex* (depending on the source geometry).

The geometry of piecewise polyhedral complexes is developed nicely and quite completely in [4], so we will not repeat this work. We do however point out a couple of highlights that we will make frequent use of.

You can define the length of a path $\ell(p) \dots$ Define a pseudo-metric on X by $d(x, x') = \inf_p \{\ell(p) \mid p \text{ connects } a \text{ to } b\}$. The triangle inequality is immediate. In the absence of no further hypotheses, this is not a metric.

thm:finitelymanyspaces

THEOREM 2.2.6 (Finitely many shapes). *If X is a piecewise polyhedral complex whose set of shapes is finite, then the induced path pseudo metric is a metric.*

Here's the idea of the proof. For details, see [4].

We end this section with a few examples that we will frequently cite.

DEFINITION 2.2.7 (Hemisphere metric). Let X be an arbitrary combinatorial 2-complex. The 1-skeleton of X can be turned into a metric space by making each edge isometric to an open unit interval. Next, consider a 2-cell R in X and let $k \in \mathbb{Z}^+$ be the length of the combinatorial path along which it is attached. We assign a metric to R so that it is isometric to the northern hemisphere of a 2-sphere in \mathbb{R}^3 of radius r , where r has been chosen so that the length of the equator is exactly k . In other words, $r = k/2\pi$. Finally, the attaching map is the natural isometry between the equator of the hemisphere and the combinatorial path in $X^{(1)}$.

In §2.2 we establish that for any combinatorial 2-complex X these local metrics on the 1-cells and 2-cells combine to turn X in a geodesic metric space with the unusual property that the inclusion map $X^{(1)} \rightarrow X$ is an isometric embedding.

THEOREM 2.2.8. *If X is a taut 2-complex, then the local hemisphere metric defines a global geodesic metric on X .*

PROOF. (the key thing to prove is that every pair of points is connected by a length minimizing path.)

(observe that the metric 1-skeleton isometrically embeds; i.e. given two points in the 1-skeleton, the shortest path between them lies in the 1-skeleton.)

(the only paths that need to be consider are those that head straight to a vertex in the boundary of the cell we're in. This reduces things to combinatorial distances.) \square

(notice that this does NOT depend on something like a finitely many cell types assumption and it's true for arbitrary taut 2-complexes)

exmp:p-eucl-2-cplx

EXAMPLE 2.2.9. Every finite 2-complexes can be converted into a piecewise Euclidean complex. The most common method for doing this is to assign to each

2-cell the metric structure of a regular Euclidean n -gon, whose sides are of unit length². As X is finite, this can be done with finitely many shapes, hence X with the path metric is a metric space. Moreover, the universal cover \tilde{X} is a geometry that $\pi_1(X)$ acts on geometrically.

For example, let X be the 2-torus with the standard cellular structure with one vertex, two edges and a single 2-cell thought of as a 4-gon. Giving the 4-gon the metric structure of a unit Euclidean square turns X into a standard flat torus, and \tilde{X} is isometric to \mathbb{R}^2 .

While the piecewise Euclidean structure described above is often useful, there is a spherical variation of it with its own utility.

PROPOSITION 2.2.10 (2-complexes). *Let X be a finite 2-complex. Then X and its universal cover \tilde{X} admit a piecewise spherical metric making them geometries and the action of $\pi_1(X)$ on \tilde{X} is geometric. Further, this can be arranged so that the embedding $\tilde{X}^{(1)} \hookrightarrow \tilde{X}$ is an isometric embedding.*

prop:p-sph-2-cplx

PROOF. Each edge of X is given the metric structure of the unit interval. Since X is a combinatorial cell complex, each 2-disk can be identified with an n -gon, where n is the number of edges hit by the attaching map (counted with multiplicity). Give such a disk the metric structure of a hemisphere of radius $n/2\pi$, where the boundary loop has been divided into n unit intervals. Since the set of shapes is finite, this describes a metric space.

Let x and y be any two points in $\tilde{X}^{(1)}$, and let $p : [0, d] \rightarrow \tilde{X}$ be a geodesic joining them. Since any geodesic is a local geodesic, if the path traced by p ever leaves the 1-skeleton and enters a 2-disk D , then its first return to the 1-skeleton must occur at the opposite point on the boundary of D . But as D has a hemispherical metric, it would have been just as efficient to have travelled around the boundary. Thus any geodesic path in \tilde{X} connecting points in the 1-skeleton can be converted to a geodesic path contained in the 1-skeleton. \square

Giving cell complexes piecewise polyhedral metrics is not a process that is restricted to dimensions 1 and 2. In fact, Theorem 2.2.6 immediately establishes that any finite, connected simplicial complex can be converted into a geometry.

PROPOSITION 2.2.11. *If Σ is any finite, connected simplicial complex, then Σ admits a piecewise Euclidean structure where the metric, restricted to any simplex, is that of a regular Euclidean simplex whose edges have length 1. The universal cover $\tilde{\Sigma}$ is then a geometry that $\pi_1(\Sigma)$ acts on geometrically.*

prop:simplicialmetric

2.2.1. Metrics and bi-Lipschitz maps.

2.2.2. Quasi-isometries.

COROLLARY 2.2.12. *Let H be a finite-index subgroup of a finitely generated group G , hence H is also finitely-generated by Theorem 2.3.14. The intrinsic metric on H (with respect to any finite generating set) is quasi-isometric to the intrinsic metric on G (with respect to any finite generating set).*

illustrate using change of finite generating set

²In the case where $n = 1$ you can use a disk whose boundary has circumference 1; for $n = 2$ use a bigon formed by two arcs of unit circles, each arc of length 1.

PROOF. We show that the inclusion map $H \hookrightarrow G$ is a quasi-isometry. Since H is a finite index subgroup, there are coset representatives such that

$$G = H \cup g_1 H \cup \cdots \cup g_n H.$$

The inclusion is then k -quasi-onto, where $k = \text{Max}\{\|g_i\|\}$.

Because all finite generating sets lead to quasi-isometric metrics on G , we may assume that the generating set for G contains the generating set for H . So $d_G(h, h') \leq d_H(h, h')$.

As the Cayley graph is locally finite, there is a universal bound on the number of elements of H that can occur in a ball of radius $(k + 1)$ about any element of G . Thus there is a universal bound \hat{k} on the H -distance between any two. Let $d_G(h, h') = n$ and let p be a path connecting h to h' of length n , passing through the vertices $\{g_0 = h, g_1, g_2, \dots, g_n = h'\}$. For each g_i pick an element h_i where $d_G(g_i, h_i) \leq k$. It follows that $d_H(h_i, h_{i+1}) \leq \hat{k}$. Thus $d_H(h, h') \leq \hat{k}d_G(h, h')$. Thus the embedding of H in G is K -bi-Lipschitz for $K = \text{Max}\{k, \hat{k}\}$. Since $K \geq k$ it then follows that the embedding is a K -quasi-isometry. \square

DEFINITION 2.2.13 (Commensurable). Commensurability is the symmetric and transitive closure of the finite-index subgroup relation. Since a finite-index subgroup H of a finite-index subgroup K of a group G is itself a finite-index subgroup of G , groups G and G' are *commensurable* iff they can be connected by a finite sequence of groups $G = G_0, H_0, G_1, H_1, \dots, H_{n-1}, G_n = G'$ where for each $i = 0, \dots, n - 1$, H_i is isomorphic both to a finite-index subgroup of G_i and to a finite-index subgroup of G_{i+1} .

weakly commensurable

COROLLARY 2.2.14 (Commensurable implies quasi-isometric). *Let G and G' be commensurable groups. If G is finitely-generated then so is G' and, moreover, the intrinsic metrics on G and G' are quasi-isometric.*

REMARK 2.2.15. In general the commensurability relation is strictly weaker than being quasi-isometric. For example, consider two closed hyperbolic 3-manifolds with irrationally related volumes.

PROPOSITION 2.2.16 (Homotopy equivalence implies quasi-isometry). *If $f : X \rightarrow Y$ is a homotopy equivalence between compact connected cell complexes, then any lift $\tilde{f} : \tilde{X} \rightarrow \tilde{Y}$ is a quasi-isometry.*

PROOF. Pick g such that $fg \cong \mathbf{1}$ and $gf \cong \mathbf{1}$ and pick a lift of g so that $\tilde{f}\tilde{g}$ fixes a base point. (Look at the lifts of the paths traced out by the homotopy from fg to the identity) \square

The following is [7, Theorem 11.1]; you can essentially find it in a paper by Milnor published in JDG in 1968 where he discusses growth in groups and manifolds. (include the Charney-Crisp extension)

THEOREM 2.2.17 (Milnor-Švarc Theorem). *If a group G acts geometrically on two geometries X and Y then X and Y are quasi-isometric.*

PROOF. Since quasi-isometry is an equivalence relation, it suffices to prove that X is quasi-isometric to G with the word metric relative to some finite generating

set S . Fix a compact fundamental domain D for the action of G on X , let U be an ϵ -neighborhood of D for some $\epsilon > 0$, that is

$$U = \{x \in X \mid d(x, D) < \epsilon\}$$

and for convenience, pick

$$S = \{g \in G \mid g \cdot U \cap U \neq \emptyset\}.$$

Define a function $\xi : G \rightarrow X$ by choosing a point $x \in D$ and setting $\xi(g) = g \cdot x \in X$. Since the action of G on X is cocompact, the map ξ is quasi-onto. In fact, if δ is the diameter of U , that is $\delta = \text{Max}\{d(x, y) \mid x, y \in U\}$, then ξ is δ -quasi-onto.

To establish that ξ is bi-Lipschitz, pick a geodesic in $\Gamma = \text{CAYLEY}(G, S)$ connecting g to h . This geodesic passes through a sequence of vertices of Γ giving an associated sequence of elements in G , $\{g = g_0, g_1, g_2, \dots, g_n = h\}$, where $n = d(g, h)$ and $g_{i-1}s_i = g_i$ for some $s_i \in S$. As $s_i U \cap U \neq \emptyset$, the sequence of elements describes a chain of copies of U in X , $\{gU = g_0U, g_1U, \dots, g_nU = hU\}$, where $g_iU \cap g_{i+1}U \neq \emptyset$. The triangle inequality implies that $d(g_i \cdot x, g_{i+1} \cdot x) \leq 2\delta$ and further uses of the triangle inequality show that $d(g \cdot x, h \cdot x) \leq 2\delta \cdot d(g, h)$ (see Figure 1). In other words, $d(\xi(g), \xi(h)) \leq 2\delta \cdot d(g, h)$.




FIGURE 1. The proof of the Milnor-Švarc Theorem

Let $\gamma : [0, d] \rightarrow X$ be a geodesic segment connecting $g \cdot x$ to $h \cdot x$, where $d = d(g \cdot x, h \cdot x)$. Divide the domain into $m = \lceil (d + 1)/\epsilon \rceil$ equal pieces to form smaller geodesics γ_i , each of length $< \epsilon$. Since the length of each of these geodesic segments is $< \epsilon$, each is contained in a copy of U . Thus we have a chain of copies of U , $\{g_0U, g_1U, g_2U, \dots, g_mU\}$ where $g_{i-1}U \cap g_iU \neq \emptyset$. Thus for each i , $1 \leq i \leq m$, there is an $s_i \in S$ such that $g_{i-1}s_i = g_i$. Further, since g_0U contains $g \cdot x$, $g_0U \cap gU \neq \emptyset$, so either $g_0 = g$ or there is some $s_0 \in S$ such that $g_0s_0 = g$. Similarly either $g_m = h$ or $g_ms_{m+1} = h$ for some $s_{m+1} \in S$. Thus

$$d(g, h) \leq m + 2 < \epsilon^{-1}d(\xi(g), \xi(h)) + 3 + \lceil \epsilon^{-1} \rceil$$

hence

$$\frac{1}{\epsilon^{-1}}d(g, h) - (3 + \lceil \epsilon^{-1} \rceil) \leq d(\xi(g), \xi(h)).$$

It follows that if $K = \text{Max}\{\epsilon, 3 + \lceil \epsilon^{-1} \rceil, 2\delta\}$ then ξ is a K -quasi-isometry. \square

If G is the fundamental group of a finite cell complex, X , then the universal cover \tilde{X} admits piecewise polyhedral metrics (see 2.2.9 through 2.2.11). The action of G on \tilde{X} is then a geometric action on a geometry, hence we have the following corollary.

COROLLARY 2.2.18 (Fundamental groups and universal covers). *Let X be a compact connected metric cell complex whose fundamental group is G . Then the universal cover \tilde{X} is quasi-isometric to G with the word metric with respect to any finite generating set.*

Following Jim Cannon [cite] we call a property of a finitely generated group *geometric* if it is independent of which finite generating set is chosen. We call a property of a finitely generated group a *QI-invariant* if it is independent of quasi-isometry type.

finite-index subgroups, virtual properties, all finite groups have the same geometric and QI properties.

2.3. Geometric properties of groups

There are a number of naturally occurring ways to alter “the” word metric on a finitely generated group G . One is to switch finite generating sets: If S and S' both generate G , and there is a $g \in S'$ which is not in S , then $d_S(1, g) > d_{S'}(1, g) = 1$. Another is to keep the generating set fixed, but to assign positive weights $\omega(s)$ to each $s \in S$, where if s and s^{-1} are both contained in S , then $\omega(s) = \omega(s^{-1})$. One can then give the Cayley graph $\Gamma = \text{CAYLEY}(G, S)$ the metric where each edge associated to generator s carries the metric of the interval $[0, \omega(s)]$, and use this to describe a metric on G by restricting to the vertices. By Theorem 2.2.6 the induced path metric is a metric.³ Of course, one can also combine these two metric varying techniques.

Call any metric derived from choosing a finite generating set and assigning positive weights to the generators a *Cayley graph metric* on G . The question is to what extent these Cayley graph metrics on G are ‘equivalent’.

LEMMA 2.3.1 (Redundant generators). *Let G be a group and let S a generating set for G . If S' is a set consisting of S and exactly one additional element t , then the metrics d_S and $d_{S'}$ on G are K -bi-Lipschitz equivalent where $K = \|t\|_S$.*

PROOF. Since any path in $\text{CAYLEY}(G, S)$ is also a path in $\text{CAYLEY}(G, S')$, $d_S(g, g') \geq d_{S'}(g, g') \geq \frac{1}{K}d_{S'}(g, g')$ for all $g, g' \in G$. Conversely, let w be a word over S that represents $t \in S' \setminus S$ and that realizes $\|t\|$. Given any path p connecting g to g' in $\text{CAYLEY}(G, S')$, we can find a path connecting these vertices in $\text{CAYLEY}(G, S)$ that is at most K times as long by replacing each edge labelled t (or t^{-1}) with the path described by the word w (or w^{-1}).




FIGURE 2. The bi-Lipschitz equivalence of Cayley graphs

Thus $d_{S'}(g, g') \geq K \cdot d_S(g, g')$ for all $g, g' \in G$, establishing the bi-Lipschitz embedding. Because the map is a bijection, it is also quasi-onto. \square

³This second option is surprisingly useful at times, and mimics constructions used to study mapping class groups of surfaces by looking at the space of metrics on a given topological surface. This method of varying the metric on edges also shows up in Culler and Vogtmann’s construction of Outer Space, a geometry for $\text{Out}(\mathbb{F}_n)$.

LEMMA 2.3.2 (Rescaling). *Let G be a group with finite generating set S . Let d denote the standard word metric induced by S and let \hat{d} denote a word metric induced by some fixed assignment of positive weights to the elements of S . Then (G, d) and (G, \hat{d}) are bi-Lipschitz equivalent.*

PROOF. Let $m = \min\{\omega(s) \mid s \in S\}$ and let d_m denote the metric induced by assigning each edge in $\text{CAYLEY}(G, S)$ the metric of $[0, m]$. Then the metric d_m is just $m \cdot d$, hence d_m and d are bi-Lipschitz equivalent. By Lemma 2.1.3 it suffices to show that d_m and \hat{d} are bi-Lipschitz equivalent.

Define $K = \text{Max}\{\omega(s)/m \mid s \in S\}$ and let p be a path connecting g to g' in $\text{CAYLEY}(G, S)$. If $\ell_m(p)$ denotes the length of p with respect to the d_m metric, and $\hat{\ell}(p)$ denotes its length with respect to \hat{d} , then $\frac{1}{K}\ell_m(p) \leq \hat{\ell}(p) \leq K \cdot \ell_m(p)$. It follows that

$$\frac{1}{K}d_m(g, g') \leq \hat{d}(g, g') \leq K \cdot d_m(g, g').$$

Hence the identity map establishes the bi-Lipschitz equivalence. \square

The preceding lemmas combine to establish:

PROPOSITION 2.3.3. *If G is a group that admits a finite set of generators, then every finite generating set for G , with every rescaling, determines the exact same bi-Lipschitz equivalence class of metrics on G .*

With little more than the core definitions, one can find deep connections between properties of a geometry for a group G and topological properties of G . In this section we establish that finiteness properties of groups are geometric properties, and moreover, they are quasi-isometry invariants.

By Theorem 1.3.8 we know that a group G is finitely generated iff it acts freely and cocompactly on a connected graph. Similarly Theorem 1.3.29 shows G is finitely presented iff it acts freely and cocompactly on a 1-connected cell complex. The natural generalization of these finiteness properties is that a group is of *type \mathcal{F}_n* iff it is the fundamental group of a cell complex X with finite n -skeleton. (Thus being of type \mathcal{F}_1 is equivalent to being finitely generated, and \mathcal{F}_2 is the same as being finitely presented.) For further information on these and other finiteness properties see KEN BROWN'S BOOK.

In this section we focus on the property of being finitely presented. However, as the results in this section generalize to these \mathcal{F}_n , we state more general results, even though our arguments will be restricted to the case $n = 2$.

THEOREM 2.3.4 (Finiteness and connectivity theorem). *For $n \geq 0$, a group G acts geometrically on an $(n - 1)$ -connected geometry X iff G is of type \mathcal{F}_n .*

PROOF. The \Leftarrow direction follows immediately from the definition of \mathcal{F}_n and Proposition 2.2.11. For in this case, the group G acts freely and cocompactly on the n -skeleton of the universal cover of its $K(G, 1)$, and this universal cover can be given a metric making it into a geometry. The other direction is more subtle.

The argument for the base case, $n = 1$, was essentially given BEFORE. The idea is to take a fundamental domain $D \subset X$ and define

$$S = \{g \in G \mid gD \cap D \neq \emptyset\}.$$

As the action of G on X is geometric, S is finite. The fact that S is a set of generators for G follows from the fact that X is 0-connected.

We assume then that G is finitely generated and that G acts geometrically on a 1-connected geometry X . Let Γ be a Cayley graph for G with respect to some finite generating set. Our goal is to carefully add finitely many G -equivariant classes of 2-cells to Γ that kill $\pi_1(\Gamma)$. By the Milnor-Švarc Theorem (2.2.17), Γ and X are quasi-isometric. Let $f : \Gamma \rightarrow X$ be a quasi-isometry and let $F : X \rightarrow \Gamma$ be a quasi-inverse of f . We let K denote the quasi-isometry constant, chosen large enough so that

$$d_X(f \circ F(x), x) \leq K \text{ and } d_\Gamma(F \circ f(\gamma), \gamma) \leq K$$

for all $x \in X$ and $\gamma \in \Gamma$. Fix a constant $\epsilon > 0$ such that if v and w are adjacent vertices in Γ , then $d_X(f(v), f(w)) < \epsilon$.

Let U be a fundamental domain for the action of G on X , and define

$$V = N_\epsilon(U) = \{x \in X \mid d(x, U) < \epsilon\}.$$

Notice that if v and w are adjacent vertices in Γ , then $f(v)$ and $f(w)$ are contained in $g \cdot V$ for some $g \in G$. Conversely, let δ be sufficiently large so that if v and w are vertices in Γ where $f(v)$ and $f(w)$ are contained in the same copy of V then $d_\Gamma(v, w) \leq \delta$.

Let B_δ be the ball of radius δ centered at the identity in Γ . Let (u, v, w) be an ordered triple of vertices in B_δ , where any two or even all three vertices might be equal. Choose edge paths from u to v , v to w and w to u , each of which is reduced and contained in B_δ . There is then a circuit in B_δ given by the paths $u \rightarrow v$, $v \rightarrow w$, and $w \rightarrow u$; attach a 2-cell for every such circuit. Note in particular that this 2-cell depends on the original ordered triple (u, v, w) and the choices of paths connecting these vertices. Thus the same basic circuit in B_δ will bound multiple 2-cells as defined above. Since the action $G \curvearrowright \Gamma$ takes ordered triples of vertices to ordered triples of vertices, and paths to paths, we can compose these attaching maps with the action of G on Γ to equivariantly distribute 2-cells throughout Γ . Denote the resulting 2-complex by $\Gamma^{(2)}$.

We claim $\Gamma^{(2)}$ is a free, cocompact G -complex that is 1-connected. The action $G \curvearrowright \Gamma^{(2)}$ is essentially given by the construction. If σ is a 2-cell of $\Gamma^{(2)}$ then σ is associated to an ordered triple of vertices (u, v, w) , along with paths connecting them. An element $g \in G$ takes σ to the 2-cell associated to the triple $(g \cdot u, g \cdot v, g \cdot w)$, along with the g -images of the original paths. The action of G is free on Γ , thus to show that it is free on $\Gamma^{(2)}$ it suffices to show that $g\sigma \neq \sigma$, unless g is the identity. But for $g\sigma$ to equal σ , it would have to be the case that $gu = u$, which is impossible unless g is the identity. Since B_δ is finite, claim that $\Gamma^{(2)} \setminus G$ is finite is immediate. The main difficulty is in establishing that we have added enough 2-cells to kill the fundamental group.

Let $\phi : S^1 \rightarrow \Gamma$ represent an element of $\pi_1(\Gamma)$, where we may assume that S^1 has been subdivided and ϕ has been homotoped so that the map is simplicial. Form a map $\varphi : S^1 \rightarrow X$ where if v is a vertex of S^1 , then $\varphi(v) = f(\phi(v))$, and if v and w are adjacent vertices of S^1 then $\varphi([v, w])$ is a geodesic (necessarily contained inside a copy of U).

Because the geometry X is 1-connected, φ can be extended to a cellular map $\hat{\varphi} : D \rightarrow X$ where D is a simplicial 2-disk whose boundary has the same cellular structure as S^1 . Further, we may refine the simplicial structure of D so that the image of any simplex in D is contained in a copy of V . The disk D and map $\hat{\varphi}$ can be used to extend $\phi : S^1 \rightarrow \Gamma$ to a map $\hat{\phi} : D \rightarrow \Gamma^{(2)}$.

If v is a vertex in the interior of D , choose $\hat{\phi}(v)$ so that $f(\hat{\phi}(v))$ is in the same copy of V as $\hat{\phi}(v)$. If $[v, w]$ is an interior edge of D , map it to an edge path in Γ connecting $\hat{\phi}(v)$ to $\hat{\phi}(w)$ inside a ball of radius δ . If σ is a 2-simplex, then the image of its boundary is a circuit inside a ball of radius δ . By construction, there are multiple 2-cells filling this circuit, so we may send σ to a 2-cell whose boundary is this circuit. \square

We can now extend Theorems 1.1.4 and 1.3.8.

COROLLARY 2.3.5. *A group G is of type \mathcal{F}_n iff there is a geometric action of G on an $(n - 1)$ -connected cell complex. In particular, G is finitely generated iff there is a geometric action on a connected cell complex, and G is finitely presented iff there is a geometric action on a 1-connected cell complex.*

cor:geometricfiniteness

At this point the reader may feel a bit overburdened by the multiple characterizations of being finitely presented, yet we will add just one additional characterization. Let G be a finitely generated group with a fixed finite generating set S with its induced word metric. For any integer $k \geq 1$ let $\Delta_k(G)$ be the simplicial complex whose simplices consist of subsets of G of diameter $\leq k$. Notice that the 1-skeleton of $\Delta_1(G)$ is the Cayley graph of G with respect to S , and in general the 1-skeleton of any $\Delta_k(G)$ contains this Cayley graph of G . Thus each $\Delta_k(G)$ is a connected simplicial complex and by definition $\Delta_k(G) \subset \Delta_{k+l}(G)$. There are induced maps

$$\pi_1(\Delta_k(G)) \rightarrow \pi_1(\Delta_{k+l}(G)),$$

one for each triple $(n, k, l) \in \mathbb{N} \times \mathbb{N} \times \mathbb{N}$. (Since the complexes are connected, and nested, we have not specified a specific fixed base point.) The directed system $\{\pi_1(\Delta_k(G))\}$ is *essentially trivial* if for each $k \in \mathbb{N}$ there is an $l \in \mathbb{N}$ such that the map $\pi_1(\Delta_k(G)) \rightarrow \pi_1(\Delta_{k+l}(G))$ is trivial.

If you are familiar with the notion of a direct limit, you should note that being essentially trivial is stronger than stating the direct limit is trivial. In fact, the direct limit $\lim_{k \rightarrow \infty} \pi_1(\Delta_k(G))$ is trivial for any finitely generated group G . To see this notice that any $g \in \pi_1[\Delta_k(G)]$ can be represented by a map $\phi : \mathbb{S}^1 \rightarrow \Delta_k(G)$ whose image has diameter D . Thus $g \in \text{KER}[\pi_1(\Delta_k(G)) \rightarrow \pi_1(\Delta_{k+D}(G))]$. Being essentially trivial is stating that there is a constant l such that $\pi_1(\Delta_k(G)) = \text{KER}[\pi_1(\Delta_k(G)) \rightarrow \pi_1(\Delta_{k+l}(G))]$.

The following result is a corollary of a more general theorem, often referred to as ‘‘Brown’s Criterion,’’ established by Ken Brown in [5].

THEOREM 2.3.6. *A finitely generated group G is of type \mathcal{F}_n iff the directed systems $\{\pi_i(\Delta_k(G))\}$ are essentially trivial for $1 \leq i < n$. In particular, G is finitely presented iff the directed system $\{\pi_1(\Delta_k(G))\}$ is essentially trivial.*

thm:BrownCrit

PROOF. Again, we are only going to prove the case of $n = 2$, that is, we show that G is finitely presented iff $\{\pi_1(\Delta_k(G))\}$ is essentially trivial.

Let $\widetilde{\Delta}_k(G)$ denote the universal cover of $\Delta_k(G)$ and let G_k be the group of all homeomorphisms of $\widetilde{\Delta}_k(G)$ that cover elements of G . Notice that G acts geometrically on $\Delta_k(G)$, hence G_k acts geometrically on $\widetilde{\Delta}_k(G)$, and therefore G_k is finitely presented by Corollary 2.3.5. For each $k \in \mathbb{N}$ there is a short exact sequence

$$1 \rightarrow \pi_1(\Delta_k(G)) \rightarrow G_k \rightarrow G \rightarrow 1$$

and for each k and $l \in \mathbb{N}$ there is an induced commutative diagram

$$\begin{array}{ccccccc} 1 & \longrightarrow & \pi_1(\Delta_k(G)) & \longrightarrow & G_k & \longrightarrow & G \longrightarrow 1 \\ & & \downarrow & & \downarrow & & \parallel \\ 1 & \longrightarrow & \pi_1(\Delta_{k+l}(G)) & \longrightarrow & G_{k+l} & \longrightarrow & G \longrightarrow 1 \end{array}$$

Assume first that G is finitely presented. Then for any $k \in \mathbb{N}$, $\pi_1(\Delta_k)$ is finitely generated as a normal subgroup of G_k . That is, there is a finite collection of elements $\{a_1, \dots, a_m\} \subset G_k$ whose normal closure is $\pi_1(\Delta_k)$. Since for each a_i there is an l_i such that $a_i \in \text{KER}[\pi_1(\Delta_k) \rightarrow \pi_1(\Delta_{k+l_i})]$, we make take $l = \text{Max}\{l_i\}$ to see that

$$\{a_1, \dots, a_m\} \subset \text{KER}[\pi_1(\Delta_k(G)) \rightarrow \pi_1(\Delta_{k+l}(G))] .$$

Hence this map is trivial and thus the directed system $\{\pi_1(\Delta_k(G))\}$ is essentially trivial.

Conversely, assume the directed system is essentially trivial, hence given k there is an l such that

$$\pi_1(\Delta_k(G)) = \text{KER}[\pi_1(\Delta_k(G)) \rightarrow \pi_1(\Delta_{k+l}(G))] .$$

Since $G \simeq G_k/\pi_1(\Delta_k(G))$, the map $G_k \rightarrow G_{k+l}$ in the commutative diagram above induces a map $G \rightarrow G_{k+l}$. This map provides a section of the map $G_{k+l} \rightarrow G$, so G is a retract of the finitely presented group G_{k+l} , and hence G is itself finitely presented. \square

This was first established by Juan Alonso.

COROLLARY 2.3.7. *The property of being \mathcal{F}_n is a QI-invariant.*

PROOF. Let G and H be quasi-isometric groups where H is finitely presented. Let $\phi : G \rightarrow H$ be a quasi-isometry with QI-constant K , and let $\phi^{-1} : H \rightarrow G$ be a quasi-inverse. The map ϕ induces a map $\hat{\phi} : \Delta_k(G) \rightarrow \Delta_{Kk+K}(H)$ defined at the level of vertices by $\hat{\phi}(g) = \phi(g)$ and extended to simplices by noticing that if $d_G(g, g') \leq k$ then $d_H(\phi(g), \phi(g')) \leq Kk + K$. Since H is finitely presented, Theorem 2.3.6 implies there is a number L such that $\pi_1(\Delta_{Kk+K}(H)) = \text{KER}[\pi_1(\Delta_{Kk+K}(H)) \rightarrow \pi_1(\Delta_L(H))]$.

Define the map $\iota : \Delta_k(G) \rightarrow \Delta_{KL+K}(G)$ as the composition of the following three maps

$$\iota : \Delta_k(G) \rightarrow \Delta_{Kk+K}(H) \hookrightarrow \Delta_L(H) \rightarrow \Delta_{KL+K}(G) ,$$

where the last map is induced by the fact that if $d_H(h, h') \leq L$ then

$$d_G(\phi^{-1}(h), \phi^{-1}(h')) \leq KL + K .$$

Notice that $\iota(g) = \phi^{-1} \circ \phi(g)$, hence $d_G(g, \iota(g)) \leq K$. It follows that in $\Delta_{KL+K}(G)$, $\sigma \cup \iota(\sigma)$ is contained in a simplex for any σ . Hence the map ι is homotopic to the inclusion map $\Delta_k(G) \hookrightarrow \Delta_{KL+K}(G)$. But by its construction, ι kills the fundamental group, hence the directed system $\{\pi_1(\Delta_k(G))\}$ is essentially trivial. \square

2.3.1. Maps between cell complexes. Let (X, x) and (Y, y) be based connected cell complexes with $\pi_1(X, x) = G$ and $\pi_1(Y, y) = H$. The next few results investigate the relationship between the homotopy equivalence classes of based maps $(X, x) \rightarrow (Y, y)$ and the group homomorphisms $G \rightarrow H$. Certainly each based map induces a group homomorphism that is well-defined up to base-point preserving homotopy. The question we want to explore is when is the function $\text{hom}((X, x), (Y, y)) \rightarrow \text{hom}(G, H)$ onto and when is it one-to-one.

PROPOSITION 2.3.8 (Homomorphisms and maps). *If (X, x) is a presentation 2-complex with $\pi_1(X, x) = G$ and (Y, y) is an arbitrary based cell complex with $\pi_1(Y, y) = H$, then for every group homomorphism $h : G \rightarrow H$ there is a continuous map $f : (X, x) \rightarrow (Y, y)$ such that $f_* = h$.*

PROOF. (existence) simply build the map. We know the equivalence class of loops to which each edge should be sent. Pick one for each. Next, the boundary cycle of each 2-cell in X is sent to a null-homotopic loop in Y . Use this null-homotopy to extend the map constructed so far over this 2-cell.

(the right map) Since A generates G , the map induced by

□

LEMMA 2.3.9. *Let (X, x) and (Y, y) be based connected cell complexes with $\pi_1(X, x) = G$ and $\pi_1(Y, y) = H$. If X is 2-dimensional, then the map $\text{hom}((X, x), (Y, y)) \rightarrow \text{hom}(G, H)$ is onto. In particular, every group homomorphism $G \rightarrow H$ is induced by some based map $(X, x) \rightarrow (Y, y)$.*

PROOF. (add in)

□

(the restriction is necessary; consider the identity homomorphism when Y is the 2-skeleton of X . concretely, let Y be a solid 3-cube with all 8 vertices identified and let X be its 2-skeleton – work on this)

Redo with k -connectivity?

LEMMA 2.3.10. *Injective when X is k -dimensional and Y is k -connected.*

THEOREM 2.3.11. *Let (X, x) is a based connected 2-dimensional cell complex with $\pi_1(X, x) = G$. If X is a classifying space (i.e. if X is aspherical), then there is a bijection between the automorphisms of G and the based homotopy equivalences $(X, x) \rightarrow (X, x)$ up to base point preserving homotopy.*

Moved from Chapter 1

2.3.2. Commensurability. There are several natural constructions that start with finitely generated and/or finitely presented groups and produce additional groups that has similar finiteness properties. Our first example of such a construction involves short exact sequences of groups.

THEOREM 2.3.12 (Short exact sequences). *If $1 \rightarrow N \hookrightarrow G \twoheadrightarrow Q \rightarrow 1$ is a short exact sequence of groups, then there is a systematic way to construct a topological presentation for G from topological presentations for N and Q . As a consequence of this construction when N and Q are finitely generated, G is finitely generated, and when N and Q are finitely presented, G is finitely presented.*

PROOF. Let K_N be a topological presentation for N with 1-skeleton R_A and let K_Q be a topological presentation for Q with 1-skeleton R_B . The construction of a topological presentation for G starts with the wedge product $K = K_N \vee K_Q$ and then adds several new 2-cells, one for each ordered pair $(a, b) \in S_A \times S_B$. Roughly speaking the extra relations record how preimages of the symmetric generators of Q conjugate the images of the symmetric generators of N . For each $a \in S_A$ and $b \in S_B$ let n_a and q_b be the elements of N and Q they represent. Finally, if $i : N \hookrightarrow G$ and $s : G \twoheadrightarrow Q$ denote the maps in the short exact sequence, let $g_a = i(n_a)$ and for each $b \in S_B$ pick an element $g_b \in s^{-1}(q_b)$.

Because $i(N)$ is normal in G , the element $g_b g_a g_b^{-1}$ is in $i(N)$. Thus, there is a word $w_{a,b} \in (S_A)^*$ that represents $g_b g_a g_b^{-1}$ and we attach a 2-cell to K along the closed path described by the word $bab^{-1}(w_{a,b})^{-1}$. Let K' be the 2-complex that

There's a problem with the current proof outline: K_G should not contain a copy of K_Q because we don't know that the ses splits.

results from attaching such a 2-cell to K for each $(a, b) \in S_A \times S_B$. The claim is that the one vertex combinatorial 2-complex K' is a topological presentation of G . To see this note that when we crush the copy of K_N inside K' to a point, the copy of K_Q injects into the quotient and the cells not in K_Q are now attached along null-homotopic paths. Thus, $\pi_1(K'/K_N) \cong \pi_1(K_Q) = Q$, and by Proposition 1.2.21 the induced map $\pi_1(K') \twoheadrightarrow Q$ is onto.

This gives us a surjection $\pi_1(K') \twoheadrightarrow Q$ whose kernel is the normal subgroup generated by K_N . But $\pi_1(K_N)$ is a normal subgroup of $\pi_1(K')$. Thus we have a commutative diagram

$$\begin{array}{ccccccc} 1 & \longrightarrow & \pi_1(K_N) & \longrightarrow & \pi_1(K') & \longrightarrow & \pi_1(K_Q) & \longrightarrow & 1 \\ & & \text{iso} \downarrow & & \downarrow & & \text{iso} \downarrow & & \dots \\ 1 & \longrightarrow & N & \longrightarrow & G & \longrightarrow & Q & \longrightarrow & 1 \end{array}$$

And so $\pi_1(K') = G$. If N and Q are finitely generated, then both K_N and K_Q can be chosen to have a finite 1-skeleton, K' will have a finite 1-skeleton, and G is finitely generated. If N and Q are finitely presented, then both K_N and K_Q can be chosen to be finite 2-complexes, K' will be a finite 2-complex, and G is finitely presented. \square

Our second example relates the finiteness properties of a group to those of a finite index subgroup. Its proof uses an elementary lemma about normal subgroups.

LEMMA 2.3.13 (Creating normal subgroups). *Every finite index subgroup of G contains a finite index subgroup normal in G .*

PROOF. Let H be a subgroup of G and let G/H be the set of left cosets of H . There is a group homomorphism $\rho : G \rightarrow \text{SYM}_{G/H}$ that sends $g \in G$ to the way left multiplication by g permutes the left cosets of H . The kernel of this representation is a normal subgroup $N \triangleleft G$ with $N \subset H$ since H is the stabilizer of the identity coset. Finally, when H is finite index in G , $\text{SYM}_{G/H}$ is a finite group, and the index of N in G must also be finite. \square

The proof of Lemma 2.3.13 extends to subgroups of infinite index with only minor modifications. See Exercise 11 and Exercise 12.

THEOREM 2.3.14 (Finite index). *If H is a finite index subgroup of G , then H is finitely generated iff G is finitely generated and H is finitely presented iff G is finitely presented.*

PROOF. One direction is a straightforward consequence of the topological characterizations: if G is the fundamental group of a cell complex X with a compact i -skeleton and H is finite index in G , then the finite cover of X corresponding to H also has a compact i -skeleton. In the other direction, assume H is finitely generated or finitely presented and let N be a finite index subgroup of H that is normal in G (Lemma 2.3.13). Because H is finitely generated/finitely presented, so is N . Finally, since N is normal we have a short exact sequence of groups $1 \rightarrow N \rightarrow G \rightarrow Q \rightarrow 1$ with $Q := G/N$, and Q , being a finite group, has a finite presentation (Example 1.3.21). Theorem 2.3.12 completes the proof. \square

DEFINITION 2.3.15 (Commensurability). Two groups G and G' are called *commensurable* if there is a finite sequence of groups G_0, G_1, \dots, G_k such that $G \cong G_0$,

$G_k \cong G'$ and for each pair of adjacent groups in the sequence, one of them is isomorphic to a finite index subgroup of the other. In other words, commensurability is the most refined equivalence relation on groups in which every group is related to each of its finite index subgroups.

By Theorem 2.3.14, commensurable groups have similar finiteness properties.

COROLLARY 2.3.16 (Commensurability). *If G and G' are commensurable groups, then G is finitely generated iff G' is finitely generated and G is finitely presented iff G' is finitely presented.*

2.3.3. Notes on Sources. One of the primary sources for the material in this chapter is the wonderful article by Jim Cannon in the book “Ergodic theory, symbolic dynamics and hyperbolic space” edited by Tim Bedford, Michael Keane and Caroline Series. In particular, this is the source of the finiteness result and the interested reader can consult it for a complete proof for all n .

(Metrical viewpoint)

Stallings, Tits, Serre, Bass, Gromov, Thurston.

Presenting a group G in terms of generators and defining relations is as arbitrary a procedure as choosing a coordinate system to describe a geometric configuration. [20, p.120]

Move this quote elsewhere

Exercises

Metrics and quasi-isometries

1. In general you need finitely many shapes.
2. (Tietze transformations) Describe Tietze transformations (move to Chapter 3?)
3. Jon’s hemispherical metric is a metric even if there are infinitely many shapes. (For example a fg but not fp group.)
4. Consider the word metric on the integers with generating sets $A = \{1\}$ and $B = \{2, 3\}$. Calculate explicitly the distance function $d_B(n, m)$. By Lemma XXX it is clear that the metrics d_A and d_B are quasi-isometric. Use your explicit answer to find the optimal constants a, b, c, d so that $a * d_B(n, m) + b \leq d_A(n, m) \leq c * d_B(n, m) + d$ for all $n, m \in \mathbb{Z}$.
5. Let G be a group and let S and S' be distinct, possibly infinite, generating sets for G . Show that if the symmetric difference of S and S' is finite (i.e. there are only a finite number of generators in one set but not the other) then the identity map $\mathbf{1} : (G, d_S) \rightarrow (G, d_{S'})$ is a quasi-isometry between these metric spaces. More generally, show that this map is a quasi-isometry when the length of each element in S is uniformly bounded in the $d_{S'}$ metric and, conversely, the length of each element in S' is uniformly bounded in the d_S metric.
6. Let X be the unit torus (i.e. the space constructed from the square with side length 1 where opposite sides have been identified in an orientation preserving way). Consider the two metrics on its fundamental group \mathbb{Z}^2 derived from (1) the geodesic distance d_X in the universal cover and (2) the word metric d_A on the 1-skeleton. Prove these two metrics are quasi-isometric and find the optimal constants a, b, c, d so that $a * d_X(x, y) + b \leq d_A(x, y) \leq c * d_X(x, y) + d$ for all $x, y \in \mathbb{Z}^2$.

7. Consider the group $BS(1,2) = \langle ab = ba^2 \rangle$ and let d denote the word metric with respect to this generating set. The subgroup $H = \langle a \rangle$ has its own word intrinsic metric with respect to its generating set $\{a\}$ and an ambient metric which simply restricts the word metric on $BS(1,2)$ to the subgroup H . Calculate the distance from e to a^n in both metrics. Then prove that these two metrics on H are *not* quasi-isometric to each other. When the subgroup metric is quasi-isometric to the restriction of the ambient metric, the subgroup is said to be quasi-isometrically embedded. The subgroup H in $BS(1,2)$ is the simplest example of a subgroup which is not quasi-isometrically embedded in its ambient group.

Geometries and geometric actions

8. One can extend Theorem 2.3.4 to include a slightly lower base case: A group G acts geometrically on a proper metric space iff G is countable. Hint: pick a sequence of elements that generate larger and larger subgroups and look at the metric where the n -th element has length n .
9. A fg subgroup H of a finitely generated group G is a *quasi-retract* if the inclusion $H \hookrightarrow G$ is a quasi-isometric embedding and there is a map $\phi : G \rightarrow H$ with an associated constant K such

$$d_H(h, \phi(h)) \leq K$$

for all $h \in H$. Prove that if G is finitely presented, then so is any quasi-retract of G .

Want an example of a quasi-retract? The trefoil knot group is virtually $\mathbb{F}_2 \times \mathbb{Z}$. The \mathbb{F}_2 subgroup is a retract of $\mathbb{F}_2 \times \mathbb{Z}$. Show it's a quasi-retract of the trefoil knot group.

Check this is a good example!

10. Show that the following are commensurable classes of groups:
- fg non-abelian free groups
 - fundamental groups of hyperbolic surfaces
11. Theorem 2.3.4 describes how to construct a free action on a 1-connected complex from a geometric action on a 1-connected geometry. The infinite dihedral group is $\langle s, t \mid s^2 = t^2 = 1 \rangle$ and it acts on the real line with point stabilizers $\simeq \mathbb{Z}_2$. Convert this to a free action on a 1-connected complex.
12. Prove the “space of Cayley graph metrics” on a group G is path connected. (topology given by fixing the size of K ?)
13. Prove Lemma 2.1.3 and Proposition 2.1.5.
14. The *geometric dimension* of a group G is the minimal dimension of a $K(G,1)$. For example, a group G has geometric dimension 1 iff it is free. Prove that the geometric dimension of the free abelian group \mathbb{Z}^n is n .

Dehn's Fundamental Problems

*“The general theory of groups defined in this way [via a finite presentation], in so far as they are infinite, does not appear to be very well developed. There are above all three fundamental problems whose solution is very important and probably not possible without a penetrating study of the subject.”*¹

Max Dehn in 1911 [9]

We begin with a series of fundamental questions. Suppose a closed loop is drawn on a compact surface. Is there an easy way to tell whether or not it is trivial in the fundamental group? Notice that this question is not about some particular loop.

The question can be reformulated as follows. Is there some uniform procedure that can be set up in advance so that no matter which closed loop is selected as an input, the procedure determines, in a finite amount of time, whether or not the loop represents the identity element in the fundamental group.

In 1912 Max Dehn proved that the answer to this question is “yes” for compact surfaces but the proof is not as general as one would like. If compact surface is replaced by compact 3-manifold the answer is not currently known (comment on Grigori Perelman, Jason Manning). If we consider compact 4-manifolds the answer is a resounding “no”.

1. How do we construct the universal cover of X ?
2. Which maps $f, g : \mathbb{S}^1 \rightarrow X$ are homotopic?
3. Which complexes X, Y are homotopy equivalent?

3.1. Disc diagrams and van Kampen's lemma

DEFINITION 3.1.1 (Disc diagrams). A *disc diagram* is a finite contractible combinatorial 2-complex D together with a specified embedding into the plane \mathbb{R}^2 . If D is homeomorphic to a disc then D is called *non-singular*. Otherwise it is a *singular disc diagram*. For example, disc diagrams with cut vertices or cut edges are singular.

DEFINITION 3.1.2 (Boundary cycle). The explicit embedding into \mathbb{R}^2 is used to select an oriented boundary cycle. (say more)

(somewhere – Chapter 3 – we should probably put a Lebesgue number argument that shows that every continuous map from $\mathbb{S}^1 \rightarrow X^1$ can be homotoped

¹“Die allgemeine Theorie derartig definierter Gruppen, sofern sie unendlich sind, scheint bisher sehr wenig entwickelt zu sein. Hier sind es vor allem *drei fundamentale Probleme*, deren Lösung sehr wichtig und wohl nicht ohne eindringendes Studium der Materie möglich ist.”

to a combinatorial loop. Similarly every continuous map from $\mathbb{D}^2 \rightarrow X^2$ can be homotoped to a van Kampen diagram.)

Alternatively, a disc diagram can be defined as the complement of a (open) 2-cell inside a combinatorial 2-sphere. Since \mathbb{S}^2 is compact, the complex is finite, the complement of an open 2-cell is clearly a subcomplex, and it is not too hard to see that this complement must be contractible.

The following lemma is immediate.

LEMMA 3.1.3. *Let X be a combinatorial 2-complex, (let C be a combinatorial circle), and let $p : C \rightarrow X$ be a combinatorial loop. If there is a disc diagram D and combinatorial maps $i : C \rightarrow D$ and $q : D \rightarrow X$ such that $q \circ i = p$ then p is null-homotopic.*

Van Kampen's lemma establishes the converse.

LEMMA 3.1.4 (Van Kampen's lemma). *Let X be a combinatorial 2-complex, (let C be a combinatorial circle), and let $p : C \rightarrow X$ be a combinatorial loop. If p is null-homotopic, then there exists a disc diagram D and combinatorial maps $i : C \rightarrow D$ and $q : D \rightarrow X$ such that $q \circ i = p$.*

import the proof and the pictures from Fans-Ladders

PROOF.

□

THEOREM 3.1.5 (Constructing the universal cover). *Let X be a compact combinatorial 2-complex. The word problem for $\pi_1(X)$ is algorithmically decidable iff its universal cover is algorithmically constructible iff the Cayley graph is constructible.*

3.2. Algorithms and decidability

DEFINITION 3.2.1 (Recursively enumerable). Let U be some well-understood set and let A be some subset of U . If there exists an algorithm that only outputs elements of A and has the property that every element of A is produced as output after some finite length of time, then this algorithm proves that the members of A are *recursively enumerable*. If A and $U \setminus A$ are both recursively enumerable then membership in A is said to be *decidable*.

THEOREM 3.2.2. *Let X be a combinatorial 2-complex. The collection of null-homotopic loops in X are recursively enumerable.*

PROOF. This follows from van Kampen's Lemma. The finite list of disc diagrams with exactly v vertices, e edges and f 2-cells can be effectively enumerated. For each disc diagram there are a finite number of maps to X . Concatenating this countable collection of finite lists enumerates all disc diagrams over X . For each disc diagram over X output its boundary cycle. □

In the 1930s it came as a bit of a shock when it was discovered that there are subsets of the natural numbers whose members can be effectively listed but with an undecidable membership problem.

THEOREM 3.2.3 (Gödel). *There exists a subset $A \subset \mathbb{N}$ which is recursively enumerable, but not decidable.*

This result was then transferred to group theory by Tarski.

THEOREM 3.2.4 (Tarski). *There is a finitely presented group with an undecidable word problem. This means, in particular, that there exists a finite combinatorial 2-complex X where the collection of combinatorial loops in X that are not null-homotopic cannot be recursively enumerated.*

The theorem does not say that there is some mysterious finitely presented group G with this property. The proof given by Tarski proceeds by giving an explicit concrete example and showing that the halting problem can be encoded into the group structure of G .

filling function. Dehn functions

PROPOSITION 3.2.5 (Filling Theorem). *If G acts geometrically on a simply connected manifold M then the filling function on M is equivalent to the Dehn function on G .*

THEOREM 3.2.6 (Dehn functions and computability). *A finitely presented group G has a decidable word problem iff its Dehn function is recursive.*

PROOF. The proof of this isn't quite as hard as it might seem. If the Dehn function is a recursive function, then given a word w we first compute the value of $\delta_P(n)$ where $n = |w|$. Then we enumerate all disc diagrams with at most n 2-cells. If none of these have w as the boundary cycle, then we can conclude that w is not equal to the identity.

The converse will be left a little vague. If there is a computer program which conclusively proves that w is or is not trivial for any arbitrary input w , it must implicitly compute some recursive bound on the number of 2-cells used. In particular, the Dehn function must be recursive. \square

3.3. Dehn's algorithm for surface groups

The word problem asks whether an arbitrary word can be reduced to the unique word of length 0. Sometimes there are obvious ways to immediately shorten the length of the word under consideration. Dehn's algorithm incorporates and formalizes these obvious shortenings.

EXAMPLE 3.3.1 (Surface groups). Let S be the surface constructed by identifying the opposite edges of an octagon in an orientation preserving way (i.e. the standard complex for the relation $abcd = dcba$.) We can build the universal cover by starting with a circle octagon... (build rings and then look at an arbitrary closed loop)

DEFINITION 3.3.2 (Dehn's algorithm). Dehn's algorithm begins with a procedure where you systematically try to *locally shorten* a combinatorial loop by pushing a subpath across a single 2-cell. In particular, suppose that p contains a subpath that is also more than half of the boundary of a 2-cell (in the combinatorial metric). If you find such a subpath, homotopy this subpath across the 2-cell to create a strictly shorter closed loop. Repeat. At some point this shortening process stops at a local minimum. If the final path is trivial, then you have proved that the original loop was null-homotopic. The other possibility is that you stop at a word that cannot be locally shortened by crossing a single 2-cell. The shortening process can be applied to any combinatorial loop in any combinatorial 2-complex.

A presentation is said to satisfy *Dehn's algorithm* if one can prove that every null-homotopic closed loop can be locally shortened. This implies that a non-trivial

follow John Stillwell's book
which gives Dehn's original
argument

loop with no local shortening is guaranteed to represent a non-trivial element of the fundamental group and that our relatively straightforward shortening procedure can be used as a diagnostic to detect whether a closed path is trivial or not. Simply run then check to see whether the result was the trivial loop. If it is, the original loop represent the identity element and if it is not, then it does not.

THEOREM 3.3.3. *If a presentation satisfies Dehn's algorithm, then the group it defines has a solvable word problem.*

(reference and poach from Chuck Miller's excellent survey article [22])

3.4. Combinatorial Curvature

(the text below has been lifted from the paper with D. Wise [21]. It still needs to be blended in and simplified)

In this section, we state and prove a version of the combinatorial Gauss-Bonnet theorem, followed by two applications. It was first stated and proven for diagrams which embed in a sphere without boundary by Gersten in [12] and Pride in [24], thereby refining some earlier ideas of Lyndon's concerning (p, q) -maps (see [19]) as well as an idea of Sieradski's [27]. The Gauss-Bonnet theorem was stated for surfaces in [13]. In this article, we prove a generalization to arbitrary 2-complexes. Since first writing this article we have learned that this theorem was proven earlier for piecewise constant curvature 2-complexes by Ballmann and Buyalo [2]. The proof is the same.

DEFINITION 3.4.1 (Links and perimeters). Let X be a locally finite 2-complex and let x be a point in its 1-skeleton. The cells of X each have a natural partial metric obtained by making every 1-cell isometric to the unit interval and every n -sided 2-cell isometric to a Euclidean disc of circumference n whose boundary has been subdivided into n curves of length 1. In this metric, the set of points which are a distance equal to ϵ from x will form a finite graph. If ϵ is sufficiently small, then the graph obtained is independent of the choice of ϵ . This well-defined graph is the *link of x in X* and is denoted by $\text{Link}(x)$. If v is a 0-cell of X , then the graph $\text{Link}(v)$ is called the link of the 0-cell v . When X contains a single 0-cell v , then $\text{Link}(v)$ is the *star graph* or *coinitial graph* of the presentation X encodes. To avoid confusion, we will discuss $\text{Link}(v)$ using the language of *vertices* and *edges* and reserve the terms 0-cells and 1-cells for the 2-complex X containing v . Notice that the link of a 0-cell can be an arbitrary finite graph. In contrast, if x lies in the interior of a 1-cell e of X , then the link of x has a very particular form: $\text{Link}(x)$ will have exactly two vertices (corresponding to the two ends of e) and a finite number of edges connecting these two vertices. The number of edges in $\text{Link}(x)$ is called the *perimeter* of e and will be denoted $\mathbf{P}(e)$. The word "perimeter" is used because if each 1-cell is thought to have length 1, then this is the length of the boundary created when the 1-cell e is removed from X .

DEFINITION 3.4.2 (Corners and sides). Let X be a 2-complex, let v be a 0-cell in X , let $R \rightarrow X$ be a 2-cell in X , and let x be a point in the interior of a 1-cell e in X . If we regard the 2-cells of X as polygons, then the edges of $\text{Link}(v)$ correspond to the corners of these polygons attached to v . We will refer to a particular edge in $\text{Link}(v)$ as a *corner of R at v* if this edge comes from the polygon $R \rightarrow X$. Similarly, the edges of $\text{Link}(x)$ correspond to the sides of these polygons attached

to e , and we will refer to a particular edge in $\text{Link}(x)$ as a *side of R at e* if this edge comes from $R \rightarrow X$.

REMARK 3.4.3. It is immediate from the definition that the 2-cell $R \rightarrow X$ contributes exactly $|\partial R|$ corners at 0-cells of X and exactly $|\partial R|$ sides at 1-cells of X . Since the number of sides contributed by each polygon is the same as the number of corners, the total number of sides in X equals the total number of corners.

DEFINITION 3.4.4 (Combinatorial curvature). We say X is an *angled 2-complex* provided that every corner c of X has been assigned a real number $\angle c$ called the *angle of c* , and X is *positively angled* if all these angles are positive. If f is a 2-cell of X then the *curvature of f* is defined to be the sum of the angles assigned to its corners minus $(|\partial f| - 2)\pi$ (which is the expected Euclidean angle sum). In symbols we have

$$\kappa(f) = \left(\sum_{c \in \text{Corners}(f)} \angle c \right) - |\partial f|\pi + 2\pi$$

The curvatures of the 2-cells of X are its *2-cell curvatures*. If v is a 0-cell of X then the *curvature of v* is defined to be 2π minus $\pi \cdot \chi(\text{Link}(v))$ minus the sum of the angles assigned to corners at v . If X embeds in the plane and v is an interior 0-cell then $\text{Link}(v)$ is a circle and the curvature measures the difference between the expected Euclidean angle sum of 2π and the actual angle sum. This *0-cell curvature* equation gives the appropriate generalization of this idea to arbitrary 2-complexes. In symbols

$$\kappa(v) = 2\pi - \pi \cdot \chi(\text{Link}(v)) - \left(\sum_{c \in \text{Corners}(v)} \angle c \right)$$

REMARK 3.4.5. Let D be a positively angled diagram and let v be a 0-cell of D . If $\text{Link}(v)$ is not a complete circle (but is not empty), then $\chi(\text{Link}(v)) \geq 1$ and thus $\kappa(v) \leq \pi$. Moreover, $\kappa(v) = \pi$ if and only if $\text{Link}(v)$ is a single vertex and v is the tip of a spur. Furthermore, if $\text{Link}(v)$ is disconnected, then $\chi(\text{Link}(v)) \geq 2$ and thus $\kappa(v) \leq 0$. In this case, $\kappa(v) = 0$ if and only if the $\text{Link}(v)$ consists of exactly two isolated vertices.

We can now state and prove the combinatorial Gauss-Bonnet theorem.

THEOREM 3.4.6 (Combinatorial Gauss-Bonnet). *If X is an angled 2-complex then the sum of the 2-cell curvatures and the 0-cell curvatures is 2π times the Euler characteristic of X .*

$$(1) \quad \sum_{f \in 2\text{-cells}(X)} \kappa(f) + \sum_{v \in 0\text{-cells}(X)} \kappa(v) = 2\pi \cdot \chi(X).$$

In particular, if X is a disc diagram then this sum will be 2π and if X is an annular diagram the sum will be 0.

PROOF. For convenience we will define the following pair of constants.

$$C = \sum_{c \in \text{Corners}(X)} \angle c \quad P = \sum_{e \in 1\text{-cells}(X)} \mathbf{P}(e)$$

The proof will follow from the following two equations:

$$(2) \quad \sum_{f \in 2\text{-cells}(X)} \text{Curvature}(f) = C - \pi P + 2\pi F$$

$$(3) \quad \sum_{v \in 0\text{-cells}(X)} \text{Curvature}(v) = 2\pi V - 2\pi E + \pi P - C$$

where the letters V , E , and F represent the number of 0-cells, 1-cells, and 2-cells in X , respectively. To prove the theorem, one simply adds Equations 2 and 3, and observes that the corner sums and the perimeter sums cancel, leaving $2\pi(V - E + F)$, which is exactly 2π times the Euler characteristic of X . The remainder of the proof justifies these two equations.

In the definition of the curvature of a 2-cell, there are three terms. The first term contributes C and the last term contributes $2\pi F$ towards the sum of all 2-cell curvatures. Observe that the sum of the lengths of the boundaries of the 2-cells of X is precisely the number of sides of 2-cells in X , and this is the total number of sides at 1-cells of X which is precisely P . Thus the middle term contributes $-\pi P$ towards the sum, and Equation 2 has been established.

Similarly, in the definition of the curvature of a 0-cell there are three terms. The first term contributes $2\pi V$ and the last term contributes $-C$ towards the sum of the 0-cell curvatures. Thus to establish Equation 3 it remains to show that the sum of the Euler characteristics of the links of the 0-cells is $2E - P$. We will consider the vertices and edges in the links separately. Note that the edges in the links of the 0-cells are in one-to-one correspondence with the corners of X . Since each 2-cell has as many corners as sides, the total number of edges occurring in the links of the 0-cells is P . On the other hand, the vertices in the links of the 0-cells correspond to the ends of the 1-cells of X . Since each 1-cell of X contributes two distinct vertices to the links of the 0-cells, the total number of vertices which occur in the links is $2E$. Finally, since the Euler characteristic of a graph is the number of vertices minus the number of edges, the sum of the Euler characteristics of the links is $2E - P$. This establishes Equation 3 and completes the proof. \square

In the remainder of the section we present two quick applications of Theorem 3.4.6. The first application is perhaps the most surprising, and it is the source of essentially all of the results of small cancellation theory.

THEOREM 3.4.7. *Let D be a positively angled disc diagram. Suppose that each 2-cell and each interior 0-cell of D has nonpositive curvature. Then one of the following holds:*

1. D is trivial.
2. D is a subdivided interval.
3. There are at least three 0-cells in ∂D with positive curvature.

PROOF. By the Combinatorial Gauss-Bonnet theorem, the total curvature of D is exactly 2π . Therefore there must be some 0-cells in ∂D with positive curvature. First observe that if there is a 0-cell v in ∂D with $\kappa(v) > \pi$ then $\text{Link}(v)$ is empty, and therefore D is trivial because it is connected and so $D = v$.

The other possibility is that no boundary 0-cell has curvature larger than π . Now if there are at most two sources of positive curvature, then both of these 0-cells must have curvature exactly equal to π and all other curvatures must equal 0. Let v_0 be one of the two 0-cells of positive curvature. By Remark 3.4.5, the link of v_0 must consist of a single 0-cell. If the 0-cell v_1 at the other end of the unique 1-cell emanating from v_0 is not the other 0-cell of positive curvature, then $\kappa(v_1) = 0$, and $\text{Link}(v_1)$ must be disconnected. By Remark 3.4.5 we conclude that $\text{Link}(v_1)$ has a specific structure: it consists of two disconnected vertices, and thus there is

a unique additional 1-cell incident at v_1 . Repeating this argument, we eventually reach the other 0-cell of positive curvature and the proof is complete. \square

As a consequence, diagrams satisfying $C(p)$ - $T(q)$ for large p and q have restricted structures.

THEOREM 3.4.8. *Let D be a $C(3)$ - $T(6)$ [$C(4)$ - $T(4)$] disc diagram. Then one of the following holds:*

1. D is trivial.
2. D is a subdivided interval.
3. There are at least three 0-cells in ∂D with connected links and valence ≤ 3 [≤ 2].

PROOF. We assign an angle of $\pi/3$ [$\pi/2$] to each corner of D . Let R be a 2-cell of D . By the $C(3)$ [$C(4)$] condition and the convention on sides, R has at least 3 [respectively 4] corners and therefore has nonpositive curvature. By the $T(6)$ [$T(4)$] condition, each interior 0-cell has nonpositive curvature. Let v be a 0-cell in ∂D . If $\text{Link}(v)$ is disconnected then v will be nonpositively curved because $\chi(\text{Link}(v)) \geq 2$. On the other hand, if $\text{Link}(v)$ is connected and v has valence at least 4 [respectively 3], then v is nonpositively curved because of the way the angles are assigned. An application of Theorem 3.4.7 completes the proof. \square

(old refconv:psides)

end the import from FL

DEFINITION 3.4.9 (Angled 2-complex). The *corner* of a polygon is determined by a vertex in its boundary cycle. An *angled 2-complex* is one where every corner of every polygon has been assigned a real number called its *angle*. There is no requirement that these real numbers be positive or even non-negative. If every assigned angle is positive or non-negative, the complex is called a *positively* or *non-negatively angled 2-complex*.

DEFINITION 3.4.10 (Curvatures).

DEFINITION 3.4.11 (Small cancellation conditions).

small cancellation groups and the combinatorial Gauss-Bonnet theorem.
Do the $C(3) - T(6)$ situation.

EXAMPLE 3.4.12 (Surfaces). Let S be a compact surface. If it is not a 2-sphere or a projective plane then there is a 1-vertex triangulation of S with at least 6 corners meeting at the unique vertex. In particular, it is a $C(3) - T(6)$ complex.

(Redo the surface example using CGB and a single regular $4n$ -gon.)

Notes on sources. Martin Bridson has a very nice article [3] on the geometry of the word problem. More specifically, he details the relationship between Dehn functions and filling functions in Riemannian manifolds.

Exercises

CHAPTER 4

Hyperbolic Geometry

The Milnor-Švarc theorem (Theorem 2.2.17) firmly establishes a close connection between geometries and finitely generated groups. In this chapter is a short digression into the world of hyperbolic geometry to set the stage for the following chapter on Gromov hyperbolic groups. Keeping this chapter brief has been difficult since there are many beautiful aspects of hyperbolic geometry that are easily explained. Thus virtually all of the topics included contribute directly to our main goal: to prove that every triangle in hyperbolic space is uniformly thin. To do this we need an explicit hyperbolic metric in one of the standard models of hyperbolic space, and a working knowledge of some of its isometries. We begin with the study of circle preserving maps from the 2-sphere to itself.

In particular, the theory of Möbius transformations is introduced via circle preserving maps of the 2-sphere. We define the hyperbolic metric as the unique metric (up to rescaling) that is invariant under all Möbius transformations stabilizing a disc.

4.1. Circle-preserving maps

When a plane in \mathbb{R}^3 intersects the unit 2-sphere in more than one point, we call the result a *circle in \mathbb{S}^2* (Figure 1). The homeomorphisms of \mathbb{S}^2 that send circles to circles form a group under composition. This group contains the orthogonal group $O(3)$, the full isometry group of \mathbb{S}^2 , but there are many other homeomorphisms that preserve the circles. In this section we characterise these homeomorphisms.

The key points are:
1. Circle preserv maps
2. Classification theorem
3. Lin frac trans
4. Point Stab = Isom(\mathbb{R}^2)
5. Disc Stab = Isom(H²)
6. Invariant metrics
7. Thin triangles

FIGURE 1. A circle in a 2-sphere.

DEFINITION 4.1.1 (Stereographic projection). Use (x, y, z) for a point on the unit sphere (i.e. with $x^2 + y^2 + z^2 = 1$) and $(X, Y, 0)$ for a point in the xy -plane. Every non-horizontal line through the point $(0, 0, 1)$ contains exactly one additional point on the unit sphere and one point in the plane. This bijection is called *stereographic projection*.

PROPOSITION 4.1.2 (Circles under stereographic projection). *The stereographic projection map sends circles in \mathbb{S}^2 to lines and circles in \mathbb{R}^2 and vice versa. More specifically a circle in \mathbb{S}^2 is sent to a line iff it contains the north pole.*

PROOF. To see the equivalence between circles through the north pole and lines in the xy -plane, simply notice that both conditions are in one-to-one correspondence with the non-horizontal planes through the point $(0, 0, 1)$. More specifically, each non-horizontal plane through $(0, 0, 1)$ can be intersected with the unit sphere to create a circle on \mathbb{S}^2 or with the xy -plane to create a line and it is clear that stereographic projection provides a bijection between them.

So suppose that we have a circle on \mathbb{S}^2 that does not contain the point $(0, 0, 1)$. It is described by a pair of equations $x^2 + y^2 + z^2 = 1$ and $ax + by + cz = d$ where a, b, c, d are fixed constants and since we are assuming that this plane nontrivially intersects \mathbb{S}^2 , we know that $a^2 + b^2 + c^2 > d^2$. On the other hand, a circle in the plane is described by equations $(X - A)^2 + (Y - B)^2 = C^2$.

It's straightforward to verify the correspondence. The equations connecting them are, in one direction $a = 2A$, $b = 2B$, $c = (A^2 + B^2 - C^2 - 1)$ and $d = (A^2 + B^2 - C^2 + 1)$, and in the other $2A = a$, $2B = b$, and $2C = \sqrt{a^2 + b^2 + c^2 - d^2}$. \square

finish this and make it more explicit. In particular, use my notes from lecture to fill in the short algebraic proof

We are now able to create lots of maps $h : \mathbb{S}^2 \rightarrow \mathbb{S}^2$ that preserves circles. Simply consider $h = f \circ g \circ f^{-1}$ where $g : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ is an map from \mathbb{R}^2 to itself that sends circles to circles and lines to lines. Three easy examples are translations, rotations and dilations.

EXAMPLE 4.1.3 (Translations). The action on \mathbb{S}^2 corresponding to a translation preserves the circles through the north pole corresponding to lines parallel to the translation direction and it shifts the lines through the north pole corresponding to the perpendicular lines. The only point on \mathbb{S}^2 that is fixed is the north pole itself.

PROPOSITION 4.1.4. *The only orientation preserving circle map fixing only one point on \mathbb{S}^2 is the projection of a translation.*

If f fixes a point we might as well assume that it's the north pole. If f fixes another point we might as well assume that it's the south pole (by conjugating by a translation).

EXAMPLE 4.1.5 (Rotations). Rotations fix two points on \mathbb{S}^2 and they preserves some circles.

EXAMPLE 4.1.6 (Dilations). Dilations have a north-south dynamic with one attracting fixed point and one repelling fixed point.

We now work to show that there is a circle preserving map that sends any three points on \mathbb{S}^2 to any other three points. This is a property known as being *3-transitive*.

Using rigid motions, translations, rotations and dilations we can assume that f fixes three points.

PROPOSITION 4.1.7 (Triply transitive). *For any triple of distinct points x, y , and z there exists an orientation preserving circle preserving homeomorphism $f : \widehat{\mathbb{C}} \rightarrow \widehat{\mathbb{C}}$ with $f(x) = \infty$, $f(y) = 0$, and $f(z) = 1$. As a consequence, the group of circle preserving maps is triply transitive on $\widehat{\mathbb{C}}$.*

LEMMA 4.1.8 (Rhombi). *A rhombus contains an inscribed circle, but a non-rhombic parallelogram does not.*

COROLLARY 4.1.9 (Circles to circles). *If f is a Möbius transformation that fixes ∞ , then the action of f on the stereographic projection sends circles in \mathbb{R}^2 centered at a to circles in \mathbb{R}^2 centered at $f(a)$.*

PROOF. Let x and y be two nonantipodal points on the circle of radius r centered at a and consider the rhombus with sides ax and ay . The map f preserves parallels and sends lines to lines so the image of this rhombus under f is a parallelogram. Moreover, since f preserves incidence, the circle inscribed in the rhombus must be sent to a circle inscribed in its image, forcing the image to be a rhombus. Thus $d(f(a), f(x)) = d(f(a), f(y))$. Call this value s . The same argument shows that every point z on the circle radius r centered at a is sent to a point $f(z)$ on the circle of radius s centered at $f(a)$ (since no such z is antipodal to both x and y). \square

It turns out that we don't need to know much about f to know it completely.

LEMMA 4.1.10 (Midpoints). *If f is a Möbius transformation that fixes ∞ , a and b , then f also fixes the point halfway between a and b .*

PROOF. Let $r = d(a, b)$. By (previous stuff) the circle of radius r centered at a is sent to itself as is the circle of radius r centered at b . These two circles intersect in two points c and d and $f(\{c, d\}) = \{c, d\}$. Even if they are switched by f , the line through c and d is sent to itself by f and its point of intersection with the line through a and b must be fixed. \square

LEMMA 4.1.11 (Three points). *The only orientation-preserving circle-preserving map of S^2 that fixes three points is the identity map.*

PROOF. Without loss of generality, assume that ∞ is one of the fixed points. The argument begins as in the previous proof, but the addition of orientation preserving implies that c and d are each fixed by f . The (rhombus idea) implies that the entire lattice that they generate in \mathbb{C} is fixed. The (midpoint idea) extends this to a fixed set which is dense. Continuity completes the proof. \square

(add in corollaries that give the group structure in terms of linear fractional transformations).

As an obvious corollary to this lemma, there is at most one circle map f sending x to a , y to b and z to c .

PROPOSITION 4.1.12 (Constructing Möbius functions). *Let $f : \widehat{\mathbb{C}} \rightarrow \widehat{\mathbb{C}}$ be an orientation preserving circle preserving map. If $f(a) = 0$, $f(b) = \infty$ and $f(\infty) = c$ and a , b and c are distinct, then $f(z) = c \left(\frac{z-a}{z-b} \right)$.*

EXAMPLE 4.1.13 (Rotations). Proposition 4.1.12 can be used to quickly write down explicit transformations for the 48 rigid motions that permute the six special points $\{0, \infty, 1, -1, i, -i\}$. For example, the one-quarter rotation that fixes i and -1 and sends $0 \rightarrow i \rightarrow \infty \rightarrow (-i) \rightarrow 0$ is described by the function $f(z) = i \left(\frac{z-i}{z+i} \right)$ since $f(i) = 0$, $f(-i) = \infty$, and $f(\infty) = i$. Note for later use that after pre- and post-composing with stereographic projection, this rotation sends the upper half-space to the unit disc.

EXAMPLE 4.1.14 (Möbius functions). An orientation preserving circle preserving maps is uniquely determined by how it acts on any three distinct points.




FIGURE 2. The standard coordinates on $\widehat{\mathbb{C}} = \mathbb{C}P^1$.

REMARK 4.1.15 (Higher dimensional analogues). Most of what we have done readily extends to higher dimensions. Consider the group of motions of homeomorphisms of the S^n that send subspheres (determined by non-trivial affine intersections) to subspheres. This includes the transitive action by the orthogonal group. Thus we only need to understand the homeomorphisms that fix the north pole. Using stereographic projection we find that every similarity of \mathbb{R}^n (i.e. all compositions of orthogonal transformations, translations and dilations) lead to Möbius transformations of the n -sphere. The rhombus argument still holds...

REMARK 4.1.16 (Möbius geometry). More generally, given any two fields $\mathcal{F}_1 \subset \mathcal{F}_2$ and positive integers $k \leq l$, we can study the orbits of $\mathcal{F}_1 P^k$ inside $\mathcal{F}_2 P^l$. Circles in 2-spheres are the special case of copies of $\mathbb{R}P^1$ inside $\mathbb{C}P^1$.

4.1.1. Möbius transformations. Here is the outline of the steps used to show that Möbius transformations are the only circle preserving homeomorphisms from $\widehat{\mathbb{C}}$ to $\widehat{\mathbb{C}}$. The key result is the following:

THEOREM 4.1.17 (Three points). *The only orientation preserving, circle preserving homeomorphism from $\widehat{\mathbb{C}}$ to $\widehat{\mathbb{C}}$ that fixes 0, 1, and ∞ is the identity.*

PROOF. Let f be such a map.

1. First f preserves the intersection number between two circles.
2. Because f fixes ∞ , f sends lines to lines
3. and f sends parallel lines to parallel lines
4. parallelograms go to parallelograms
5. rhombi go to rhombi
6. circles centered at z go to circles centered at $f(z)$.
7. $\text{FIX}(f)$ is parallelogram closed
8. $\text{FIX}(f)$ is midpoint closed
9. $\text{FIX}(f)$ is dense (using o.p.)

10. $\text{FIX}(f)$ is $\bar{\mathbb{C}}$. (using continuity)

□

COROLLARY 4.1.18 (Uniquely triply transitive). *There exists a unique orientation preserving circle preserving homeomorphism that takes any three points to any other three points*

4.2. Models of hyperbolic space

(define disc in \mathbb{S}^2 as a nontrivial intersection of a half-space with \mathbb{S}^2 .)

The Möbius transformations act transitively on the discs in the unit 2-sphere.

Can we define a metric that is invariant under all of the Möbius transformations that stabilize this disc. The answer is yes.

For concreteness consider a hemisphere in \mathbb{S}^2 . Two standard choices would be the hemisphere bounded by the real axis containing i , or the hemisphere bounded by the unit circle in \mathbb{C} that contains 0. In Figure 2 these would be southern hemisphere and the front hemisphere.

bounded by the real axis where we pick the side containing i .

DEFINITION 4.2.1 (Upper half-space model). Let \mathbb{U} be the upper half-space in the complex plane and define the *hyperbolic length* of a piecewise differentiable path $f : [a, b] \rightarrow \mathbb{U}$ to be given by the integral

$$\int_f \frac{|dz|}{\text{Im}(z)}$$

The next step is to use the metric on the upper half plane to induce a metric on the unit disc.

DEFINITION 4.2.2 (Poincaré Disc Model). Let \mathbb{D} be the open unit disc in the complex plane and define the *hyperbolic length* of a piecewise differentiable path $f : [a, b] \rightarrow \mathbb{D}$ to be given by the integral

$$\int_f \frac{2}{1 - |z|^2} |dz|$$

Two comments: first, it is easy to show that these two elements of arc length end up defining the same notion of arc length in the hemisphere of \mathbb{S}^2 .

PROPOSITION 4.2.3 (Lengths). *Let h be the Möbius function that sends the upper half plane to the unit disc via a rigid rotation of the corresponding hemispheres in $\hat{\mathbb{C}}$. For every piecewise differentiable path $f : [0, 1] \rightarrow \mathbb{U}$, the length of f calculated in the upperhalf space model is equal to the length of $h(f)$ calculated in the disc model.*

PROOF. The function h is defined by the formula $h(z) = i \left(\frac{z-i}{z+i} \right)$. (the key calculations include $|z+i|^2 - |z-i|^2 = 4 \text{Im}(z)$ and $h'(z) = -2(z+i)^{-2}$). □

The length function $d_{\mathbb{D}}$ is clearly invariant under rotations around 0 and $d_{\mathbb{U}}$ is clearly invariant under translations and dilations. Together these observations show that this notion of distance is invariant under $PSL_2(\mathbb{R})$.

EXAMPLE 4.2.4. Show that the hyperbolic length of the segment from the origin to $(r, 0)$ is $\log \left(\frac{1+r}{1-r} \right) = 2 \tanh^{-1}(r)$.

Discuss (1) the upper half-space model, (2) Poincaré disc model, (3) the Klein model and (4) the hyperbolic model - as well as their relations. In all of them describe the geodesics, and in at least one, describe the distance function explicitly.

4.3. Isometries of hyperbolic space

Isometry between them, conformal, angles and geodesics, group of isometries, homogeneous space,

4.4. Triangles in hyperbolic space

	\mathbb{S}^2	\mathbb{H}^2
law of sines	$\frac{\sin a}{\sin \alpha} = \frac{\sin b}{\sin \beta} = \frac{\sin c}{\sin \gamma}$	$\frac{\sinh a}{\sin \alpha} = \frac{\sinh b}{\sin \beta} = \frac{\sinh c}{\sin \gamma}$
1st law of cosines	$\cos \gamma = \frac{\cos c - \cos a \cos b}{\sin a \sin b}$	$\cos \gamma = \frac{\cosh a \cosh b - \cosh c}{\sinh a \sinh b}$
2nd law of cosines	$\cos c = \frac{\cos \alpha \cos \beta + \cos \gamma}{\sin \alpha \sin \beta}$	$\cosh c = \frac{\cos \alpha \cos \beta + \cos \gamma}{\sin \alpha \sin \beta}$
area	$(\alpha + \beta + \gamma) - \pi$	$\pi - (\alpha + \beta + \gamma)$

REMARK 4.4.1 (Labeling conventions). The conventions for labeling a (hyperbolic) triangle are as follows. Capital letters for vertices, lower case for side lengths and Greek for angle measures. If a vertex is labeled A , then the size of the angle at A is called α and the length of the opposite side is called a . A fully labeled triangle is show in Figure 3.




FIGURE 3. Conventions for labeling a hyperbolic triangle.

LEMMA 4.4.2 (Law of sines). *If ABC is an arbitrary hyperbolic triangle, then*

$$\frac{\sinh a}{\sin \alpha} = \frac{\sinh b}{\sin \beta} = \frac{\sinh c}{\sin \gamma}$$

LEMMA 4.4.3 (Law of cosines; first version). *If ABC is an arbitrary hyperbolic triangle, then*

$$\cos \gamma = \frac{\cosh a \cosh b - \cosh c}{\sinh a \sinh b}$$


DEFINITION 4.4.4 (Thin triangles). A triangle ABC is called k -thin if every point on any side of the triangle is at most k units from some point on one of the other two sides of the triangle.

THEOREM 4.4.5. *Every triangle in hyperbolic n -space is $\log(1 + \sqrt{2})$ -thin.*

PROOF. Let ABC be an arbitrary hyperbolic triangle and let x be a point on the AB side of the triangle. In the Poincaré disc model we can use the isometries of hyperbolic space to arrange the picture so that x lies at the origin, and the side AB lies along the real axis. The ideal points at either end of this geodesic are 1 and -1 . Without loss of generality we can assume that C has a non-negative imaginary part and a non-positive real part. Under these conditions we can argue that the line $y = x$ intersects line segment AC before it intersects the line connecting 1 and i . (argue this)

The latter distance is $\log(1 + \sqrt{2})$. To see this notice that the line connecting 1 and i is an arc of the unique circle that intersects the unit circle at right angles. In this case the circle is obvious. It is centered at $1 + i$ and has radius 1 . Thus the point y of intersection is (Euclidean) distance $\sqrt{2} - 1$ from the origin. The estimate now follows from the Example above. \square

look up whether this really is δ -thin or δ -slim, or one of the other equivalent notions.



4.4.1. $PSL_2(\mathbb{Z})$. (put most of this into the exercises) Farey fractions, Ford circles, Apollonius, binary quadratic forms (Conway's book [8])

REMARK 4.4.6 (Comment on history).

4.4.2. **Historical Notes.** As most undergraduates learn at some point, Euclid begins the first book of the Elements with a list of "common notions" that starts as follows.

1. A straight line may be drawn from any point to any other point.
2. A finite straight line may be extended continuously in a straight line
3. A circle may be drawn with any center and any radius
4. All right angles are equal

5. If a straight line falling on two straight lines makes the interior angles on the same side less than two right angles, the two straight lines, if extended indefinitely, meet on the side where the angles are less than two right angles.

Euclid's "Fifth postulate", as it came to be called, clearly stands out as a notion different in quality from those preceding it.

The existence of non-Euclidean geometries came as a shock to many 19th century mathematicians, but in retrospect it seems quite natural. In this chapter we review some of the standard results about the hyperbolic plane and its isometry group.

4.4.3. Annotated Bibliography. The material for this chapter is quite classical. We have used and referred to the following works.

1. Jim Anderson "Hyperbolic Geometry" \square
2. Riccardo Benedetti and Carlo Petronio "Lectures on Hyperbolic Geometry" \square
3. John Ratcliffe "Foundations of Hyperbolic Manifolds"
4. David Mumford, Caroline Series, David Wright "Indra's Pearls"

CHAPTER 5

Gromov's Hyperbolic Groups

(The basic model for this chapter should be the good concise discussion in Chapter *III.Γ* of Bridson-Haefliger [4], pp.399–437 or the section on hyperbolic groups in Roe [25].)

(From the outline.tex file)

The key points to highlight are:

- Opening gambit: properties of hyperbolic manifolds.
- 1-skeleton of a hyperbolic manifolds, quasi-geodesic triangles are thin. quasi-geodesics are uniformly close to geodesics.
- Basic definitions.
- Spaces first: δ -hyperbolic space.
- thin triangles, 4-point (quasi double max), [slim triangles. incenter, tripod, bigons, Gromov inner product one, etc. in exercises?]
- Examples. variably curved manifolds. universal cover a wedge of hyperbolic manifolds. $c(3)$ - $t(7)$, twisted figure 8 knot.
- Quasi-geos fellow travel.
- exponential divergence of geodesics.
- Nicening δ -hyperbolic spaces, Rips complex.
- Gromov hyperbolic groups. defined as groups acting geometrically on δ -hyp spaces. Then apply Rips construction to preimage of a basepoint to get a geometric action on the Rips complex.
- Cor: hyp- δ nice finiteness properties.
- Linear IP = Hyperbolic. (subquadratic δ -hyp)
- Dehn's algorithm = Hyperbolic
- no $Z + Z$

5.0.4. Exercise:

- Prove the various definitions are equivalent.
- R-trees as connected 0-hyperbolic spaces.
- Tropical stuff
- Twisted figure 8 knot doesn't embed in a 3-manifold.

5.0.5. Maybe stuff.

- Finitely many cone types
- Free subgroups (g^n, h^n)
- automaticity.
- elementary hyperbolic
- quasiconvexity of Z subgroups.
- translation length

5.1. δ -hyperbolic spaces

The most basic objects are δ -hyperbolic spaces. Recall that a *geodesic* is a length minimizing curve and that a *geodesic metric space* is a metric space in which every pair of points is connected by a geodesic.

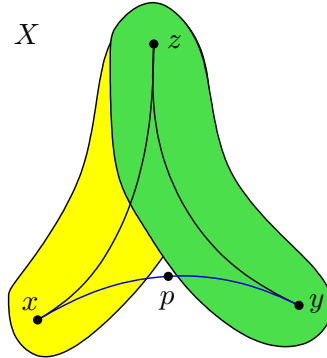


FIGURE 1. A triangle which is not quite δ -thin.

DEFINITION 5.1.1 (δ -hyperbolic). A geodesic metric space X is δ -hyperbolic if there is a fixed $\delta \geq 0$ such that for all points $x, y, z \in X$ and for all geodesics connecting x, y , and z and for all points p on the chosen geodesic connecting x to y , the distance from p to the union of the other two geodesics is at most δ .

A space is called *hyperbolic* if it is δ -hyperbolic for some value of δ .

(Remark on the dangerous aspects of this terminology. Mention \mathbb{H}^3 minus horoballs, complement of the figure 8 knot, etc.)

5.2. Quasi-geodesics

DEFINITION 5.2.1 (Quasi-geodesics). A *quasi-geodesic* is an isometrically embedding of an interval of the reals.

As we noted in the earlier section on quasi-isometries, in a typical metric space such as \mathbb{R}^2 , a quasi-isometric embedding of an interval need not look anything like a geodesic between its endpoints.

EXAMPLE 5.2.2 (Quasi-geodesics are necessarily close). If α and β are quasi-geodesics that start and end at the same endpoints, then there

The following is an easy consequence of δ -hyperbolicity.

PROPOSITION 5.2.3. *Let X be a δ -hyperbolic geodesic metric space, let α be a geodesic from x to y , and let β be an arbitrary rectifiable curve from x to y . If p is a point on α , and $\ell(\beta) \leq 2^{N+1}$, then $d(p, \beta) \leq N\delta + 1$.*

PROOF. The proof is an easy induction. When $N = 0$, $\ell(\beta) \leq 2$. Since α is a geodesic, its length is at most 2 and p is within 1 unit of one of the endpoints. On the other hand, if $2^N < \ell(\beta) \leq 2^{N+1}$, then by drawing a geodesic triangle connecting x, y and the point z half-way along the curve β , we can find a point q on one of these sides that is at most δ from p . Applying the induction hypothesis completes the proof. \square

COROLLARY 5.2.4. “Circles” in δ -hyperbolic spaces have circumferences that are exponential functions of their radii.

THEOREM 5.2.5 (Stability of quasi-geodesics). *If X is a δ -hyperbolic space then there exists an M depending only on δ and K such that every pair of coterminous K -quasi-geodesics are within M of each other.*

Lead up to this.

THEOREM 5.2.6 (Stability of hyperbolicity). *If X is δ -hyperbolic and X' is quasi-isometric to X , then X' is δ' -hyperbolic for some δ' .*

PROOF. Take a geodesic triangle in X' , pull it back to a quasi-geodesic triangle in X , relate it to a geodesic triangle in X , apply the δ -hyperbolic condition and then push everything forward again to X' . A careful accounting of the constants involved shows that X' is δ' -hyperbolic. \square

We are now ready to define the notion of a hyperbolic group.

DEFINITION 5.2.7 (Gromov hyperbolic groups). A finitely generated group G is called *Gromov hyperbolic* or *word hyperbolic* (or simply *hyperbolic* when there is no danger of confusion) if its intrinsic geometry is hyperbolic

COROLLARY 5.2.8. *If G and G' are quasi-isometric groups and G is Gromov hyperbolic, then G' is Gromov hyperbolic.*

LEMMA 5.2.9. *If M^n is a closed and compact hyperbolic n -manifold then $\pi_1(M)$ is a Gromov hyperbolic group.*

REMARK 5.2.10 (Hyperbolic groups and hyperbolic manifolds with boundary). It is an unfortunate accident of terminology that the fundamental group of a compact hyperbolic manifold with boundary need not be a Gromov hyperbolic group. In particular, it is relatively easy to construct a hyperbolic 3-manifold M with a torus boundary component. In this situation, $\pi_1(\partial M) = \mathbb{Z} \times \mathbb{Z}$ injects into $\pi_1(M)$, proving that $\pi_1(M)$ is not a Gromov hyperbolic group.

(an example of this phenomenon is the complement of the figure 8 knot)

5.3. Equivalent definitions

Basic Definitions and their equivalence. Plus key properties. Dehn’s algorithm.

DEFINITION 5.3.1 (Inradius).

DEFINITION 5.3.2 (Tripods).

DEFINITION 5.3.3 (Double maximum).

DEFINITION 5.3.4 (Roughly linear). Let X be a metric space and let ϵ be a fixed constant. If $d_X(x, z) + d_X(z, y) - d_X(x, y) \leq \epsilon$, then z is ϵ -quasi-between x and y and x, z and y are *roughly linear*. Note that in a geodesic metric space, z is 0-quasi-between x and y iff there is a geodesic from x to y through z .

DEFINITION 5.3.5 (Vertex hyperbolic). (discuss the difference between being δ -hyperbolic and having a δ -hyperbolic vertex metric. So long as there is a cocompact group action, the two notions are equivalent.)

PROPOSITION 5.3.6. *If X is a geodesic metric space which is δ -hyperbolic and L is a subset of X which is ϵ -quasi-onto, then the induced metric on L is $\delta + 2\epsilon$ roughly thin.*

DEFINITION 5.3.7 (Maximum almost occurs twice). Let X be a geodesic metric space. It has the δ double maximum property if for all quadruples of points, x, y, z and w , the following inequality is satisfied:

$$d_X(x, y) + d_X(z, w) \leq \max\{d_X(x, z) + d_X(y, w), d_X(x, w) + d_X(y, z)\} + 2\delta$$

(discuss the situation in trees and then convert it to δ -hyperbolic spaces) (comment on the relation with tropical geometry).

(where should this go?)

THEOREM 5.3.8 (Characterization theorem). *A group G acts geometrically on hyperbolic 3-space \mathbb{H}^3 if and only if its Cayley graph is quasi-isometric with \mathbb{H}^3 . (Theorem 11.4 from [7] by Cannon and Cooper)*

5.4. Rips' complex

DEFINITION 5.4.1 (Discrete metric spaces). A *discrete metric space* is any metric space where the induced topology is discrete. I.e. for every $x \in X$ there is an ϵ such that $B_\epsilon(x)$ only contains x . Examples include the 0-skeleton of any metric cell complex in the induced metric.

(there are minor restrictions that need to be made, but these can be handled in the definition of a metric cell complex).

DEFINITION 5.4.2 (Rips complex). Let L be a discrete metric space and let $d > 0$ be a fixed constant. The simplicial complex $\text{RIPS}_d(L)$, known as the *Rips complex*, is defined as follows. Start with a 0-skeleton labeled by the points of L . Next add an edge between distinct vertices iff the distance between the corresponding points in X is at most d . In other words, if we let u and v denote distinct elements of L inside X as well as the corresponding vertices in the simplicial complex we are constructing, then the vertices u and v are connected by an edge iff $d_X(u, v) \leq d$. The last step is to attach all possible simplices to this graph so that the final result is a flag complex. (say more and/or define flag elsewhere)

THEOREM 5.4.3 (Contractibility). *If X is a δ -hyperbolic metric space and L is an ϵ -quasi-onto discrete subset of X , then for all $d \geq 4\delta + 2\epsilon$, the simplicial complex $\text{RIPS}_d(L)$ is contractible.*

PROOF. Fix $d \geq 4\delta + 2\epsilon$ and let $Y = \text{RIPS}_d(L)$. By Whitehead's theorem (Theorem A.4.8) it is sufficient to show that every map $\mathbb{S}^n \rightarrow Y$ is null-homotopic. Notice that if C is any compact space, any map $f : C \rightarrow Y$ can be factored as $C \rightarrow K \hookrightarrow Y$ where K is a finite subcomplex containing the image and the second factor is the inclusion map. In particular, \mathbb{S}^n is compact, and it is thus sufficient to show that for every finite subcomplex K in Y , the inclusion map is null-homotopic. As a base case of an induction, notice that for every pair of vertices u and v in K , $d_X(u, v) \leq d$, then K is contained in a simplex of Y , and the inclusion map is thus null-homotopic.

The strategy at this point is to use the δ -hyperbolicity to construct explicit homotopies from the inclusion map $K \rightarrow Y$ (with K finite) to another map whose

image is contained in a proper subcomplex of K . It will be clear from the construction that this procedure can be applied systematically to reduce the vertex diameter of the image of K to the point where the base case can complete the proof.

(somewhere assume that K is the full subcomplex on this vertex set)

Let K be a finite simplicial complex and let x be one of its vertices. Consider the value $m = \max\{d_X(x, v) | v \in K^{(0)}\}$, and let y be a vertex of K that achieves this maximal distance. If $m \leq d/2$ then every pair of vertices in K are at most d apart and we are in the base case. Thus we may assume that $d_X(x, y) \geq d/2$. Consider a geodesic in X connecting x and y and let z be a point in L that is at most ϵ from the point on the geodesic distance $d/2$ from y and $m - d/2$ from x . See Figure XXXX. We define a map from K to K which fixes the 0-skeleton of K except that it sends y to z . The key claim is that this map on the vertices can be extended to a simplicial map. The only concern involves the simplices containing y . For example, if there is an edge connecting w and y , then there needs to be an edge connecting w and z . Thus we need to know that $d_X(x, y) \leq d$ implies $d_X(w, z) \leq d$. Conversely, this is sufficient since K is a flag complex.

Using the quasi-double maximum version of δ -hyperbolicity, we know that $d_X(w, z) + m \leq \max\{d + m - d/2 + \epsilon, m + d/2 + \epsilon\} + 2\delta = m + d/2 + \epsilon + 2\delta \leq m + d$.

(Final step, mumble something about how this map is homotopic to the identity map.) \square

some consequences

Mention that group actions on X that preserve L extend to group actions on $\text{RIPS}_d(L)$.

5.5. Finite subgroups

As a warm-up do finite groups acting on trees and finite groups acting on Euclidean space.

LEMMA 5.5.1. *If G is a finite group acting by isometries on a tree X then the $X^G \neq \emptyset$ where X^G denotes the set of global fixed points.*

PROOF. Let x be a point in X and consider the smallest subtree of X containing the orbit of x under the action of G . This is a finite tree T . Since G is acting by isometries, it is a bijection on the vertex set and sends leaves to leaves. Thus the subtree T' of T obtained by removing all the leaves of T (if non-empty) is also preserved by the action of G . Continuing in this way we find a tree T with no internal vertices that is preserved by the action of G . This T is either a single vertex or a single edge. Done. \square

LEMMA 5.5.2. *If G is a finite group acting by isometries on $X = \mathbb{R}^n$, then $X^G \neq \emptyset$.*

PROOF. (Find the set of best centers and show there's only one.) \square

With these two proofs as preamble, it is perhaps surprising that a finite group acting on a Rips complex need not have a global fixed point. The first step can

Argue that finite groups stabilize some bounded orbit. (maybe allude to Bob Olivers strange theorem [23]).

On the other hand, stabilizers are clearly finite.

5.5.1. Final remarks. Reasons for studying negatively curved groups (from [7] p.319)

1. In a certain formal sense, negatively curved groups are the most common among finitely presented groups
2. Among 2 and 3 manifolds, the negatively curved manifolds are far and away the most common.
3. Computationally the negatively curved groups exhibit behavior just one step more complicated than free products and HNN extensions.
4. Negatively curved groups are surprisingly stable under natural constructions.

EXAMPLE 5.5.3 (Virtually Free Groups). A finitely generated group G is called *virtually free* if it contains a finite index subgroup which is free.

From [7]

The task of geometric group theory is to study geometric group actions and group actions that are nearly geometric in some sense. The goal is to determine which groups can act on which geometries and to determine the combinatorial properties of the groups which can so act.

Ends and Boundaries

When the real line is considered intuitively, most observers would agree that it has two distinct “ends” whereas something qualitatively different is going on with the portion of the complex plane far away from the origin. The technical concept of an end of a space is a way to formalize this intuition.

Use the second half of Chapter III.Γ of Bridson-Haefliger [4] as a reference

6.1. The Space of Ends

Opening gambit: number of ends of R , R^2 , and a trivalent tree.
The key points to highlight are:

- Ends. Categorical definition
- Proper rays. Quotient definition
- (Freudenthal-Hopf) Every group has 0,1,2 or infy ends.
- 0 iff finite
- 2 iff virt. Z
- mention stallings, and remark that one-ended is where its at.
- The space of ends has a topology. (its perfect)
- QI - i homeo.

DEFINITION 6.1.1 (Ends via projective limits). If X is a metric space, then let $U(X)$ denote its set of unbounded path components. Given a metric space X , the sets $U(X \setminus C)$ where C is an arbitrary compact subspace form a projective system. The ends of X are the projective limit of this system. Alternatively, (give the proper ray definition)

DEFINITION 6.1.2 (Proper rays). A map $f : X \rightarrow Y$ is *proper* if the inverse image of each compact set $S \subset Y$ is a compact subset of X .

DEFINITION 6.1.3 (Ends via equivalence relations). The ends of X can also be defined using an equivalence relation on the collection of all proper rays in X . Call two rays *equivalent with respect to C* if they both eventually remain in the same unbounded component of $X \setminus C$. Two rays are *equivalent* if and only if they are equivalent with respect to every possible compact subset C .

(should we prove these two definitions are the same?)

Notice that if $C' \supset C$ then there is a well-defined map $U(X \setminus C') \rightarrow U(X \setminus C)$. These are the maps in the diagram with respect to which we taking the projective limit. (awkwardly said)

THEOREM 6.1.4 (Freudenthal-Hopf). *Every finitely generated group has either 0, 1, 2, or infinitely many ends.*

Categorical definition
Proper rays. Quotient definition
(Freudenthal-Hopf) Every group has 0,1,2 or infy ends.
0 iff finite
2 iff virt. Z
mention stallings, and remark that one-ended is where its at.
The space of ends has a topology. (its perfect)
QI - i homeo.
Projective limits.

PROOF. Notice that as larger and larger compact subsets C of X are removed, the cardinality of $U(X \setminus C)$ can only increase and the number of ends is the supremum of these increasing values. Thus it suffices to prove the following. If there exists a compact set C such that $X \setminus C$ has at least $k \geq 3$ unbounded components, then there exists a compact set $C' \supset C$ such that $X \setminus C'$ has at least $k' > k$ unbounded components. (add in the standard move things around proof) \square

LEMMA 6.1.5. *Every quasi-isometry $f : X \rightarrow Y$ between geodesic metric spaces induces a homeomorphism $f_* : \text{ENDS}(X) \rightarrow \text{ENDS}(Y)$. Moreover, the map $QI(X) \rightarrow \text{Homeo}(\text{Ends}(X))$ is a group homomorphism.*

As a corollary we have the following.

PROPOSITION 6.1.6. *The number of ends is a geometry property.*

6.2. The Boundary at Infinity

- Boundaries.
- Geodesics rays in R^2
- Gromov-Hausdorff distance
- geo / equiv = boundary
- same for H^2 . Note that in the hyp case, we can use quasi-geo.
- Still works. δ -hyperbolic. Still works.
- connected components of the boundary = ends.
- Remark on Mostow rigidity.

Various definitions: boundaries of hyperbolic groups,
(look at the AIM conference description for ideas)

Exercises

- π_1 at infinty. (R^2)
- f.g. x f.g. is one-ended.
- Bestvina-Kapovich-Kleiner [Inv02]. Embeddings. $K_{3,3}$ and $F_2 \times F_2$.

6.2.1. Maybe stuff. Stallings infnty ends - i splits. (torsion-free version only?)

Splittings and Quasiconvexity

(I've simply joined together old ch 8 and 9 and tried out some section titles. Nothing is where it should be yet)

One of the main goals of this chapter is to introduce enough of the ideas and results so that afterwards, Serre's book on trees [26] should be readable.

7.1. Actions on trees

Stallings's theorem and actions on \mathbb{Z} -trees seemed like a topic which was something that wasn't so elementary as to be presumable, but so basic that it seemed strange to leave it to a case study. Plus, if we include a short chapter on this there are many opportunities to build on it later. Although the definitions do not seem inherently metric, the content of the theorem is that whether a group has a non-trivial splitting only depends on the purely geometry notion of the number of ends that its geometry possesses.

HNN extensions, amalgamated free products, Britton's lemma, actions of groups on (\mathbb{Z} -)trees, splittings, graphs of groups, Stallings theorem, complexes of groups

7.2. Scott and Wall's approach

7.3. Bass-Serre Theory

7.4. Amalgamations and Quasiconvexity

relating the metric on an amalgamated free product to the metrics on all sides. Quasiconvexity, the \mathbb{Z} 's inside \mathbb{Z}^2 , gen. set dependent in general, not so in a hyper. group.

(Higman, Neumann, Neumann)

Epilogue: Where to go from here

All of mathematics is a tale about groups.

Henri Poincaré ¹

Moved from the intro

Some examples include the book by Martin Bridson and André Haefliger on the geometry of nonpositively curved spaces [4], the six-author work on automatic groups [10] and, of course the original long article by Misha Gromov [14] where many key results about hyperbolic groups were outlined for the first time. More recent examples include a book by Mike Davis on Coxeter groups and one by Ross Geoghegan on topological methods in group theory.

(Detailed pointers to other sources of information. Sort of a bibliographic essay.)

[USE the S-book stuff!!!]

7.4.1. Notes on History. Moved from Ch.1: Here's a quasi-random list of names/bits of history that I might want to work in as I go.

Names: (Topological / combinatorial viewpoint)

Prehistory: Niels Abel, Evariste Galois, Dyck, Arthur Cayley,

Founders: Max Dehn, Henri Poincaré

First steps: Tietze, Nielsen, Reidemeister, Schreier, van Kampen,

Coxeter, Baumslag, Miller, Boone

Grushko, Adyan, Novikov

Early group theory books: Burnside [6], Hall [15], Kurosh [17, 18], Magnus-Karrass-Solitar [20], Lyndon-Schupp [19], Stillwell [29].

¹The context for this quote is provided by Hawkins: "Thus when Lie visited Paris in 1882, he found Poincaré already convinced of the importance of group theoretic ideas in geometry - and throughout mathematics. According to Lie's report to Klein in a letter from Paris, Poincaré had explained to Lie that all of mathematics was a tale about groups. Lie indicated that Poincaré did not know the Erlanger Programm, and so he described it to him. There is no evidence, however, that Poincaré ever studied the Programm or was particularly influenced by it since he himself had already arrived at a group theoretic interpretation of geometry." [Hawkins 1984, p. 447-448]

APPENDIX A

Algebraic Topology

This appendix is a brief review of basic algebraic topology. The idea is to make explicit the foundations on which geometric group theory is built and to establish standard notation and terminology.

A.1. Cell complexes and Euler characteristics

The notion of a cell complex is flexible enough to construct complicated spaces, but restrictive enough to avoid pathological examples such as the topologist's sine curve or the Hawaiian earring. The most basic cell complexes are the simplicial ones.

DEFINITION A.1.1 (Simplicial complexes). An *abstract simplicial complex* is a collection \mathcal{S} of finite subsets of a fixed set V such that $\tau \subset \sigma \in \mathcal{S}$ implies $\tau \in \mathcal{S}$. The elements of V are called *vertices* and the elements of \mathcal{S} are called *simplices* because of the shapes they produce in the geometric realization. Let U be a real vector space with a basis whose elements are indexed by V . To each $\sigma \in \mathcal{S}$ we associate the subset of U formed by all nonnegative linear combinations

$$\sum_{v \in \sigma} \lambda_v v \quad \text{with} \quad \sum_{v \in \sigma} \lambda_v = 1.$$

If σ has n elements, then this set is an ordinary $(n-1)$ -simplex. (For $n = 0, 1, 2$ and 3 , an n -simplex is a point, an interval, a triangle and a tetrahedron.) The union of the simplices associated to each $\sigma \in \mathcal{S}$ is the (topological) *geometric realization*. For convenience we use \mathcal{S} for both the abstract simplicial complex and for its geometric realization, and we use σ for both a finite subset of V and the topological simplex it contributes to \mathcal{S} .

Because of their concrete description, simplicial complexes are nice to work with, but there are situations where they are unnaturally restrictive. A more flexible construction involves iteratively attaching cells.

DEFINITION A.1.2 (Attaching spaces along subspaces). If X_0 and X_1 are topological spaces, A is a subspace of X_1 and $f : A \rightarrow X_0$ is a continuous map, then we can form a quotient of $X_0 \sqcup X_1$ by identifying each point $a \in A$ with its image $f(a) \in X_0$. The resulting space X is denoted $X_0 \sqcup_f X_1$ and it is described as the space X_0 with X_1 *attached along A via f* . See Figure 1.

DEFINITION A.1.3 (Cell complexes). The notion of a *cell complex* or *CW complex* (terms we use interchangeably) is defined inductively, dimension by dimension. A *0-dimensional cell complex* is an arbitrary set of points called *0-cells* with the discrete topology. An *n -dimensional cell complex* or *n -complex* X is constructed by




FIGURE 1. A schematic representation of spaces and maps used to construct $X = X_0 \sqcup_f X_1$.

attaching a disjoint union of n -discs along their boundary spheres to an already constructed $(n-1)$ -dimensional cell complex X^{n-1} . In particular, let $E^n = \coprod \mathbb{D}^n$ be a disjoint union of n -discs and for each n -disc fix a continuous map $f : \partial\mathbb{D}^n \rightarrow X^{n-1}$, called the *attaching map*. There is then an induced map $F : \partial E^n \rightarrow X^{n-1}$ and the complex $X = X^{n-1} \sqcup_F E^n$ is an n -dimensional cell complex. For any $j < n$, the space X^j embeds into X and thus X^j can be viewed as a subspace of X ; it is called the *j -skeleton of X* . Since unadorned superscripts often indicate dimension, we use $X^{(j)}$ to denote the j -skeleton of a cell complex X .

The interiors of the n -discs map homeomorphically into X and these images are called the *n -cells of X* . Since the points of X can be partitioned into $X^{(n-1)}$ and the n -cells of X , by induction, the set X can be viewed as a disjoint union of its j -cells, $0 \leq j \leq n$. For convenience, we often refer to 0-cells and 1-cells as *vertices* and *edges* respectively, and 1-complexes as *graphs*. A cell complex is *finite* if it has only finitely many cells.

Infinite dimensional cell complexes can also be constructed. Given cell complexes $X^0 \subset X^1 \subset \dots \subset X^k \subset \dots$ where each X^k is a k -dimensional cell complex constructed by attaching k -discs along their boundary to the previous complex in the list, we let X denote the union of these nested spaces and declare $U \subset X$ to be an open subset of X iff $U \cap X^k$ is open in X^k for all $k \geq 0$.

REMARK A.1.4 (Dimension -1). The inductive construction described above could actually have started one step earlier by declaring the empty topological space to be a (-1) -dimensional cell complex X^{-1} . The 0-dimensional cell complexes are constructed by attaching a disjoint union of 0-discs along their boundary spheres to this (-1) -dimensional cell complex. Because \mathbb{D}^0 is the entire space \mathbb{R}^0 , it is open in the topology of \mathbb{R}^0 and thus its boundary is empty. This is completely consistent with the idea that $\partial\mathbb{D}^0 = \mathbb{S}^{-1}$ since \mathbb{S}^{-1} , by definition, is the set of vectors in \mathbb{R}^0 of length 1, which is, once again, empty. As a consequence, $E^0 = \coprod \mathbb{D}^0$ is a set of points with the discrete topology, ∂E^0 is the empty set, and the only choice we have for $F : \partial E^0 \rightarrow X^{-1}$ is the empty map between empty spaces. The resulting space $X = X^{-1} \sqcup_F E^0$ is then a set of points with the discrete topology. This convention is often useful. For example, one can define the k -cells of X as the images of the interiors of the k -discs under their attaching maps with no need to single out the 0-cells of X for separate treatment.

A *subcomplex* of a cell complex X is a union of j -cells that is closed in the topology of X . The various skeleta are obvious examples of subcomplexes, but there are many others. The fact that cell complexes are well-behaved is illustrated by the following theorem:

THEOREM A.1.5 (Cell complex properties). *Every cell complex X is normal and Hausdorff. It is connected iff it is path-connected iff its 1-skeleton is connected. It is compact iff it has only finitely many cells. And every compact subspace of X is contained in some finite subcomplex.*

As a consequence of Theorem A.1.5, the image of any k -disc, being compact, is contained in some finite subcomplex. Historically, this property was called *closure-finite* since the finite subcomplex contains the closure of the corresponding k -cell. Cell complexes were originally defined in a way that relied heavily on the *closure-finite* property and the use of the *weak topology* on the union. Hence the name *CW complex*.

DEFINITION A.1.6 (Euler characteristics). Let X be a finite cell complex and let c_i denote the number of i -cells that X contains. The (ordinary) *Euler characteristic* of X is equal to $\sum_{i \geq 0} (-1)^i c_i$ and it is denoted $\chi(X)$. The *reduced Euler characteristic* of X is a slight modification of the Euler characteristic where we consider the empty set as a (-1) -dimensional cell of X . When viewed in this way $c_{-1} = 1$ and the alternating sum over the cells of X yields $\sum_{i \geq -1} (-1)^i c_i = \chi(X) - 1$. The reduced Euler characteristic is denoted $\tilde{\chi}(X)$. Despite the redundancy, it is useful to have both $\chi(X)$ and $\tilde{\chi}(X)$ available. For example, $\tilde{\chi}(\mathbb{S}^n) = (-1)^n$ and $\chi(X \times Y) = \chi(X) \times \chi(Y)$. Neither pattern can be stated as cleanly in the other notation.

There is great freedom in the definition of a cell complex, as the nature of the attaching maps is not very restrictive. In particular, it is not the case that every cell complex is homeomorphic to a simplicial complex (Exercise 4). In this book we often restrict ourselves to a simpler situation where the spaces are always homeomorphic to simplicial complexes (Exercise 5).

DEFINITION A.1.7 (Cellular maps and combinatorial complexes). A map $Y \rightarrow X$ between cell complexes is *cellular* if its restriction to each cell of Y is a homeomorphism onto a cell of X . A cell complex X is *combinatorial* if a cell structure can be imposed on the domain of each attaching map of each k -cell of X so that the result is a cellular map between cell complexes. In the literature, combinatorial cell complexes are also known as regular cell complexes.

A.2. Fundamental groups and van Kampen's theorem

Next we shift our attention from spaces to maps.

DEFINITION A.2.1 (Homotopic maps). Two maps $g, h : X \rightarrow Y$ are *homotopic* if there is a map $F : X \times I \rightarrow Y$ (a *homotopy*) such that $g = f_0$ and $h = f_1$ where $f_t : X \rightarrow Y$ is the map defined by the equation $f_t(x) = F(x, t)$. We write $g \cong h$ when g and h are homotopic maps. When $g : X \rightarrow Y$ is homotopic to a constant map (i.e. a map whose image is a single point of Y), then g is *null-homotopic*. If A is a subspace of X and there is a homotopy $F : X \times I \rightarrow Y$ such that $F(a, s) = F(a, t)$ for all $s, t \in I$, then g and h are *homotopic relative to A* .

Recall that a *based space* is a pair (X, x) where X is a topological space and x is a point of X and a *based map* is a map from (Y, y) to (X, x) is a map $f : Y \rightarrow X$ with $f(y) = x$. Such a map is denoted $f : (Y, y) \rightarrow (X, x)$.

DEFINITION A.2.2 (Fundamental groups). The *fundamental group* of a cell complex X based at a point x is the set of equivalence classes of paths in X that start and end at x where two paths are considered equivalent if they are homotopic relative to their endpoints. The multiplication of two such classes is defined by taking the equivalence class of the concatenation of representatives. If x and \hat{x} are two points in the same connected component of X , then any path connecting them induces an isomorphism $\pi_1(X, x) \approx \pi_1(X, \hat{x})$. We should note, however, that the exact isomorphism usually depends on the choice of a connecting path. Since fundamental groups of connected cell complexes are well-defined up to isomorphism, basepoints are occasionally suppressed.

One of the key properties of this construction is its functoriality.

PROPOSITION A.2.3 (Functoriality). *There is a functor from the category of based topological spaces to the category of groups such that the image of (X, x) is $\pi_1(X, x)$ and the image of the map $f : (X, x) \rightarrow (Y, y)$ is the group homomorphism $f_* : \pi_1(X, x) \rightarrow \pi_1(Y, y)$. In particular, if $f = gh$ as based maps, then the corresponding group homomorphisms satisfy $f_* = g_*h_*$, and if f is the identity map then f_* is the identity group homomorphism.*

To illustrate the benefits of functoriality, consider a retraction onto a subspace. Let A be a subspace of X and let $i : A \rightarrow X$ be the inclusion map. Recall that a map $r : X \rightarrow A$ is called a *retraction* if $ri = \mathbf{1}_A$ and it is a *deformation retraction* if, in addition, ir is homotopic to $\mathbf{1}_X$ relative to the subspace A .

PROPOSITION A.2.4 (Retractions and fundamental groups). *If A is a connected subspace of a connected space X , $i : A \rightarrow X$ is the inclusion map and $r : X \rightarrow A$ is a retraction, then r_* is surjective and i_* is injective. In particular, $\pi_1(A, a)$ can be viewed as a subgroup of $\pi_1(X, i(a))$.*

PROOF. Pick $a \in A$. By Proposition A.2.3, $r_*i_* = \mathbf{1}_G$ where $G = \pi_1(A, a)$. The rest follows from the fact that $\mathbf{1}_G$ is a bijection. \square

add connecting text

PROPOSITION A.2.5. *For every subcomplex A of a cell complex X there is a small open neighborhood N of A such that N deformation retracts to A . In particular, N is homotopy equivalent to A .*

PROPOSITION A.2.6. *Let X be a cell complex. Then every map $f : \mathbb{S}^1 \rightarrow X$ can be homotoped to a map $\hat{f} : \mathbb{S}^1 \rightarrow X^{(1)}$. In particular, if $\iota : X^{(1)} \hookrightarrow X$ is the inclusion of the 1-skeleton into X , then the induced map $\iota_* : \pi_1(X^{(1)}) \rightarrow \pi_1(X)$ is a surjection.*

DEFINITION A.2.7 (Wedge products). Let $\{(X_\alpha, x_\alpha)\}$ be a collection of based spaces. The *wedge product* of this collection is the quotient of their disjoint union in which all of the base points have been identified:

$$\bigvee X_\alpha = \prod_{\alpha} X_\alpha / \{x_\alpha \sim x_\beta\}.$$

The resulting space is denoted $\bigvee_{\alpha} X_\alpha$ or $X \vee Y$ when only two spaces are involved.

If X is a cell complex with basepoint x , and X can be expressed as a union of subcomplexes A_α , each of which contains x , then there is a map

$$\phi : \bigvee A_i \rightarrow X$$




FIGURE 2. A decomposition into subcomplexes induces a map from the wedge product of the subcomplexes

defined by making ϕ an isomorphism when restricted to any A_i .

THEOREM A.2.8 (van Kampen’s Theorem). *Let X be a cell complex that can be expressed as a union of path-connected subcomplexes $X = \bigcup A_\alpha$, where for all pairs of distinct indices $A_\alpha \cap A_\beta = C$, for a fixed, path-connected subcomplex C . The resulting map from the wedge product of the pieces, $\phi : \bigvee A_\alpha \rightarrow X$ induces a surjection on the level of fundamental groups: $\phi_* : \pi_1(\bigvee A_\alpha) \twoheadrightarrow \pi_1(X)$. Further, if for each index α we let ι_α denote the induced map $\pi_1(C) \rightarrow \pi_1(A_\alpha)$, then the kernel of ϕ_* is the normal subgroup generated by $\{\iota_\alpha(c)\iota_\beta(c^{-1}) \mid c \in \pi_1(C)\}$.*

COROLLARY A.2.9. *Let X be a cell complex and let $\iota : X^{(2)} \hookrightarrow X$ be the inclusion of its 2-skeleton. Then the induced map $\iota_* : \pi_1(X^{(2)}) \rightarrow \pi_1(X)$ is an isomorphism.*

PROOF. In higher dimensions, when you attach cells it is along 1-connected subspaces, so the kernel is trivial. \square

A.3. Group actions and covering spaces

In the prologue we began our study of the fundamental group of the complement of the trefoil knot, $G \approx \pi_1(\mathbb{S}^3 \setminus K)$, by forming a cell complex \mathcal{D} with $G \approx \pi_1(\mathcal{D})$. In order to study the structure of G we needed to understand not just the structure of \mathcal{D} , but how G acts on $\tilde{\mathcal{D}}$.

DEFINITION A.3.1 (Group actions). A *left action* of a group G on a mathematical structure X is a group homomorphism from G to $\text{AUT}(X)$, the group of all invertible structure preserving maps under function composition. Thus, if X is a topological space, $\text{AUT}(X)$ is the group of all homeomorphisms from X to itself. More explicitly, a left group action of G on X is a function $a : G \times X \rightarrow X$ such that (1) for each $g \in G$, the restriction $g \cdot : X \rightarrow X$ defined by $g \cdot (x) = a(g, x)$ is a homeomorphism from X to itself, (2) $g \cdot (h \cdot (x)) = (gh) \cdot (x)$ for all $g, h \in G$ and for all $x \in X$, and (3) the identity element of G restricts to the identity homeomorphism.

Left group actions are denoted $G \curvearrowright X$, which is read as “ G acts on X ”. The word “left” is usually suppressed since the sidedness of the action is implied by the way that functions are denoted. In order to define a *right action* of G on X we would need to use algebraist notation (i.e. we would have to write $(x)f$ instead of

$f(x)$ to describe the function f applied to the point x). The few occasions where algebraist notation for functions and right group actions are needed are clearly indicated.

DEFINITION A.3.2 (Proper group actions). Let $G \curvearrowright X$, where X is a topological space. The group G is acting *properly discontinuously* on X if for every point $x \in X$ there is a neighborhood U of x such that $\{g \in G \mid g(U) \cap U \neq \emptyset\}$ is finite. The action is *free* if the open set U can always be chosen so that this set contains only the identity element of G . The *stabilizer* of a point $x \in X$ is the subgroup $\text{Stab}(x) = \{g \in G \mid g(x) = x\}$, and a group action is *proper* if all of the point stabilizers are finite.

REMARK A.3.3 (Group actions and categories). The notion of a group action depends on the category used to define $\text{AUT}(X)$. Consider the group $\text{AUT}(\mathbb{S}^1)$. When \mathbb{S}^1 is viewed as a cell complex, the natural maps are cellular maps and $\text{AUT}(\mathbb{S}^1)$ is either a cyclic group of order 2 (in the one 0-cell case) or a finite dihedral group whose order depends on the number of 0-cells in \mathbb{S}^1 ; when \mathbb{S}^1 is viewed as a metric space, the natural maps to are isometries and $\text{AUT}(\mathbb{S}^1)$ becomes the Lie group $O(2)$; and when \mathbb{S}^1 is viewed purely as a topological space, $\text{AUT}(\mathbb{S}^1)$ contains all homeomorphisms from \mathbb{S}^1 to itself, which is quite a large group.

Cite Caligari?

Suppose $\cdot : G \times X \rightarrow X$ is a left group action of a group G on a space X . Because group elements are invertible, every map $g \cdot : X \rightarrow X$ is necessarily one-to-one and onto. Moreover, when X has any additional structure (such as a cell structure, or an orientation on its 1-skeleton, etc.), we shall assume that the action of G preserves this additional structure. In the case of a cell structure, this means that each map $g \cdot$ induces a bijection from the i -cells of X to the i -cells of X .

DEFINITION A.3.4 (Quotients). Given an action $G \curvearrowright X$, the *quotient* of the action is the quotient space formed by identifying $g \cdot x$ with x for each $x \in X$ and $g \in G$. It is denoted $G \backslash X$. A *fundamental domain* for an action $G \curvearrowright X$ is a path connected, closed subset $\mathcal{F} \subset X$ such that $G \cdot \mathcal{F} = X$ with no proper subset of \mathcal{F} satisfying these conditions. When X is a cell complex one can always find a fundamental domain that is a subcomplex, but this is not required. Note that given a fundamental domain \mathcal{F} there is an induced surjection $\mathcal{F} \rightarrow G \backslash X$. A group action $G \curvearrowright X$ is *cocompact* if $G \backslash X$ is compact, or equivalently, if there is a compact fundamental domain.

PROPOSITION A.3.5 (Free actions have quotients). *If $G \curvearrowright X$ is a free left action of a group G on a cell complex X where, by convention, the action respects the cell structure, then there is a well-defined cell structure on its quotient $G \backslash X$.*

A.3.1. Covering spaces. A map $f : Y \rightarrow X$ between path-connected topological spaces X and Y is called a *covering map* when for every $x \in X$ there exists an open set U containing x such that $f^{-1}(U)$ can be written as a disjoint union of open sets U_α where f restricted to each U_α is a homeomorphism. When $f : Y \rightarrow X$ is a covering map then Y is called a *cover of X* . A covering map must be a local homeomorphism, but in general this is not sufficient (Exercise 8). For cell complexes, however, the two concepts are equivalent.

PROPOSITION A.3.6 (Recognizing covers). *If X and Y are connected cell complexes, then $f : Y \rightarrow X$ is a covering map if and only if f is a local homeomorphism.*

If $f : Y \rightarrow X$, $g : Z \rightarrow X$, and $h : Z \rightarrow Y$ are maps such that $f \circ h = g$, then h is called a *lift of g through f* . When f and g are based maps, then we additionally require h to be a based map taking the base point of Z to the basepoint of Y . The definition of a cover is designed to facilitate the creation of lifts.

THEOREM A.3.7 (Map lifting). *Let (X, x) , (Y, y) and (Z, z) be path connected based spaces, let $f : (Y, y) \rightarrow (X, x)$ be a cover and let $g : (Z, z) \rightarrow (X, x)$ be an arbitrary map. When Z is a cell complex there exists a based map $h : (Z, z) \rightarrow (Y, y)$ such that $f \circ h = g$ iff $g_*(\pi_1(Z, z)) \subset f_*(\pi_1(Y, y))$. Moreover, when such a map exists, it is unique.*

Special cases of Theorem A.3.7 have their own names. When Z is a 1-cell, it is called *path lifting* and when Z is a 2-cell it is called *homotopy lifting*. In both cases the condition is trivially satisfied since $g_*(\pi_1(Z, z))$ is the trivial subgroup of $\pi_1(X, x)$. Homotopy lifting is used to show that if f is cover then f_* is injective.

PROPOSITION A.3.8 (Covers and subgroups). *If $f : Y \rightarrow X$ is a covering with $f(y) = x$, then $f_* : \pi_1(Y, y) \rightarrow \pi_1(X, x)$ is an injection. In particular, the fundamental group of Y at y can be viewed as a subgroup of the fundamental group of X at x .*

Let $f : Y \rightarrow X$ be a covering and let $f(y) = x$. The right stabilizers of f (i.e. the maps $g : Y \rightarrow Y$ such that $f \circ g = f$), are called *deck transformations* and they form a *group of deck transformations* under composition. When the group of deck transformations of f acts transitively on the preimages of x , then f is called a *regular covering* and Y is a *regular cover of X* . Regular covers correspond to normal subgroups.

PROPOSITION A.3.9 (Regular covers and normal subgroups). *If $f : Y \rightarrow X$ is a covering with $f(y) = x$, then Y is a regular cover of X iff $f_*(\pi_1(Y, y))$ is a normal subgroup of $\pi_1(X, x)$. Moreover, when Y is a regular cover of X the quotient of $\pi_1(X, x)$ by $f_*(\pi_1(Y, y))$ is isomorphic to the group of deck transformations.*

If $f : Y \rightarrow X$ is a covering, X and Y are connected spaces, and Y is simply connected, then Y is called the *universal cover of X* . An easy application of Theorem A.3.7 shows that universal covers are unique (up to the natural notion of equivalence defined by lifts in both directions whose compositions are identity maps).

THEOREM A.3.10 (Fundamental theorem of covering spaces). *If X is connected topological space that has a universal cover \tilde{X} , then there is a natural bijection between the connected covers of X and the subgroups of $\pi_1(X, x)$.*

(indicate the proof since this uses the quotient by the H -action defined earlier)
Cell complexes, as usual, are extremely well behaved.

PROPOSITION A.3.11 (Recognizing universal covers). *Every connected cell complex has a universal cover. Moreover, if X and Y are connected cell complexes, then Y is the universal cover of X iff Y is simply connected and there exists a local homeomorphism $f : Y \rightarrow X$.*

A.4. Homotopy invariants and Whitehead's theorem

(homotopy type, contractibility, n -connected)

DEFINITION A.4.1 (Homotopy equivalences). A map $f : X \rightarrow Y$ is a *homotopy equivalence* if there exists a map $g : Y \rightarrow X$ such that both compositions are homotopic to the appropriate identity map. In symbols this requires $fg \cong \mathbf{1}_Y$ and $gf \cong \mathbf{1}_X$. Two spaces X and Y are *homotopy equivalent* and have the same *homotopy type* if there exists a homotopy equivalence $f : X \rightarrow Y$. A *homotopy invariant* of a space X is something defined using X where the resulting answer or object depends only on the homotopy type of X .

PROPOSITION A.4.2 (Fundamental groups are homotopy invariants). *If $f : X \rightarrow Y$ is a homotopy equivalence and $f(x) = y$, then $f_* : \pi_1(X, x) \rightarrow \pi_1(Y, y)$ is an isomorphism. In particular, connected spaces with the same homotopy type have isomorphic fundamental groups.*

REMARK A.4.3. It is a basic result from algebraic topology that the Euler characteristic of a finite cell complex only depends on its topology and not on details of its cellular structure. Using cellular homology the alternating sum of the c_i is easily seen to be equal to the alternating sum of the betti numbers of X . But since all homology theories agree on finite cell complexes, and singular homology is insensitive to the cell structure of X , the Euler characteristic only depends on the topology of X .

THEOREM A.4.4 (Invariance of $\chi(X)$). *Euler characteristic is a homotopy invariant. If X and Y are homotopy equivalent spaces and $\chi(X)$ and $\chi(Y)$ can be defined, then $\chi(X) = \chi(Y)$.*

Homology and cohomology are also homotopy invariants, and we will on occasion make use of them. However we do not review their definitions and basic properties as topics such as group cohomology are not a central focus of this book.

There are two common ways to modify a cell complex without changing its homotopy type. One is to collapse a contractible subcomplex and the other is to replace an attaching map with an alternate map homotopic to it. This section is devoted to an application of the first; discussion of the second is postponed until Section 1.1. For a proof of the following results see Chapter 0 in [16].

THEOREM A.4.5 (Collapsing contractible subcomplexes). *If A is a contractible subcomplex of a cell complex X , then the quotient map $X \rightarrow X/A$ is a homotopy equivalence.*

THEOREM A.4.6 (Modifying the attaching maps). *If A is a subcomplex of a cell complex X_1 and $f, g : A \rightarrow X_0$ are homotopic maps, then the spaces $X_0 \sqcup_f X_1$ and $X_0 \sqcup_g X_1$ are homotopy equivalent.*

THEOREM A.4.7 (Contractibility). *If X is a connected topological space, then the following conditions are equivalent.*

1. X has the homotopy type of a point (i.e. X is contractible)
2. the identity map $\mathbf{1} : X \rightarrow X$ is null-homotopic
3. every map $Y \rightarrow X$ is null-homotopic

A space satisfying these conditions is said to be contractible, and contractibility is a homotopy invariant.

PROOF. Exercise 13. □

Theorem A.4.7 is true for arbitrary topological spaces. For connected *cell complexes*, it is sufficient to show that every map $Y \rightarrow X$ where Y is compact is null-homotopic. In fact, it is sufficient to show that for every $n \geq 0$ and each map $\mathbb{S}^n \rightarrow X$ is null-homotopic. That this is implied by the above is clear. That it is sufficient to show contractibility is a part of a nontrivial theorem due to J.H.C. Whitehead.

THEOREM A.4.8 (Whitehead's theorem). *A cell complex X is contractible iff for every $n \geq 0$, each map $\mathbb{S}^n \rightarrow X$ is null-homotopic.*

PROPOSITION A.4.9. *The nested union of n -connected cell complexes is n -connected. More specifically, if $A_0 \subset A_1 \subset \cdots \subset A_k \subset \cdots$ is a nested sequence of n -connected subcomplexes of a cell complex X and $A = \cup_{k \geq 0} A_k$, then A itself is n -connected. As a consequence, the nested union of contractible cell complexes is contractible.*

PROOF. Any map $f : \mathbb{S}^m \rightarrow A$ with $m \leq n$ is contained in a finite subcomplex B . Since each cell of B is contained in some A_i and there are only finitely many cells in B , all of B is contained in some A_i . The fact that A_i is n -connected now implies that f is homotopic to a constant map inside $A_i \subset A$. Thus A is n -connected. The final assertion follows by Theorem A.4.8. \square

There is a family of homotopy invariant properties that sits between being connected and being contractible. The most common is being *simply connected*, that is being path connected and having trivial fundamental group (although some do not require simply connected spaces to be path connected).

DEFINITION A.4.10 (Connectivity). A topological space X is *n -connected* if for all $k \leq n$, each map $\mathbb{S}^k \rightarrow X$ is null-homotopic. Being 0-connected is the same as path connected, and 1-connected is the same as *simply connected*.

REMARK A.4.11 ($\pi_n(X, x)$). For those familiar with the definition of the higher homotopy groups, $\pi_n(X, x)$, it is easy to prove that our definition of n -connectivity is equivalent to the condition that $\pi_i(X, x)$ is trivial for $i \leq n$. See Exercise XXX. (add an exercise where we hint how to make a tail and let it wiggle.)

A.5. Classifying spaces and Hurewicz's theorem

In the prologue we illustrated how one can understand certain facts about certain groups G via their actions on contractible complexes \tilde{X} . In that particular case it was helpful that the action was free and the complex was contractible. An *Eilenberg-MacLane space* for a group G is a cell complex whose fundamental group is G and whose universal cover is contractible. Such a space is also referred to as a $K(G, 1)$ and as a *classifying space* for G .

THEOREM A.5.1 (Eilenberg-MacLane spaces). *For every group G there exists a connected cell complex X whose universal cover is contractible and whose fundamental group is G . Moreover, if X and Y have contractible universal covers and isomorphic fundamental groups, then X and Y are homotopy equivalent.*

The proof that any two $K(G, 1)$ s are homotopy equivalent (*Hurewicz's Theorem*) is a bit too long of a distraction for us. (See Theorem 1B.8 in [16].) The existence claim can be viewed as a topological variation of Cayley's Theorem. Cayley's Theorem states that every group can be faithfully represented as a group of

permutations. The proof constructs an action of G on its own elements via left multiplication. The proof we use below extends this action of G on its own elements, and in the end yields a faithful representation of G as a group of deck transformations of a contractible topological space.

PROOF OF EXISTENCE. To prove the existence of $K(\pi, 1)$ s, start with a vertex set of the form $G \times \mathbb{N}$ and think of the second coordinate as describing the “column” to which the vertex belongs. Extend this vertex set to a simplicial complex by declaring that any finite set of vertices drawn from distinct columns forms the vertex set of a simplex. Call this complex X and notice that it is just the countable join of discrete sets of vertices, each of cardinality $|G|$. In particular, if $|G| = |H|$, then the simplicial complexes built from G and H are the same. Standard tools from algebraic topology, like the Künneth formulas, prove that X is contractible.

Since G acts on itself by left multiplication (Cayley’s Theorem), it also acts on $G \times \mathbb{N}$ by left multiplication applied to the first coordinate. This action preserves columns and any n -tuple of vertices coming from distinct columns will be taken to another n -tuple of vertices coming from (the same) distinct columns. As the action is free when restricted to any column, the action of G on X is also free. Thus the quotient $G \backslash X$ is a $K(G, 1)$. \square

These facts enable one to apply homotopy invariants in the study of groups. We say that a homotopy invariant assertion is true of a group G iff it is true of any (and thus every) Eilenberg-MacLane space for G . In particular, one can declare the homology and cohomology groups of a group to be the homology and cohomology groups of any $K(G, 1)$.

DEFINITION A.5.2 (Finite type). A group G is of *finite type* if it admits a finite $K(G, 1)$ -complex. Equivalently, a group G is of finite type if there is a free, cocompact action of G on a finite dimensional, contractible cell complex.

DEFINITION A.5.3 (Euler characteristics of groups). If G is a group of finite type, then the *Euler characteristic* of G is Euler characteristic of any finite $K(G, 1)$. (If you happen to have a non-finite $K(G, 1)$, when G is in fact of finite type, then the Euler characteristic can still be computed by taking the alternating sum of the betti numbers, which are homotopy invariants.) For example, the fundamental group G of the complement of the trefoil knot has a $K(G, 1)$ described in the prologue. This complex has two vertices, five edges, and three faces, hence $\chi(G) = 0$.

PROPOSITION A.5.4. *Let G be a group with a finite $K(G, 1)$, X . If H is a finite index subgroup of G , then $\chi(H)$ exists, and*

$$\chi(H) = [G : H] \cdot \chi(G)$$

PROOF. The cover \overline{X} of X whose fundamental group is H is a $K(H, 1)$. Since it is a $[G : H]$ -fold cover, if X contains c_i i -cells, then \overline{X} contains $\overline{c}_i = [G : H] \cdot c_i$ i -cells. Thus

$$\chi(\overline{X}) = \sum (-1)^i \overline{c}_i = [G : H] \sum (-1)^i c_i = [G : H] \chi(X).$$

\square

We leave the following corollary as a (fun) exercise.

COROLLARY A.5.5. *Let S_g be the closed orientable surface of genus g , and fix two integers g and h , both greater than 1. Then $\pi_1(S_g)$ is a finite index subgroup of $\pi_1(S_h)$ if and only if $g - 1$ is a multiple of $h - 1$.*

PROPOSITION A.5.6. *Let G be a group of finite type. Then G contains no non-trivial finite subgroup.*

PROOF. Let X be a finite $K(G, 1)$, of dimension d , and let \tilde{X} be its universal cover. Assume to the contrary that G has a non-trivial finite subgroup, and therefore that G has a subgroup isomorphic to a finite cyclic group \mathbb{Z}_n . Since G acts freely on \tilde{X} , $\mathbb{Z}_n \curvearrowright \tilde{X}$. It follows that $H_i(\mathbb{Z}_n) = H_i(\mathbb{Z}_n \backslash \tilde{X})$ and in particular, $H_i(\mathbb{Z}_n) = 0$ for all $i > d$. But $H_{2j+1}(\mathbb{Z}_n) \approx \mathbb{Z}_n$ for all $j \geq 0$. \square

Avoiding groups with torsion is often overly restrictive. For example, consider the group G of isometries of the Euclidean plane, generated by reflections in the sides of an equilateral triangle. This action is not free, but it is cocompact and proper. As we will see, such actions are often more than sufficient for one to derive deep facts about the group.

Exercises

Cell complexes

1. Let X be a cell complex, let x and y be 0-cells of X and let A be a connected finite subcomplex containing x and y with a minimum number of cell. Prove that A is the image of an embedded interval $f : I \rightarrow X$ starting at x and ending at y .
2. Let X be a 1-complex that contains a 1-disc where both endpoints are attached to the same 0-cell. Use a retraction to show that X is not simply-connected.
3. Let X be a 1-complex and let $f : \mathbb{S}^1 \rightarrow X$ be an embedding. Show that X is not simply-connected by collapsing all but one 1-cell of the image of f and applying the previous exercise.
4. Let $f : [0, 1] \rightarrow [0, 1]$ be an infinitely oscillating function like that shown on the left in Figure 3. Use this function as part of an attaching map (as indicated on the right in Figure 3) to create a 2-complex with three vertices, three edges, and a single 2-cell. Show that this “shower curtain complex” is not homeomorphic to any simplicial complex.




FIGURE 3. The “shower curtain complex” is not homeomorphic to any simplicial complex.

5. Prove that every combinatorial cell complex is homeomorphic to a simplicial complex.
6. (Classifying compact surfaces) (sketch out how to classify compact surfaces) Here are the main steps.
 - a. show that you only need a single 2-cell.
 - b. make all moebius edges adjacent
 - c. isolate crossing annular edges
 - d. remove noncrossed annular edges
 - e. eliminate the mixed case
 - f. use the abelianizations to distinguish the remaining cases
7. Prove Corollary A.5.5.
8. Let Y be the image of the map $g : \mathbb{R} \rightarrow \mathbb{R}^3$ defined by $g(t) = (\cos t, \sin t, e^t)$ and turn Y into a path connected topological space by giving it the subspace topology. There is a map $f : Y \rightarrow \mathbb{S}^1$ that comes from projecting onto the first two coordinates. Prove that f is a local homeomorphism but not a covering map.
9. Show that every one-relator Artin group is a torus knot group. Which torus knot groups arise in this way?

Group actions and covering spaces

10. Add an exercise that covers the Hawaiian earring. (Figure 4) (fix: do not have universal covers. Identifying a single point in a simply connected metric space with a single point in another simply connected metric space does not need to result in a simply connected space. This space, which is the union of the circles centered at $(0, 1/n)$ and tangent to the x -axis...)




FIGURE 4. Hawaiian earring.

11. (Normal subgroups) Let H be a subgroup of G , let $A := G/H$ be the set of left H -cosets, let $\kappa := |A|$ be the index of H in G , and let $\kappa!$ denotes the size of SYM_A (the bijections $A \rightarrow A$ under composition). Prove that there is a normal subgroup N of G contained in H whose index in G is at most $\kappa!$.
12. (Infinite Index) Let A be a set and let $G = \text{SYM}_A$ be the group of all permutations (i.e. bijections) $f : A \rightarrow A$ under function composition. Choose an element $a \in A$ and let H be the subgroup of permutations that fix a . Prove that index of H in G is $\kappa = |A|$ and that the only normal subgroup of G in H is the trivial subgroup (index $\kappa!$).

Homotopy invariants and Whitehead's theorem

13. (Contractibility) Prove that the characterizations of contractibility list in Theorem A.4.7 are equivalent.

APPENDIX B

Hints

This appendix collects hints to the exercises.

Chapter 1

1. Find an explicit procedure that contracts arbitrary closed paths in $Y^{(1)}$.
3. Make each edge in X isometric to a unit interval $[0, 1]$.
4. Suppose $f : R_A \rightarrow R_B$ is a homotopy equivalence. Prove that $f(R_A)$ is contained in a finite subcomplex of R_B , and then use a retraction to show that f_* is not onto.
16. Find elements in \mathbb{Z}^n that cannot be part of a basis of \mathbb{Z}^n as a free \mathbb{Z} -module and then use Exercise 15.
27. Consider a non-trivial surface with boundary.
29. Find the 1-skeleton first.

(re)move. Consider the following method of converting a factorization of g into its free product normal form. Let g be a non-trivial element of $\pi_1(X, x)$ and let $f : I \rightarrow X^{(1)}$ be an immersed loop based at x that represents g (Proposition 1.1.2). Since every 1-cell in X belongs to a unique X_α and these subcomplexes overlap only at x , the maximal connected nontrivial subintervals of I sent into a single X_α partition f into a finite concatenation of loops f_1, f_2, \dots, f_k , all based at x , such that consecutive f_i represent elements g_i in distinct groups $\pi_1(X_\alpha, x_\alpha)$. Finally, if any of the elements g_i are trivial, then the corresponding paths f_i can be excised from f and the new path still representing g but with strictly fewer concatenations. Repeating this procedure (reparsing the path into maximal subpaths and then eliminating subpaths representing trivial elements) eventually results in a finite product of non-trivial elements satisfying the lemma since the number of concatenations strictly decreases each time this process is carried out.

Appendix A

12. Characterize the permutations in the left coset fH and in the subgroup fHf^{-1} . Then show that these conjugates of H have trivial intersection.

Bibliography

- [1] M. Aschbacher. *Finite group theory*, volume 10 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, second edition, 2000.
- [2] W. Ballmann and S. Buyalo. Nonpositively curved metrics on 2-polyhedra. *Math. Z.*, 222(1):97–134, 1996.
- [3] Martin R. Bridson. The geometry of the word problem. In *Invitations to geometry and topology*, volume 7 of *Oxf. Grad. Texts Math.*, pages 29–91. Oxford Univ. Press, Oxford, 2002.
- [4] Martin R. Bridson and André Haefliger. *Metric spaces of non-positive curvature*, volume 319 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 1999.
- [5] Kenneth S. Brown. Finiteness properties of groups. In *Proceedings of the Northwestern conference on cohomology of groups (Evanston, Ill., 1985)*, volume 44, pages 45–75, 1987.
- [6] W. Burnside. *Theory of groups of finite order*. Dover Publications Inc., New York, 1955. 2d ed.
- [7] James W. Cannon. The theory of negatively curved spaces and groups. In *Ergodic theory, symbolic dynamics, and hyperbolic spaces (Trieste, 1989)*, Oxford Sci. Publ., pages 315–369. Oxford Univ. Press, New York, 1991.
- [8] John H. Conway. *The sensual (quadratic) form*, volume 26 of *Carus Mathematical Monographs*. Mathematical Association of America, Washington, DC, 1997. With the assistance of Francis Y. C. Fung.
- [9] M. Dehn. Über unendliche diskontinuierliche Gruppen. *Math. Ann.*, 71(1):116–144, 1911.
- [10] David B. A. Epstein, James W. Cannon, Derek F. Holt, Silvio V. F. Levy, Michael S. Paterson, and William P. Thurston. *Word processing in groups*. Jones and Bartlett Publishers, Boston, MA, 1992.
- [11] Benson Farb and Lee Mosher. A rigidity theorem for the solvable Baumslag-Solitar groups. *Invent. Math.*, 131(2):419–451, 1998. With an appendix by Daryl Cooper.
- [12] S. M. Gersten. Reducible diagrams and equations over groups. In *Essays in group theory*, pages 15–73. Springer, New York-Berlin, 1987.
- [13] S. M. Gersten and H. B. Short. Small cancellation theory and automatic groups. *Invent. Math.*, 102(2):305–334, 1990.
- [14] M. Gromov. Hyperbolic groups. In *Essays in group theory*, volume 8 of *Math. Sci. Res. Inst. Publ.*, pages 75–263. Springer, New York, 1987.
- [15] Marshall Hall, Jr. *The theory of groups*. The Macmillan Co., New York, N.Y., 1959.
- [16] Allen Hatcher. *Algebraic topology*. Cambridge University Press, Cambridge, 2002.
- [17] A. G. Kurosh. *The theory of groups. Vol. I*. Chelsea Publishing Co., New York, N.Y., 1955. Translated and edited by K. A. Hirsch.
- [18] A. G. Kurosh. *The theory of groups*. Chelsea Publishing Company, New York, N.Y., 1956. Translated from the Russian and edited by K. A. Hirsch.
- [19] Roger C. Lyndon and Paul E. Schupp. *Combinatorial group theory*. Classics in Mathematics. Springer-Verlag, Berlin, 2001. Reprint of the 1977 edition.
- [20] Wilhelm Magnus, Abraham Karrass, and Donald Solitar. *Combinatorial group theory*. Dover Publications Inc., New York, revised edition, 1976. Presentations of groups in terms of generators and relations.
- [21] Jonathan P. McCammond and Daniel T. Wise. Fans and ladders in small cancellation theory. *Proc. London Math. Soc. (3)*, 84(3):599–644, 2002.
- [22] Charles F. Miller, III. Decision problems for groups—survey and reflections. In *Algorithms and classification in combinatorial group theory (Berkeley, CA, 1989)*, volume 23 of *Math. Sci. Res. Inst. Publ.*, pages 1–59. Springer, New York, 1992.

- [23] Robert Oliver. Fixed-point sets of group actions on finite acyclic complexes. *Comment. Math. Helv.*, 50:155–177, 1975.
- [24] S. J. Pride. Small cancellation conditions satisfied by one-relator groups. *Math. Z.*, 184(2):283–286, 1983.
- [25] John Roe. *Lectures on coarse geometry*, volume 31 of *University Lecture Series*. American Mathematical Society, Providence, RI, 2003.
- [26] Jean-Pierre Serre. *Trees*. Springer Monographs in Mathematics. Springer-Verlag, Berlin, 2003. Translated from the French original by John Stillwell, Corrected 2nd printing of the 1980 English translation.
- [27] A. J. Sieradski. A coloring test for asphericity. *Quart. J. Math. Oxford Ser. (2)*, 34(133):97–106, 1983.
- [28] Richard P. Stanley. *Enumerative combinatorics. Vol. 1*, volume 49 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 1997. With a foreword by Gian-Carlo Rota, Corrected reprint of the 1986 original.
- [29] John Stillwell. *Classical topology and combinatorial group theory*, volume 72 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, second edition, 1993.
- [30] K. Whyte. The large scale geometry of the higher Baumslag-Solitar groups. *Geom. Funct. Anal.*, 11(6):1327–1343, 2001.