Geometric Group Theory

Shing Tak Lam

Lent 2023*

Contents

Lecture 1

1 Combinatorial group theory

Combinatorial group theory is a sibling field to Geometric group theory. Both study infinite discrete groups.

*Based on lectures by Henry Wilton. Last updated March 13, 2024.

1.1 Free groups and presentations

Let $A = \{a_1, a_2, \dots\}$ be an *alphabet*. A group F is *free on* A if

1. There is a map of sets $A \to F$,
2. for any group G, and a map of sets $A \to G$, there exists a unique group homomorphism $F \to G$ such that

$$\begin{array}{ccc} & & F \\ & \nearrow & \downarrow \\ A & \longrightarrow & G \end{array}$$

   commutes.

This is a *universal property*. As usual, this means that F is unique up to unique isomorphism. This shows that F is determined by A, so we may write $F = F(A)$.

However (as usual with definitions by universal property), we don't know if $F(A)$ exists. We'll show this two different ways.

1. Topologically: Let

$$X = \bigvee_{a \in A} S^1$$

By the Seifert-van Kempen theorem,

$$\pi_1(X) \cong F(A)$$

2. Combinatorially: Let

$$A^* = \{\text{words in } A \sqcup A^{-1}\}$$

where $A^{-1} = \{a_1^{-1}, \dots, a_n^{-1}\}$. For example,

$$1 = \varnothing, aa, aa^{-1}, aba^{-1}b^{-1}, a^{100}ba^{-100}b, \dots$$

A word is *reducible* if it contains aa^{-1} or $a^{-1}a$ as a subword for any $a \in A$. Otherwise, it is *reduced*. We can now define

$$F(A) = \{w \in A^* \text{ reduced}\}$$

The group operation is concatenation, followed by reduction. For example,

$$(abab^{-1})(b^2a) = abab^{-1}b^2a = ababa$$

Note that reduction terminates as each reduction step reduces the length. We won't check that this is well defined or associative. The identity element is 1, inverses is clear.

A *presentation* consists of an alphabet A, which we will call *generators*, and a set $R \subseteq F(A)$, which we will call *relations*, and we write

$$\langle A \mid R \rangle = \langle a_1, a_2, \cdots \mid r_1, r_2, \dots \rangle$$

This *presents* a group

$$G = \frac{F(A)}{\langle\langle R \rangle\rangle}$$

where $\langle\langle R \rangle\rangle$ is the *normal closure* of R, i.e. the smallest normal subgroup of $F(A)$ containing R.

Example 1.1.1

Some examples of presentations which we have seen:

$$\begin{aligned} \langle a \mid a^n \rangle &\cong \mathbb{Z}/n\mathbb{Z} \\ \langle r, s \mid r^n, s^2, srsr \rangle &\cong D_{2n} \\ \langle A \mid \rangle &\cong F(A) \\ \langle a, b \mid aba^{-1}b^{-1} \rangle &\cong \mathbb{Z}^2 \\ \langle a_1, \dots, a_g, b_1, \dots, b_g \mid a_1b_1a_1^{-1}b_1^{-1} \cdots a_gb_ga_g^{-1}b_g^{-1} \rangle &\cong \pi_1(\Sigma_g) \end{aligned}$$

where Σ_g is the compact orientable surface of genus g.

As we see, presentations arise when we write down fundamental groups of spaces. In fact, all groups arise this way.

Corollary 1.1.2 (of Seifert-van Kampen). For

$$G = \langle a_1, a_2, \cdots \mid r_1, r_2, \ldots \rangle$$

there exists a space X with $\pi_1(X) = G$.

Proof. First, start of with a wedge of circles, one for each a_i. Also consider a disjoint union of discs, one for each r_i. Attach the i-cell along its boundary, which is a loop in the wedge of circles. □

This is called the *presentation complex* of G. For example, if we have $\mathbb{Z}^2 = \left\langle a, b \mid aba^{-1}b^{-1} \right\rangle$, then $X = \mathrm{T}^2$ is the 2-torus.

In 1911, Max Dehn posed the following problems.

1. (Word problem) Given $w \in A^*$, determine whether or not $w = 1$ in $G = \langle A \mid R \rangle$. That is, whether or not $w \in \langle\langle R \rangle\rangle \subseteq F(A)$.
2. (Conjugacy problem) Given $G = \langle A \mid R \rangle$, $u, v \in A^*$, determine whether or not u is conjugate to v in G.
3. (Isomorphism problem) Given $G = \langle A \mid R \rangle$, $H = \langle B \mid S \rangle$, determine whether or not G is isomorphic to H.

Lecture 2

Remark 1.1.3. The conjugacy problem is stronger than the word problem, since $w = 1$ if and only if w is conjugate to 1.

Dehn was motivated by topology, but the problems asks for algorithms. We will often solve them using geometry.

All three were *un*solved in the 1950s, as all three problems are algorithmically undecidable. Nevertheless, positive solutions are known for many reasonable classes of groups.

Let $A = \{a_1, \ldots, a_n\}$ be a finite alphabet.

Example 1.1.4 (word problem in free groups)

Let $w \in A^*$. If w is reduced, then $w = 1$ if and only if $w = \varnothing$. Otherwise, w contains a subword aa^{-1} for some $a \in A \cup A^{-1}$. Cancelling aa^{-1} produces a word $w' \in A^*$, such that $w = w'$ in $F(A)$, and $\ell(w') = \ell(w) - 2 < \ell(w)$. This terminates after finitely many steps.

We can also solve the conjugacy problem for free groups.

Definition 1.1.5

There is a natural action of $\mathbb{Z}$ on A^* permuting words. That is,

$$1 \cdot a_1 \cdots a_k = a_2 a_3 \cdots a_k a_1$$

The elements of $\mathbb{Z} \cdot w$ are called the *cyclic conjugates* of w. Note that all $u \in \mathbb{Z}w$ are conjugate to w. The orbits $\mathbb{Z} \backslash A^*$ are called *cyclic words*. A word is *cyclically reduced* if every cyclic conjugate is reduced.

Example 1.1.6

aba^{-1} is reduced, but not cyclically reduced as $ba^{-1}a$ is not reduced.

Note that if w is reduced, but not cyclically reduced, then

$$w = aw'a^{-1}$$

for some $a \in A \cup A^{-1}$. Note that w' is both conjugate to w and shorter than w. Hence after finitely many iterations, we can assume that w is cyclically reduced.

Lemma 1.1.7 (conjugacy problem in free groups)**.** If $u, v \in F(A)$ are cyclically reduced, then u is conjugate to v if and only if the corresponding cyclic words are the same.

Proof. If u, v have the same cyclic words, then v is a cyclic conjugate of u, which we have seen is a conjugate of u.

Conversely, suppose $u = gvg^{-1}$. By induction on $\ell(g)$, we can assume $g = a \in A \cup A^{-1}$. From this it follows that either $v = a^{-1}v'$, or $v = v'a$, as v is cyclically reduced. That is, $u = v'a^{-1}$ or $u = av'$. In both cases, they are cyclic conjugates. □

1.2 Historical case study

Let's briefly think about the state of topology in the early 20th century. Poincaré knew that homology classifies the compact two-dimensional surfaces. This motivated the

Conjecture 1.2.1 (Poincaré conjecture, version 1)**.** Let M be a compact connected 3-manifold, with

$$\mathrm{H}_*(M) = \begin{cases} \mathbb{Z} & * = 0, 3 \\ 0 & \text{otherwise} \end{cases}$$

Then M is homeomorphic to S^3.

Such a 3-manifold M is called a *homology sphere*.

Theorem 1.2.2 (Poincaré, 1904)**.** There exists a three dimensional homology sphere P, with

$$\pi_1(P) \twoheadrightarrow A_5$$

The moral is that: homology is not enough, we need use π_1 as well.

Conjecture 1.2.3 (Poincaré conjecture, version 2)**.** Let M be a compact connected 3-manifold, with $\pi_1(M) = 1$. Then M is homeomorphic to S^3.

This was proven by Perelman in 2003. Returning to the original conjecture, in 1910 Dehn wanted to construct more homology spheres.

Theorem 1.2.4 (Dehn, 1910)**.** There are infinitely many non-homeomorphic 3-dimensional homology spheres.

Remark 1.2.5. The isomorphism problem is exactly what is needed to distinguish these manifolds.

Lecture 3

Here is Dehn's construction. Consider the trefoil knot $K \subseteq \mathbb{R}^3 \subseteq S^3$. Let $N(K)$ be a regular open neighbourhood of K, i.e. a thickening of K. Let $N = S^3 \setminus N(K)$. This is a compact 3-manifold with boundart $\partial N \cong \mathrm{T}^2$.

Computing,

$$\mathrm{H}_*(N) = \begin{cases} \mathbb{Z} & * = 0, 1 \\ 0 & \text{otherwise} \end{cases}$$

and

$$\pi_1(N) \cong \langle x, y, z \mid x^2 = y^3 = z \rangle$$

This follows from the fact that K is a torus knot, and we can compute the π_1 of the complement of a torus knot. Note that the abelianisation map

$$\pi_1(N) \to \mathrm{H}_1(N) \cong \mathbb{Z}$$

sends z to 6, x to 3 and y to 2. Moreover, the boundary torus T has $\pi_1(T) = \mathbb{Z}^2$, generated by xy and z. Under the abelianisation map, xy is sent to 5.

We can glue a solid torus $D^2 \times S^1$ to N, by a homeomorphism on the boundary. Let λ, μ be the corresponding loops on T^2. The resulting manifold $M_\phi = N \cup_\phi T$ is closed. $\pi_1(M_\phi)$ depends on $g = \phi_*(\mu)$. By Seifert-van Kapmen,

$$\pi_1(M_\phi) = \frac{\pi_1(N)}{\langle\langle g \rangle\rangle}$$

Similarly,

$$\mathrm{H}_1(M_\phi) = \frac{\mathbb{Z}}{\langle [g] \in \mathrm{H}^1(N) \rangle}$$

To produce a homology sphere, we need to choose ϕ such that $g = \phi_*(\mu) = 1$ in $\mathrm{H}^1(N)$.

If $g = (xy)^a z^b$, then in $\mathrm{H}^1(N)$ this is mapped to $5a + 6b \in \mathbb{Z}$. Choose $a = 6n + 5$, and $b = -(5n + 4)$ for $n \in \mathbb{Z}$. He constructs ϕ_n such that

$$\phi_n(\mu) = g_n = (xy)^{6n+5} z^{-(5n+4)}$$

His family of manifolds

$$D_n = N \cup_{\phi_n} U$$

has

$$\pi_1(D_n) = \left\langle x, y, z \mid x^2 = y^2 = z, (xy)^{6n+5} = z^{5n+4} \right\rangle$$

The remaining challenge is to prove the groups $G_n = \pi_1(D_n)$ for $n \geq 0$ are pairwise non-isomorphic.

This is the isomorphism problem! In particular,

$$g_n = g_m \implies g_n \text{ and } g_m \text{ are conjugate} \implies G_n \cong G_m$$

So we also need to solve the word and conjugacy problem in $\pi_1(N)$.

1.3 van Kampen diagrams

Definition 1.3.1

A map $f : Y \to X$ of cell complexes is called *combinatorial* if for all $k \in \mathbb{Z}_{\geq 0}$, and every k-cell e^k of Y, f maps the interior $\mathrm{Int}(e^k)$ homeomorphically to the interior of a k-cell of X.

Consider a presentation $G = \langle a_i \mid r_j \rangle$, and the associated presentation complex X.

Definition 1.3.2

A *(singular) disc diagram* is a compact contractible 2-complex D, with an embedding $D \hookrightarrow \mathbb{R}^2$. A disc diagram D is *over* X if it is equipped with a combinatorial map $D \to X$.

Example 1.3.3

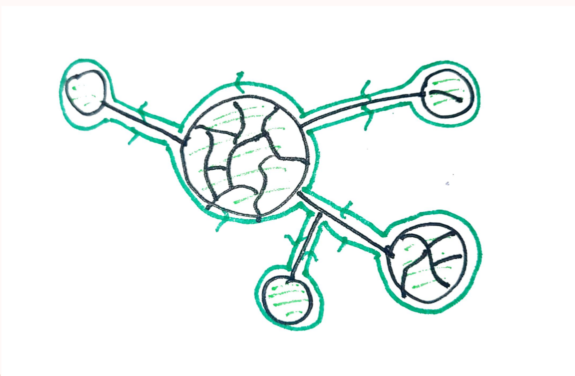

Recall that X is given by a wedge of circles, with discs glued on for each relation. So the 1-cells correspond to generators, 2-cells go to relations (or cyclic conjugates, or inverses). With this:

- every oriented 1-cell of D is labelled with some $a_i \in A$,
- so that each 2-cell has boundary which is a cyclic conjugate of some $r_j^{\pm 1}$.

Associated to each disc diagram D, we have a *boundary cycle*, which reads a (cyclic[1]) word $w \in A^*$, which reduces to an element $w' \in \langle\langle r_j \rangle\rangle \leq F(A)$. To see this, D is contractible.

D is a *van Kampen diagram for* w.

Lecture 4

Lemma 1.3.4 (van Kampen)**.** If $w \in \langle\langle R \rangle\rangle$, then there exists a van Kampen diagram for w.

Proof. Suppose $w \in \langle\langle R \rangle\rangle$. Then w can be written as

$$w = \prod_{i=1}^{k} h_i r_i^{\pm 1} h_i$$

in $F(A)$, where $h_i \in F(A)$, and $r_i \in R$. Now build a *lollipop diagram* D_0, which has boundary word w_0, which is equal to w in $F(A)$, but may not be reduced.

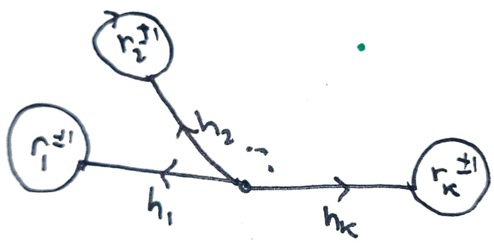

If w_0 is reduced, $w = w_0$, and so we are done. Otherwise, w_0 contains a cancelling pair, so

$$w_0 = \cdots a a^{-1} \cdots$$

for some $a \in A \cup A^{-1}$.

We can see that e_1, e_2 share a vertex. There are two cases to consider:

1. if the *origin* of e_1 is the *terminus* of e_2, then the diagram D_0 is a wedge $D_1 \vee D'$.

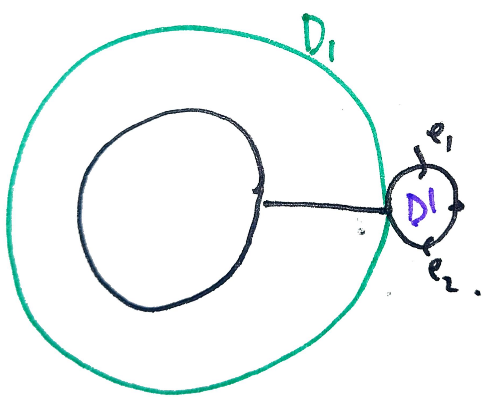

   Then D_1 is a van Kampen diagram for w_1, which is the result of cancelling a and a^{-1}.

2. if the origin of e_1 is distinct from the terminus of e_2, then we can fold the edges to get D_1,

[1] or a word once we choose a base point.

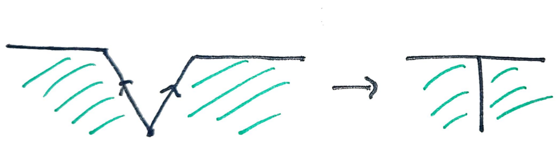

which has boundary word w_1 as above.

In either case, $w_1 = \partial D_1$ is obtained from w_0 by cancelling a pair. Therefore, we may proceeed by induction, and after finitely many repetitions, we construct a van Kampen diagram D_n such that $w_n = \partial D_n$ is reduced, and $w_n = w$ in $F(A)$. Thus, $w_n = w$ as words, and so D_n is a van Kampen diagram for w. ☐

Remark 1.3.5. The minimal number of 2-cells in a van Kampen diagram of w is the minimal number of k, such that w can be written as

$$w = \prod_{i=1}^{k} h_i r_i^{\pm 1} h_i^{-1}$$

This is called the *area* of w.

Example 1.3.6

Let $G = \mathbb{Z}^2 = \langle a, b \mid [a, b] \rangle$. Consider $w = a^n b^n a^{-n} b^{-n}$. This has van Kampen diagram

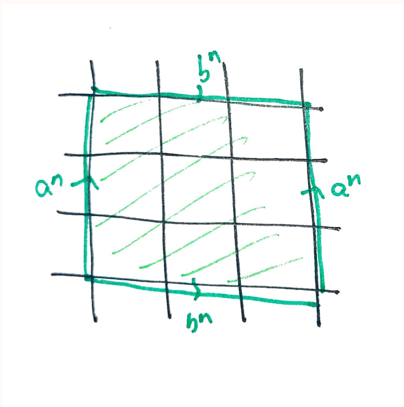

In this case, $\text{Area}(D) = n^2$. We will show D is minimal, and so $\text{Area}(w) = n^2$.

Definition 1.3.7

Let $\mathcal{P} = \langle A \mid R \rangle$ be a finite presentation of a group G. Define

$$\begin{aligned} \delta_{\mathcal{P}} : \mathbb{N} &\to \mathbb{N} \\ \ell &\mapsto \max_{w \in \langle\langle R \rangle\rangle, \ell(w) = d} \text{Area}(w) \end{aligned}$$

This is called the *Dehn function*.

Remark 1.3.8. The word problem in $\mathcal{P}$ if and only if $\delta_{\mathcal{P}}$ is computable.

2 Basics of geometric group theory

2.1 Cayley graphs

A *graph* is a 1-dimensional cell complex. Throughout, let G be a group, with finite generating set $S \subseteq G$.

Definition 2.1.1

The *Cayley graph* $\mathrm{Cay}_S(G)$ is defined as follows:

- vertices $V(\mathrm{Cay}_S(G)) = G$,
- edges $E(\mathrm{Cay}_S(G))$ correspond bijectively with $G \times S$. That is, we have an edge $g \to gs$.

Example 2.1.2

The trivial group given by $1 = \langle a, b \mid a, b \rangle$ has Cayley graph

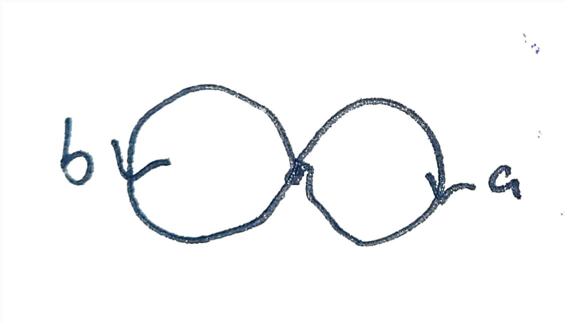

Example 2.1.3

$S_3 = \left\langle r, s \mid srsr, r^3, s^2 \right\rangle$ has Cayley graph

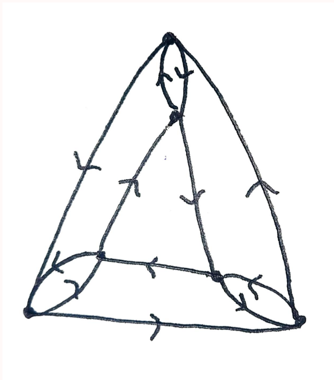

Example 2.1.4

$\mathbb{Z} = \langle 1 \rangle$ as Cayley graph

On the other hand, $\mathbb{Z} = \langle 2, 3 \rangle$ has Cayley graph

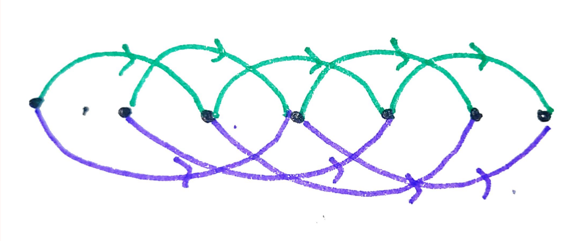

Note that the action of G on itself on the *left* extends to an action of G on $\mathrm{Cay}_S(G)$, sending an edge $h \to hs$ to $gh \to ghs$. Note the right action does not work, because of our definition of the Cayley graph.

Remark 2.1.5. The action of G on $\mathrm{Cay}_S(G)$ is free. That is, for all $x \in \mathrm{Cay}_S(G)$, $\mathrm{Stab}_G(x) = 1$.

Proposition 2.1.6. Let $G = \langle S \mid R \rangle$, and let X be the corresponding presentation complex. Then there exists a G-equivariant isomorphism of graphs

$$\mathrm{Cay}_S(G) \cong \widetilde{X}_{(1)}$$

with the 1-skeleton of the universal cover $\widetilde{X}$ of X[a].

[a] and not the other way around.

Lecture 5

Proof. Consider the natural free action of $G = \pi_1(X)$ on $\widetilde{X}$, by deck transformations. The action is by combinatorial endomorphisms. In particular, it preserves the 1-skeleton. So we have a free action of G on $\widetilde{X}_{(1)}$, which sends vertices to vertices and edges to edges.

The action of G on $\widetilde{X}_{(0)}$ is free, and as X has only one vertex, transitive. Therefore, choosing a base vertex $\widetilde{v}_0$, the orbit-stabiliser theorem provides a G-equivariant bijection

$$G \to \mathrm{Orb}(\widetilde{v}_0) = \widetilde{X}_{(0)}$$

sending g to $g \cdot \widetilde{v}_0$. So this matches up the vertices as required.

Next, let us match up the edges. For each $s \in S$, let e_s be the corresponding edge of X. Let $\widetilde{e}_s$ be the unique lift of e_s to $\widetilde{X}$, beginning at $\widetilde{v}_0$. By the definition of the action of G on $\widetilde{X}$, e_s ends at $s \cdot \widetilde{v}_0$. Now an arbitrary edge $\widetilde{e}$ of $\widetilde{X}$ maps to some e_s, under the covering map. Since egdes of X correspond to G-orbits of edges in $\widetilde{X}$, it folloes that $\widetilde{e} = g \cdot \widetilde{e}_s$ for some $s \in S$. That is, $\widetilde{e}$ is the edge from $g \cdot \widetilde{v}_0$ to $gs \cdot \widetilde{v}_0$. So it corresponds to an edge from g to gs.

This shows that the G-equivariant map $G \to \widetilde{X}_{(0)}$ extends to a G-equivariant isomorphism of graphs as claimed. □

The next proposition deepens the relationship between generating sets and path connectedness.

Proposition 2.1.7. Let $\widetilde{X}$ be a path connected topological space, and suppose that G acts on $\widetilde{X}$ by homeomorphisms. If $U \subseteq \widetilde{X}$ is an open subset, such that $G \cdot U = \widetilde{X}$, then the set

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

generates G.

Proof. Fix a base point $\widetilde{x}_0 \in U$. Now for $g \in G$, let $\gamma : [0, 1] \to \widetilde{X}$ be a path from $\widetilde{x}_0$ to $g \cdot \widetilde{x}_0$. The set $\{\gamma^{-1}(h \cdot U) \mid h \in G\}$ is an open cover of $[0, 1]$. So it has a finite subcover, $\{\gamma^{-1}(U_1), \ldots, \gamma^{-1}(U_n)\}$, where $U_i = g_i \cdot U$. We may choose the indices so that

- $\widetilde{x}_0 \in U_1$,
- $\gamma^{-1}(U_i) \cap \gamma^{-1}(U_{i+1}) \neq \varnothing$ for all i,

- $g \cdot \widetilde{x}_0 \in U_n$.

Note that the g_i need not be unique. By definition, $x_0 \in U \cap g_1 \cdot U$, and so $g_1 \in S$. Similarly, if $t_i \in \gamma^{-1}(U_i) \cap \gamma^{-1}(U_{i+1})$, then $x_i = \gamma(t_i) \in g_i \cdot U \cap g_{i+1} \cdot U$. Thus,

$$g_i^{-1} \cdot x_i \in U \cap g_i^{-1} g_{i+1} \cdot U$$

and so $s_i = g_i^{-1} g_{i+1} \in S$. Thus, $g_n = s_{n-1} \cdots s_2 s_1 g_1$, is a finite product of elements of S. Finally, $g^{-1} g_n \in S$ similarly to the above, so $g \in \langle S \rangle$ as required. □

Example 2.1.8

Let $\Gamma \subseteq \mathrm{Isom}(\mathbb{R}^2)$ be the symmetry group of the standard tiling of the plane by equilateral triangles. Let U be a thickened triangle. Using the proposition, we obtain a finite generating set of Γ. In particular, Γ is generated by the reflections in the sides of a single triangle.

In particular, this is not a covering space action, as it is not free.

Definition 2.1.9

An action of G on $\widetilde{X}$ by homeomorphisms is *properly discontinuous* if for every compact $K \subseteq \widetilde{X}$, the set

$$\{g \cdot K \cap K\}$$

is finite.

The action is *cocompact* if there exists $K \subseteq \widetilde{X}$ compact, such that

$$G \cdot K = \widetilde{X}$$

$\widetilde{X}$ is *locally compact* if for every neighbourhood U of x, there exists an open neighbourhood $V \subseteq U$ of x, such that $\overline{V} \subseteq U$ is compact.

Corollary 2.1.10. If G acts on $\widetilde{X}$ properly discontinuously and cocompactly, and $\widetilde{X}$ is path connected and locally compact, then G is finitely generated.

Lecture 6

Proof. Let $K \subseteq \widetilde{X}$ be compact, such that $G \cdot K = \widetilde{X}$. By local compactness, we may find an open U such that $K \subseteq U$, and $\overline{U}$ is compact. In particular, $G \cdot U = \widetilde{X}$, and the set

$$S = \{g \in G \mid g \cdot U \cap U \neq \varnothing\} \subseteq \{g \mid g \cdot \overline{U} \cap \overline{U} \neq \varnothing\}$$

But the right hand side is finite, so S is finite. By the proposition, S generates G. □

Corollary 2.1.11. If X is compact, locally compact and has a universal cover $\widetilde{X}$, then $\pi_1(X)$ is finitely generated.

Proof. Exercise. Sheet 1 question 10. □

2.2 The Schwarz-Milnor lemma

Cayley graphs are not just combinatorial. They admit a natural metric, called the *word metric*.

Definition 2.2.1 (word metric)
Let S generate G. Define

$$\ell_S(G) = \min\{n \mid g = \prod_{i=1}^{n} s_i^{\pm 1}, s_i \in S\}$$

This defines a metric

$$d_S(g,h) = \ell_S(g^{-1}h)$$

called the *word metric* associated to S.

The word metric is invariant under the *left* G action on itself. That is,

$$d_S(\gamma g, \gamma h) = d_S(g,h)$$

However, it is, in general, not right invariant.

Example 2.2.2
$G = \mathbb{Z}^2 = \langle a \rangle \oplus \langle b \rangle$. Then the word metric is just the ℓ_1-metric.

Remark 2.2.3. The word metric extends naturally to a left invariant metric on $\mathrm{Cay}_S(G)$, in which the interior of each edge is locally isometric to $(0,1)$. That is, the path metric.

Lemma 2.2.4. Suppose S, T are finite generating sets for G. Then there exists constants $C, C'' \geq 1$ such that

$$\frac{1}{C} d_T \leq d_S \leq C' d_T$$

Proof. Let $C = \max_{s \in S} \ell_T(s)$. Then for any $g \in G$,

$$\ell_T(g) \leq C \ell_S(g)$$

by induction. □

That is, for finitely generated groups, the word metric is well defined, up to bi-Lipschitz equivalence.

Definition 2.2.5 (quasi-isometry)
A function[a] $f: X \to Y$ between metric spaces is a *quasi-isometric embedding* if there are constants $C \leq 1$, $D \geq 0$, such that

$$\frac{1}{C} d(x,x') - D \leq d(f(x), f(x')) \leq C d(x,x') + D$$

for all $x, x' \in X$.

If in addition, there exists a constant K such that for every $y \in Y$, there exists $x \in X$ such that $\mathrm{d}(f(x), y) \leq K$, then f is called a quasi-isometry, and we write $X \overset{\text{qi}}{\sim} Y$.

[a] It does not have to continuous.

Remark 2.2.6. On examples sheet 1, we have that $\overset{\text{qi}}{\sim}$ is an equivalence relation.

Example 2.2.7
Every bounded metric space is quasi-isometric to a point.

Definition 2.2.8 (proper)
A metric space X is *proper* if closed balls in X are compact.

Definition 2.2.9 (geodesic)
A *geodesic* in X is an isometric embedding $\gamma : [a, b] \to X$. The metric space X is *geodesic* if every pair of points is joined by a geodesic.

Theorem 2.2.10 (Schwarz-Milnor). Suppose X is a proper geodesic metric space. Let G acts on X properly discontinuously and cocompactly by isometries. Then G is finitely generated, and

$$X \overset{\text{qi}}{\sim} (G, d_S)$$

for any finite generating set S of G.

Proof. Fix a base point $x_0 \in X$. Let $B = \overline{B}(x_0, K) \subseteq X$ be a closed ball, such that $G \cdot B = X$. By properness and proper discontinuity, the set

$$\{g \in G \mid d(x_0, g(x_0)) \leq 3K\}$$

is finite. Therefore, there exists $\varepsilon > 0$, such that

$$d(x_0, g(x_0)) < 2K + \varepsilon \iff d(x_0, g(x_0)) \leq 2K$$

Moreover, in this case, $gB \cap B \neq \varnothing$.

Lecture 7

If $U = B(x_0, K + \varepsilon/2)$, then

$$S = \{g \cdot U \cap U \neq \varnothing\} = \{g \cdot B \cap B \neq \varnothing\}$$

Since B is compact, S is finite, since the action is properly discontinuous. But S is a generating set for G.

Since the word metric for any two finite generating sets are bi-Lipschitz, we may prove the result for the S above. Consider the map $f : G \to X$, $f(g) = g \cdot x_0$. We claim that this is a quasi-isometry. f is quasi-surjective, since $G \cdot B = X$. It remains to prove that f is a quasi-isometric embedding. That is, we want upper and lower bounds on $d(x_0, g \cdot x_0)$ in terms of $\ell_S(g)$.

For the upper bound, take $C = \max_{s \in S} d(x_0, s \cdot x_0)$. Then

$$d(x_0, g \cdot x_0) \leq C \ell_S(g)$$

for any $g \in G$, using the triangle inequality.

For the lower bound, consider a geodesic $\gamma : [0, d(x_0, g \cdot x_0)] \to X$ from x_0 to $g \cdot x_0$. Choose a dissection of $[0, d(x_0, g \cdot x_0)]$

$$0 = t_0 < t_1 < \cdots < t_n < t_{n+1} = d(x_0, g \cdot x_0)$$

with

$$\frac{\varepsilon}{2} \leq |t_i - t_{i+1}| < \varepsilon$$

for $0 \leq i \leq n - 1$. Note we can make it so that $|t_n - t_{n+1}| < \varepsilon$, but we may not have the lower bound.

Since $G \cdot B$ covers X, for $1 \leq i \leq n$, there exists $g_i \in G$ such that $\gamma(t_i) \in g_i \cdot B$. Set $g_0 = 1$ and $g_{n+1} = g$. Then $\gamma(t_i) \in g_i \cdot B$ for all i. For each i,

$$d(g_i(x_0), g_{i+1}(x_0)) < 2K + \varepsilon$$

by the triangle inequality. Therefore, $g_i^{-1} g_{i+1} \in \{h \cdot U \cap U \neq \varnothing\} = S$. Hence $\ell_S(g) \leq n + 1$. Furthermore,

$$|t_i - t_{i-1}| \geq \varepsilon/2$$

for all $1 \leq i \leq n$, so $d(x_0, g \cdot x_0) \geq n\varepsilon/2$. Combining these,

$$\ell_S(g) \leq n + 1 \leq \frac{2}{\varepsilon} d(x_0, g \cdot x_0) + 1$$

We can rearrange this to get the lower bound. □

Example 2.2.11

Recall the two Cayley graphs of $\mathbb{Z}$, with generating sets $\{1\}$, and $\{2, 3\}$ respectively.

The Schwartz-Milnor says that these are both quasi-isometric to $\mathbb{Z}$ with an appropriate word metric. So they are quasi-isometric. More generally, for any finitely generated group G, the Cayley graphs of any two finite generating sets are quasi-isometric.

Corollary 2.2.12. If G is finitely generated, H is a subgroup with finite index in G. Then H is finitely generated, and H is quasi-isometric to G.

Proof. H acts on $\mathrm{Cay}_S(G)$. The action is cocompact as H has finite index. So it satisfies the Schwartz-Milnor lemma. □

Example 2.2.13

Let Σ_2 be the closed orientable surface of genus 2, and $G = \pi_1(\Sigma_2)$. Choose a Riemannian metric g on Σ_2 of constant curvature -1.

This pulls back to a Riemannian metric on its universal cover $\widetilde{\Sigma}_2$. By a classical theorem of differential geometry, $\widetilde{\Sigma}_2$ is isometric to the hyperbolic plane $\mathbb{H}^2$. Moreover, the action of G on the $\mathbb{H}^2$ by isometries, and properly discontinuously. The action is cocompact as the quotient is Σ_2, which is compact. So by the Schwartz-Milnor lemma, $\pi_1(\Sigma_2)$ is quasi-isometric to $\mathbb{H}^2$.

3 Case study - Free groups

Let $A = \{a_1, \ldots, a_n\}$. We will write $F_n = F(A_n)$. The *Cayley tree* is the infinite $2n$-valent tree $T_n = \mathrm{Cay}_A(F_n)$.

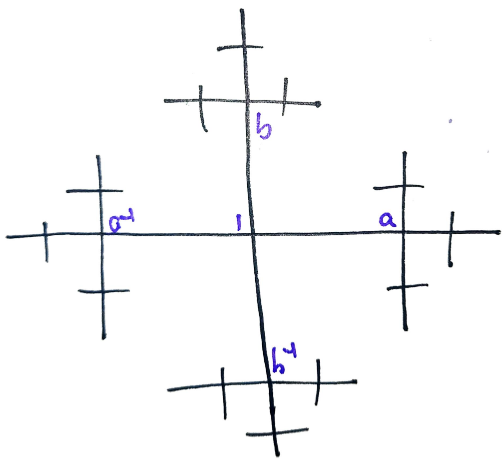

Lecture 8

In particular, every vertex looks the same. F_n acts freely on T_n. The quotient X_n is the wedge of n circles. So we recover $F_n = \pi_1(X_n)$, and T_n is the universal cover of X_n.

We can translate our *combinatorial* arguments about F_n, about geometric properties of T_n.

Words

A word $w \in A^*$ is equivalent to an *edge path*, which is a map $w : I \to X_n$, where I is an interval.

For example, consider the word $w = a^2bb^{-1}a^{-1}$. The edge path is:

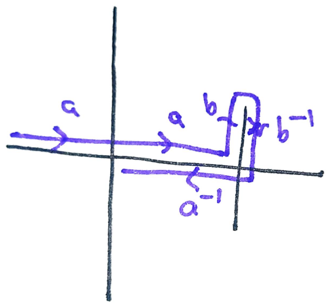

An edge path in X_n lifts to a unique edge path in T_n, based at 1. Conversely, each such path in T_n projects to a path in X_n.

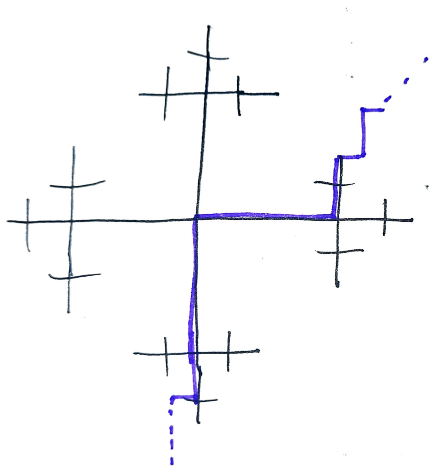

Reduced words

A word $w \in A^*$ is reduced if and only if the corresponding edge path $w : I \to X_n$ is locally injective. In turn, this holds if and only if the corresponding edge path $w : I \to T_n$ is locally injective. This is because an edge path can only fail to be locally injective at a vertex.

Clearly, the shortest path in T_n from 1 to $g \in F_n$ is injective. In particular, locally injective. So every element of F_n is represented by a reduced word.

The fact that this representative is unique follows from the next lemma.

Lemma 3.0.1. If T is a tree, and $\gamma : I \to T$ is a locally injective (edge) path, then γ is injective.

Proof. Suppose not. Let $\gamma : [a, b] \to T$ be the shortest counterexample. In particular, $\gamma(a) = \gamma(b)$, and γ is injective on (a, b). So γ descends to an injective map $S^1 \to T$. But T is a tree. Contradiction. □

Similarly, if $g \in F_n$ is shortest such that g is represented by distinct reduced words w_1, w_2, then we get an embedding $S^1 \hookrightarrow T$. Hence the reduced word is unique.

For $g \in F_n$, write $[1, g]$ for the unique injective edge path from 1 to g.

Cyclically reduced words

So far, implicitly we have chosen base points. Each (nontrivial) word $w \in A^*$ also defines a (based) edge loop, by gluing together the end points of the interval. So we have a map $S^1 \to X_n$. If we forget the base point of S^1, then two elements $u, v \in A^*$ determine the same cyclic word if and only if they represent the same edge loop $S^1 \to X_n$.

Now a word $w \in A^*$ is cyclically reduced if and only if the corresponding map $S^1 \to X_n$ is locally injective. From lifting theory, we have a lift $\widetilde{w} : \mathbb{R} \to T_n$ as below

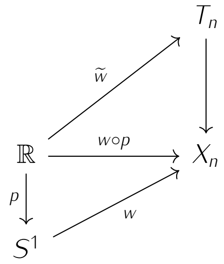

For example, if $w = ab^2$. the lift is

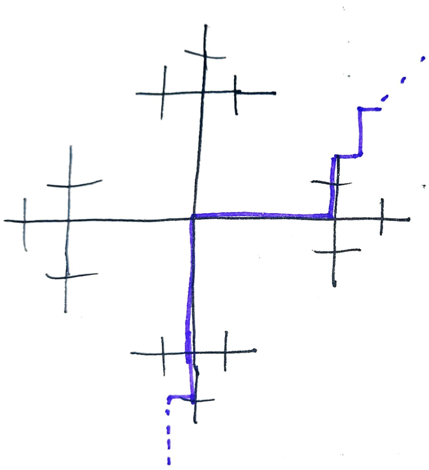

In particular, since w is locally injective, $\widetilde{w}$ is as well, and so it is injective, by the lemma. The image of $\widetilde{w}$ is called the *axis* of w.

By the definition of the action of F_n on T_n, w when thought of as a deck tranformation of T_n, preserves its axis. Note that w translates Axis(w) by $\ell(w)$. This is called the *translation length of* w, denoted as $\tau(w)$.

Lecture 9

A geometric solution to the conjugacy problem follows from:

Lemma 3.0.2. Let $u, v \in F_n$ be cyclically reduced. If u and v are conjugate, then there exists $g \in F_n$, such that

$$\ell(g) \leq \frac{1}{2}(\tau(u) + \tau(v))$$

and $u = gvg^{-1}$.

The conjugacy problem follows, as the lemma tells us that we only need to check $u = gvg^{-1}$ for finitely many g, and each of these can be checked using the word problem.

Remark 3.0.3. The statement is existence, it does not hold for all choice of g. In particular, $C(v)$ is infinite, as it contains v^k for all $k \in \mathbb{Z}$, and the length of gv^k is unbounded as $k \to \infty$.

In fact, the set of conjugators is the double coset $\langle u \rangle g \langle v \rangle$.

Proof. Suppose $u = gvg^{-1}$, with $\ell(g)$ minimal. Then

(i) If $u \in [1, g]$, then $g = uh$ for some h, and there is no cancellation. Moreover, $u = hvh^{-1}$, and if $h \neq g$, then $\ell(h) < \ell(g)$. Contradiction.

(ii) If $v \in [1, g^{-1}]$ is strictly between 1 and g as above, then $\ell(g)$ wasn't minimal.

Now consider the convex hull of $\{1, g, u, gv\}$.

For this, there are three (non-degenerate) different combinatorial types for the convex hull. The first case is

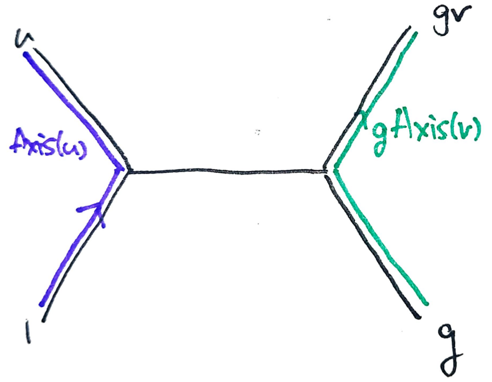

By the minimality in (i), $\ell(\alpha) > 0$. Similarly, $\ell(\beta) > 0$. On this diagram, we have Axis(u) and $g \cdot$ Axis(v) = Axis(gvg^{-1}). But $u = gvg^{-1}$. Contradiction (we will assume the middle length is non-zero for now).

The second case is

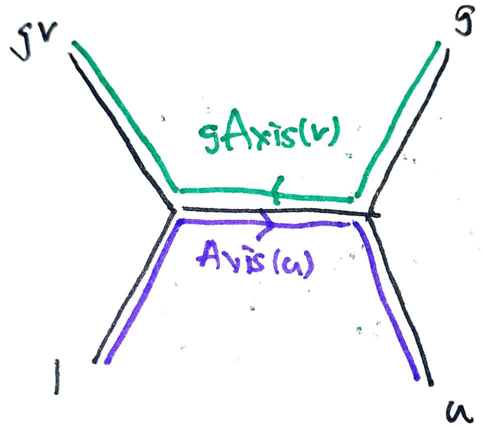

The axes are labelled. But they translate in opposite directions. Contradiction (again, we assume the middle length is non-zero).

The third case is:

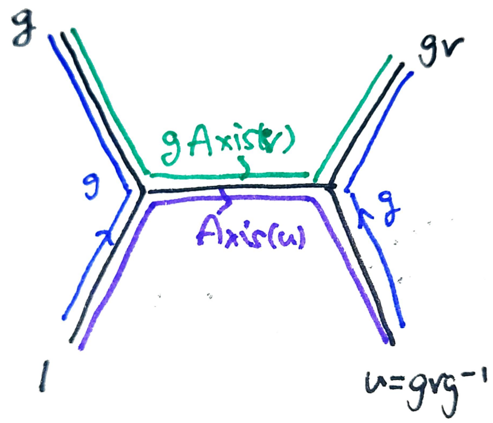

If the middle length is λ, then

$$\tau(u) + \tau(v) = 2\ell(g) + 2\lambda \geq 2\ell(g)$$

□

Subgroups of free groups

Proposition 3.0.4. If X is a (connected) graph, then $\pi_1(X)$ is free.

Proof when X has countably many cells. Let $T \subseteq X$ be a maximal tree, and let $\{e_1, e_2, \dots\}$ be the edges in X and not T. Let $X_N = T \cup \{e_1, \dots, e_N\}$. With this,

$$X = \bigcup_{n \geq 1} X_n$$

Pick a base vertex $v_0 \in T$. For each e_i, let α_i be the illustrated loop.

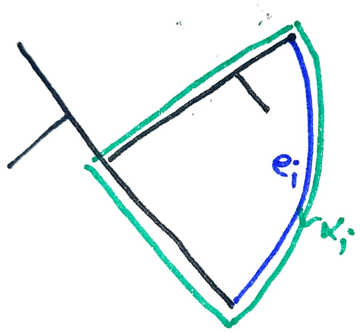

Lecture 10

Note

$$X_{n+1} = X_n \cup e_{n+1} = X_n \cup_{Y_{n+1}} (S^1 \cup I)$$

By Seifert-van Kampen,

$$\pi_1(X_{n+1}) = \pi_1(X_n) * \langle \alpha_{n+1} \rangle$$

Thus, by induction, $\pi_1(X_n)$ is free for all n, and generated by $\alpha_1, \dots, \alpha_n$. When X is countably infinite, note that every (edge) loop $\gamma \subseteq X$ is contained in X_n for some n. Thus, $\pi_1(X)$ is generated by $\{\alpha_1, \alpha_2, \dots\}$.

By the universal property of free groups, we have a surjection

$$\eta : F_\infty = \langle \alpha_1, \dots \rangle \to \pi_1(X)$$

Suppose γ is a loop representing an element of $\ker(\eta)$. As before, γ is contained in X_n for some n. So γ is in the kernel of the map

$$\langle \alpha_1, \dots, \alpha_n \rangle \to \pi_1(X_n) \to \pi_1(X)$$

The first map is an isomorphism, so $\gamma \in \ker(\pi_1(X_n) \to \pi_1(X))$.

Since X_n is a retract[2] of X, every loop which is null-homotopic in X, is null-homotopic in X_n. So $\gamma = 1$ in $\pi_1(X_n) = \langle \alpha_1, \dots, \alpha_n \rangle \leq F_\infty$. □

Corollary 3.0.5. If G acts on a tree T freely, then G is free.

Proof. The action of G on T is a *covering space action*. Since T is simply connected, $X = G \backslash T$ is a graph, with universal cover T, and $G = \pi_1(X)$ is free. □

Corollary 3.0.6 (Nielsen-Schreier)**.** Any subgroup of $H \leq F_n$ is free.

Proof. Let $T = T_n$ be the Cayley tree of F_n. Then F_n acts on T freely, and so H acts freely on T. By the previous corollary, H is free. □

[2] i.e. the inclusion $X_n \to X$ has a left inverse $r : X \to X_n$

Remark 3.0.7. The choice of generating set comes from the choice of a maximal tree in the proposition.

4 Bass-Serre theory

We will study groups acting on trees, not necessarily freely. We will also see how to glue groups together, or cut groups into pieces.

4.1 Amalgamated free products

Definition 4.1.1 (pushout)

A commutative diagram of groups

$$\begin{array}{ccc} C & \xrightarrow{i} & A \\ {\scriptstyle j}\big\downarrow & & \big\downarrow{\scriptstyle k} \\ B & \xrightarrow[\ell]{} & \Gamma \end{array}$$

is a *pushout* if for any group G, and homomorphisms $A \to G$, $B \to G$, there exists a unique homomorphism making the diagram

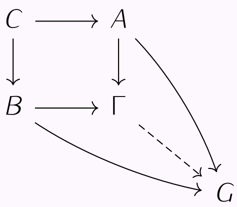

commute.

In this case, Γ is unique up to unique isomorphism, and therefore we may write $\Gamma = A \bigsqcup_C B$.

Theorem 4.1.2 (Seifert-van Kampen for cell complexes)**.** Suppose $K, L \subseteq X$ are subcomplexes, such that $X = K \cup L$. Suppose $K, L, K \cap L$ are all path connected. Then

$$\pi_1(X) = \pi_1(K) \bigsqcup_{\pi_1(K\cap L)} \pi_1(L)$$

Proof omitted. □

Note we use $\sqcup$ as it is a coproduct.

Proposition 4.1.3. Suppose $A = \langle S_A \mid R_A \rangle$, $B = \langle S_B \mid R_B \rangle$, $C = \langle \Sigma \mid \ldots \rangle$. Let i, j be represented by maps $\hat{i} : \Sigma \to F(S_A)$, $\hat{j} : \Sigma \to F(S_B)$. Then

$$A \bigsqcup_C B = \left\langle S_A, S_B \;\middle|\; R_A, R_B, \{\hat{i}(\sigma)\hat{j}(\sigma)^{-1} \mid \sigma \in \Sigma\} \right\rangle$$

Proof. Exercise. □

Example 4.1.4

If B is trivial, then

$$A \bigsqcup_C 1 = A / \langle\langle i(C) \rangle\rangle$$

Definition 4.1.5 ((amalgamated) free product)

If the maps i, j in the definition of a pushout are injective, then we write $\Gamma = A *_C B$, and call Γ the *amalgamated free product* of A and B over C.

In particular, if $C = 1$, we write $\Gamma = A * B$, and we call this the *free product* of A and B.

Theorem 4.1.6 (Britton's lemma). The *vertex group* A (or B) injects into $G = A *_C B$.

Remark 4.1.7. This is not true for pushouts. For example, $\mathbb{Z}/2 \coprod_{\mathbb{Z}} \mathbb{Z}/3 = 1$.

Lecture 11

To prove the theorem, we will construct a *graph of spaces* X, such that $G = \pi_1(X)$.

diagram

Let X_A be a presentation complex for A, and X_B be a presentation complex for B. As before, let Σ be a generating set for C. For each $\sigma \in \Sigma$, let α_σ be a based edge loop in X_A, representing $i(\sigma)$. Similarly, let β_σ be a based edge loop in X_B, representing $j(\sigma)$. To build this space:

1. Let X_A, X_B be the presentation complexes, with their based points.
2. Add in an edge t from the base point of X_A to the base point of X_B.
3. For each $\sigma \in \Sigma$, consider the following "rectangular" 2-cell **diagram** with gluing pattern $t\beta_\sigma^{-1}t^{-1}\alpha_\sigma$. Attach these to the diagram.

Call the resulting space X. By construction (and the Seifert-van Kampen theorem), $\pi_1(X) = G = A \coprod_C B$.

Proof. Suppose $g \in A$ mapsto $1 \in G = A *_C B$. Then g represented by a (based) loop γ in X_A, which is null-homotopic in X.

By van Kampen's lemma[3], γ bounds a singular disc diagram $D \to X$. Because the edge t appears in each rectabgle, and nowhere else, the rectangular 2-cells in D are arranged in strips, which we call t-corridors.

diagram

Since the boundary word is γ, which is contained in X_A. Therefore, we can't have any t on the boundary, so all of the t-corridors are annuli. Look at an *inner most disc* D_0 bounded by a t-corridor.

diagram

Since D_0 is contained in a t-corridor, it is contained in X_A or X_B. Without loss of generality (proof is symmetric), $D_0 \subseteq X_A$. Going around the t-corridor, we get a cyclic word δ in $\Sigma \cup \Sigma^{-1}$. In particular, $i(\delta)$ is the inner loop, $j(\delta)$ is the outer loop. But $i(\delta)$ bounds a disc D_0, and so it is contractible. So $i(\delta) = 1$. But i is injective, so $\delta = 1$. So $j(\delta) = 1$ in B.

By van Kampen's lemma, $j(\delta)$ has a van Kampen diagram $D_B \to X_B$. In particular, this has no t-corridors, and the same boundary as D_0 with its surrounding t-corridor. So we can remove D_0 and its surrounding t-corridor, and replace it with D_B.

diagram

This is now a van Kampen diagram, with one less t-corridor. Iterating, we can remove all of the t-corridors. But then we obtain a disc diagram Δ for γ with cells in X_A only. So $\Delta \to X_A$, and so $\gamma = 1$ in $\pi_1(X_A) = A$. $\square$

Example 4.1.8

For a closed orientable surface Σ, we can cut along a curve γ to get

$$\pi_1(\Sigma) = \pi_1(\Sigma_A) *_{\mathbb{Z}} \pi_1(\Sigma_B)$$

What happens if we cut along a non-separating curve?

Lecture 12

[3] Yes this isn't a presentation complex, it still applies.

4.2 Higman-Neumann-Neumann extensions

Definition 4.2.1 (HNN pushout)

Suppose $i, j : H \to G$ are group homomorphisms. The *HNN pushout* is the quotient

$$G\bigsqcup_{H} = \frac{G * \langle t \rangle}{\langle\langle ti(h)t^{-1}j(h) \mid h \in H \rangle\rangle}$$

The t is called the *stable letter*.

That is, we force $i(h)$ and $j(h)$ to be conjugate for all $h \in H$.

Theorem 4.2.2 (Seifert-van Kampen for non-separating decompositions)**.** Suppose Y is a connected cell complex, and $i, j : Z \hookrightarrow Y$ are two inclusion maps, with disjoint image. Define

$$X = Y\bigcup_{Z} = \frac{Y}{i(z) \sim j(z)}$$

for the result of gluing Y to itself by identifying $i(Z)$ with $j(Z)$. Then

$$\pi_1(X) \cong \pi_1(Y) \bigsqcup_{\pi_1(Z)}$$

Proof. Deferred. □

Remark 4.2.3. Suppose G has presentation $\langle a_1, \dots, a_m, t \mid r_1, \dots, r_n, p_1tq_1t^{-1}, \dots, p_\ell tq_\ell t^{-1} \rangle$, where the r_i do not involve t. Define $A = \langle a_1, \dots, a_m \mid r_1, \dots, r_n \rangle$, and define maps $i, j : F_\ell \to A$ by $i(x_k) = p_k$ and $j(x_k) = q_k$, then

$$G = A\bigsqcup_{F_\ell}$$

Definition 4.2.4 (HNN extension)

If $G = A\bigsqcup_{B}$, and the maps $B \to A$ are injective, then G is called an *HNN extension*, and we write $G = A *_B$.

Example 4.2.5

Consider $\pi_1(T^2) = \mathbb{Z}^2 = \langle a, b \mid aba^{-1}b^{-1} \rangle$.

diagram

Cut along the non-separating curve a, we get a cylinder.

diagram

The π_1 of the cylinder is $\mathbb{Z} = \langle a, c = ba'b^{-1} \mid ac^{-1} \rangle$. Consider the maps $i, j : \mathbb{Z} = \langle z \rangle \to \mathbb{Z}$, given by $i(z) = a$ and $j(z) = c$. The resulting HNN extension has presentation

$$\langle a, c, t \mid ac^{-1}, tat^{-1}c^{-1} \rangle \cong \langle a, t \mid tat^{-1}a^{-1} \rangle \cong \mathbb{Z}^2$$

Example 4.2.6

Now consider Σ_2, surface of genus 2. Here

$$\pi_1(\Sigma_2) = \langle a_1, b_1, a_2, b_2 \mid a_1b_1a_1^{-1}b_1^{-1}a_2b_2a_2^{-1}b_2^{-1} \rangle$$

We can cut along the non-separating curve a_1, and the stable letter is b_1. So we have an HNN extension.

Example 4.2.7 (Baumslag-Solitar groups)
Define
$$\mathrm{BS}(m,n) = \langle a, b \mid ba^m b^{-1} a^n \rangle$$
where $m, n \in \mathbb{Z}$ are not zero. Notice these are HNN extensions of $\mathbb{Z}$ over $\mathbb{Z}$, where we conjugate $m\mathbb{Z}$ with $n\mathbb{Z}$.

Theorem 4.2.8 (Britton's lemma for HNN extensions)**.** The vertex group A embeds into $A \underset{C}{*}$.

Proof. The same proof as for $A \underset{C}{*} B$ works. Build a graph of spaces, and apply the method of t-corridors. □

4.2.1 Sample applications of HNN extensions

- there exists an infinite group with exactly two conjugacy classes,
- there exists a *non-Hopfian* finitely presented group. That is, there exists a map $f : G \twoheadrightarrow G$ with $\ker(f) \neq 1$. In fact, $G = \mathrm{BS}(2,3)$ works,
- there exists an infinite finitely generated simple group [Higman],
- every countable group embeds into a group with two generators [HNN],
- there exists a group with an unsolvable word problem.

What about cutting surfaces along *multi-curves*? For example,
diagram

4.3 Graph of groups

For example, with the above decomposition, we have the graph
diagram

Lecture 13

First, we should *carefully* define directed (or oriented) graphs.

Definition 4.3.1 (oriented graph)
An *(oriented) graph* Γ consists of a pair of sets $V = V_\Gamma, E = E_\Gamma$. V is the set of vertices, and E is the set of egdes. We have two maps
$$\iota = \iota_\Gamma, \tau = \tau_\Gamma : E \to V$$
We call ι the *origin map*, and τ the *terminus map*.
The *realisation* of Γ is $|\Gamma|$, the 1-dimensional cell complex given by the above data.

Often we will abuse notation and not distinguish between Γ and $|\Gamma|$.

Example 4.3.2
For example, we have
diagram

Definition 4.3.3
A *graph of groups* $\mathcal{G}$ consists of:

- a graph Γ,

- assignments

$$\begin{aligned} V &\to \text{Groups} \\ v &\mapsto G_v \end{aligned}$$

  and

$$\begin{aligned} E &\to \text{Groups} \\ e &\mapsto G_e \end{aligned}$$

- injective homomorphisms

$$\iota_e : G_e \to G_{\iota(e)} \quad \text{and} \quad \tau_e : G_e \to G_{\tau(e)}$$

Example 4.3.4

Continuing with the example as above,

$$\begin{aligned} G_u &= \pi_1(\Sigma_1) \\ G_v &= \pi_1(\Sigma_2) \\ G_w &= \pi_1(\Sigma_3) \end{aligned}$$

The maps are given by the inclusions of $\pi_1(S^1) \hookrightarrow \pi_1(\Sigma_1)$.

Definition 4.3.5

Let $\mathcal{G}$ be a graph of groups, with connected underlying graph Γ. Let $T \subseteq \Gamma$ be a spanning tree. The *fundamental group of* $\mathcal{G}$ *with respect to* T, $\pi_1(\mathcal{G}, T)$ is defined as follows:

$$\frac{\left(\mathop{\ast}_{v \in V} G_v\right) * F(E_\Gamma)}{\langle\langle \{t_e \iota_e(h) t_e^{-1} \tau_e(h)^{-1} \mid e \in E, h \in G_e\} \cup \{t_e \mid e \in T\} \rangle\rangle}$$

where $F(E_\Gamma) = \langle t_e \mid e \in E \rangle$.

Example 4.3.6

diagram

In this case, the spanning tree is e, and

$$\pi_1(\mathcal{G}, T) = G_u \mathop{\ast}_{G_e} G_v$$

Now if we have

diagram

Then

$$\pi_1(\mathcal{G}, T) = G_u \mathop{\ast}_{G_e}$$

Theorem 4.3.7 (Seifert-van Kampen for graphs of groups). Let Γ be a graph. For each vertex $v \in V$, $e \in E$, let X_v, X_e be connected cell complexes, and let $\iota_e : X_e \to X_{\iota(e)}, \tau_e : X_e \to X_{\tau(e)}$ be inclusions of subcomplexes, or equivalently, injective cellular maps. Moreover, assume that the maps induce injections on π_1.

Let

$$X = \frac{\bigsqcup_{v \in V} X_v}{\iota_e(x) \sim \tau_e(x)}$$

Setting $G_v = \pi_1(X_v)$, $G_e = \pi_1(X_e)$ and so on, defines a graph of groups $\mathcal{G}$. Then

$$\pi_1(X) \cong \pi_1(\mathcal{G}, T)$$

for any spanning tree T.

Proof idea when Γ *is finite.* Induct on the number of edges of Γ, and the two Seifert-van Kampen theorems we have seen. □

Remark 4.3.8. It follows (for example by taking the spaces to be presentation complexes), that $\pi_1(\mathcal{G}, T)$ does not depend, up to isomorphism, on T. Thus, we will write $\pi_1(\mathcal{G})$.

4.3.1 Quotients

Suppose G acts on a tree T (or any graph). That is, G acts on V_T and on E_T, so that

$$\iota(g \cdot \tilde{e}) = g \cdot \iota(\tilde{e}) \quad \text{and} \quad \tau(g \cdot \tilde{e}) = g \cdot \tau(\tilde{e})$$

There is a natural quotient graph $\Gamma = G\backslash T$. In this case,

$$\begin{aligned} V_\Gamma &= G\backslash V_T \\ E_\Gamma &= G\backslash E_T \\ \iota_\Gamma(G \cdot \tilde{e}) &= G \cdot \iota(\tilde{e}) \\ \tau_\Gamma(G \cdot \tilde{e}) &= G \cdot \tau(\tilde{e}) \end{aligned}$$

Furthermore, Γ is naturally a graph of groups. Let $v = G\tilde{v} \in V_\Gamma$. Set $G_v = \mathrm{Stab}_G(\tilde{v})$. This is well defined, up to conjugation in G. Similarly, if $e = G \cdot \tilde{e}$, then $G_e = \mathrm{Stab}_G(\tilde{e})$.

Suppose $\iota(e) = v$. So $G \cdot \iota(\tilde{e}) = G\tilde{v}$. So we may choose $\tilde{e}$, such that $\iota(\tilde{e}) = \tilde{v}$. Now $G_e = \mathrm{Stab}_G(\tilde{e}) \subseteq \mathrm{Stab}_G(\tilde{v}) = G_v$. So the map is the inclusion map, which is injective.

Let ι_e be the inclusion homomorphism $G_e \to G_v$.

Remark 4.3.9. ι_e is well defined, up to conjugation in G_v.

Define τ_e similarly.

Example 4.3.10

Let $\mathbb{Z} = \langle t \rangle$ act on $\mathbb{R}$, considered as a graph

diagram

and t is translation by 1. The quotient is $\mathbb{Z}\backslash\mathbb{R} = S^1$. The associated graph of groups is

diagram

So $\mathbb{Z}$ is an HNN extension of 1 by itself.

Example 4.3.11

Let $D_\infty = \langle s, t \mid s^2, t^2 \rangle$ act on $\mathbb{R}$. The graph is the same as the above. s acts by reflection in 0, and t acts by reflection in 1. In this case, $D_\infty\backslash\mathbb{R}$ is the graph

diagram

and we have an associated graph of groups

diagram

So $D_\infty = (\mathbb{Z}/2\mathbb{Z}) * (\mathbb{Z}/2\mathbb{Z})$.

Lecture 14

4.4 Bass-Serre tree

The main theorem of the subject is due to Serre, although we adopt a topological approach, due to Scott and Wall.

Theorem 4.4.1 (Serre, the fundamental theorem of Bass-Serre theory)**.** Let $\mathcal{G}$ be a graph of groups, with connected underlying graph Γ. Let $G = \pi_1(\mathcal{G})$. Then G acts on a tree T, such that

$$\mathcal{G} \cong G \backslash T$$

T is called the *Bass-Serre tree* of $\mathcal{G}$.

Remark 4.4.2. Letting G act on a tree T is equivalent to cutting G into pieces. The theorem says that $\mathcal{G}$ has a "universal cover" T, on which $G = \pi_1(\mathcal{G})$ acts, and we recover $\mathcal{G}$ as the quotient.

Sketch proof. Using presentation complexes, build a "graph of spaces" $\mathcal{X}$ corresponding to $\mathcal{G}$.

diagram

For each $\bullet$, let $X_\bullet$ be a presentation complex for $G_\bullet$. Then build X as follows

diagram

For each "edge space", take a product with the interval $[-1, 1]$. We can use the homomorphism of groups to glue the ends of the cylinder to the appropriate vertex spaces. This is the data for $\mathcal{X}$, and X is the resulting space.

Let $\widetilde{X}$ be the universal cover of X. It looks something like

diagram

The result is a graph of spaces $\widetilde{\mathcal{X}}$, where each vertex space $\widetilde{X}_{\widetilde{v}}$ is the universal cover of some X_v, and so on. The edge space is $[-1, 1] \times \widetilde{X}_{\widetilde{e}}$, where $\widetilde{X}_{\widetilde{e}}$ is the universal cover of X_e. Let $\widetilde{\Gamma}$ be the underlying graph of $\widetilde{\mathcal{X}}$. Now note that $\widetilde{X}$ retracts onto $\widetilde{\Gamma}$, by crushing all of the edge and vertex spaces to their base points. That is, we have maps

$$\iota : \widetilde{\Gamma} \hookrightarrow \widetilde{X} \quad \text{and} \quad r : \widetilde{X} \to \widetilde{\Gamma}$$

such that $r \circ \iota \simeq \mathrm{id}$. So $\iota_* : \pi_1(\widetilde{\Gamma}) \to \pi_1(\widetilde{X})$ is injective. But $\widetilde{X}$ is a universal cover, so simply connected. Hence $\pi_1(\widetilde{\Gamma})$ is simply connected. But a simply connected graph is a tree, so $\widetilde{\Gamma}$ is a tree.

Set $T = \widetilde{\Gamma}$. □

Proposition 4.4.3. Let G act on T with quotient $\mathcal{G}$. Then

(i) there exists a G-equivariant bijection

$$V_T \leftrightarrow \bigsqcup_{v \in V_\Gamma} G/G_v$$

(ii) there exists a G-equivariant bijection

$$E_T \leftrightarrow \bigsqcup_{e \in E_\Gamma} G/G_e$$

(iii) for any $\widetilde{v} \in V_T$, mapping to $v \in V_\Gamma$, the set of edges of T incident at $\widetilde{v}$ is G-equivariantly bijective with

$$\left(\bigsqcup_{\iota(e)=v} G_v/\iota_e(G_e) \right) \sqcup \left(\bigsqcup_{\tau(e)=v} G_v/\tau_e(G_e) \right)$$

Proof. For (i), choose orbit representatives $\widetilde{v} \in G \cdot \widetilde{v} = v \in V_\Gamma$. Orbit stabiliser says that the map $G \to G \cdot \widetilde{v}$ defines a G-equivariant bijection $G/G_v \to G \cdot \widetilde{v}$.

For (ii), let G act on the set of edges. For (iii), let $\mathrm{Stab}_G(\widetilde{v})$ act on the set of incident edges. □

Remark 4.4.4. In particular, T is determined by the algebraic data of $\mathcal{G}$, and so it is unique.

Example 4.4.5

For

diagram

we have Bass-Serre tree

diagram

Example 4.4.6

For

diagram

we have Bass-Serre tree

diagram

Example 4.4.7

Here, $F_2 = \pi_1(\mathcal{G}) = \mathbb{Z} \underset{1}{*} \mathbb{Z}$, and the graph of groups is

diagram

The Bass-Serre tree is

diagram

which is the tree with countably infinite valence at each vertex.

Example 4.4.8

On the other hand, we have another graph of groups

diagram

with Bass-Serre tree

diagram

with is the usual Cayley tree.

Lecture 15

How do stable letters $t_e \in \pi_1(\mathcal{G})$ act on T? Choose a maximal tree M in Γ. The action of G in T also depends on a choice of lift $\widetilde{M} \subseteq T$, where we lift by the quotient map $T \to \Gamma$.

For example, when D_∞ acts on $\mathbb{R}$, the Bass-Serre tree is

diagram

and the lift of a maximal tree is

diagram

The choice of $\widetilde{M}$ determine choices of lifts of vertices $\widetilde{v} \in T$ mapping to $v \in \Gamma$. For each edge $e \in E_\Gamma$ not contained in M, choose a lift $\widetilde{e}$ such that $\iota(\widetilde{e}) = \widetilde{\iota(e)}$. The action of t_e on T is determined by the fact that:

$$t_e \widetilde{\tau(e)} = \tau(\widetilde{e})$$

Most importantly, we can understand elements of $G = \pi_1(\mathcal{G}, M)$ via reduced words.

Definition 4.4.9 (loop)

Fix a base vertex $v_0 \in V_\Gamma$. Consider an element

$$w = g_0 t_1^{\pm 1} \cdots g_{k=1} t_k^{\pm 1} g_k \in \left(\mathop{\ast}_{v \in V_\Gamma} G_v \right) * F(E_\Gamma)$$

where $g_i \in G_{v_i}$, and $t_i = t_{e_i}$ is the corresponding stable letter. Then w is a *(based) loop* if:

(i) $v_0 = v_k$, which is also the base vertex we fixed at the start of the definition.

(ii) the path $e_1^{\pm}1 \cdots e_k^{\pm}1$ is a loop in Γ based at v_0,

(iii) "if it goes" $t_i g_i$, then $v_i = \tau(e_i)$. On the other hand, "if it goes" $t_i^{-1} g_i$, then $v_i = \iota(e_i)$.

Recall the relations in $\pi_1(\mathcal{G})$ say that

$$t_e \iota_e(G_e) t_e^{-1} = \tau_e(G_e)$$

Definition 4.4.10 (pinch)
A sub-path of a loop is called a *pinch* if it is of the form:

(i) $t_e \iota_e(h) t_e^{-1}$ for $h \in G_e$, or

(ii) $t_e^{-1} \tau_e(h) t_e$ for $h \in G_e$.

Remark 4.4.11. Loops should be thought of as defining paths in the Bass-Serre tree.

A pinch corresponds to when the path double backs on itself. A based loop without pinches is called *reduced*.

Theorem 4.4.12 (normal form for graphs of groups)**.** Let $\mathcal{G}$ be a graph of groups. Then

(i) every element $g \in \pi_1(\mathcal{G})$ is represented by a based loop γ,

(ii) if γ is reduced, then g is non-trivial.

Remark 4.4.13 (about the proof)**.** (i) The unique path $[\tilde{v}_0, g\tilde{v}_0]$ defines a loop representing g,

(ii) reduced loops correspond to locally injective paths in T, which are globally injective. Hence $g\tilde{v}_0 \neq \tilde{v}_0$.

5 Property FA

Suppose G acts on a tree. A *global fixed point* $p \in T$ for G is a point $x_0 \in T$ such that $\text{Stab}(x_0) = G$. We say G acts *trivially* on T if there is a global fixed point.

Example 5.0.1
Let $\mathbb{Z}$ act on the tree T
diagram
The central point is a global fixed point. The quotient is
diagram

If G acts on some tree non-trivially, then we say that G splits. Otherwise, we say that G has *property FA*. Here is a result from examples sheet 2:

Lemma 5.0.2. If ϕ is an isometry of a tree T, then either:

(i) ϕ fixes a point, or

(ii) ϕ translates a line a positive distance.

In (i), ϕ is *elliptic*, and in (ii), ϕ is *hyperbolic*.

Remark 5.0.3. If the order of ϕ is finite, then ϕ is elliptic.

Lemma 5.0.4. Suppose $\phi, \psi \in \text{Isom}(T)$ are both elliptic, $\text{Fix}(\phi) \cap \text{Fix}(\psi) = \varnothing$, then $\phi \circ \psi$ is hyperbolic.

Proof. Note that $\text{Fix}(\phi)$ and $\text{Fix}(\psi)$ are connected subtrees of T. Let $[x, y]$ be the unique path from $\text{Fix}(\phi)$ to $\text{Fix}(\psi)$.

Let $I = [x, y] \cup [\psi^{-1}x, \psi^{-1}y]$. Note $\psi^{-1}[x, y]$ is the path from $\psi^{-1}\,\text{Fix}(\phi)$ to $\text{Fix}(\psi)$.

diagram

Now note that $I \cap \phi\psi I = \{x\}$, and so repeating this, we have a line

$$\bigcup_{n \in \mathbb{Z}} (\phi\psi)^n I$$

which is preserved by $\phi\psi$. In fact, the line is translated by $2d(x, y)$. Thus, $\phi\psi$ is hyperbolic. □

Lecture 16

Next, we need a version of the Helly property.

Lemma 5.0.5 (Helly property for trees). Suppose T is a tree, $T_1, \dots, T_n$ are subtrees. If $T_i \cap T_j \neq \varnothing$ for every i, j. Then

$$\bigcap_{i=1}^{n} T_i \neq \varnothing$$

Proof. We induct on n. $n = 1, 2$ are trivial. Let $T' = T_{n-1} \cap T_n$.

Claim 5.0.6. $T' \cap T_i \neq \varnothing$ for all $i < n - 1$.

Once we show the claim, we are done by induction.

Proof of claim. Suppose not.

diagram

Then we get a non-trivial cycle in T. Contradiction. □

□

Theorem 5.0.7 (criterion for FA). Let G be a group, and suppose $S = \{s_1, \dots, s_n\}$ is a generating set. If

(i) s_i has finite order for all i,

(ii) for all i, j, either $s_i s_j$ or $s_j s_i$ has finite order.

Then G has property FA.

Proof. Suppose G acts on a tree T. Let $T_i = \text{Fix}(s_i)$. Since s_i has finite order, T_i is non-empty. Since at least one of $s_i s_j, s_j s_i$ has finite order, $T_i \cap T_j$ is non-empty for all i, j. Hence by the Helly property,

$$\bigcap_{i=1}^{n} T_i \neq \varnothing$$

But this is the set of global fixed points of G. □

Example 5.0.8

Let Γ be the group generated by the reflections in the sides of an equilateral triangle, say reflections r_ℓ, r_m, r_n, where $r_\bullet$ is reflection in the line $\bullet$.

So $\Gamma = \langle r_\ell, r_m, r_n \rangle \leq \text{Isom}(\mathbb{R}^2)$. Note that $r_\ell^2, r_m^2, r_n^2 = 1$. Composition of two reflection is a rotation of order 3. So Γ has property FA (but it is infinite).

On sheet 3, Dehn's examples also have property FA. The corresponding 3-manifolds are "non-Haken".

6 Fuchsian groups

6.1 Hyperbolic geometry

Let $\mathbb{H}^2$ denote the hyperbolic plane. Recall we have the disc model and the upper half plane model, both contained in $\mathbb{C}$.

diagram

which have metrics

$$\frac{4|dz|^2}{(1-|z|^2)^2} \text{ and } \frac{|dz|^2}{|\text{Im}(z)|^2}$$

respectively. The geodesics in $\mathbb{H}^2$ (with both models) are lines, or arcs of circles which intersect the boundary orthogonally.

We will write $\ell^+ = \{iy \mid y > 0\}$ in the upper half plane. In this case, if $s > t$, then

$$d(is, it) = \int_t^s \frac{dy}{y} = \log\left(\frac{s}{t}\right)$$

One more useful fact is a special case of the Gauss-Bonnet theorem.

Proposition 6.1.1 (Gauss-Bonnet for triangles)**.** if $\Delta \subseteq \mathbb{H}^2$ is a geodesic triangle, with interior angles α, β, γ, then

$$\text{Area}(\Delta) = \pi - (\alpha + \beta + \gamma)$$

In particular, $\alpha + \beta + \gamma < \pi$.

Corollary 6.1.2. If $P \subseteq \mathbb{H}^2$ is a geodesic n-gon, with interior angles α_i, then

$$\text{Area}(P) = (n-2)\pi - \sum_i \alpha_i$$

Recall that $\text{Isom}^+(\mathbb{H}^2) \cong \text{PSL}(2, \mathbb{R})$, acting on the upper half plane model by Möbius transformations.

Definition 6.1.3 (Fuchsian group)

If $\Gamma \leq \text{PSL}(2, \mathbb{R})$ is a subgroup which acts properly discontinuously on $\mathbb{H}^2$, then Γ is called a *Fuchsian group*.

We can also think of them as the discrete subgroups of $\text{PSL}(2, \mathbb{R})$. Some basic facts of $\text{PSL}(2, \mathbb{R})$:

Proposition 6.1.4.

(i) The action of $\text{PSL}(2, \mathbb{R})$ on $\mathbb{H}$ extends continuously to $\overline{\mathbb{H}}^2$, which is $\mathbb{H}^2 \cup \partial\mathbb{H}^2$.

(ii) $\text{PSL}(2, \mathbb{R})$ is transitive on triples of distinct points on $\mathbb{R} \cup \{\infty\}$,

(iii) if $\phi \in \text{PSL}(2, \mathbb{R})$ and fixes any three distinct points in $\overline{\mathbb{H}}^2$, then $\phi = \text{id}$.

Corollary 6.1.5 (classification of (orientation preserving) isometries of $\mathbb{H}^2$)**.** Suppose $\phi \in \text{Isom}^+(\mathbb{H}^2)$. Then one of the following holds:

(i) ϕ fixes a point in $\mathbb{H}^2$, which is unique unless $\phi = \text{id}$.

(ii) ϕ fixes a unique point in $\partial\mathbb{H}^2$,

(iii) ϕ preserves a unique geodesic in $\mathbb{H}^2$, which it translates a positive distance.

In (i), ϕ is elliptic, in (ii), ϕ is parabolic, and in (iii), ϕ is hyperbolic.

Remark 6.1.6. If Γ is a Fuchsian group, ϕ is elliptic, then ϕ must have finite order.

Lecture 17

Proof. Recall that $\phi : \mathbb{H}^2 \to \mathbb{H}^2$ extends continuously to a homeomorphism $\phi : \overline{\mathbb{H}}^2 \to \overline{\mathbb{H}}^2$. By Brouwer's fixed point theorem, $\mathrm{Fix}(\phi) \subseteq \overline{\mathbb{H}}^2$ is non-empty. We saw that if ϕ has at least three fixed points, then $\phi = \mathrm{id}$, so we can consider it case by case.

1. $\mathrm{Fix}(\phi) = \{\xi\} \subseteq \overline{\mathbb{H}}^2$.
   (a) If $\xi \in \mathbb{H}^2$, then ϕ is elliptic.
   (b) If $\xi \in \partial\mathbb{H}^2$, then ϕ is parabolic.
2. $\mathrm{Fix}(\phi) = \{\xi^+, \xi^-\}$.
   (a) If $\xi^+ \in \mathbb{H}^2$ (without loss of generality), we get a unique geodesic from ξ^+ to ξ^-. ϕ preserves the geodesic $[\xi^+, \xi^-)$. But then there is (at least) thre fixed points. So $\phi = \mathrm{id}$.
   (b) If $\xi^+, \xi^- \in \partial\mathbb{H}^2$, then again we have a uique geodesic from ξ^+ to ξ^-, and ϕ preserves it. Since ϕ has two fixed points, ϕ must act on the geodesic by a translation by a positive distance.

□

When ϕ is hyperbolic, we call the geodesic it preserves its *axis*.

6.2 Examples of Fuchsian groups

Recall $\Gamma \leq \mathrm{Isom}^+(\mathbb{H}^2)$ is *Fuchsian* if the action of Γ on $\mathbb{H}^2$ is properly discontinuous. In particular, for all $x \in \mathbb{H}^2$, $\mathrm{Stab}_\Gamma(x)$ is finite.

Lets start with some easy examples.

Example 6.2.1

Consider the disc model. The metric is radially symmetric, and so all rotations about 0 are isometries. In particular,

$$z \mapsto e^{2\pi i/n} z$$

is an isometry, generates $\mathbb{Z}/n\mathbb{Z} \leq \mathrm{Isom}^(\mathbb{H}^2)$.

In fact, any elliptic isometry is conjugate to this one.

Example 6.2.2

Now consider the upper half plane model. Consider the map $z \mapsto \lambda z$, for any $\lambda \in \mathbb{R}_{>1}$. This is an element of $\mathrm{Isom}^+(\mathbb{H}^2)$. The axis is ℓ^+. This gives $\mathbb{Z} \cong \langle \phi \rangle \leq \mathrm{Isom}^+(\mathbb{H}^2)$.

In fact, any elliptic isometry is conjugate to this one.

Example 6.2.3

Define $\psi(z) = z + 1$. This is an isometry of $\mathbb{H}^2$. This gives a parabolic isometry, where the fixed point is ∞. This gives $\mathbb{Z} \cong \langle \psi \rangle \leq \mathrm{Isom}^+(\mathbb{H}^2)$.

In fact, any parabolic isometry is conjugate to this one.

These examples are called *elementary*. There's one more elementary example

Example 6.2.4

Consider upper half plane. Let s_1 be rotation by π about i, and s_2 be rotation by π about $2i$. Then we get

$$\langle s_1, s_2 \rangle \cong D_\infty$$

Example 6.2.5

Let Σ_g be a closed orientable surface of genus g, with $g \geq 2$. In this case, $\widetilde{\Sigma}_g$ is isometric to $\mathbb{H}^2$. Then $\pi_1(\widetilde{\Sigma}_g)$ is Fuchsian.

Definition 6.2.6

Let $p, q, r \in \mathbb{Z}_{\geq 1}$. The (p, q, r)-*triangle group* is defined by the presentation

$$\Gamma(p, q, r) = \langle a, b, c \mid a^p, b^q, c^r, abc \rangle = \langle a, b \mid a^p, b^q, (ab)^{-r} \rangle$$

From our criterion for FA, $\Gamma(p, q, r)$ has property FA. Thus, it does not split, and so we can't use the techniques we have developed so far.

Is $\Gamma(p, q, r)$ non-trivial? infinite? and so on?

Example 6.2.7

$\Gamma(2, 3, 1) = 1$.

However, many interesting examples arise from Poincaré's polygon theorem.

Theorem 6.2.8 (Poincaré's polygon theorem). If $p^{-1} + q^{-1} + r^{-1} < 1$, then $\Gamma(p, q, r)$ is an infinite Fuchsian group.

Remark 6.2.9. The converse is morally true. That is, the other cases are all finite or non-Fuchsian.

Proof. We start with a geodesic triangle $\Delta \subseteq \mathbb{H}^2$ with interior angles $\pi/p, \pi/q, \pi/r$.

Lecture 18

diagram

Let α denote rotation about u, with angle $2\pi/p$; β about v, with angle $2\pi/q$ and γ about w, with angle $2\pi/r$. Note all of these are *anticlockwise*.

Let $G = \langle \alpha, \beta, \gamma \rangle \leq \mathrm{Isom}^+(\mathbb{H}^2)$. Clearly $\alpha^p = \beta^q = \gamma^r = 1$. Next, we show $\alpha\beta\gamma = 1$.

diagram

We see that $\beta(w) = \alpha^{-1}(w) = w'$. Hence $\alpha\beta\gamma(w) = \alpha\beta(w) = w$. Similarly, $\gamma(u) = \beta^{-1}(u)$, and so $\alpha\beta\gamma(u) = \alpha\beta\beta^{-1}(u) = \alpha(u) = u$. Hence by the classification of orientation preserving isometries of $\mathbb{H}^2$, it fixes two distinct points in $\mathbb{H}^2$ and so it is trivial.

Hence we have a surjective homomorphism $f : \Gamma(p, q, r) \twoheadrightarrow G$, sending a to α and so on. We will show that f is an isomorphism. Let r_ℓ denote reflection in the line ℓ, and $Q = \Delta \cup r_\ell(\Delta)$.

diagram

Define

$$\widetilde{Q} = \frac{\Gamma \times Q}{\sim}$$

where $\sim$ is the relation given by $(gc, x) \sim (g, c(x))$ for $x \in m$, and $(gb, y) \sim (g, b(y))$ for $y \in n'$. Next, define the *development map*

$$F : \widetilde{Q} \to \mathbb{H}^2$$
$$F(g, x) = f(g)x$$

Note $\widetilde{Q}$ is a complete geodesic metric space, via the path metric, and F is a local isometry, sending (sufficiently small) open balls in $\widetilde{Q}$ isometrically to small open balls in $\mathbb{H}^2$. In fact, F is an isometric embedding.

Indeed, if $x, y \in \widetilde{Q}$, and $[x, y]$ is a geodesic, then $F([x, y])$ is a local geodesic[4] from $F(x)$ to $F(y)$. But local geodesics in $\mathbb{H}^2$ are global geodesics. So $d(F(x), F(y)) = d(x, y)$. Next, we prove that F is surjective. $\text{im}(F)$ is open, since it sends small open balls to small open balls. On the other hand, $\widetilde{Q}$ is complete, and hence so is $\text{im}(F)$. But complete subsets of a metric space are closed, and so $\text{im}(F)$ is closed. Thus, by connectedness, F is an isometry.

So $\widetilde{Q}$ is isometric to $\mathbb{H}^2$, and the action of Γ on $\widetilde{Q}$ is properly discontinuous by construction, so Γ is Fuchsian. Since Q is compact, and F us surjective, Γ must be infinite. □

> **Remark 6.2.10.** It follows from the construction of $\widetilde{Q}$ that only $\Gamma \cdot u, \Gamma \cdot v, \Gamma \cdot w$ has non-trivial stabiliser. Moreover, $\text{Stab}(u) = \langle a \rangle$, $\text{Stab}(v) = \langle b \rangle$ and $\text{Stab}(w) = \langle c \rangle$.
>
> Q is called a *fundamental domain* for the action of Γ on $\mathbb{H}^2$.

6.3 Centres and Dehn's examples

> **Lemma 6.3.1.** Suppose $1/p + 1/q + 1/r < 1$. If $g \in \Gamma(p, q, r)$, and the order of g is finite, then g is in the conjugate of one of $\langle a \rangle, \langle b \rangle, \langle c \rangle$.

Proof. We saw that finite order elements of $\mathbb{H}^2$ fix a point in $\mathbb{H}^2$. If $g \neq 1$, then the fixed point z must be in the orbit of one of u, v, w. Say (without loss of generality) $z = hu$. So $ghu = hu$, and so $h^{-1}gh \in \text{Stab}(u) = \langle a \rangle$. □

> **Proposition 6.3.2.** If Γ is a non-elementary Fuchsian group, then $Z(\Gamma) = 1$.

Proof. Suppose $\gamma \in Z(\Gamma) \setminus 1$. Consider $\text{Fix}(\gamma) \subseteq \overline{\mathbb{H}}^2$. Note that for $g \in \Gamma, x \in \text{Fix}(\gamma)$, $gx = g\gamma x = \gamma g x$, and so $gx \in \text{Fix}(\gamma)$.

Lecture 19

Now we need to do some case analysis:

- if γ is elliptic, then $\text{Fix}(\gamma) = \{x\} \subseteq \mathbb{H}^2$. Without loss of generality, $x = 0$ in the disc model $\mathbb{D} \subseteq \mathbb{C}$. From this,
$$\text{Stab}_{\text{Isom}^+(\mathbb{H}^2)}(0) = \{z \mapsto e^{i\theta} z\}$$
By proper discontinuity, Γ is a subgroup of the above, and so it is a finite cyclic group.
- if γ is parabolic, then without loss of generality $\text{Fix}(\gamma) = \{\infty\}$ in the upper half plane model. A direct computation shows that
$$\text{Stab}_{\text{Isom}^+(\mathbb{H}^2)}(\infty) = \{z \mapsto az + b\}$$
For γ to be the only fixed point, necessarily $a = 1$, and so $\gamma(z) = z + c$ for some $c \in \mathbb{R}$ non-zero. But g commutes with γ only if $a = 1$, and so
$$\Gamma \leq \{z \mapsto z + b \mid b \in \mathbb{R}\}$$
Any discrete subgroup of $\mathbb{R}$ is isomorphic to $\mathbb{Z}$.
- if γ is hyperbolic, without loss of generlity $\text{Fix}(\gamma) = \{0, \infty\}$ in the upper half plane model. So Γ acts by isometries $\text{Axis}(\gamma) = \ell^+$, and so $\Gamma \cong \mathbb{Z}$ or D_∞ by proper discontinuity.

□

We can now analyse Dehn's examples. Recall

$$G_n = \left\langle x, y, z \mid x^2 = y^3 = z, (xy)^{6n+5} = z^{5n+4} \right\rangle$$

[4] i.e. locally it is a geodesic.

for $n \geq 0$. Note that $z \in Z(G_n)$. Let $\Gamma_n = \frac{G_n}{\langle z \rangle} = \left\langle x, y \mid x^2, y^3, (xy)^{6n+5} \right\rangle = \Gamma(2, 3, 6n+5)$. Therefore, Γ_n is a Fuchsian triangle group if $n \geq 1$, and so $Z(\Gamma_n) = 1$. Hence $Z(G_n) = \langle z \rangle$. Therefore, if $\phi : G_m \to G_n$ is an isomorphism, then $\phi(Z(G_m)) = Z(G_n)$, and so

$$\Gamma_m = \frac{G_m}{Z(G_m)} \cong \frac{G_n}{Z(G_n)} = \Gamma_n$$

But the order of torsion elements in Γ_n are the divisors of $2, 3, 6n+5$. Hence if $\Gamma_m \cong \Gamma_n$, we must have that $m = n$. We have proven:

Theorem 6.3.3 (Dehn). There are infinitely many non-homeomorphic 3-dimensional homology spheres.

7 Hyperbolic groups

The goal is to define a notion of coarse hyperbolic geometry. This is something which looks like hyperbolic geometry that is invariant under quasi-isometry.

7.1 Hyperbolic metric spaces

Let X be a geodesic metric space. A *geodesic triangle* is a triple of geodesics

$$\Delta = [x, y] \cup [y, z] \cup [z, x]$$

For $A \subseteq X$, let

$$N_\delta(A) = \{y \in X \mid \exists x \in A, d(x, y) \leq \delta\} = \bigcup_{x \in A} B_\delta(x)$$

be its (closed) δ-neighbourhood.

Definition 7.1.1

Let $\delta \geq 0$. A geodesic triangle Δ is δ-*slim* if the δ-neighbourhood of any two sides cover the third side. So

$$[x, y] \subseteq N_\delta([x, z] \cup [y, z])$$

and so on.

Definition 7.1.2

X is called δ-*hyperbolic* if every geodesic triangle $\Delta \subseteq X$ is δ-slim. We also say X is *Gromov-hyperbolic*, or *hyperbolic*.

Example 7.1.3

If $\operatorname{diam}(X) = \delta$, then X is δ-hyperbolic.

Example 7.1.4

If X is a tree, then X is 0-hyperbolic.

Example 7.1.5 (non-example)

Euclidean space is not Gromov-hyperbolic.

Example 7.1.6
$\mathbb{H}^2$ is hyperbolic. To see this, Δ is δ-slim, where δ is the radius of the largest semicircle which we can inscribe in Δ.

Let $A(r)$ be the area of a circle of radius r in $\mathbb{H}^2$. But now

$$\frac{1}{2}A(\delta) \leq \text{Area}(\Delta) < \pi$$

Since $A(\delta) \to \infty$ as $\delta \to \infty$, we see that $\mathbb{H}^2$ is δ-hyperbolic for sufficiently large δ.

7.2 The Mostow-Morse lemma

The goal is to prove that Gromov-hyperbolicity is a quasi-isometry invariant.

Lecture 20

Definition 7.2.1 (quasigeodesic)
A path $\gamma : [a, b] \to X$ is a (λ, ε)-*quasigeodesic* if γ is a (λ, ε)-quasi isometric embedding. That is,

$$\frac{1}{\lambda}|s - t| - \varepsilon \leq d(\gamma(s), \gamma(t)) \leq \lambda|s - t| + \varepsilon$$

Definition 7.2.2 (Hausdoff distance)
Let $A, B \subseteq X$ be nonempty subsets of a metric space X. Let

$$N_C(A) = \bigcup_{a \in A} B_c(a) = \{x \in X \mid \exists a \in A, d(x, a) \leq c\}$$

The *Hausdorff distance* is

$$d_{\text{Haus}}(A, B) = \inf\{c > 0 \mid A \subseteq N_c(B) \text{ and } B \subseteq N_c(A)\}$$

Definition 7.2.3 (length)
Let $\gamma : [a, b] \to X$ be a path. The *length* of γ is

$$\ell(\gamma) = \sup_{\mathcal{D}} \sum_{i=1}^{n} d(\gamma(t_{i-1}), \gamma(t_i))$$

where $\mathcal{D}$ ranges over all dissections

$$a = t_0 < t_1 < \cdots < t_n = b$$

Lemma 7.2.4. For any $\lambda \geq 1, \varepsilon \geq 0$, there are $\lambda' \geq 1, \varepsilon' \geq 1$, such that for any geodesic metric space X, and any (λ, ε)-quasigeodesic $\alpha : [a, b] \to X$, there exists a continuous (λ', ε')-quasigeodesic $\alpha' : [a, b] \to X$, such that

(i) $\alpha'(a) = \alpha(a), \alpha'(b) = \alpha(b)$,

(ii) $d_{\text{Haus}}(\text{im}(\alpha), \text{im}(\alpha')) \leq \lambda + \varepsilon$,

(iii) $\ell(\alpha'|_{[s,t]}) \leq \lambda' d(\alpha'(s), \alpha'(t)) + \varepsilon$, for all $a \leq s \leq t \leq b$.

Proof. Let $I = \{a, b\} \cup (a, b) \cap \mathbb{Z}$. Define α' by setting $\alpha'(t) = \alpha(t)$ for all $t \in I$, and then interpolating using a (reparametrised) geodesic between points of I.

Continuity is clear, and so is (i). (ii) is easy. The fact that α' is a quasi-geodesic and (iii) follow from easy, but tedious calculations. □

Lemma 7.2.5. Let X be a δ-hyperbolic metric space. Suppose $\beta : [a, b] \to X$ is a geodesic, $\alpha : [a, b] \to X$ is a continuous path, with $\alpha(a) = \beta(a)$, $\alpha(b) = \beta(b)$. Then

$$d(\beta(t), \mathrm{im}(\alpha)) \leq \delta \lfloor \log_2(\ell(\alpha)) \rfloor + 1$$

Proof. Let

$$N = \lfloor \log_2(\ell(\alpha)) \rfloor$$

The proof proceeds by induction on N. If $N \leq 0$ then $\ell(\alpha) \leq 1$ and we are done.

Consider the geodesic triangle with vertices

$$\alpha(a), \alpha(b), \alpha\left(\frac{a+b}{2}\right)$$

Since X is δ-hyperbolic, $\beta(t)$ has distance at most δ from one of the other edges of the triangle. Call the corresponding half of α α', and the geodesic β'. Now

$$\ell(\alpha') = \frac{\ell(\alpha)}{2} \implies \lfloor \log_2(\ell(\alpha')) \rfloor = N - 1$$

and we have a point $\beta'(t')$ such that $d(\beta(t), \beta'(t')) \leq \delta$. By inductive hypothesis,

$$\begin{aligned} d(\beta(t), \mathrm{im}(\alpha)) &\leq d(\beta(t), \mathrm{im}(\alpha')) \\ &\leq d(\beta(t), \beta'(t')) + d(\beta'(t'), \mathrm{im}(\alpha')) \\ &\leq \delta + \delta(N-1) + 1 \\ &= \delta N + 1 \end{aligned}$$

as required. □

We are now ready for the main result of this section:

Theorem 7.2.6 (Mostow-Morse lemma)**.** Let X be a (geodesic) δ-hyperbolic space. Let $\alpha : [a', b'] \to X$ be a (λ, ε)-quasigeodesic, and $\beta : [a, b] \to X$ a geodesic, with

$$\beta(a) = \alpha(a') \quad \text{and} \quad \beta(b) = \alpha(b')$$

Then there exists a constant $C = C(\lambda, \varepsilon, \delta)$, such that

$$d_{\mathrm{Haus}}(\mathrm{im}(\alpha), \mathrm{im}(\beta)) \leq C$$

Proof. We may replace α by the result of lemma 7.2.4. In particular, α is continuous, and

$$\ell\left(\alpha_{[s,t]}\right) \leq \lambda|s - t| + \varepsilon$$

for $a \leq s \leq t \leq b$. We need to bound

$$C_1 = \inf\{C \mid \mathrm{im}(\beta) \subseteq N_C(\mathrm{im}(\alpha))\} \quad \text{and} \quad C_2 = \inf\{C \mid \mathrm{im}(\alpha) \subseteq N_C(\mathrm{im}(\beta))\}$$

We'll first bound C_1. For this, we'll need to bound

$$d(\beta(t), \mathrm{im}(\alpha)) = \inf_{t' \in [a,b]} d(\beta(t), \alpha(t'))$$

Let $C = \sup_{t \in [a,b]} d(\beta(t), \mathrm{im}(\alpha))$. Since $[a', b']$ is compact, it is realised at some $\beta(t)$. Let

$$r = \max\{a, t - 2C\} \quad \text{and} \quad s = \min\{b, t + 2C\}$$

Define the path γ by going from $\beta(r)$ to the closest point $\alpha(r')$ on α, following α until the closest point $\alpha(s')$ to $\beta(s)$, and then going to $\beta(s)$. Then

$$\begin{aligned}\ell(\gamma) &\leq 2C + \ell(\alpha|_{[r',s']}) \\ &\leq 2C + \lambda d(\alpha(r'), \alpha(s')) + \varepsilon \\ &\leq 6\lambda C + 2C + \varepsilon\end{aligned}$$

Lecture 21

On the other hand, the lemma above shows that

$$C \leq \delta \lfloor \log_2(\ell(\gamma)) \rfloor + 1$$

Thus,

$$C \leq \delta \lfloor \log_2(6\lambda C + 2C + \varepsilon) \rfloor + 1$$

Since the left hand side is linear, and the right hand side is logarithmic in C, there is an upper bound on C, which only depends on δ, λ and ε.

Next, we need to bound C_2, i.e. we need to bound $d(\alpha(t), \text{im}(\beta))$. Let $[s', r'] \subseteq [a', b']$ be maximal such that $\alpha|_{(s',r')}$ lies outside of $N_C(\text{im}(\beta))$. Here, C is the constant from above. By continuity, there exists $t \in [a, b]$, and $s \in [a', s']$, $r \in [r', b']$ such that

$$d(\beta(t), \alpha(s)), d(\beta(t), \alpha(r)) \leq C$$

as the interval is connected. Thus, $d(\alpha(r), \alpha(s)) \leq 2C$. Hence

$$\ell(\alpha|_{[s',r']}) \leq \ell(\alpha|_{[s,r]}) \leq \lambda d(\alpha(s), \alpha(r)) + \varepsilon \leq 2\lambda C + \varepsilon$$

Hence every point on α is at most $2C\lambda + C + \varepsilon$ from $\text{im}(\beta)$. □

Corollary 7.2.7. Let X, Y be geodesic metric spaces. If X is δ-hyperbolic, and X is quasi-isometric to Y, then Y is δ'-hyperbolic for some δ'.

Proof. Let $f : X \to Y, g : Y \to X$ be (λ, ε)-quasi-isometries, such that

$$d(f(g(y)), y) \leq \varepsilon \quad \text{and} \quad d(g(f(x)), x) \leq \varepsilon$$

Consider a geodesic triangle

$$[y_1, y_2] \cup [y_2, y_3] \cup [y_3, y_1] \subseteq Y$$

Consider $y \in [y_1, y_2]$. By the Mostow-Morse lemma, there exists $x \in [g(y_1), g(y_2)]$ such that

$$d(x, g(y)) \leq C$$

Since X is δ-hyperbolic, there exists (without loss of generality) $x' \in [g(y_2), g(y_3)]$ such that $d(x, x') \leq \delta$. By the Mostow-Morse lemma again, there exists $y' \in [y_2, y_3]$ such that

$$d(x', g(y')) \leq C$$

In summary,

$$d(g(y), g(y')) \leq 2C + \delta$$

and so

$$d(f(g(y)), f(g(y'))) \leq \lambda(2C + \delta) + \varepsilon$$

and thus

$$d(y, y') \leq \lambda(2C + \delta) + 3\varepsilon$$

The right hand side is a function of δ, λ and ε only. □

Example 7.2.8

Let $G = \pi_1(\Sigma_2)$. This has presentation

$$\langle a_1, b_1, a_2, b_2 \mid [a_1, b_1][a_2, b_2] \rangle$$

By the Schwarz-Milnor lemma, $\text{Cay}(G) \overset{\text{qi}}{\sim} \mathbb{H}^2$, which is Gromov hyperbolic, and so $\text{Cay}(G)$ is Gromov hyperbolic.

7.3 Hyperbolic groups

Using the previous corollary, the following properties of a group G are all equivalent.

1. G has a finite generating set S, such that $\mathrm{Cay}(G, S)$ is Gromov hyperbolic.
2. G is finitely generated, and for any finite generating set S, $\mathrm{Cay}(G, S)$ is Gromov hyperbolic.
3. G acts properly discontinuously and cocompactly by isometries on some proper geodesic Gromov hyperbolic metric space X.
4. Every proper geodesic metric space X on which G acts properly discontinuously and compactly is Gromov hyperbolic.

Definition 7.3.1

G is *(word) hyperbolic* if any of the above hold.

Lecture 22

Example 7.3.2

If G is finite, then $\mathrm{Cay}_S(G)$ is bounded, and so hyperbolic.

Example 7.3.3

If $G = F_m$, then the standard generating set gives $\mathrm{Cay}_S(G)$ which is a tree. Recall that trees are 0-hyperbolic.

Example 7.3.4

If $\mathbb{Z}^2$ was hyperbolic, then $\mathbb{R}^2$ would be Gromov hyperbolic, which it is not.

Example 7.3.5

For $g \geq 2$, let Σ_g be the closed oriented surface of genus g. Let $G\pi_1(\Sigma_g)$. Then G acts on $\mathbb{H}^2$ properly discontinuously, cocompactly by isometries. Thus, G is hyperbolic.

Remark 7.3.6. Sometimes authors say a group acts on a space *geometrically* if it acts properly discontinuously and cocompactly by isometries.

Example 7.3.7

$\pi_1(M)$ is hyperbolic if M is any closed Riemannian manifold with negative sectional curvature.

Example 7.3.8

$\mathrm{SL}_2(\mathbb{Z}) \cong \mathbb{Z}/4 \underset{\mathbb{Z}/2}{*} \mathbb{Z}/6$. The Bass-Serre tree is an infinite 3-valent tree T, and $\mathrm{SL}_2(\mathbb{Z})$ acts geometrically on T, so $\mathrm{SL}_2(\mathbb{Z})$ is hyperbolic.

Example 7.3.9 (random finitely presented groups)
If

$$G = \langle a_1, \dots, a_m \mid r_1, \dots, r_n \rangle$$

is "chosen at random", then G is infinite and hyperbolic.

7.4 Local geodesics

Our goal is to solve the word problem in hyperbolic groups. The key ingredient is a "local to global" statement, about geodesics in hyperbolic metric spaces.

Definition 7.4.1
A path γ in a metric space X is a c-*local geodesic* if $d(\gamma(s), \gamma(t)) = |s - t|$ whenever $|s - t| \leq c$.

Lemma 7.4.2. Let X be a δ-hyperbolic metric space. If $\alpha : [a, b] \to X$ is a 10δ-local geodesic, then

$$\operatorname{im}(\alpha) \subseteq N_{2\delta}([\alpha(a), \alpha(b)])$$

for any geodesic $[\alpha(a), \alpha(b)]$.

Proof. Let

$$C = \sup_{t \in [a,b]} d(\alpha(t), [\alpha(a), \alpha(b)])$$

Say it is realised at $t_0 \in [a, b]$. Let $r = \max\{a, t_0 - 5\delta\}, s = \min\{b, t_0 + 5\delta\}$.

Let $x, y, z \in [\alpha(a), \alpha(b)]$ be the closest points to $\alpha(r), \alpha(s), \alpha(t_0)$ respectively. Then $d(x, \alpha(r)), d(y, \alpha(r)) \leq C$, and $d(\alpha(t_0), z) = C$. Consider the quadrilateral with vertices $\alpha(r), \alpha(s), x, y$.

Note we can subdivide it into two triangles. and so any point p on $\alpha([r, s])$ is within distance 2δ of one of the other three sides. Apply this to $p = \alpha(t_0)$. Suppose there is a point $w \in [\alpha(r), x]$, such that $d(\alpha(t_0), w) \leq 2\delta$. But then

$$d(\alpha(r), w) \geq d(\alpha(r), \alpha(t_0)) - d(\alpha(t_0), w) \geq 5\delta - 2\delta = 3\delta$$

In this case,

$$\begin{aligned} d(\alpha(t_0), x) &\leq 2\delta + d(w, x) \\ &< 3\delta + d(w, x) \\ &\leq d(\alpha(r), x) \\ &\leq C \end{aligned}$$

But this contradicts $d(\alpha(t_0), [\alpha(a), \alpha(b)]) = C$. Therefore, $\alpha(t_0)$ is not within 2δ of $[\alpha(r), x]$. By symmetry, it is not within 2δ of $[\alpha(s), y]$. Thus, it is within 2δ of $[x, y]$. With this, $C \leq 2\delta$. □

Remark 7.4.3. This is a coarse analogue of the fact that local geodesics in trees are global geodesics.

A consequence of this is key to solving the word problem in hyperbolic groups.

Lemma 7.4.4 (shortcuts in hyperbolic spaces)**.** Let X be δ-hyperbolic. Any loop $\alpha : [a, b] \to X$ such that $\ell(\alpha) > 4\delta$ contains $a \leq s < t \leq b$, such that

$$d(\alpha(s), \alpha(t)) < \ell(\alpha|_{[s,t]}) \leq 10\delta \qquad (*)$$

Proof. Unless $(*)$ is satisfied, then α is a 10δ-local geodesic. By the previous lemma,

$$\operatorname{im}(\alpha) \subseteq N_{2\delta}([\alpha(a), \alpha(b)]) = B_{2\delta}(\alpha(a))$$

Since α is a 10δ local geodesic, and $\operatorname{diam}(B_{2\delta}(\alpha(a))) \leq 4\delta$, it follows that $\ell(\alpha) \leq 4\delta$. □

Lecture 23

7.5 Dehn's algorithm

We will solve the word problem for all hyperbolic groups, using an algorithm that Dehn exhibited for hyperbolic surface groups, in 1912.

Theorem 7.5.1 (relations in hyperbolic groups). Let G be a hyperbolic group, and S a finite generating set. For every non-trivial edge loop α in $\mathrm{Cay}_S(G)$, there is an edge loop γ of length at most 20δ, such that

$$\ell(\alpha\beta\gamma\beta^{-1}) < \ell(\alpha)$$

for some choice of path β from 1 to a point on γ.

Proof. If $\ell(\alpha) \leq 20\delta$, then we can take $\gamma = \alpha^{-1}$. Then $\alpha\gamma$ is homotopic to the constant loop, and so $\ell(\alpha\gamma) = 0 < \ell(\alpha)$.

Otherwise, from the previous lemma, let

$$\beta = (\alpha|_{[t,b]})^{-1}$$
$$\gamma = (\alpha|_{[s,t]})^{-1} \cdot [\alpha(s), \alpha(t)]$$

Then $\ell(\gamma) < 20\delta$, and $\alpha\beta\gamma\beta^{-1}$ is homotopic to

$$\alpha|_{[a,s]} \cdot [\alpha(s), \alpha(t)] \cdot \alpha|_{[t,b]}$$

which has length less than $\ell(\alpha)$. □

Corollary 7.5.2 (Gromov). Hyperbolic groups are finitely presented.

Proof. Let S be a finite generating set for a hyperbolic group G. Consider $\mathrm{Cay}_S(G)$. This is δ-hyperbolic for some δ. Let

$$R = \{\text{edge loops in } \mathrm{Cay}_S(G) \text{ based at 1 with length at most } 20\delta\}$$

This is a finite set, with size at most $(2|S|)^{20\delta}$ say. We claim that $\langle S \mid R\rangle$ is a presentation for G. To see this, by the theorem, and induction on length, every relation is a product of conjugates of elements of R. □

Corollary 7.5.3 (Dehn, Gromov). Let G be a hyperbolic group. The word problem in G is solvable.

Proof. Consider the presentation $G = \langle S \mid R\rangle$, constructed in the previous corollary. Let $w \in F(S)$. The theorem tells us that if w represents the trivial element in G, then there is a cyclic conjugate w' of w, and $r \in R$, such that $\ell(w'r) < \ell(w)$. To see this, let $\alpha = w$ and let $w' = \beta^{-1}\alpha\beta$, $r = \gamma$. Since w has finitely many cyclic conjugates, and R is finite, we have finitely many combinations of (w', r) to check. If we fine one such combination, then we replace w with $w'r$ and repeat.

On the other hand, if we cannot find (w', r), then it must be the case that w did not represent a loop.

Since $\ell(w'r) < \ell(w)$, this process has to terminate, either showing that w is not a loop, or when $w'r$ is the trivial element. □

Remark 7.5.4. A presentation in the corollary is called a *Dehn presentation*. That is, a presentation $\langle S \mid R\rangle$, such that for any non-trivial word w, with $w = 1$ in G, there exists $h \in G, r^{\pm 1} \in R$ such that

$$\ell(whrh^{-1}) < \ell(w)$$

It turns out a group G has a Dehn presentation if and only if G is hyperbolic.

Lecture 24

8 *Outlook, further topics, open problems*

Random groups

Fix a generating set $S = \{a_1, \dots, a_m\}$. Fix $n \geq 1$, choose a subset

$$\{r_1, \dots, r_n\} \subseteq F(S)$$

uniformly at random such that $\ell(r_i) = \ell$ for all i. Consider the resulting group

$$G = \langle a_1, \dots, a_m \mid r_1, \dots, r_n \rangle$$

For any property P of groups, we can look at

$$\mathbb{P}(G \in P)$$

which depends on m, n, ℓ. We say that a *random group has property P* if

$$\mathbb{P}(G \in P) \to 1$$

as $\ell \to \infty$.

Theorem 8.0.1 (Gromov, Ol'shanski)**.** For $m \geq 2$, a random group is infinite and hyperbolic.

Subgroups

One of the most important open problems concern subgroups of hyperbolic groups.

Conjecture 8.0.2 (surface subgroup)**.** Unless G is virtually[a] free, if G is word hyperbolic, then there exists a surface Σ_g of genus $g \geq 2$, such that $\pi_1(\Sigma_g) \leq G$.

[a] has a finite index subgroup which is

This has been proven in a special case by Kahn-Markovich, when $G = \pi_1(M^3)$ for M a compact 3-manifold.

Representations and residual finiteness

A group G is *linear* if it is a subgroup of $\mathrm{GL}(n, \mathbb{C})$ for some n. That is, it has a faithful representation over $\mathbb{C}$.

Theorem 8.0.3 (M. Kapovich)**.** There is a hyperbolic group which is not linear.

But a weaker property is also important.

Definition 8.0.4

A group G is *residually finite* if for any $g \in G$ non-trivial, there exists a homomorphism $f : G \to Q$ finite, such that $f(g) \neq 1$.

All finitely generated linear groups are residually finite. Then it is an open question whether every hyperbolic group is residually finite. Recent progress includes

Theorem 8.0.5 (Olivier-Wise, Agol)**.** Random groups are residually finite. In fact, they are linear.

Boundaries

Recall that $\partial \mathbb{H}^2 = S^1$.

Definition 8.0.6

Let X be a proper hyperbolic metric space. A *geodesic ray* is an isometric embedding $\gamma : [0, \infty) \to X$. We say that $\gamma_1 \sim \gamma_2$ if there exists $C \geq 0$ such that

$$d(\gamma_1(t), \gamma_2(t)) \leq C$$

for all t.

The *Gromov boundary* of X is defined to be

$$\partial_\infty X = \frac{\{\text{geodesic rays in X}\}}{\sim}$$

Remark 8.0.7. $\partial_\infty X$ admits a natural boundary, so that $\partial_\infty X$ and $X \cup \partial_\infty X$ are compact.

A quasi-isometry $f : X \to Y$ induces a homeomorphism $\partial_\infty X \to \partial_\infty Y$. Thus, for a hyperbolic group G, we may define

$$\partial_\infty G = \partial_\infty \operatorname{Cay}_S(G)$$

Example 8.0.8

If G is a cocompact Fuchsian group (e.g. $\pi_1(\Sigma_g)$ and triangle groups), then G is quasi-isometric to $\mathbb{H}^2$, and so $\partial_\infty G = \partial \mathbb{H}^2 = S^1$.

Theorem 8.0.9. If G is hyperbolic and $\partial_\infty G \cong S^1$, then G is virtually Fuchsian.

Conjecture 8.0.10 (Cannon). If G is hyperbolic, and $\partial_\infty G \cong S^2$, then G is virtually $\pi_1(M)$ for M a 3-manifold.

Non-positive curvature

Definition 8.0.11

Suppose X is a geodesic metric space. Each geodesic triangle in X has a well defined (up to isometry) *comparison triangle* $\overline{\Delta} \subseteq \mathbb{R}^n$. That is, it is a triangle with the same side lengths as Δ. Let $f : \overline{\Delta} \to \Delta$ be the natural map.

X is *CAT(0)* if $d(x, y) \geq d(f(x), f(y))$ for all $x, y \in \overline{\Delta}$.

One question: Does every hyperbolic group act geometrically on a CAT(0) space?