Introduction to
Modern Mathematics
ALM 33, pp. 125–144

© *Higher Education Press*
and International Press
Beijing–Boston

Geometric Group Theory: An Introduction*

Jason Behrstock†

Abstract

These lectures are designed to provide a brief introduction to geometric group theory for advanced undergraduates and beginning graduate students. The first lecture provides some foundational definitions and an introduction to thinking about groups geometrically. Each of the remaining lectures are dedicated to a classical theorem in the field. The choice of topics is designed to illustrate the power of the geometric viewpoint in proving algebraic properties of groups.

AMS Classification Numbers: 20F65, 57M07 , 20F67, 20F69.

Keywords and Phrases: Geometric group theory.

Overview

A central aspect of geometric group theory is that finitely generated groups can be treated as geometric objects and that this view point allows one to prove results that are otherwise difficult or inaccessible through purely algebraic means.

The topics for the five lectures will be:

(1) Groups as geometric objects
(2) Curves on surfaces and the word problem
(3) Geometry of finite presentability
(4) Groups and their ends
(5) How many groups are there?

Students should be familiar (or familiarize themselves) with basic notions in group theory and topology such as: free groups, presenting a group by generators and relators, normal subgroups, fundamental groups, and covering spaces. For instance, for group theory see: "An introduction to the theory of groups", by Rotman. For topology see: "Algebraic Topology", by Hatcher (Chapters 0 and 1), available at: `http://www.math.cornell.edu/hatcher/AT/ATpage.html`

*Lecture notes from 2013 Summer School on Modern Mathematics, Beijing, China.

†Lehman College and The Graduate Center, City University of New York, U.S.A. Email: jason.behrstock@lehman.cuny.edu

For further introductions to various topics studied in geometric group theory, one can browse the following surveys, each of which has some easily accessible topics and several more advanced ones:

B. Bowditch, A course on geometric group theory. Download available: `http://www.warwick.ac.uk/masgak/papers/bhb-ggtcourse.pdf`

C. Drutu and M. Kapovich, Lectures on Geometric Group Theory. Download available: `http://www.math.ucdavis.edu/kapovich/EPR/ggt.pdf`

J. Meier, Groups, Graphs and Trees. Cambridge University Press.

H. Short (edit.): MSRI Notes on Hyperbolic groups. Download available: `http://www.cmi.univ-mrs.fr/hamish/Papers/MSRInotes2004.pdf`

M. Gromov, Asymptotic invariants of infinite groups. LMS lecture note series v. 182 (1993).

1 Groups as geometric objects

Let G be a group with a finite generating set S. We will always assume that the identity element, e, is not in S and that the generating set is *symmetric*, i.e., if $x \in S$ then $x^{-1} \in S$.

There are many interesting examples of finitely generated groups, below are a few basic examples to keep in mind.

- $\mathbb{Z}$ under addition. This is isomorphic to $\langle a \rangle$ under multiplication, where each element is a^n for some $n \in \mathbb{Z}$.
- $\mathbb{Z}/5\mathbb{Z}$ under addition. This is isomorphic to $\langle a; a^5 = 1 \rangle$ under multiplication, so that each element is a^n for some $n \in \{0,1,2,3,4\}$.
- $\mathbb{Z} \times \mathbb{Z} = \mathbb{Z}^2 \cong \{(n,m); n,m \in \mathbb{Z}\}$. The *free abelian group on two generators.* Equivalently, $\langle a,b; aba^{-1}b^{-1} = 1 \rangle$ so that each element is $a^n b^m$ for some $n,m \in \mathbb{Z}$.
- $\mathbb{F}_2 \cong \langle a,b \rangle$. The *free group on two generators.* Any word is of the form $a^{n_1}b^{n_2}a^{n_3}\dots b^{n_k}$ for $n_1,\dots,n_k \in \mathbb{Z}$ and $k \in \mathbb{N}$. The multiplication here is given by concatenating words.
- $\pi_1(S_g)$. The *fundamental group of S_g*, where S_g denotes the closed orientable surface of genus g. Example: $\pi_1(S_2) = \langle a,b,c,d; aba^{-1}b^{-1}cdc^{-1}d^{-1} \rangle$.
- $\mathcal{MCG}(S_g) \cong \mathrm{Out}(\pi_1(S_g)) \cong \mathrm{Homeo}(S_g)/\mathrm{Homeo}_0(S_g)$. The *mapping class group of S_g*.

For a finitely generated group, G, we define the *Cayley graph of G with respect to the generating set S* as the graph whose vertex and edge sets consist of the following.

- Vertices $\Longleftrightarrow$ one for each element of G
- Edges $\Longleftrightarrow$ one connecting each pair $a,b \in G$ for which there exists $c \in S$ satisfying $a = b \cdot c$

The Cayley graph of $\mathbb{F}_2$ with its canonical generating set is the regular 4–valent tree. In the case of $\mathbb{Z}^2$, there is a natural set of generators consisting of the elements $(0,1)$ and $(1,0)$, and their inverses; with respect to these generators, we obtain a Cayley graph whose underlying set is the collection of horizontal and

vertical lines in $\mathbb{R}^2$ through the integer x and y axes; i.e., this Cayley graph looks like graph paper. Further examples appear in Figures 1 and 2.

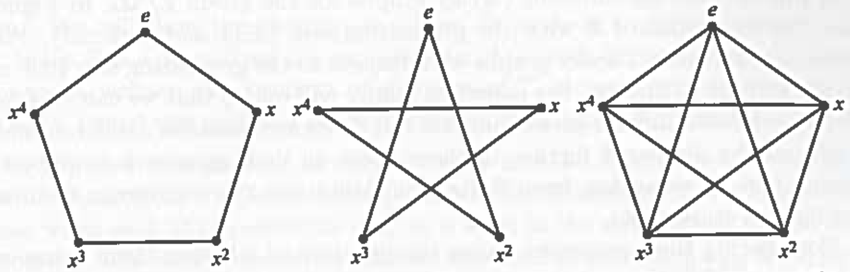

(a) with generating set $\{x^{\pm1}\}$ (b) with generating set $\{x^{\pm2}\}$ (c) with generating set $\{x^{\pm3}\}$

Figure 1. Cayley graphs of $\mathbb{Z}/5\mathbb{Z} = \langle x; x^5 \rangle$ with different generating sets. Note that all line segments should be considered to have equal length.

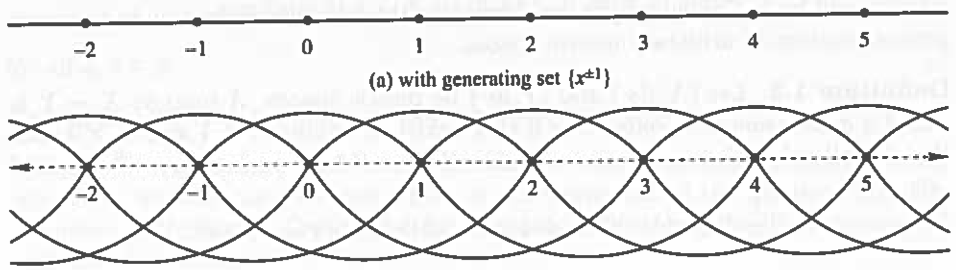

(a) with generating set $\{x^{\pm1}\}$

(b) with generating set $\{\pm2, \pm3\}$

Figure 2. Cayley graphs of $\mathbb{Z}$ with different generating sets. Note that all arcs between vertices are of equal length.

Associated to any finitely generated group, we may associate the *word metric on G with respect to the finite generating set S* as follows: For each $x,y \in G$ define:

$$d_{G,S}(x,y) = \|x^{-1}y\|$$

where the norm on the right side is defined to be the smallest number of letters (with multiplicity) from the set S needed to represent the group element $x^{-1}y$.

Remark 1.1. It is easily checked that this metric is *left-invariant*, that is, multiplication on the left is an isometry of the metric. More precisely, for each $a,b,g \in G$ we have $d_{G,S}(ga,gb) = d_{G,S}(a,b)$

Exercise 1.2. *Prove the following are equivalent:*

(1) $d_{G,S}(x,y) = n$
(2) $x = y \cdot w$ *where w is a word of length n*
(3) *Consider the Cayley graph of G with respect to S and metrize this graph by making each edge isometric to $[0,1]$. In this graph the shortest length of a path connecting x to y is n.*

As one can see by generalizing the above examples, any given group may have many distinct Cayley graphs, depending on the choice of finite generating set. In Figure 1 we see different Cayley graphs for the group $\mathbb{Z}/5\mathbb{Z}$. In Figure 2, we see Cayley graphs of $\mathbb{Z}$ with the generating sets $\{\pm 1\}$ and $\{\pm 2, \pm 3\}$. As an exercise, you can draw Cayley graphs with respect to the generating sets $\{\pm 1, \pm 3\}$, $\{\pm 1, \pm 2, \pm 3\}$, or ... Indeed, the attentive reader will notice that we can even build Cayley graphs with infinite generating set, but those are beasts of a different nature and will not be discussed further in these notes as their geometry is quite a bit different than those arising from finite generating sets (we encourage the reader to justify this statement).

Considering these examples raises the question of whether there is any relation between different word metrics on the same group. Contrary to what might be your expectation from these examples, it turns out that the answer is "Yes!" when given an appropriate formulation. The formulation which we consider uses the following definition which makes precise the notion that two Cayley graphs for the same group "look roughly the same from far away." Although our main interest will be in Cayley graphs, we formulate the following definition in the more general context of arbitrary metric spaces.

Definition 1.3. Let (X, d_X) and (Y, d_Y) be metric spaces. A map $\phi\colon X \to Y$ is called a *quasi-isometric embedding* if there exist constants $K \geqslant 1$ and $C \geqslant 0$ such that for all $a, b \in X$:

$$\frac{1}{K} d_X(a,b) - C \leqslant d_Y(\phi(a), \phi(b)) \leqslant K d_X(a,b) + C.$$

This map is called a *quasi-isometry* if, additionally, every point of Y lies in the C–neighborhood of the image of ϕ.

Definition 1.4. Let $f\colon X \to Y$ be a quasi-isometry. We say a quasi-isometry $g\colon Y \to X$ is a *quasi-inverse* to f if the maps satisfy the following:

$$d_X(g \circ f, 1_X) < \infty \text{ and } d_Y(f \circ g, 1_Y) < \infty$$

where 1_X denotes the identity map on X and distance is given by the supremum over all points in the relevant space.

The following exercise is a good way to get used to working with quasi-isometries, even though it will not be needed explicitly in these lectures.

Exercise 1.5. *Given a quasi-isometry between metric spaces, $f\colon X \to Y$, prove that f has a quasi-inverse.*

We will now prove that for any finitely generated group the word metric is unique up to quasi-isometry, i.e., any two word metrics on a fixed finitely generated group are quasi-isometric.

Notation 1.6. *When the group is clear we simplify the notation $d_{G,S}$ to d_S or even d when this will not cause ambiguity. Also, for $w \in G$ we use the notation $\|w\|_S = d_S(e, w)$.*

Proposition 1.7. *If S, T are two finite generating sets for a group, G, then (G, d_S) and (G, d_T) are quasi-isometric.*

Proof. We will show the identity map from d_S to d_T is a quasi-isometry.

Since T is a generating set, any element of G, and, in particular, any element of S can be written as a product of elements of T. Accordingly we can set $K_1 = \max_{s \in S} d_T(e, s)$ and this will be a positive integer. Now, since S generates, any $w \in G$ can be factored as a product of elements of S and the shortest way to do so is a word in elements of S of length $\|w\|_S$. Given this factorization as elements of S, we can write each of the elements of S as a word in the elements of T with length at most K_1. We have thus shown that for any $w \in G$ we have $\|w\|_T \leqslant K_1 \|w\|_S$.

Setting $K_2 = \max_{t \in T} d_S(e, t)$, the argument above with the roles of S and T reversed shows that for any $w \in G$ we have $\|w\|_S \leqslant K_2 \|w\|_T$.

Taking $K = \max\{K_1, K_2\}$ and $C = 0$, we then have:

$$\frac{1}{K} d_S(a,b) - C \leqslant d_T(a,b) \leqslant K d_S(a,b) + C$$

for all $a, b \in S$. □

Example 1.8. *Any two finite groups are quasi-isometric.*

Milnor–Švarc Lemma and its applications. After recording a few preliminary definitions we will state a result which is sometimes called the *Fundamental Observation of Geometric Group Theory*. The proof of this result will be postponed until the third lecture.

Definition 1.9. A *geodesic metric space* is a metric space X with the property that for any two points $x, y \in X$ there exists an isometric embedding, $g\colon [0, d(x,y)] \to X$, of an interval into X which satisfies $g(0) = x$ and $g(d(x,y)) = y$.

We say that a metric space X is *proper* if closed balls in X are compact. Note that any proper metric space is locally compact and complete.

A discrete group, G, acting by homeomorphisms on a locally compact topological space is said to act *properly* if for any compact subset $K \subseteq X$, the set

$$\{g \in G\colon gK \cap K \neq \varnothing\}$$

is finite. If the quotient of a space by a group action is compact, we say the group action is *cocompact*

Exercise 1.10. *Prove that a Cayley graph of a finitely generated group is a proper geodesic metric space. Further, we noted previously that a group acts on itself isometrically by left multiplication; check that in doing so G acts on its Cayley graph properly and cocompactly.*

The following is a quintessential geometric group theory result: starting with no hypothesis on a group except that it acts in a reasonable way on a space with some mild geometric/topological hypotheses, we then use geometry to deduce an algebraic result, which in this case is that the group in question is finitely generated.

Lemma 1.11 (Milnor–Švarc). *Let G be a group acting, on the left, properly and isometrically on a proper geodesic metric space, X. If $G\backslash X$ is compact, then G is finitely generated and is quasi-isometric to X.*

We say two groups are *commensurable* if they "differ by a finite subgroup"; more precisely this means that one can be obtained by the other by passage to a finite index subgroup and/or by quotienting by a finite normal subgroup. Below are two immediate corollaries of the Milnor–Švarc Lemma which show that *commensurable* groups are quasi-isometric.

Corollary 1.12. *Let G be a finitely generated group and G' a subgroup of finite index. Then G' is also finitely generated and is quasi-isometric to G.*

Corollary 1.13. *Let G be a finitely generated group and N a finite normal subgroup of G, then G is quasi-isometric to $N\backslash G$.*

The next corollary is used so frequently, that it is itself sometimes referred to as the Milnor–Švarc Lemma.

Corollary 1.14. *Let M be a compact Riemannian manifold, then $\pi_1(M)$ is quasi-isometric to $\widetilde{M}$, the universal covering space of M.*

The next two examples are immediate consequences of the above corollary; we single them out for attention because they are examples which we will consider again in later lectures.

Example 1.15. *Let T^2 be the torus. Its fundamental group $\pi_1(T^2) \cong \mathbb{Z}^2$ is quasi-isometric to $\mathbb{R}^2$.*

Example 1.16. *Let S_g be the closed orientable surface of genus $g \geqslant 2$, then $\pi_1(S_g)$ is quasi-isometric to the hyperbolic plane, $\mathbb{H}^2$.*

To verify the second example, choose a Riemannian metric of constant curvature -1 on S_g and consider the associated action of $\pi_1(S_g)$ on $\mathbb{H}^2$.

Remark 1.17. Generalizing this example, we note that in each dimension, all compact manifolds of constant curvature -1 have quasi-isometric fundamental groups, since they are each quasi-isometric to $\mathbb{H}^n$. Using the (highly non-trivial) result that there exists hyperbolic 3–dimensional manifolds whose volumes are not rationally related, from such examples one can obtain many quasi-isometric groups which are not commensurable.

Exercise 1.18. *For each $m, n \geqslant 2$, the groups $\mathbb{F}_m$ and $\mathbb{F}_n$ are commensurable. In particular, for all $m, n \geqslant 2$ the groups $\mathbb{F}_m$ and $\mathbb{F}_n$ are quasi-isometric.*

Hint. This follows from the algebraic fact that for all $k \geqslant 2$ the group $\mathbb{F}_k$ is a subgroup of $\mathbb{F}_2$. We sketch a topological proof of this algebraic fact. Consider $\mathbb{F}_2 = \pi_1(S^1 \vee S^1)$, where $S^1 \vee S^1$ denotes the wedge of two circles. Each $(k-1)$-sheeted cover of $S^1 \vee S^1$ is a graph with $(k-1)$ vertices and $2(k-1)$ edges. Thus the Euler characteristic of any such cover is $1-k$ and hence its fundamental group is $\mathbb{F}_k$. It follows that $\mathbb{F}_k$ is an index $k-1$ subgroup of $\mathbb{F}_2$ and hence that $\mathbb{F}_k$ is quasi-isometric to $\mathbb{F}_2$. □

Exercise 1.19. *Any two infinite regular trees whose valence at each vertex is at least 3 are quasi-isometric. The case where the valence is a multiple of 4 is a consequence of the previous exercise about F_n. Find a geometric argument that treats all cases and thus gives an alternate proof of the above exercise as well.*

2 Curves on surfaces and the word problem

Problem 2.1 (Contractibility problem). *Let S_g denote the closed orientable surface of genus g and let γ denote a closed curve on S_g. Can γ be contracted to a point?*

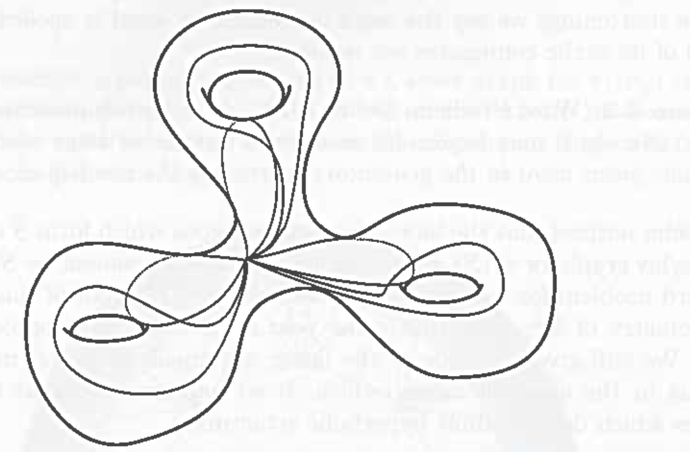

Figure 3. A curve, γ, on the closed orientable genus three surface, S_3.

We start with an observation used by Dehn in the early 1900's and perhaps known earlier. Given a surface S, we write $\widetilde{S}$ to denote the universal cover of S and $\widetilde{\gamma}$ to refer to a lift of γ to $\widetilde{S}$.

Observation. *γ is a contractible closed curve in $S_g \Leftrightarrow \widetilde{\gamma}$ is a closed curve in $\widetilde{S}_g$.*

Proof of observation. ($\Rightarrow$) Since γ can be contracted to a point, we can do so fixing a basepoint. Such a contraction then lifts to a contraction of $\widetilde{\gamma}$ fixing both endpoints. Hence these two endpoints of $\widetilde{\gamma}$ must coincide and thus we have that $\widetilde{\gamma}$ is a closed curve in $\widetilde{S}_g$.

($\Leftarrow$) Since the universal cover of a closed surface is homeomorphic to either $\mathbb{S}^2$ or $\mathbb{R}^2$, we know that if $\widetilde{\gamma}$ is a closed curve in $\widetilde{S}_g$ then $\widetilde{\gamma}$ is contractible in $\widetilde{S}_g$. Viewing $\widetilde{S}_g$ as being tessellated by polygons given by translations of a fundamental domain, we see that by contracting $\widetilde{\gamma}$ through one polygon in $\widetilde{S}_g$ at a time, that this contraction in the cover projects to a contraction of γ in S_g. □

Problem 2.2 (Alternative formulation of Contractibility Problem). *In $\pi_1(S_g)$, is $[\gamma] = 1$? Here $[\gamma]$ denotes the homotopy class of the curve γ.*

Recall that $\pi_1(S_g) = \langle a_1, \ldots, a_g, b_1, \ldots, b_g; a_1b_1a_1^{-1}b_1^{-1}, \ldots, a_gb_ga_g^{-1}b_g^{-1}\rangle$. Formally, we can think of a group as the quotient of the free group (*generators*) by the normal subgroup generated by the *relators*. When the set of generators is finite we say the group is *finitely generated*, when the set of relators is also finite, then we say the group is *finitely presented*; by an abuse of grammar, it is common parlance to also use these phrases in the case when there exists such a presentation, even if a particular presentation is not indicated. Given any *word*[1], w, in the generators (and their inverses), we can consider the associated element in the group, which we denote $[w]$. Note that there are many words which all have the same associated group element. A word may contain an "obvious shortenings" if it admits a subword of the form $a^{-1}a$ or aa^{-1} where a is a generator, since the word obtained by removing this subword still represents the same group element. If there are no obvious shortenings we say the word is *reduced*. A word is *cyclically reduced* if it and all of its cyclic conjugates are reduced.

Problem 2.3 (Word Problem; Dehn, 1912). *In a finitely presented group, give a method whereby it may be decided in a finite number of steps whether or not any particular given word in the generators represents the identity element.*

Dehn noticed that the labelled net of polygons which form $\widetilde{S}$ can also be seen as a Cayley graph for $\pi_1(S)$. Thus the contractibility problem for S is equivalent to the word problem for $\pi_1(S)$. This led to Dehn's 1912 proof of this problem using the geometry of $\mathbb{H}^2$; later that same year he gave a more topological/algebraic proof. We will give a version of the latter argument after first mentioning what happens in the simplest cases (which, it so happens, are also the only closed surfaces which do not admit hyperbolic structures).

Example 2.4. *The two-dimensional sphere, $\mathbb{S}^2$, which is the closed oriented surface of genus 0, satisfies: $\pi_1(\mathbb{S}^2) = \{1\}$. Thus every word in this group is already equal to 1 and the word problem admits a trivial solution.*

The two-dimensional torus $\mathbb{T}^2$, which is the closed oriented surface of genus 1, satisfies: $\pi_1(\mathbb{T}^2) = \langle a, b; aba^{-1}b^{-1}\rangle \cong \mathbb{Z}\times\mathbb{Z}$, so for any word we can take as a "normal form" the unique representative of the form a^mb^n. One can arrive at this normal form by switching the order of adjacent letters $a^{\pm1}$ and $b^{\pm1}$. The solvability of the word problem then follows from this and the fact that $a^mb^n = 1 \Leftrightarrow m=n=0$.

It remains to discuss the case of closed orientable surfaces with $g \geqslant 2$.

Remark 2.5. Note that there exists a Cayley graph for the fundamental group of a closed orientable surfaces with $g \geqslant 2$ which consists of a planar net of $4g$-gons with $4g$ at each vertex. We will work with this particular Cayley graph, which will be described in more detail below.

Lemma 2.6. *If γ is a closed path in the Cayley graph of $\pi_1(S_g)$ with $g \geqslant 2$, then γ has either a "spike" (see Figure 4) or a subpath consisting of more than half the edges in the boundary of a polygon, in succession.*

[1] A *word* is a finite sequence of elements from a given alphabet.

Proof. We assume γ does not contain any spikes; note that this is equivalent to assuming that γ represents a cyclically reduced word. It thus suffices to show that there is a subpath consisting of more than half the edges in the boundary of a polygon.

Figure 4. A "spike" is a subpath that consists of an edge followed by the same edge in the opposite direction.

We now describe a planar embedding of a Cayley graph for $\pi_1(S_g)$ (see Figure 5). Fix a collection $C_1, C_2, \ldots, C_n, \ldots$ of nested circles with increasing radii. Each circle will be subdivided into circumferential edges, and there will be edges from C_i to C_{i+1}, called *radial edges*.

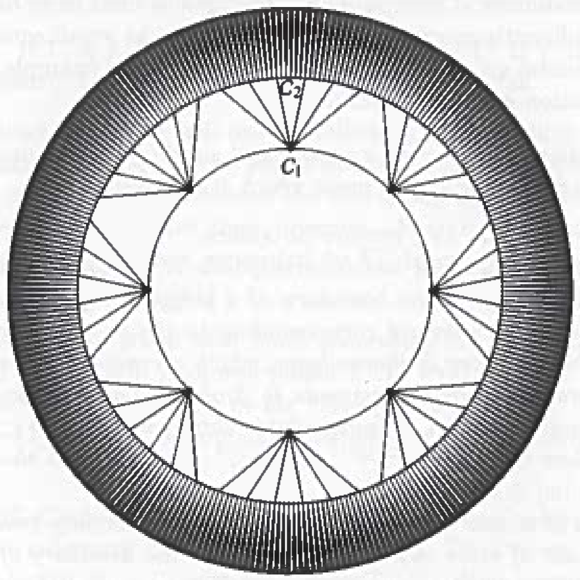

Figure 5. A portion of the Cayley graph of $\pi_1(S_2)$.

- Divide C_1 into $4g$ arcs to form the first $4g$-gon.
- Each vertex on C_1 will emit $4g-2$ radial edges to C_2 thereby dividing C_2 into $4g(4g-2)$ arcs.
- Subdivide each arc in C_2 so that every region between C_1 and C_2 is a $4g$-gon (i.e., insert either $4g-3$ or $4g-4$ vertices depending on whether the polygon contains an arc of C_1 or not).

- Repeat the second and third step to construct radial edges from C_2 to C_3, from C_3 to C_4, etc.

By Remark 2.5, this is a (topological embedding of a) Cayley graph for $\pi_1(S_g)$ into the plane. A fact which we will not use, but which we note for the reader, is that this embedding strongly distorts distances.

Consider a closed path γ in this graph starting at a vertex of C_1. Let C_k be the outermost circumferential portion reached by γ; we may assume $k > 1$ since otherwise γ must travel around the full polygon C_1. The curve γ must travel across a radial edge e from C_{k-1} to C_k, at which point, since it is not a spike, it must travel along C_k, and at last return C_{k-1} via a radial edge from C_k to C_{k-1}. Now e and the initial segment of γ following e determine a $4g$-gon. Along this $4g$-gon, γ must travel along either $4g-1$ edges (when this $4g$-gon has no edge on C_{k-1}) or $4g-2$ edges (when this $4g$-gon has one edge on C_{k-1}); this follows from the assumption that no spike exists and that C_k is the largest circle that γ reaches. Since $g \geqslant 2$ we have that $4g-2 > \frac{4g}{2}$ and thus we have found a polygon more than half of whose boundary is traversed by a subpath of γ. □

The following is known as *Dehn's algorithm*; it was originally formulated in the context of fundamental groups of surfaces, as it is used here, but it can also be applied in a significantly more general setting known as *small cancellation groups*, of which fundamental groups of surfaces are an important example. We will return to small cancellation groups in Lecture 5.

Dehn's Algorithm. *Start with a non-trivial word w in the generators of $\pi_1(S)$, and let $[w] \in \pi_1(S)$ denote the element which it represents.*

(1) *Cyclically reduce w, i.e., remove all spikes.*
(2) *If w contains a subword which corresponds to a subpath consisting of more than half the edges in the boundary of a polygon, replace this subword with the inverse of the subword corresponding to the subpath consisting of those edges in the boundary of that polygon which are not in the previous subword.*
Algebraically this corresponds to finding a subword $w_1 \subset w$ and r, a cyclic permutation of a defining relator satisfying $w_1w_2 = r$ and $2|w_1| > |r|$. In w we then replace w_1 by $r^{-1}w_1 = w_2^{-1}$ which gives a shorter word which represents $[w] \in \pi_1(S)$.
(3) *Repeat the first two steps until there are no longer any spikes nor subpaths which consist of more than half the edges in the boundary of a polygon (i.e., there is neither an "obvious shortenings" nor a subword consisting of more than half of a relator).*

Theorem 2.7 (Dehn). *There exists a solution to the Word Problem for the fundamental group of a closed orientable surfaces.*

Proof. In genus 0 and 1 we have already provided a solution in Example 2.4. Now we consider a surface S with genus greater than 1.

Fix a non-trivial word w in the generators of $\pi_1(S)$ and let γ be the corresponding curve in a Cayley graph for $\pi_1(S)$. We can remove all spikes without changing the class in $\pi_1(S)$, so without loss of generality we may assume that our word w is reduced. We now proceed to step two in Dehn's Algorithm and search w for subwords which are more than half a relator.

Now, if γ is a closed curve (or, equivalently, if $[w] = 1$) this path contains a subpath which travels along more than half of a relator, by Lemma 2.6. Hence, if there is no such subword, then $[w] \neq 1 \in \pi_1(S)$. On the other hand, if there exists such a subpath, then we continue following Dehn's Algorithm. If we start Dehn's Algorithm with a closed curve, each iteration produces a shorter closed curve in the Cayley graph (or equivalently a shorter word w satisfying $[w] = 1$) and thus in a finite number of steps w will be reduced to the trivial word, otherwise we will have a word that can't be shortened, certifying non-triviality. □

Exercise 2.8. *Using the techniques of this Lecture, give a solution to the Word Problem for closed non-orientable surfaces.*

3 Geometry of finite presentability

In this lecture we will develop some geometry related to finite presentability and also provide a proof of the Milnor–Švarc Lemma, which was stated in Lecture 1.

Theorem 3.1. *If Γ is a finitely presented group and Γ' is a finitely generated group quasi-isometric to Γ, then Γ' must be finitely presented.*

To prove this theorem we will use the following lemma which gives a topological condition characterizing when a set of words forms a complete set of relation for a given group. We let $\mathbb{F}(A)$ denote the free group generated on the alphabet given by the elements of A and their inverses. Also, recall that given a subset $\mathcal{R} \subset \mathbb{F}(A)$ the *normal subgroup generated by $\mathcal{R}$*, denoted $\langle\langle \mathcal{R} \rangle\rangle$, is defined to be $\langle \{g^{-1}rg : r^{\pm 1} \in \mathcal{R} \text{ and } g \in \mathbb{F}(A)\} \rangle$

Lemma 3.2. *Let Γ be a group with finite generating set A and let $\mathcal{R}$ be a subset of the kernel of the natural homomorphism $\mathbb{F}(A) \to \Gamma$. Build a 2-complex by attaching 2-cells to each edge loop in the Cayley graph of (Γ, A) which is labelled by a reduced word of $\mathcal{R}$. This 2-complex is simply-connected if and only if $\langle\langle \mathcal{R} \rangle\rangle =$* $\ker(\mathbb{F}(A) \twoheadrightarrow \Gamma)$.

Sketch of proof. $\mathrm{Cayley}(\Gamma, A)$ is the quotient of the tree $\mathrm{Cayley}(\mathbb{F}(A), A)$ by the free action of the subgroup given by $\ker(\mathbb{F}(A) \twoheadrightarrow \Gamma)$. Thus there exists a natural identification of $\ker(\mathbb{F}(A) \twoheadrightarrow \Gamma)$ with $\pi_1(\mathrm{Cayley}(\Gamma, A))$. Now apply the Seifert–van Kampen theorem from algebraic topology. □

Proposition 3.3. *A group admits a finite presentation if and only if it acts properly discontinuously and cocompactly by isometries on a simply-connected geodesic space.*

Sketch of proof. ($\Leftarrow$): This direction is a bit harder and uses different tools then we need for the rest of the lectures, so we will omit it.

($\Rightarrow$): Metrize the simply-connected 2-complex from Lemma 3.2 as a piecewise Euclidean complex where all edges have length 1 and all 2-cells are regular polygons. The action on the Cayley graph then extends to the desired action on the complex. □

Now we are ready to prove the main theorem of this lecture that finite presentability is preserved by quasi-isometries:

Sketch of proof of Theorem 3.1. Fix Γ, Γ' where $\Gamma = \langle A; \mathcal{R}\rangle$ is finitely presented and Γ' is quasi-isometric to Γ. Let ℓ be the length of the longest word in $\mathcal{R}$. By Lemma 3.2, the 2–complex $\mathcal{K}$, obtained by attaching 2–cells to the Cayley graph for (Γ, A) along all edge loops of length $\leqslant \ell$, is simply-connected. Note that to produce the complex $\mathcal{K}$ the collection of 2–cells added to the Cayley graph is locally finite and all of uniformly bounded diameter and hence $\mathcal{K}$ is quasi-isometric to Γ.

To show Γ' admits a finite presentation, by the above Proposition, it suffices to find a simply-connected 2–complex $\mathcal{K}'$ on which the group acts properly discontinuously and cocompactly by isometries. To do this we start with a fixed Cayley graph for Γ' and add to this a locally finite collection of 2–cells to yield a simply-connected 2–complex $\mathcal{K}'$. In particular, we add a 2–cell to each loop in the Cayley graph for Γ' which is of length at most M, for some sufficient large M. It remains to show that this complex which we obtain is simply-connected. Note that since Γ' is quasi-isometric to Γ there exist a quasi-isometry $\phi\colon \mathcal{K}' \to \mathcal{K}$.

Now, we want to show that $\mathcal{K}'$ is simply-connected, i.e., that any loop in $\mathcal{K}'$ is the boundary of a disk in $\mathcal{K}'$. We will sketch the argument here under the additional assumption that ϕ and ϕ^{-1} are continuous; we leave it as an exercise for the careful reader to prove that the *a priori* non-continuous quasi-isometry ϕ can be assumed to be continuous.

Consider a loop $\gamma \subset \mathcal{K}'$ and note that $\phi(\gamma)$ is a loop in $\mathcal{K}'$ (this is using the continuity of ϕ which you proved could be assumed). Now, since $\mathcal{K}'$ is simply-connected, there exists a disk $D \subset \mathcal{K}'$ whose boundary is $\phi(\gamma)$. Then map this disk back to $\mathcal{K}$ by the inverse of the quasi-isometry, which we have also assumed to be continuous. This map of a disk into $\mathcal{K}'$ yields a filling for the original loop in $\mathcal{K}'$. □

By keeping count of the number of cells of the 2–complex in the above proof one can show:

Corollary 3.4 (Alonso). *Among finitely presented groups, the property of having a solvable word problem is preserved by quasi-isometries.*

We now give a proof of a result introduced in the first Lecture:

Lemma 3.5 (Milnor–Švarc). *Let G be a group acting, on the left, properly and isometrically on a proper geodesic metric space, X. If $G\backslash X$ is compact, then G is finitely generated and is quasi-isometric to X.*

Proof. Let $\pi\colon X \to G\backslash X$ be the canonical projection. Endow $G\backslash X$ with the metric $d(p,q) = \inf\{d(x,y)\colon x \in \pi^{-1}(p), y \in \pi^{-1}(q)\}$. Since $G\backslash X$ is compact its diameter, $R = \sup\{d(p,q)\colon p,q \in G\backslash X\}$, is finite.

Fix a base point $x_0 \in X$ and set $B = \{x \in X\colon d(x_0,x) < R\}$. Note that since X is proper B is compact. The sets $\{\gamma B : \gamma \in G\}$ form a covering of X. Define

$$S = \{\gamma \in G\colon \gamma \neq e \text{ and } \gamma B \cap B \neq \varnothing\}$$

noting that $S = S^{-1}$, and, since the action is proper, $|S| < \infty$.

Set $r = \inf_{\gamma \in G-(S\cup\{e\})} d(B, \gamma B)$. Without loss of generality, we assume $G - (S \cup \{e\})$ is non-empty, since otherwise the result follows immediately since G would be finite and X compact.

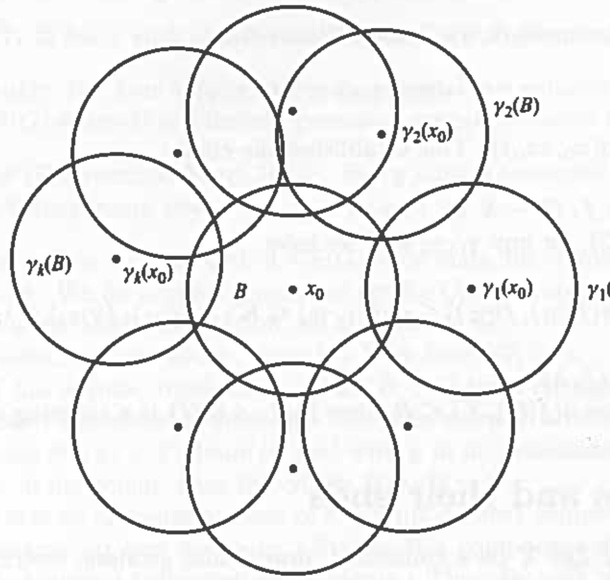

Figure 6. B and γB with $\gamma \in S$.

Claim (1). $r > 0$.

Proof. Choose $\gamma' \in G - (S \cup \{e\})$ and set $r' = d(B, \gamma' B)$. By the definition of S and the compactness of B, we have that $r' > 0$. Now consider the set $\{\gamma \in G - (S \cup \{e\})\colon d(B, \gamma B) \leqslant r'\}$ which is non-empty (since it contains γ') and finite (since G acts properly). Thus r is actually the minimum over a finite set of positive numbers. This establishes the claim. ■

Claim (2). *S generates G and there exists $\lambda > 1$ such that for all $\gamma \in G$ we have*

$$\frac{1}{\lambda} d(x_0, \gamma x_0) \leqslant d_S(e, \gamma) \leqslant \frac{1}{r} d(x_0, \gamma x_0) + 1.$$

Proof. Fix $\gamma \in G$. Let k be the smallest integer such that $d(x_0, \gamma x_0) < k \cdot r + R$. Since X is geodesic, there exist $x_1, x_2, \ldots, x_k, x_{k+1} = \gamma x_0$ such that $d(x_0, x_1) \leqslant R$ and $d(x_i, x_{i+1}) \leqslant r$ for all $i \in \{1, 2, \ldots, k\}$.

Since $\{gB\}_{g\in G}$ covers X, we can choose $\gamma_1, \gamma_2, \dots, \gamma_{k+1} \in G$ such that $x_i \in \gamma_i B$ for $1 \leqslant i \leqslant k+1$ with $\gamma_1 = e$ and $\gamma_{k+1} = \gamma$. Set $s_i = \gamma_i^{-1}\gamma_{i+1}$. Note that we then have that $\gamma = s_1 s_2 \dots s_k$. We will now show that for each i we have $s_i \in S \cup \{e\}$. Since G acts on the left isometrically on X, we have $d(\gamma_i^{-1}x_i, \gamma_i^{-1}x_{i+1}) < r$ where $\gamma_i^{-1}x_i \in B$ and $\gamma_i^{-1}x_{i+1} = s_i\gamma_{i+1}^{-1}x_{i+1} \in s_i B$. Since $r = \inf_{\gamma \in G-(S\cup\{e\})} d(B, \gamma B)$ this then implies that $\gamma_i \in S \cup \{e\}$.

Hence $d_S(e,\gamma) \leqslant k$. Since we chose k to satisfy $(k-1)r + R \leqslant d(x_0, \gamma x_0)$, we have:

$$d_S(e,\gamma) \leqslant \frac{1}{r} d(x_0, \gamma x_0) + 1 - \frac{R}{r}.$$

By induction on $d_S(e,\gamma)$, it then follows that:

$$d(x_0, \gamma x_0) \leqslant \lambda d_S(e,\gamma)$$

where $\lambda = \sup_{s\in S}\{d(x_0, s x_0)\}$. This establishes the claim. ■

Now define $f\colon G \to X$ by $f\colon \gamma \mapsto \gamma x_0$.

By claim (2), for any $\gamma_1, \gamma_2 \in G$ we have

$$\frac{1}{K} d(f(\gamma_1), f(\gamma_2)) \leqslant d_S(\gamma_1, \gamma_2) \leqslant K \cdot d(f(\gamma_1), f(\gamma_2)) + 1$$

where $K = \max\{\frac{1}{r}, \lambda\}$.

And we have $d(f(G), X) \leqslant R$ since $\{\gamma B : \gamma \in G\}$ is a covering of X. □

4 Groups and their ends

Definition 4.1. Let X be a connected, proper and geodesic metric space (e.g., a locally finite graph). The *number of ends* of X, denoted $e(X)$, is defined to be the limit as $n \to \infty$ of the number of infinite diameter connected components of $X - B_n(x_0)$ where x_0 is a fixed point in X and $B_n(x_0)$ is the ball of radius n about x_0.

Exercise 4.2. *Prove that for any space X as above this limit either exists or is infinite and that the value of the limit is independent of the choice of $x_0 \in X$.*

Exercise 4.3. *If two connected, proper, geodesic metric spaces are quasi-isometric, then they have the same number of ends.*

Hint. Consider a path connecting two points which stays away from some large ball in the space X. Let $\phi\colon X \to Y$ be a quasi-isometry. Show that the image of the path (which need not be a path since quasi-isometries need not be continuous), can be "connected up" to form a path disjoint from a (possible smaller, but still fairly large) ball in Y. □

Definition 4.4. The *number of ends* of a finitely generated group is defined to be the number of ends of a(ny) Cayley graph with respect to some finite generating set.

Example 4.5. *Consider $\mathbb{F}_2 = \langle a, b\rangle$ and the Cayley graph of this group with respect to the generators $\{a, b\}$. For any non-negative integer n, removing a ball of radius n leaves 4×3^n unbounded connected components. Thus $\mathbb{F}_2$ has infinitely many ends.*

The poof of the next Proposition is left as an easy exercise.

Proposition 4.6. *A finitely generated group has zero ends if and only if it is finite.*

Considering Cayley graphs with respect to their standard generating set, we easily obtain the following two examples.

Example 4.7. *$\mathbb{Z}$ has 2 ends and, for each integer $n \geqslant 2$, the group $\mathbb{Z}^n$ has 1 end.*

Surprisingly, the four simple examples given above exhaust all possibilities for values of $e(G)$ where G is a finitely generated group, as shown by the following:

Theorem 4.8 (Freudenthal–Hopf; 1944). *Every finitely generated group has either* 0, 1, 2, *or infinitely many ends.*

Proof. Suppose G is a group with $3 \leqslant e(G) < \infty$ ends, for simplicity of notation we set $e(G) = k$. We fix a finite generating set for G and a corresponding Cayley graph, which as an abuse of notation, we also denote G.

By the exercise given above, since $k > 0$ we have $|G| = \infty$.

Since G has k ends, there exists a ball $B \subset G$ large enough so that $G - B$ has k unbounded connected components. Further, since G is infinite, there exists $g \in G$ satisfying $d(e, g) > 2 \cdot \mathrm{diam}(B)$ and with g in an unbounded component of $G - B$. Note, in particular, that this yields $B \cap gB = \varnothing$.

Since g acts as an isometry, one of the k unbounded connected components of $G - B$ contains gB and removing gB from this component divides it into at least $k - 1$ unbounded connected components. Thus $B \cup gB$ is a compact set in G whose removal leaves at least $2(k-1)$ unbounded components. Hence the space obtained by removing the ball of radius $2d(e, g)$ from G has at least $2(k-1)$ unbounded connected components. Since $k > 3$ we then have a contradiction, since by hypothesis $e(G) = k$ and whence $k = e(G) \geqslant 2(k-1) > k$. □

The following theorem, first proven by Wall in 1967, yields a straightforward way to recognize exactly which groups have two ends. There are a number of proofs of this result many of which are elementary. I like the following argument, the details of which appear in the book listed above by John Meier.

Theorem 4.9 (Wall). *A group G has two ends if and only if G contains a finite index subgroup isomorphic to $\mathbb{Z}$.*

Proof. ($\Leftarrow$): This follows from the fact that $\mathbb{Z}$ has two ends and that the number of ends is preserved by commensurability since it is preserved by quasi-isometries.

($\Rightarrow$): Recall the definition of symmetric difference, $A \triangle B = (A\cup B) - (A \cap B)$, and two easy consequences, whose proofs are left as exercises:

(1) For all A, B, D we have $A \triangle B = (A \triangle D) \triangle (B \triangle D)$;

(2) For all X and $A, B \subset X$ we have $A \triangle B = (X - A) \triangle (X - B)$.

Let $F \subset G$ be a finite graph for which $G - F$ has exactly two unbounded components. We will use the notation F^c to denote the complement of F in G. Without loss of generality we may assume F^c has exactly two connected components (we do this by, if necessary, adding all bounded components to F). The set F divides G into two unbounded components; we think of these as "left" and "right" sides and call the vertices of one of these sides R.

For any $g \in G$ we can consider the translates $g \cdot R = \{g \cdot r : r \in R\}$ and $g \cdot R^c = \{g \cdot l : l \notin R\}$. Note: $g \cdot R^c = (g \cdot R)^c$.

To continue the proof, we need the following two lemmas. When we speak of the cardiality of a subset of G we will mean the number of vertices in that subset.

Lemma 4.10. *Because G has two ends, either $R \triangle gR$ or $(R \triangle gR)^c$ is finite.*

Proof. $G - (F \cup gF)$ contains 2 unbounded components. Each of R and $g \cdot R$ union at most finitely many additional points (i.e., some of F, $g \cdot F$, or the bounded components) is equal to exactly one of these components.

Hence, if both R and $g \cdot R$ contain the same unbounded component, then $R \triangle gR$ is finite; otherwise, $(R \triangle gR)^c$ is finite. □

We leave the proof of the next Lemma as an exercise.

Lemma 4.11. *The subset of elements $H = \{g \in G : R \triangle gR\}$ is finite and forms a subgroup which is of index at most 2 in G.*

Let $h \in H$ satisfy $hF \cap F = \varnothing$, then either $hF \subset R$ or $hF \subset R^c$. Hence in this case we have either

$$R \cap hR^c = \varnothing \text{ and } R^c \cap hR \neq \varnothing$$

or

$$R \cap hR^c \neq \varnothing \text{ and } R^c \cap hR = \varnothing.$$

We think of the first case above as "translating to the left," as the image of R^c under h is disjoint from R and the image of R under h intersects R^c, similarly, think of the second case as "translating to the right."

We now formulate and prove the following proposition which will provide the key step in proving the theorem:

Proposition 4.12. *Let G and H be as above, then there exists a homomorphism $\phi : H \to \mathbb{Z}$ whose kernel is finite.*

Proof. Let $h \in H$, so that $R \Delta hR$ is finite. Thus $|R \cap hR^c| < \infty$ and $|R^c \cap hR| < \infty$.

We define the map $\phi : H \to \mathbb{Z}$ by setting:

$$\phi(h) = |R \cap hR^c| - |R^c \cap hR|.$$

We begin by the following straightforward observation for any $g, h \in H$:

$$R \cap hR^c = (R \cap hR^c \cap hgR^c) \sqcup \underbrace{(R \cap hR^c \cap hgR)}_{\text{A}}$$

$$\text{and } R^c \cap hR = \underbrace{(R^c \cap hR \cap hgR^c)}_{\text{B}} \sqcup (R^c \cap hR \cap hgR).$$

Hence

$$\phi(h) = |R \cap hR^c \cap hgR^c| + \underbrace{|R \cap hR^c \cap hgR|}_{\text{A}}$$
$$- \underbrace{|R^c \cap hR \cap hgR^c|}_{\text{B}} - |R^c \cap hR \cap hgR|.$$

The first equality here is by definition of ϕ, the second is since the h acts on G as an isometry: $\phi(g) = |R \cap gR^c| - |R^c \cap gR| = |hR \cap hgR^c| - |hR^c \cap hgR|$. Since each of these terms can be written as the sum of two terms, one measuring the intersection with R and the other with R^c, we can rewrite $\phi(g)$ as:

$$\phi(g) = |R \cap hR \cap hgR^c| + \underbrace{|R^c \cap hR \cap hgR^c|}_{\text{B}}$$
$$- \underbrace{|R \cap hR^c \cap hgR|}_{\text{A}} - |R^c \cap hR^c \cap hgR|.$$

Thus, we see that in $\phi(h) + \phi(g)$ the A and B terms are paired with opposite signs and thus we have: $\phi(h) + \phi(g) = |R \cap hgR^c| - |R^c \cap hgR| = \phi(hg)$, i.e., ϕ is a homomorphism from H to $\mathbb{Z}$.

Since $|F| < \infty$, there exists only finitely many $h \in H$ satisfying $hF \cap F \neq \varnothing$. Since Lemma 4.11 shows that $gF \cap F = \varnothing$ implies $\phi(g) \neq 0$, we know an element can only be in the kernal if it is in F. Then, since F is finite, we are done. □

To finish the proof of Wall's Theorem we now just need to observe that the previous Proposition yields a finite index subgroup $H \subset G$ and a homomorphism $\phi : H \to \mathbb{Z}$ with finite kernel. If $h \in H$ satisfies $\phi(h) \neq 0$, then h generates a finite index subgroup isomorphic to $\mathbb{Z}$, which is also a finite index subgroup in G. □

Groups with infinitely many ends also admit a particularly nice characterization, which we now state. This result is one of the first major theorems in modern geometric group theory and its proof is beyond the scope of these lectures.

Theorem 4.13 (Stallings Ends Theorem). *Let G be a torsion-free[2] finitely generated group. Then G has infinitely many ends if and only if G can be written as the free product of a pair of groups H and K which are both non-trivial and not both $\mathbb{Z}/2\mathbb{Z}$.*

5 How many groups are there?

Theorem 5.1 (B. Neumann). *There exists uncountably many finitely generated groups up to isomorphism.*

Neumann's original proof of this theorem in 1937, as well as most subsequent proofs rely on delicate arguments using torsion. In this Lecture we discuss an alternate approach, namely, we will describe the following stronger and purely geometric result, which has the Neumann Theorem as an immediate corollary:

[2] The torsion-free hypothesis can be removed; if this is done the phrase "free product" must be replaced by "amalgamation along a finite group," a phrase which might be unfamiliar to readers of these notes, which is why we took the special case in the main text.

Theorem 5.2 (Bowditch). *There exists uncountably many quasi-isometry classes of torsion-free small-cancellation 2-generator groups.*

We will begin by defining the small-cancellation condition. To do this we will fix a finite presentation $\langle A; \mathcal{R}\rangle$ with the property that for each relator $r \in \mathcal{R}$, the set $\mathbb{R}$ contains r^{-1} and also all cyclic conjugates of r. (Note that to any finite presentation, we can add finitely many additional relators to put it in this form.)

Definition 5.3. A cyclically-reduced, non-trivial word $w \in \mathbb{F}(A)$ is called a *piece* if there exist $r_1, r_2 \in \mathcal{R} \subset \mathbb{F}(A)$ such that $r_1 = wr_1'$ and $r_2 = wr_2'$ (i.e., if there exists two distinct relations for which w is an initial subword of both). We now define two of the most commonly used *small-cancellation conditions*:

We say $\langle A; \mathcal{R}\rangle$ satisfies the $C(p)$*-condition* if no element of $\mathcal{R}$ is a product of fewer than p pieces. We say $\langle A; \mathcal{R}\rangle$ satisfies the $C'(\frac{1}{p})$*-condition* if each piece w occurring in $r \in \mathcal{R}$ satisfies $p \cdot \ell(w) < \ell(r)$.

Remark 5.4. $C'(\frac{1}{p})$-condition implies $C(p)$-condition.

Example 5.5. *In earlier lectures we discussed $\pi_1(S_g)$ with the presentation*

$$\langle a_1, \ldots, a_g, b_1, \ldots, b_g; \prod_{i=1}^{g} a_g b_g a_g^{-1} b_g^{-1}\rangle.$$

To study pieces we need to modify this presentation by adding additional relators corresponding to each of the $4g-1$ other cyclic conjugates of this relator (for instance, $b_1 a_1^{-1} b_1^{-1} \cdots a_g b_g a_g^{-1} b_g^{-1} a_1^{-1}$, etc) as well as the inverse of this relator (namely, the word $\prod_{i=0}^{g-1} b_{g-i} a_{g-i} b_{g-i}^{-1} a_{g-i}^{-1}$) and each of its cyclic conjugates. In this presentation we see that a maximal piece is a single letter, so this presentation satisfies $C(4g)$ and $C'(\frac{1}{4g-1})$.

Exercise 5.6. $\mathbb{Z}^3 = \langle a, b, c; aba^{-1}b^{-1}, aca^{-1}c^{-1}, bcb^{-1}c^{-1}\rangle$. *Write the presentation which contains all inverses and cyclic conjugates of each relator. Then show that this presentation satisfies $C(4)$.*

A powerful tool in working with small-cancellation groups is a result called Greendlinger's Lemma which provides that in a group which satisfies a sufficiently strong small-cancellation condition, for instance $C'(\frac{1}{6})$, then any word representing the identity element contains a subword which is more than half of either a relator or a cyclic conjugate of a relator.

Hence, any $C'(\frac{1}{6})$-presented group admits a *Dehn presentation*: a presentation $\langle A; \mathcal{R}\rangle$ such that for any reduced non-trivial word $w \in \mathbb{F}(A)$ which represents the identity in $\langle A; \mathcal{R}\rangle$ there exist a relator $r = r_1 r_2 \in \mathcal{R}$ with $w = w_1 r_1 w_2$ and $\ell(r_1) > \ell(r_2)$ (and thus w can be represented by the shorter word $w = w_1 r_2^{-1} w_2$).

Upshot: the word problem is solvable in any group with a Dehn presentation by applying *Dehn's Algorithm*, as we saw in Lecture 1.

Definition 5.7. For Bowditch's Theorem 5.2 we use *small-cancellation* to mean $C'(\frac{1}{6})$, in particular, if r is a common subword of two relators r_1 and r_2 then either $r_1 = r_2$ or $6|r| < \min\{|r_1|, |r_2|\}$.

Let $\mathcal{P}(\mathbb{N})$ denote the power set of $\mathbb{N}$. Given $F, F' \in \mathcal{P}(\mathbb{N})$, we write $F \sim F'$ if their symmetric difference $F \triangle F'$ is finite. This is an equivalence relation on an uncountable set with each equivalence class countable. Thus there exist uncountably many equivalence classes, indeed the number of distict equivalence classes is the cardinality of $2^{\mathbb{N}}$.

Definition 5.8. Given $F \in \mathcal{P}(\mathbb{N})$, define $S(F) = \{2^{2^n} : n \in F\} \subset \mathbb{N}$. For any $p \in \mathbb{N}$, define

$$w_p = a(a^p b^p)^7 \in \mathbb{F}(a, b).$$

Given $S \subset \mathbb{N}$, define the group $\Gamma(S) = \langle a, b; \{w_p : p \in S\}\rangle$.

Proposition 5.9. *For any $F \in \mathcal{P}(\mathbb{N})$, the group $\Gamma(S(F))$ is $C'(\frac{1}{6})$.*

Proof. The length of w_p is $14p+1$. For any $p' \neq p$, the longest piece which w_p and $w_{p'}$ share is $a^p b^p$ (here we assume, without loss of generality, that $p < p'$) and hence has length $2p$, and $6 \times 2p < 14p + 1$. □

Remark 5.10. A result in small-cancellation theory states that any torsion must arise in the cyclic symmetry of a relator (e.g., a group has an order two element if and only if it has a relator of the form w^2). Hence, $\Gamma(S(F))$ is torsion-free.

Theorem 5.11. *If $F, F' \in \mathcal{P}(\mathbb{N})$ and $\Gamma(S(F))$ is quasi-isometric to $\Gamma(S(F'))$, then $F \sim F'$.*

Sketch of proof. We are only going to give the idea of the argument here. The two main ingredients in the argument are the fact that the group is $C'(\frac{1}{6})$ and the super-exponential growth of length of relators.

In a group with a presentation satisfying $C'(\frac{1}{6})$, the lengths of relators determine the size of "holes" in the Cayley graph, i.e., long loops in the graph for which there is no shortcut. The sizes of these holes geometrically determine a subset of $\mathbb{N}$; a quasi-isometry of the group only changes these sizes by a bounded multiplicative and additive amount. Hence, arranging these holes to grow super-exponentially ensures that from the quasi-isometry type we can recover, up to finite error, the original set of sizes, as we shall see in the following lemma. □

Let $A \subset \mathbb{N}$ and $q \in \mathbb{N}$, define $q \cdot A = \{qa : a \in A\}$.

Lemma 5.12. *Let $F, F' \subset \mathbb{N}$ for which $q \cdot S(F)$ and $q \cdot S(F')$ satisfy the property that there exists $k \in \mathbb{N}$ such that for each $m \in q \cdot S(F)$ satisfying $m > k$ there exists $m' \in q \cdot S(F')$ satisfying $\frac{m}{k} \leqslant m' \leqslant km$. Then $F \sim F'$.*

Proof. If not, then there exist $n \in F \setminus F'$, such that $2^{2^{n-1}} > k$. Since $2^{2^{n-1}} q \in q \cdot S(F)$, there exists $m \in F'$ such that $\frac{2^{2^n} q}{k} \leqslant 2^{2^m} q \leqslant 2^{2^n} kq$. Thus $2^{2^{n-1}} < \frac{2^{2^n}}{k} \leqslant 2^{2^m} \leqslant 2^{2^n} k < 2^{2^{n+1}}$. Hence $m = n$ which contradicts our assumption that $n \notin F'$. □