Contents

1 Foundations of geometric group theory

1.1 Word metrics and Cayley graphs

Let S be a finite generating set for a group G. We assume S is *symmetric*, i.e., if $s \in S$, then $s^{-1} \in S$.

Definition 1.1. A *word* in S is a concatenation of elements of S: $W = s_1 s_2 \dots s_n$ where $s_i \in S$. We let $\|W\|$ denote the length of the word, so here $\|W\| = n$. If two words W, U are letter-for-letter equivalent, we write $W \equiv U$. If W, U represent the same element of the group G, we write $W =_G U$. If $g \in G$, write $|g|_S$ for the length of the shortest word in S that represents g in G. Given $g, h \in G$, let $d_S(g,h) = |g^{-1}h|_S$. The metric d_S is the *word metric* on G with respect to S.

Definition 1.2. The *Cayley graph* of G with respect to S, which we denote by $\Gamma(G,S)$ is the graph whose vertex set is $\{g \in G\}$ with a (directed) edge from a vertex g to a vertex h whenever $h = gs$ for some $s \in S$.

Each edge has an orientation and a label. But if $gs = h$, then $g = hs^{-1}$ as well, so there is also a directed edge from h to g labeled by s^{-1}. This is rather cumbersome, so we typically don't draw the edge labeled by s^{-1}. We just keep in mind that if we traverse an edge labeled by s going in the opposite direction as its orientation, we think of that edge as labeled by s^{-1}.

Given a path p in $\Gamma(G,S)$, let $\mathbf{Lab}(p)$ be its label in $\Gamma(G,S)$. We denote the initial point of p by p_- and the terminal point of p by p_+.

Remark 1.3. d_S is a metric on the vertex set of $\Gamma(G,S)$. Identifying each edge with the unit interval $[0,1]$ extends this metric to the whole graph.

Exercises 1.4. *Draw the following Cayley graphs.*

1. $\Gamma(\mathbb{Z}/4\mathbb{Z},\{1\})$
2. $\Gamma(\mathbb{Z},\{1\})$
3. $\Gamma(\mathbb{Z},\{2,3\})$
4. $\Gamma(\mathbb{Z}^2,\{a=(1,0),b=(0,1)\})$
5. $\Gamma(\mathbb{F}_2,\{a,b\})$*, where* $\mathbb{F}_2=\langle a,b\rangle$ *(i.e., the group with two generators and no relations)*
6. $\Gamma(BS(1,2),\{a,t\})$*, where* $BS(1,2)=\langle a,t \mid tat^{-1}=a^2\rangle$.

1.2 Metric spaces

Let (X,d) denote a metric space X with its metric d. If the metric is clear, we may just write X.

If $x\in X$ and $r\geq 0$, then $B_R(x)=\{y\in X \mid d(x,y)\leq r\}$ is the (closed) ball of radius r about x. If $A\subseteq X$, then we write $\mathcal{N}_r(A)=\{y\in X \mid d(y,a)\leq r\}=A^{+r}$ for the (closed) r–neighborhood of A in X.

Definition 1.5. A *path* is X is a continuous map $p\colon [a,b]\to X$ for some $[a,b]\subseteq\mathbb{R}$. By an abuse of notation, we also call the image of this map p. Then $p_-=p(a)$ and $p_+=p(b)$. A *ray* is a continuous map $p\colon [a,\infty)\to X$, and a *bi-infinite path* is a continuous map $p\colon (-\infty,\infty)\to X$.

Given a path p, the *length* of p is

$$\ell(p) := \sup_{a\leq t_1\leq t_2\leq\cdots\leq t_n\leq b}\sum_{i=1}^{n-1} d(p(t_i),p(t_{i+1})), \tag{1}$$

where the supremum is taken over all n and all choices of $t_1,\ldots,t_n$. This supremum is often infinite. We typically consider *rectifiable* paths, which are those with $\ell(p)<\infty$.

Example 1.6 (From BH I.1 Ex. 1.19). Let $X=[0,1]$, and let $0=t_0<t_1<\cdots<t_n<\cdots$ be an infinite sequence in $[0,1]$ converging to 1. let $c\colon [0,1]\to X$ be any path such that $c(0)=0$, $c(t_n)=\sum_{k=1}^{n}\frac{(-1)^{k+1}}{k}$. Then c is not rectifiable because its length is bounded below by the harmonic series $\sum_{i=1}^{\infty}\frac{1}{k}$, which diverges to infinity.

Definition 1.7. A path is *geodesic* if $\ell(p)=d(p_-,p_+)$. A ray is *geodesic* if for all $a\leq s\leq t<\infty$, $\ell(p|_{[s,t]})=d(p(s),p(t))$, and similarly for a bi-infinite path. A metric space is *geodesic* if every pair of points can be connected by a geodesic.

A weaker condition one can put on a metric space is that it is a *length space*, which means that it is a path-connected space such that for all $x,y\in X$, $d(x,y)=\inf\{\ell(p)\mid p_-=x,p_+=y\}$. For example, consider a space with two points, x and y, with countably many edges between them indexed by the positive integers so that the n-th edge has length $1+\frac{1}{n}$. Here, $d(x,y)=\inf(1+frac1n)$, but there is no path of length 1 between them. Most statements in this class will hold for length spaces, but the proofs tend to be a bit more technical, so we will mostly work with geodesic spaces.

Definition 1.8. Let (X, d_X) and (Y, d_Y) be metric spaces. A map $f\colon X \to Y$ is:

- an *isometry* if for all $x, y \in X$,
$$d_Y(f(x), f(y)) = d_X(x, y)$$
and f is surjective;
- a *biLipschitz equivalence* if there exists $\lambda \geq 1$ such that for all $x, y \in X$,
$$\frac{1}{\lambda} d_X(x, y) \leq d_Y(f(x), f(y)) \leq \lambda d_X(x, y)$$
and f is surjective;
- a (λ, c)–*quasi-isometry* if there exists $\lambda \geq 1$ and $c \geq 0$ such that for all $x, y \in X$
$$\frac{1}{\lambda} d_X(x, y) - c \leq d_Y(f(x), f(y)) \leq \lambda d_X(x, y) + c$$
and $f(X)^{+c} = Y$, i.e., f is *coarsely surjective.*

We write $X \sim_{Lip} Y$ or $X \sim_{qi} Y$ if there exists a biLipschitz equivalence or (λ, c)–quasi-isometry from X to Y, respectively.

If the assumption of (coarse) surjectivity is dropped, then f is a(n) isometric, biLipschitz, or quasi-isometric *embedding*, respectively. Note that quasi-isometries are not required to be continuous, though it will follow from the definitions that isometries and biLipschitz equivalences are continuous.

Exercises 1.9.

1. *Show that $\sim_{Lip}$ and $\sim_{qi}$ are equivalence relations.*
2. *Let X, Y be bounded metric spaces. Show that $X \sim_{qi} Y$.*
3. *Show that $\Gamma(\mathbb{Z}, \{1\})$ is quasi-isometric to the real line $\mathbb{R}$.*
4. *Show that $\Gamma(\mathbb{Z}, \{1\}) \sim_{qi} \Gamma(\mathbb{Z}, \{2, 3\})$.*
5. *In general, show that if $S, T \subseteq G$ are two finite generating sets, then $(G, d_S) \sim_{Lip} (G, d_T)$, and so $\Gamma(G, S) \sim_{qi} \Gamma(G, T)$.*

1.3 Group actions

Let G act on a metric space X, that is, there exists a map $\varphi\colon G \times X \to X$ such that $\varphi(gh, x) = \varphi(g, \varphi(h, x))$. We write $G \curvearrowright X$, and, to save space, we often drop the φ notation simply write gx (or $g.x$ or $g \cdot x$) for $\varphi(g, x)$. In this class, we will ALWAYS assume the action is by isometries, that is, for all $g \in G$, $d_X(x, y) = d_X(gx, gy)$.

There is a bijection between actions of G on X and homomorphims $\rho\colon G \to \mathrm{Isom}(X)$, where $\mathrm{Isom}(X)$ is the group of isometries of the space X.

Definition 1.10. The action is *faithful* if ρ is injective, i.e., for all $g \in G$, there exists $x \in X$ such that $gx \neq x$. The action is *free* if for all $x \in X$, $\mathrm{Stab}_G(x) = \{1\}$, where $\mathrm{Stab}_G(x) = \{g \in G \mid gx = x\}$, i.e., for all $x \in X$ for all $g \in G$, $gx \neq x$.

The action is *proper* if for all bounded subsets $B \subseteq X$,
$$\#\{g \in G \mid gB \cap B \neq \emptyset\} < \infty.$$
This is also called *metric properness.* There is a topological version where bounded is replaced by compact.

The action is *cobounded* (resp. *cocompact*) if there exists a cobounded (resp. cocompact) subset $K \subseteq X$ such that
$$\bigcup_{g \in G} gK = X.$$

Group actions can be used to prove algebraic results. Here is just one example:

Theorem 1.11. *Suppose G acts freely on $\mathbb{R}^n$ (where $\mathbb{R}^n$ has the Euclidean metric). Then G is torsion free.*

Proof. Exercise. $\square$

[Hint: Use the fact that any finite set of points $\{x_1, \ldots, x_n\}$ in $\mathbb{R}^n$ has a unique *centroid*, that is, a point $w \in \mathbb{R}^n$ that minimizes the distance $\sum d_{\mathbb{R}^n}(w, x_i)$.]

The following is one of the the foundational results of geometric group theory. A *proper* metric space is one in which closed balls are compact.

Theorem 1.12 (Milnor–Swarz Lemma)**.** *Let G be a group acting properly and coboundedly on a propoer geodesic metric space X. Then G has a finite generating set S and $\Gamma(G,S) \sim_{qi} X$.*

Proof. Fix $x \in X$. Choose a bounded subset K of X so that $\bigcup_{g \in G} gK = X$. Since K is bounded, there exists a closed ball B of radius D such that $K \subseteq B$. Thus we have $\bigcup_{g \in G} gB = X$. Since X is proper, B is compact. Let

$$S = \{g \in G \setminus \{1\} \mid B \cap gB \neq \emptyset\}.$$

Since $G \curvearrowright X$ is proper, the set S is finite, and $S = S^{-1}$ by the definition of S.

Consider $\inf\{d(B, gB) \mid g \in G \setminus (S \cup \{1\})\}$. For some $g \in G \setminus (S \cup \{1\})$, the distance $d(B, gB)$ is a positive constant R, by the definition of S. The set H of elements of G such that $d(B, gB) \leq R$ is contained in

$$\{g \in G \mid gB(x, D+R) \cap B(x, D+R)\}$$

and hence is finite. Since

$$\inf\{d(B, gB) \mid g \in G \setminus (S \cup \{1\})\} = \inf\{d(B, gB) \mid g \in H \setminus (S \cup \{1\})\},$$

and the latter is taken over a finite set, we have that the infimum is realized by some $h_0 \in H \setminus (S \cup \{1\})$. For this h_0, we have that $d(B, hB) = 2L$ for some positive constant L. By definition, if $d(B, gB) < 2L$, then $g \in S$.

We first show that S generates G. Let $g \in G$ and consider a geodesic $[x, gx]$ in X. let $k = \left\lfloor \frac{d(x,gx)}{L} \right\rfloor$. Then there exists a finite sequence a points along the geodesic $x = x_0, x_1, \ldots x_{k+1} = gx$ such that $d(x_i, x_{i+1}) \leq L$ for every $i \in \{0, \ldots, k\}$. For each such i, let $h_i \in G$ be such that $x_i \in h_i B$. We take $h_0 = 1$ and $h_{k+1} = g$. As $d(B, h_i^{-1} h_{i+1} B) = d(h_i B, h_{i+1} B) \leq d(x_i, x_{i+1}) = L$, it follows that $h_i^{-1} h_{i+1} \in S$. Also,

$$h_0 (h_0^{-1} h_1)(h_1^{-1} h_2) \cdots (h_k^{-1} h_{k+1}) = h_{k+1} = g,$$

and so $g \in \langle S \rangle$. Therefore the finite set S generates G.

It remains to show that $\Gamma(G,S)$ is quasi-isometric to X. let $f\colon G \to X$ be defined by $f(g) = gx$ (so f is the orbit map). Since $G \curvearrowright X$ is cobounded, this map is $2D$–coarsely surjective.

The argument above shows that

$$|g|_S \leq k + 1 \leq \frac{1}{L} d(x, gx) + 1.$$

For the other half of the inequality, let $|g|_S = m$, and let $w = s_1 s_2 \ldots s_m$ be a word in S such that $w =_G g$.

By the triangle inequality,

$$\begin{aligned} d(x, gx) &= d(x, s_1 s_2 \dots s_m) \\ &\le d(x, s_1 x) + d(s_1 x, s_1 s_2 x) + \cdots + d(s_1 \dots s_{m-1} x, s_1 \dots s_m x) \\ &= \sum_{i=1}^{m} d(x, s_i x) \\ &\le 2Dm \\ &= 2D|g|_S. \end{aligned}$$

Thus, we have for any $g \in G$:

$$L|g|_s - L \le d(x, gx) \le 2D|g|_S.$$

Note that since both the word metric and the metric on X are left-invariant (under the action of G), this suffices to prove the statement for any $h, g \in G$, rather than the pair $1, g \in G$ for which we have proven the result. □

Corollary 1.13. *The following follow immediately from the Milnor–Schwarz Lemma:*

1. *If G is finitely generated and H is a finite-index subgroup of G, then $H \sim_{qi} G$.*
2. *If $N \trianglelefteq G$ is a finite normal subgroup and G/N is finitely generated, then G is finitely generated and $G \sim_{qi} G/N$.*
3. *If M is a closed Riemannian manifold and $\tilde{M}$ is its universal cover, then $\pi_1(M) \sim_{qi} \tilde{M}$.*
4. *If G is a connected Lie group with left-invariant Riemannian metric and $\Gamma \le G$ is a uniform lattice, then Γ is finitely generated and $\Gamma \sim_{qi} G$.*

2 Geometry of hyperbolic metric spaces

Let X be a metric space. Given two points $x, y \in X$, let $[x, y]$ be a geodesic from X to Y.

2.1 Four equivalent definitions of hyperbolicity

Assume X is geodesic. For any geodesic triangle T in X with vertices x, y, z the following hold.

1. There exist points $i_z \in [x, y], i_y \in [x, z], i_x \in [y, z]$ such that $a := d(x, i_z) = d(x, i_y)$, $b := d(y, i_z) = d(y, i_x)$ and $c := d(z, i_x) = d(z, i_y)$. See Figure 1. The points i_x, i_y, i_z are called *internal points* of T.
2. Let T' be a *comparison tripod*, that is, a graph with 3 vertices of valence one labeled v_x, v_y, and v_z, and a central vertex of valence 3 labeled o, so that the lengths of the edges from v_x, v_y, and v_z to o are a, b, and c, respectively. The map $T \to T'$ which sends x, y, z to v_x, v_y, v_z respectively, extends uniquely to a map $\chi \colon T \to T'$ which restricts to an isometry on each side of T. This map sends i_x, i_y, i_z to o, and every other point in T' except v_x, v_y, and v_z will have a preimage in T consisting of exactly two points.

Definition 2.1 (Thin triangles)**.** Let $\delta \ge 0$. A geodesic metric spaces satisfies $\mathrm{Hyp}_1(\delta)$ if for any geodesic triangle T in X and any $a, b \in T$ which map to the same point in T', we have $d(a, b) \le \delta$. The triangle T is called *δ–thin*.

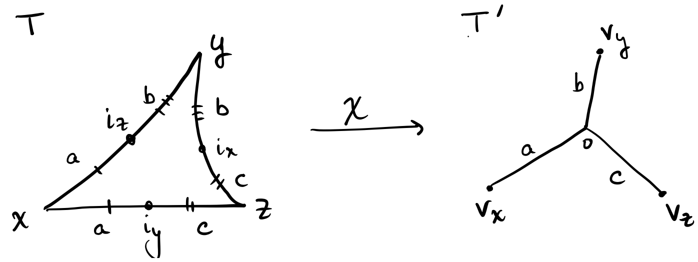

Figure 1: A geodesic triangle T in X and its comparison tripod T'.

Definition 2.2 (Insize)**.** Let $\delta \geq 0$. A geodesic metric spaces satisfies $\mathrm{Hyp}_2(\delta)$ if for any geodesic triangle T in X, $\mathrm{diam}(\{i_x, i_y, i_z\}) \leq \delta$. The quantity $\mathrm{diam}(\{i_x, i_y, i_z\})$ is called the *insize* of the triangle T.

The following is the most common definition of a (geodesic) hyperbolic metric space. It is attributed to Rips, and is sometimes called the Rips condition.

Definition 2.3 (Slim triangles)**.** Let $\delta \geq 0$. A geodesic metric spaces satisfies $\mathrm{Hyp}_3(\delta)$ if for any geodesic triangle T in X, we have $[x, y] \subseteq \mathcal{N}_\delta([x, z] \cup [y, z])$. Equivalently, given any point $a \in [x, y]$, there exists a point $b \in [x, z] \cup [y, z]$ such that $d(a, b) \leq \delta$.

Now let X be a (not necessarily geodesic) metric space and $x_0 \in X$.

Definition 2.4. The *Gromov product* of $x, y \in X$ with respect to x_0 is

$$(x \mid y)_{x_0} = \frac{1}{2}(d(x_0, x) + d(x_0, y) - d(x, y)).$$

If X is a tree, then $(x \mid y)_{x_0}$ measures the distance from x_0 to the geodesic $[x, y]$ (draw a picture of this!). Intuitively, $(x \mid y)_{x_0}$ will in general measure how long geodesics from x_0 to x and from x_0 to y "fellow travel." But note that we did not assume X was geodesic. Another way to think about the Gromov product is that it measures the failure of the triangle inequality to be an equality: in general we have that $d(x_0, x) + d(x_0, y) \geq d(x, y)$, and we call $d(x_0, x) + d(x_0, y) - d(x, y) \geq 0$ the *defect*. The quantity $(x \mid y)_{x_0}$ is half the defect.

The following is Gromov's original definition of a hyperbolic metric space.

Definition 2.5. Let $\delta \geq 0$. A metric spaces satisfies $\mathrm{Hyp}_4(\delta)$ if for any four points $x, y, z, w \in X$,

$$(x \mid y)_w \geq \min\{(x \mid z)_w, (y \mid z)_w\} - \delta.$$

Exercises 2.6. *If X is a tree, then X satisfies Hyp$_4(0)$. (We discussed this in class, but see if you can reconstruct the reasoning!)*

We can reformulate $\mathrm{Hyp}_4(\delta)$ as a more symmetric 4–*point condition*:

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

If you consider $x, y, z,$ and w to be the vertices of a tetrahedron, then the three sums in (2) are the sums of the lengths of the pairs of opposite sides. In a tree, a tetrahedron is a subtree with 6 vertices: the vertices x, y, z, w have valence 1, and there are two vertices with valence three. In such a tetrahedron, you can check that there is a tie for the two largest sums of opposite sides, and so $d(x, w) + d(y, z) \leq$

$\max\{d(x,y)+d(z,w), d(x,z)+d(y,w)$. The 4–point condition makes this coarse by requiring only that there is a *coarse* tie for the two largest sums.

Up to modifying δ by a constant multiple, all four conditions $\mathrm{Hyp}_i(\delta)$ are equivalent:

Theorem 2.7. *For any geodesic metric space X, any $\delta \geq 0$, and any geodesic triangle T, the following hold:*

(1) T is δ–thin $\implies$ T is δ–slim $\implies$ the insize of T is at most 2δ $\implies$ T is 2δ–thin; and

(2) the insize of T is at most δ $\implies$ Hyp$_4(\delta)$ $\implies$ the insize of T is at most 6δ.

Proof. We first prove (1). It is clear that the first implication holds.

For the second implication, let $x, y, z \in X$, and consider a geodesic triangle T. Suppose T is δ–slim. Then there exists a point p on (without loss of generality) $[x,y]$ at distance at most δ from i_x. By the triangle inequality, $|d(y,p) - d(y,i_x)| \leq \delta$. As $d(y,i_x) = d(y,i_z)$, we have $d(p,i_z) \leq \delta$. Thus $d(i_x,i_z) \leq 2\delta$. Similarly for the other pairs.

For the third implication, suppose p lies on $[y,z]$ is such that $d(y,p) < d(y,i_x)$. Then the fiber of $\chi\colon T \to T'$ containing p is $\{p,q\}$, where $q \in [y,x]$ and $d(y,p) = d(y,q)$. We will show that $d(p,q) \leq 2\delta$ by building a geodesic triangle with p,q as internal points.

let $c\colon [0,1] \to X$ be a monotone parametrization of $[y,z]$, and for each $t \in [0,1]$, consider a geodesic triangle $T_t = T(x,y,c(t))$, two of whose sides are $[y,x]$ and $c([0,t])$. The interal point of T_t on $c([0,t])$ varies continuously (though maybe not monotonically) as a function of t. At $t=0$, it is y, while at $t=1$ it is i_x, so for some $t \in [0,1]$ it is p. Since $d(y,p) = d(y,q)$, q must also be an internal point for T_t. As we assume the insize of all triangles is at most 2δ, we have $d(p,q) \leq 2\delta$.

We now prove (2). For the first implication, assume that the insize of any geodesic triangle in X is at most δ. Given $x,y,z,w \in X$, we may assume without loss of generality that $S := d(x,z) + d(y,w) \leq M := d(x,y) + d(z,w) \leq L := d(x,w) + d(y,z)$. To show the 4–point condition holds, we must show that $L \leq M + 2\delta$. Let $T = T(x,w,y)$ and $T' = T'(x,w,z)$ be geodesic triangles, and denote their internal points by i_x, i_w, i_y and i'_x, i'_w, i'_z, respectively. Consider the path from y to z which passes through i_x, i_y, i'_z, and i'_w, in that order. By the triangle inequality, we have

$$\begin{aligned} d(y,z) &\leq d(y,i_x) + d(i_x,i_y) + d(i_y,i'_z) + d(i'_z,i'_w) + d(i'_w,z) \\ &\leq d(y,i_x) + \delta + d(i_y,i'_z) + \delta + d(i'_w,z) \\ &= d(y,i_w) + d(i_y,i'_z) + d(z,i'_x) + 2\delta. \end{aligned}$$

We also have that

$$d(x,w) = d(x,i_y) + d(i'_z,w) - d(i'_z,i_y) = d(x,i_w) + d(i'_x,w) + d(i'_z,i_y).$$

Thus

$$\begin{aligned} L &\leq d(y,i_w) + d(i_y,i'_z) + d(z,i'_x) + 2\delta + d(x,i_w) + d(i'_x,w) + d(i'_z,i_y) \\ &= (d(x,i_w) + d(i_w,y)) + (d(w,i'_x) + d(z,i'_x)) + 2\delta \\ &= M + 2\delta, \end{aligned}$$

as desired.

For the second implication, assume that X satisfies $\mathrm{Hyp}_4(\delta)$. Consider a triangle $T = T(x,y,z)$. We will apply the four-point condition to the points x, y, z, and $i_x \in [y,z]$. The largest sum of pairs of distances in the four-point condition will be $d(x,i_x) + d(y,z)$. To see this, notice that $2(d(x,i_x) + d(y,z)) = (d(x,i_x) + d(i_x,y)) + (d(x,i_x) + d(i_x,z)) + d(y,z)$ is greater than the perimeter P of T, while the other sums are $d(z,i_x) + d(x,y) = d(y,i_x) + d(x,z) = P/2$. Therefore the four-point condition implies that $d(x,i_x) + d(y,z) \leq$

$P/2+2\delta$. Since $d(y,z)+d(x,i_z)$ is also equal to $P/2$, we conclude that $|d(x,i_x)-d(x,i_z)| \leq 2\delta$. Similarly, we have $|d(z,i_x)-d(z,i_z)| \leq 2\delta$.

We now consider the four points $\{x,z,i_x,i_z\}$. In this case the three sums of pairs of distances are $d(x,z)=d(x,i_z)+d(z,i_x)$, $d(i_x,i_z)+d(x,z)$, and $d(x,i_x)+d(z,i_z)$. By the previous paragraph, the third sum is at most $d(x,i_z)+d(z,i_x)+4\delta = d(x,z)+4\delta$. Thus we see that the largest two pairs of sums are the last two listed above. Applying the four-point condition, we have $d(i_x,i_z) \leq 6\delta$. Analogous arguments yield that $d(i_x,i_y) \leq 6\delta$ and $d(i_y,i_z) \leq 6\delta$, as well, completing the proof. □

Definition 2.8. We say that a metric space X is δ–hyperbolic (or Gromov hyperbolic) if it satisfies $\mathrm{Hyp}_4(\delta)$. If X is a geodesic metric space, we will also assume that δ has been chosen so that $\mathrm{Hyp}_1(\delta)$–$\mathrm{Hyp}_3(\delta)$ also hold.

Note: In a geodesic δ–hyperbolic metric space, the geometric interpretation of the Gromov product is coarsely the same as in the case of a tree: $(x \mid y)_{x_0}$ and $d(x_0,[x,y])$ will differ by at most δ.

Example 2.9. Examples of hyperbolic metric spaces:

1. If X is a bounded diameter metric space, then X is δ–hyperbolic with $\delta = \mathrm{diam}(X)$.

2. If X is a simplicial tree, then X is 0–hyperbolic.

3. The classical hyperbolic space $\mathbb{H}^2$ is δ–hyperbolic. (Exercise: Find δ! Hint: first show that it suffices to consider ideal triangles – those with all vertices at infinity. Then use the (isometric) action of $PSL_2(\mathbb{R})$ on $\mathbb{H}^2$ to show that it suffices to consider the ideal triangle with vertices at $-1, 1$, and ∞, in the upper half plane model.)

4. $\mathbb{H}^n$ for $n \geq 3$.

5. A space X is an *ultrametric space* if the following strong version of the triangle inequality holds for all $x,y,z \in X$: $d(x,y) \leq \max\{d(x,z), d(y,z)\}$. (In other words, a (geodesic) space is an ultrametric space is there is always a tie for the two longest sides of a triangle.) Ultrametric spaces often come up in relation to number theory; for example, the p–adic numbers $\mathbb{Q}_p$ with the p–adic metric space is an ultrametric space. All ultrametric spaces are 0–hyperbolic (this can be shown using the fourth definition of hyperbolicity).

6. An $\mathbb{R}$–tree is a generalization of a simplicial tree which can be defined as a space in which every geodesic triangle is a tripod. $\mathbb{R}$–trees are 0–hyperbolic.

7. Non-example: $\mathbb{R}^2$ (and more generally $\mathbb{R}^n$) is not hyperbolic for any δ. In fact, more is true: no δ–hyperbolic space can contain an isometrically copy of $\mathbb{R}^2$ (called a *flat*), or even a quasi-isometrically copy (called a *quasi-flat*).

Exercises 2.10. *Let X be a δ–hyperbolic metric space and $P = p_1 \dots p_n$ be a geodesic n–gon in X with $n \geq 3$. Let a be a point on p_i for some $1 \leq i \leq n$. Prove that there exists $j \neq i$ and $b \in p_j$ such that $d(a,b) \leq (n-2)\delta$. (In fact, $n-2$ can be replaced by $\log_2(n)$.)*

2.2 Quasigeodesic stability

We'd like to define a hyperbolic group as a finitely generated group so that every Cayley graph is δ–hyperbolic for some δ. However, this is a cumbersome (and impractical) definition, as we'd have to check it for *every single* finite generating set of the group. What would be nice is if the property of a metric space being δ–hyperbolic was a quasi-isometry invariant (in some sense which we'll make precise later). The most straightforward definitions of hyperbolicity involve geodesic triangles, and the problem is when you only look at things up to quasi-isometry, you lose control of how geodesics behave. In particular, the image of a geodesic under a quasi-isometry is not necessarily a geodesic any more. This motivates the notion of a *quasi-geodesic.*

Definition 2.11. A (λ, ϵ)–*quasi-geodesic* in a metric space X is a (λ, ϵ)–quasi-isometric embedding $c\colon I \to X$, where $I \subset \mathbb{R}$ is connected. More explicitly, we have

$$\frac{1}{\lambda}|t - t'| - \epsilon \leq d(c(t), c(t')) \leq \lambda|t - t'| + \epsilon$$

for all $t, t' \in I$. If $I = [a, b]$, then $c(a)$ and $c(b)$ are called the *endpoints* of c. If $I = [0, \infty)$, then c is called a *quasi-geodesic ray.*

We typically abuse notation and call the image of c the quasi-geodesic c. Note that since quasi-isometries are not necessarily continuous, neither are quasi-geodesics. In particular, a quasi-geodesic could just be a discrete set of points. However, it is shown in Bridson–Haefliger [**?**, Lemma III.H.1.11] that in any geodesic metric space, all quasi-geodesics $c\colon [a, b] \to X$ are within Hausdorff distance at most $\lambda + \epsilon$ of a *tame* quasi-geodesic $c'\colon [a, b] \to X$, that, is, a continuous quasi-geodesic with the same endpoints that satisfies the following bound on length: for any $t, t' \in [a, b]$, we have

$$\ell(c'|_{[t,t']}) \leq \lambda d(c'(t), c'(t')) + \epsilon,$$

where $\ell(c'|_{[t,t']})$ denotes the length of the subpath of c from $c'(t)$ to $c'(t')$.

In practice, this length bound is often very useful, and so when we work in geodesic metric spaces, we often consider only tame quasi-geodesics.

(The *Hausdorff distance* between closed subsets A and B of a metric space X is the infimum of all ϵ such that $A \subseteq B^{+\epsilon}$ and $B \subseteq A^{+\epsilon}$. We write $d_{Haus}(A, B)$.)

In general, quasi-geodesics in metric spaces can look nothing like geodesics. For example, the map $[0, \infty) \to \mathbb{R}^2$ given in polar coordinates by $t \mapsto (t, \log(1 + t))$ is a quasi-geodesic ray.

Lemma 2.12. *Let X be a δ–hyperbolic geodesic metric space. Let c be a continuous (rectifiable) path in X. If $[p, q]$ is a geodesic connecting the endpoints of c, then for every point $x \in [p, q]$, there exists a point $y \in c$ such that*

$$d(x, b) \leq \delta|\log_2(\ell(c))| + 1.$$

Proof. Exercise! □

There are two very important properties of quasi-geodesics in hyperbolic metric spaces: a local-to-global property, and the Morse lemma. We begin with the Morse lemma.

Theorem 2.13 (Morse Lemma; Quasi-geoesic stability)**.** *For all $\delta \geq 0$, $\lambda \geq 1$, and $\varepsilon \geq 0$, there exists a constant $R(\delta, \lambda, \varepsilon) \geq 0$ such that the following holds. If X is a δ–hyperbolic metric space, c is a (λ, ε)–quasi-geodesic in X, and $[p, q]$ is a geodesic connecting the endpoints of c, then*

$$d_{Haus}(c, [p, q]) \leq R.$$

Before giving the proof, we note several consequences of the Morse lemma.

Corollary 2.14. *For every $\lambda \geq 1$, $\varepsilon \geq 0$, and $\delta \geq 0$, there is a constant $M = M(\delta, \lambda, \varepsilon)$ such that a geodesic metric space X is δ–hyperbolic if and only if every (λ, ε)–quasigeodesic triangle in X is M–slim.*

Corollary 2.15. *Let X, X' be geodesic metric spaces and $f\colon X \to X'$ be a (λ, ε)–quasi-isometric embedding. If X' is δ–hyperbolic, then X is δ'–hyperbolic for some $\delta' = \delta'(\delta, \lambda, \varepsilon)$.*

Exercises 2.16. *Give careful proofs of both corollaries.*

A special case of the above corollary says that if $X \sim_{qi} X'$ and X is δ–hyperbolic, then X' is δ'–hyperbolic. This is what we mean by hyperbolicity being a quasi-isometry invariant.

We now prove the Morse Lemma.

Proof of the Morse Lemma. We may assume c is tame. Let $D = \sup\{d(x, c) \mid x \in [p, q]\}$, and let $x_0 \in [p, q]$ be the point where the supremum is attained (the point x_0 exists because the interval $[p, q]$ is compact). Then the open ball of radius D about x_0 is disjoint from c.

Let $y \in [p, x_0]$ be such that $d(y, x_0) = 2D$, and similarly for $z \in [x_0, q]$ (if such points do not exist, we take $y = p$ or $z = q$). Fix $y', z' \in c$ with $d(y, y') \leq D$ and $d(z, z') \leq D$. Choose geodesics $[y, y']$ and $[z, z']$, and let γ be the path from y to z formed by concatenating $[y, y']$, the subpath of c from y' to z', and $[z', z]$. Then by construction, γ lies outside of $B(x_0, D)$.

We have

$$\ell(\gamma) \leq 2D + \ell(c|_{[y',z']}) \leq 2D + \lambda d(y', z') + \varepsilon \leq 2D + \lambda \cdot 6D + \varepsilon.$$

Since $d(x_0, \gamma) = D$, we have $D \leq \delta |\log_2(\ell(\gamma))| + 1$, and so

$$D - 1 \leq \delta \log(2D + 6D\lambda + \varepsilon) + 1. \tag{3}$$

Since the left-hand side is linear in D while the right-hand side is logarithmic in D, (3) gives an upper bound on D depending only on δ, λ, and ε. Fix such an upper bound D_0, so that $[p, q] \subseteq \mathcal{N}_{D_0}(c)$.

We will now show that $c \subseteq \mathcal{N}_{R'}([p, q])$, where $R' = D_0(1 + \lambda) + \varepsilon/2$.

Consider a maximal subpath c' of c which lies outside the D_0–neighborhood of $[p, q]$. The closed D–neighborhood of $c|_{[p,c'_-]}$ and of $c|_{[c'_+,q]}$ are both closed sets which collectively cover $[p, q]$. Thus there exists a point w in the intersection of the two closed neighborhoods and $[p, q]$. In other words, there exists $w \in [p, q]$, $t \in c|_{[p,c'_-]}$, and $t' \in c|_{[c'_+,q]}$ such that $d(w, t) \leq D_0$ and $d(w, t') \leq D_0$. In particular, $d(t, t') \leq 2D_0$. Note that c' is a subpath of $c|_{[t,t']}$, and since c is a quasigeodesic, we have $\ell(c') \leq \ell(c|_{[t,t']}) \leq \lambda \cdot 2D_0 + \varepsilon$. Therefore starting from any point on c', we can follow c and reach either t or t' in a distance of at most $\lambda D_0 + \varepsilon/2$. From there we can reach $w \in [p, q]$ in an additional distance of D_0, and so we conclude that $c' \subseteq \mathcal{N}_{R'}([p, q])$. Applying this same argument to all such maximal subpaths, we conclude that $c \subseteq \mathcal{N}_{R'}([p, q])$. □

Definition 2.17. A finitely generated group is *hyperbolic* if there exists a finite generating set X such that $\Gamma(G, S)$ is hyperbolic.

Note that by Corollary 2.15, the "there exists" in the above definition is equivalent to "for all". Moreover, by the Milnor–Schwartz Lemma (Theorem 1.12), an equivalent definition is: A finitely generated group is *hyperbolic* if it admits a proper, cobounded action on a proper, geodesic, hyperbolic metric space.

Examples 2.18. The following groups are hyperbolic.

1. Finite groups
2. $\mathbb{F}_n$ for all $n \geq 1$
3. If M is a closed hyperbolic manifold, then $\pi_1(M)$ is hyperbolic. In particular, if S is an orientable surface of genus g, then $\pi_1(S)$ is hyperbolic if and only if $g \geq 2$.
4. $\mathbb{Z}^n$ is hyperbolic if and only if $n = 1$.

Recall that a finitely generated group is quasi-isometric to any finite-index subgroup. This motivates the following definition of "virtually".

Definition 2.19. Let P be a property of groups. Then we say G is *vitually* P if there is a finite index subgroup $G_0 \leq G$ such that G_0 is P.

For example, finite groups are virtually trivial!

Example 2.20. A virtually hyperbolic group is hyperbolic. For example, the group $PSL_2(\mathbb{Z}) = (\mathbb{Z}/3\mathbb{Z}) * (\mathbb{Z}/2\mathbb{Z})$ is virtually free and therefore hyperbolic.

Recall that we showed (as an exercise assigned in class) that $\Gamma(BS(1,2),\{a,t\})$ is not δ–hyperbolic for any δ. Therefore, the Baumslag-Solitar group $BS(1,2)$ is *not* a hyperbolic group. Similarly, $BS(1,n)$ is not a hyperbolic group.

We next describe a huge class of examples of hyperbolic groups. To do so, we need to first formally define a group presentation.

Given a set S, we denote by $F(S)$ the free group on S. Elements of $F(S)$ are equivalences classes of words over the alphabet $S^{\pm 1}$: a *word* is a finite sequence $a_1 \dots a_n$ where $a_i \in A^{\pm 1}$. Two words are *equivalent* if you can pass from one to the other by inserting or deleting words of the type ss^{-1}; for example, $s_1 s_2$ is equivalent to $s_1 s_1 s_1^{-1} s_2$. A words $a_1 \dots a_n$ is *reduced* if $a_i \neq a_{i-1}^{-1}$ for all i. There is a unique reduced word in each equivalence class. If $S = \{s_1, \dots, s_n\}$, we often write $\mathbb{F}_n$ (or $\mathbf{F}_n$ or F_n).

Let G be a group and R a subset of G. The *normal closure* $\langle\langle R\rangle\rangle$ of R in G is the smallest normal subgroup of G containing R. Equivalently,

$$\langle\langle R\rangle\rangle = \left\{ \prod_{i=1}^{k} g_i r_i g_i^{-1} \;\middle|\; k \geq 0,\, g_i \in G,\, r_i \in R^{\pm 1} \right\}.$$

Given a set S and $R \subseteq F(A)$, we say that

$$\langle S \mid R \rangle$$

is a *presentation* for the group $G \simeq F(S)/\langle\langle R\rangle\rangle$. The set S is the set of *generators* of G and R is the set of *relations*. The presentation is *finite* if both S and R are finite sets, and G is said to be *finitely presented* if it admits such a presentation.

Example 2.21. A "random" group is hyperbolic. To be precise: fix $m \geq 2$ and $n \geq 1$. We will temporarily fix $\ell \geq 0$. Consider groups of the form

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

such that r_i are all *cyclically reduced* words of length ℓ (i.e., reduced words such that the first letter is not the inverse of the last letter). Put the uniform probability distribution of the set of all such groups. This defines a group-valued random variable Γ_ℓ. For a property P, we say that a *random group is* P if

$$\lim_{\ell \to \infty} \mathbb{P}(\Gamma_\ell \text{ is } P) = 1$$

for all n, m. This is called the *few-relator model* of random groups.

Theorem 2.22 (Gromov). *A random group is infinite and hyperbolic.*

This is not the only way to define a "random group" (we call the different definitions different *models* of random groups). Generally speaking, any way you define a random group so that you don't always end up with finite groups should yield a hyperbolic group.

A hyperbolic group is *non-elementary* if its Cayley graph with respect to some (equivalently, any) finite generating set is infinite and not quasi-isometric to $\mathbb{R}$. (We call spaces quasi-isometric to $\mathbb{R}$ *quasi-lines*.)

We turn our attention to the local-to-global property. Recall that in differential geometry, geodesics are defined locally. Our geodesics, on the other hand, are a global notion. The following gives an analogous version of locally defined geodesics.

Definition 2.23. Let X be a metric space. A path $p\colon [a,b] \to X$ is a *k–local geodesic* if $d(p(t), p(t')) = |t-t'|$ for all $t, t' \in [a,b]$ with $|t - t'| \leq k$.

In other words, every subpath of p of length at most k is a geodesic.

On a general Riemannian manifold, local geodesics are not necessarily global geodesics, since, for example, on a sphere we can wrap around many times. However, that is not the case for hyperbolic spaces.

Lemma 2.24 (Local-to-global property; k–local geodesics are quasigeodesics)**.** *Let X be a δ–hyperbolic geodesic metric space, and $p\colon [a,b] \to X$ a k–local geodesic with $k > 8\delta$. Then:*

(1) $p \subseteq \mathcal{N}_{2\delta}([p_-, p_+])$;

(2) $[x,y] \subseteq \mathcal{N}_{3\delta}(p)$; and

(3) p is a (λ, ε)–quasigeodesic where $\varepsilon = 2\delta$ and $\lambda = \frac{k+4\delta}{k-4\delta}$.

We'll skip most of the proof and only prove (1); see Bridson–Haefliger [**?**, III.H.1.13] for the proofs of the other parts.

Proof of Lemma 2.24(1). Let $x = p(t)$ maximize the distance from p to $[p_-, p_+]$. We will find a point on $[p_- p_+]$ at distance at most 2δ from x.

Choose points $y = p(t - k/2)$ and $z = p(t + k/2)$. If $t - k/2 < a$, we let $y = p_-$, and similarly for z. Let y' and z' be the closest points on $[p_-, p_+]$ to y and z, respectively. Fix geodesics $[y', y]$ and $[z', z]$. Concatenating these geodesics with the subpath of $[p_-, p_+]$ between y' and z' (which we call $[y', z']$) and the subpath of p between y and z forms a *geodesic* quadrilateral. Note that this subpath of p is a geodesic because $|t + k/2 - (t - k/2)| \leq k$ and p is a k–local geodesic.

Since X is δ–hyperbolic, there exists a point w on $[y', y] \cup [y', z'] \cup [z', z]$ such that $d(x, w) \leq 2\delta$. If w is in $[y', z']$ then we are done, so assume without loss of generality that $w \in [y', y]$.

If $y = p_-$, then we must have $y' = y$, and thus $w = y \in [p_-, p_+]$ and we are done. Thus we can assume that $y = p(t - k/2)$. We have $d(y, x) = k/2 > 4\delta$ (where the equality follows from the fact that p restricts to a geodesic from y to x). By the triangle inequality, we must have $d(y, w) > 2\delta$, and so $d(y, w) > d(x, w)$. However,

$$d(x, y') \leq d(x, w) + d(w, y') < d(y, w) + d(w, y') = d(y, y').$$

Therefore,

$$d(x, [p_-, p_+]) < d(x, y') \leq d(y, y') = d(y, [p_-, p_+]),$$

which contradicts our choice of x. □

This local-to-global property has some very powerful consequences, some of which we will see when we discuss decision problems in Section 3. One straightforward corollary is the following:

Corollary 2.25. *If X is a δ–hyperbolic geodesic metric space and p a k–local geodesic for some $k > 8\delta$, then either p is constant or $p_- \neq p_+$.*

Exercises 2.26. *Prove the corollary.*

3 Decision Problems

Max Dehn formulated three classic decision problems in 1912: the word problem, the conjugacy problem, and the isomorphism problem. These are Dehn's original formulations of the problems. In each, the group is given by a finite presentation.

The Word Problem: An element of the group is given as a product of generators. One is required to give a method whereby it may be decided in a finite number of steps whether this element is the identity or not.

The Conjugacy Problem: Any two elements s and t of the group are given. A method is sought for deciding the question of whether s and t can be transformed into each other, i.e., whether there is an element u of the group satisfying

$$s = utu^{-1}.$$

The Isomorphism Problem: Given two groups, one is to decide whether they are isomorphic or not (and further, whether a given correspondence between the generators of one group and the elements of the other is an isomorphism or not).

3.1 Group presentations and Dehn functions

We'll always assume our set of relators are symmetric, so that $r^{-1} \in R$ whenever $r \in R$; this is solely for convenience.

Suppose $\langle S \mid R\rangle$ is a presentation for a group G and W is a word in S (by this we mean that $W \in F(S)$). Then $W =_G 1$ if and only if there exist $r_1, \dots, r_k \in R$ and $f_1, \dots, f_k \in F(S)$ such that

$$W =_{F(S)} f_1^{-1} r_1 f_1 \cdots f_k^{-1} r_k f_k. \tag{4}$$

In other words, W is represents the trivial word in G if and only if it differs from the form in (4) by the insertion and or deletion of pairs of the form ss^{-1} where $s \in S$. Note that (4) does not say that $W \equiv f_1^{-1} r_1 f_1 \cdots f_k^{-1} r_k f_k$, just that the two sides of the equation represent the same word in the free group $F(S)$.

We will encode this geometrically. Let Δ be a finite, connected, simply connected, planar 2–complex in which every edge is oriented and labeled by an element of S. If e is an edge of Δ with label s and $\bar{e}$ is the same edge with the opposite orientation, then the label of $\bar{e}$ is s^{-1}. Labels are paths in Δ are defined in the same way as in Cayley graphs. If Π is a 2–cell of Δ, then $\mathbf{Lab}(\partial\Pi)$ is the word obtained by choosing some point $v \in \partial\Pi$ and reading the label of the path $\partial\Pi$ starting and ending at v, traversed counterclockwise. A different choice of basepoint results in a cyclic permutation of the word $\mathbf{Lab}(\partial\Pi)$, so we only consider this word to be defined up to cyclic permutation. We define $\mathbf{Lab}(\partial\Delta)$ similarly. The complex Δ is called a *van Kampen diagram* over $\langle S \mid R\rangle$ if for every 2–cell Π of Δ, (a cyclic permutation of) $\mathbf{Lab}(\partial\Pi)$ belongs to R. By induction on the number of 2–cells in Δ, it can be shown that $\mathbf{Lab}(\partial\Delta) =_G 1$ (exercise!). The 2–cells of Δ are called *faces*.

Lemma 3.1 (Van Kampen Lemma). *Suppose $\langle S \mid R\rangle$ is a presentation for a group G and W is a word in S. Then $W =_G 1$ if and only if there exists a van Kampen diagram Δ over $\langle S \mid R\rangle$ such that* $\mathbf{Lab}(\partial\Delta) \equiv W$.

Proof. We will prove the forward direction. The backward direction is left as an exercise. Suppose $W =_G 1$. Then there exist $r_1, \dots, r_k \in R$ and $f_1, \dots, f_k \in F(S)$ such that

$$W =_{F(S)} f_1^{-1} r_1 f_1 \cdots f_k^{-1} r_k f_k.$$

We build a van Kampen diagram out of the following pieces: for each $f_i^{-1} r_i f_i$, build a "lollipop" consisting of an edge labeled by f_i and a loop labeled by r_i so that the vertex of the loop is identified with the initial vertex of the edge labeled by f_i. Now, glue these lollipops together by identifying all the terminal vertices the edges labeled by f_i to form a "wedge of lollipops." Call this space Δ'.

However, since (4) is not a letter-for-letter equivalence, it may be the case that $\mathbf{Lab}(\partial\Delta') \not\equiv W$. By construction, $\mathbf{Lab}(\partial\Delta') \equiv f_1^{-1} r_1 f_1 \cdots f_k^{-1} r_k f_k$, and this word and W are equal *as words in* $F(S)$. Recall that this means that this word and W differ by a finite sequence of moves consisting of adding and deleting subwords of the form ss^{-1}. We use this sequence of moves to modify Δ' in the following way.

Say $f_1 = f_1's$ and $f_2 = f_2's$ for some $s \in S$. Then $f_1 f_2^{-1} = f_1' s s^{-1} (f_2')^{-1}$. Then we may replace the pair of edges labeled $f_1 f_2$ in Δ' with the tripod whose three legs have labels f_1', s, and f_2'. In other words, at the vertex between f_1 and f_2, there are two incoming edges with the label s, and we replace these two incoming edges with a single edge labeled by s. (Note that in the path $f_1 f_2$, one of these edges is being traversed in the opposite direction, which is where the s_1^{-1} comes from.) You can think of this as "folding" the two edges labeled by s and the edge labeled by s^{-1} together into a single edge labeled by s. This gives a new van Kampen diagram**, and after performing this operation a finite number of times, this procedure results in a van Kampen diagram Δ with $\mathbf{Lab}(\partial\Delta) \equiv W$.

**This is not quite true. Van Kampen stated: "the two 1–cells can be brought into coincidence by a deformation without any other changes in the complex." However, there is some subtlety here. The key point is that a van Kampen diagram is a *planar* diagram, and one can imagine folding edges in this way may result in the introduction of spheres. It is possible to get around this issue, but, in certain situations, one must modify the complex Δ' by more than just folding edges together – in particular, any spheres created must be removed. This can be done in such a way that the number of faces does not increase, the boundary label of each remaining face is either unchanged or changed only by a free reduction (of the type ss^{-1}) and the boundary label of Δ' is unchanged or changed only by a free reduction. We'll draw some pictures of this in class, and the following notes by Hamish Short have a good, detailed description of the process (see page 13): `https://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.210.3797&rep=rep1&type=pdf`. □

Exercises 3.2. *Suppose Δ is a van Kampen diagram over $\langle S \mid R\rangle$ and p is a closed (combinatorial) path in Δ. Prove that* $\mathbf{Lab}(p) =_G 1$.

From this exercise, we see that if you fix a vertex $v \in \Delta$, there is a well-defined, label-preserving map from the 1–skeleton of Δ to $\Gamma(G, S)$ which sends v to 1.

Example 3.3 (From Short's notes, p.10). Consider the group $G = \langle x, y \mid xyx^{-1}y^{-1}\rangle$ and the word $W = x^3yx^{-1}y^{-1}x^{-2}y^2x^{-2}y^{-2}x^2$. Figure 2 shows a van Kampen diagram over $\langle a, b, c \mid [a, b]\rangle$ whose boundary has label W, and the same diagram for W deconstructed as $x^2rx^{-2} \cdot yx^{-1}rxy^{-1} \cdot yx^{-2}rx^2y^{-1} \cdot x^{-1}rx \cdot x^{-2}rx^2$, where $= xyx^1y^{-1}$.

Definition 3.4. Given a van Kampen diagram Δ, let $\mathrm{Area}(\Delta)$ be the number of 2–cells of Δ. Fixing $G = \langle S \mid R\rangle$ and a word W in S, let

$$\mathrm{Area}(W) = \min\{\mathrm{Area}(\Delta) \mid \Delta \text{ is a van Kampen diagram over } \langle S \mid R\rangle \text{ and } \mathbf{Lab}(\partial\Delta) \equiv W\}.$$

Equivalently, $\mathrm{Area}(W)$ is the minimum k in the product (4). The *Dehn function* of a finite presentation $\mathcal{P} = \langle S \mid R\rangle$ is the function $\delta_G \colon \mathbb{N} \to \mathbb{N}$ defined by

$$\delta_{\mathcal{P}}(n) = \max\{\mathrm{Area}(W) \mid \|W\| = n,\, W =_G 1\}.$$

So far, we have only defined the Dehn function of a presentation. We would like to define the Dehn function of a group. To do so, we first introduce the following equivalence.

Definition 3.5. If $g, h \colon \mathbb{N} \to [0, \infty)$, write $g \preceq h$ if there exists K such that for all n,

$$g(n) \leq Kh(Kn + K) + Kn + K.$$

We say g and h are *equivalent*, and write $g \simeq h$, if $g \preceq h$ and $h \preceq g$.

While, in general, the Dehn function of a group is not well-defined because it can depend on the choice of presentation, we will see that in fact all such Dehn functions are equivalent in the above sense. Note that in the above definition of the equivalence $\simeq$, the linear term Kn is necessary: If $G = \{1\}$ is the trivial group, then $G = \langle s \mid s\rangle$, and $\mathrm{Area}(s^n) = n$.

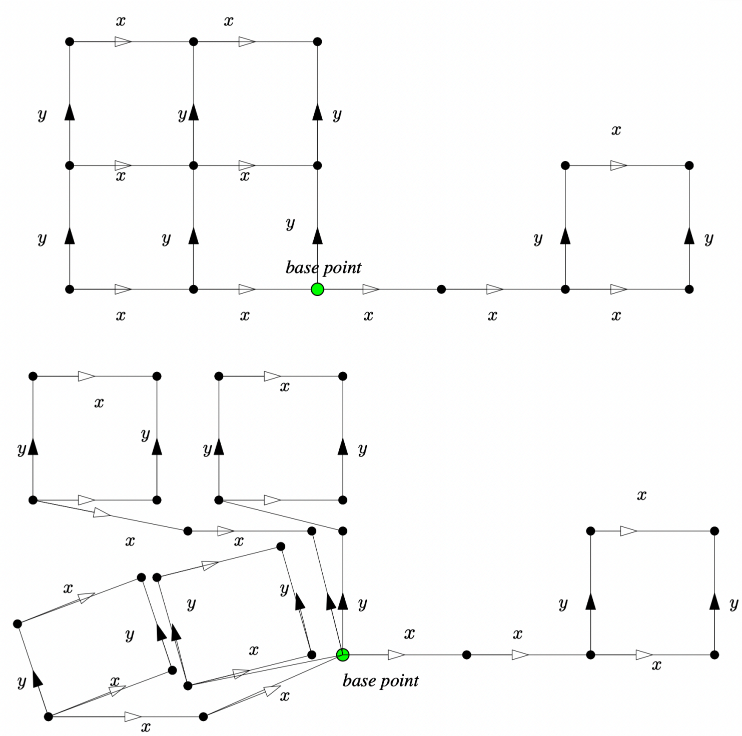

Figure 2: Example 3.3.

Exercises 3.6. *1. Show that $\simeq$ (introduced in class) is an equivalence relation on functions.*

2. Show that $f_1(n) = 1$, $f_2(n) = \log n$, and $f_3(n) = n$ are all equivalent.

3. Show that two polynomials of degree at least 2 are equivalent if and only if they have the same degree.

4. Show that 2^n and 3^n are equivalent.

Lemma 3.7. *If $G = \langle S \mid R\rangle$ and H is generated by a finite set T such that $\Gamma(G,S)$ is quasi-isometric to $\Gamma(H,T)$, then there is a finite presentation $H = \langle T \mid R'\rangle$ such that the Dehn functions with respect to $\langle S \mid R\rangle$ and $\langle T \mid R'\rangle$ are equivalent.*

Proof. Let $\delta_{\mathcal{P}}$ be the Dehn function associated to the given presentation of G. Let $M = \max\{\|r\| \mid r \in R\}$, and let $f\colon \Gamma(H,T) \to \Gamma(G,S)$ be a (λ, c)–quasi-isometry. Let p be a closed (combinatorial) path in $\Gamma(H,T)$ with vertices $v_1, v_2, \ldots, v_n, v_{n+1} = v_2$. Let q be the closed path in $\Gamma(G,S)$ formed by concatenating geodesics $f(v_i)$ to $f(v_{i+1})$. Since f is a (λ, c)–quasi-isometry, each such geodesic has length at most $\lambda + c$, and so $\ell(q) \leq (\lambda + c)n$. Since q is a closed path, we have $\mathbf{Lab}(q) =_G 1$, and so there exists a van Kampen diagram Δ with $\mathbf{Lab}(\partial\Delta) = \mathbf{Lab}(q)$. Choose Δ so that $\mathrm{Area}(\Delta) \leq \delta_{\mathcal{P}}((\lambda + c)n)$.

We identify the 1–skeleton $\Delta^{(1)}$ of Δ with its image in $\Gamma(G,S)$ under the natural map which sends $\partial\Delta$ to q.

We will build a map $g\colon \Delta^{(1)} \to \Gamma(H,T)$ as follows. For each interior vertex $v \in \Delta^{(1)}$, choose a vertex $u \in \Gamma(H,T)$ such that $d_S(f(u), v) \leq c$. Define $g(v) = u$. For each exterior vertex $v \in \Delta^{(1)}$, v lies on a geodesic $f(v_i)$ to $f(v_{i+1})$. If v is closer to $f(v_i)$, define $g(v) = v_i$, and otherwise define $g(v) = v_{i+1}$. Whenever $u, v \in \Delta^{(1)}$ are connected by an edge, join $g(u)$ and $g(v)$ be a geodesic in $\Gamma(H,T)$. Notice that for such $u, v \in \Delta^{(1)}$, we have $d_S(f(g(u)), f(g(v)) \leq c + 1 + (\lambda + c)/2$ (if exactly one of u, v is exterior) and $d_S(f(g(u)), f(g(v))) \leq 2c + 1$ (if both u, v are interior). Since f is a (λ, c)–quasi-isometry, this implies that in these two cases

$$d_T(g(u), g(v)) \leq \lambda(3c + 1 + (\lambda + c)/2). \tag{5}$$

If u and v are both exterior vertices that are adjacent, then $d_T(g(u), g(v)) \leq 2$ and so (5) still holds.

Therefore, if Π is a face of Δ, then there is a closed loop in $g(\Delta^{(1)}) \subseteq \Gamma(H, T)$ corresponding to the image of $\partial\Pi$ with length at most

$$\lambda(3c + 1 + (\lambda + c)/2)\ell(\partial\Pi) \leq \lambda(3c + 1 + (\lambda + c)/2)M.$$

Let $R' = \{r \in F(T) \mid \|r\| \leq \lambda(3c + 1 + (\lambda + c)/2)M \text{ and } r =_H 1\}$. Then $|R'| < \infty$, since T is a finite set. From the above discussion, we have that there exists a van Kampen diagram Δ' whose 1–skeleton is $g(\Delta^{(1)})$, and each 2–cell is labeled by an element of R'. Thus Δ' is a van Kampen diagram over $\langle T \mid R' \rangle$ and $\mathbf{Lab}(\partial\Delta') \equiv W$. Therefore $\mathcal{Q} = \langle T \mid R' \rangle$ is a finite presentation for H, and in particular H is finitely presented.

Furthermore,

$$\mathrm{Area}(W) \leq \mathrm{Area}(\Delta') = \mathrm{Area}(\Delta) \leq \delta_{\mathcal{P}}((\lambda + c)n).$$

Since W is an arbitrary trivial word of length n, we see that

$$\delta_{\mathcal{Q}}(n) \leq \delta_{\mathcal{P}}((\lambda + c)n).$$

Reversing the roles of G and H in the proof gives the other inequality (with possibly different constants). Therefore $\delta_H \simeq \delta_G$. □

The following corollaries are immediate from the above theorem.

Corollary 3.8. *The property of being finitely presented is a quasi-isometry invariant among finitely generated groups.*

Corollary 3.9. *If $\mathcal{P}$ and $\mathcal{Q}$ are two finite presentations for a group G, then $\delta_{\mathcal{P}} \simeq \delta_{\mathcal{Q}}$.*

In light of this, we define δ_G, which we call the *Dehn function of* G, to be $\delta_{\mathcal{P}}$ for any presentation $\mathcal{P}$ of G, and understand that this function is only well-defined up to $\simeq$. In particular, if $\delta_{\mathcal{P}}$ is bounded above by a linear (quadratic, exponential, etc.) function for some finite presentation of G, then the same is true for all finite presentations of G.

Exercises 3.10. *Prove that a finite group has at most linear Dehn function.*

Exercises 3.11. *Prove that a finitely generated abelian group has at most quadratic Dehn function, and that the Dehn function of $\mathbb{Z}^2$ is equivalent to n^2.*

Definition 3.12. If a group G has a presentation $\mathcal{P}$ such that $\delta_{\mathcal{P}}$ is bounded above by a linear (quadratic, exponential, etc.) Dehn function, we say G *satisfies a linear (quadratic, exponential, etc.) isoperimetric inequality.*

3.2 The word problem

We will first focus on the word problem. The Dehn function of a presentation encodes whether or not the word problem is solvable.

Theorem 3.13. *The word problem for a presentation $\mathcal{P}$ is solvable if and only if the Dehn function $\delta_{\mathcal{P}}$ is computable.*

Proof. For the backwards direction, suppose $\delta_{\mathcal{P}}$ is computable. Note that if $W =_G 1$ with associated van Kampen diagram Δ, then we can write W as

$$W =_{F(S)} f_1^{-1} r_1 f_1 \cdots f_k^{-1} r_k f_k,$$

where each f_i has length at most $\|W\|$. (Why?) Since S is finite, there are only finitely many words of length $\|W\|$. Thus, there are only finitely many words of the form $f_i r_i f_i^{-1}$. Since $\delta_{\mathcal{P}}$ is computable, the area of Δ is computable, and so W is a product of at most $\delta_{\mathcal{P}}(\|W\|)$ words of the form $f_i r_i f_i^{-1}$. There are (computably) finitely many such products, and hence, by trying all these combinations, we can determine if $W =_G 1$.

For the forwards direction, assume that the word problem is solvable for the presentation $\mathcal{P}$. For a given value of n, we can thus find all words of length at most n which are equal to the identity in G. Then for each such word, we calculate its area, and this allows us to compute the Dehn function of $\mathcal{P}$. $\square$

The following is a difficult theorem:

Theorem 3.14 (Novikov–Boone Theorem, 1955 & 1958)**.** *There is finite presentation for which the word problem is not solvable.*

(To see an explicit presentation of such a group, in fact, a group with 2 generators, see pg. 5 of `https://eprint.iacr.org/2014/528.pdf`.)

Given a group presentation $\langle S \mid R \rangle$ and word W in S, *Dehn's agorithm* is the following procedure. First, freely reduce W (that is, delete all pairs ss^{-1} appearing in W). If this produces the empty word, the algorithm stops. Otherwise, W is freely reduced and non-empty. In this case, search W for subwords U such that U is also a subword of a (cyclic shift of a) relator $r \in R$ and $\|U\| > \frac{1}{2}\|r\|$. If no such subword exists, the algorithm stops. If such a U exists, then there is a (possibly empty) word V (the complement of U in r) such that $UV^{-1} =_G 1$ and $\|V\| < \|U\|$. In this case, the algorithm replaces U with V and repeats.

If the presentation $\langle S \mid R \rangle$ is finite, then Dehn's algorithm terminates after finitely many steps and outputs either a non-trivial word or the trivial word.

Definition 3.15. Let $\langle S \mid R \rangle$ be a finite presentation for a group G. The presentation is a *Dehn presentation* if the following hold:

1. There is a set of words $u_1, \ldots, u_k, v_1, \ldots, v_k$ in $F(S)$, and the set of relators are of the form $r_i = u_i v_i^{-1}$.
2. For each i, $\|v_i\| \leq \|u_i\|$
3. For any nonempty reduced word $w \in F(S)$, with $w =_G 1$, at least one of u_i or u_i^{-1} must appear as a subword.

Note that this third condition is very strong, and implies, in particular that applying the relation r_i (which amounts to replacing the subword u_i with v_i) is guaranteed to shorten the word. In general, given a presentation and a word that represents the identity in the group, it may be the case that in order to find a shorter word representing the same group element, one must first apply a relation that makes the word *longer*, and only then can one find a relation which will make the word shorter. This third condition also guarantees that if $\langle S \mid R \rangle$ is a Dehn presentation, then Dehn's algorithm described above is a solution to the word problem. In other words, given a word W, if Dehn's algorithm will return the trivial word exactly when $W =_G 1$.

Exercises 3.16. *Find a group presentation that does not have a Dehn presentation.*

Exercises 3.17. *Suppose $\langle S \mid R \rangle$ is a Dehn presentation for a group G. Prove that G satisfies a linear isoperimetric inequality.*

Theorem 3.18. *For any finitely generated group G, the following are equivalent:*

(1) G is hyperbolic;

(2) G has a Dehn presentation $\langle S \mid R\rangle$;

(3) G is finitely presented and satisfies a linear isoperimetric inequality.

Proof. Let G be generated by a finite set S, and suppose G is δ–hyperbolic.

(1) $\implies$ (2): First, fix any $K > 8\delta$, and consider the finite set of freely reduced words in $F(S)$ of length at most K. Now, check which of the words on the list represent the same group element in G (we can do this because there are only finitely many words). Let u_i be the non-geodesic spelling of this word and let v_i be a geodesic spelling of this same word, so that $\|v_i\| < \|u_i\|$. Let

$$R = \{u_i v_i^{-1}\}.$$

We claim that $\langle S \mid R\rangle$ is a Dehn presentation for G. The first two items are clear by construction. For the third, let $W \in F(S)$ be a non-empty word with $\|W\| =_G 1$. Let p be a path labeled by W in $\Gamma(G,S)$ with $p_- = p_+ = 1$. If $\ell(p) > 8\delta$, then p is not an 8δ–local geodesic by Corollary 2.25. Therefore, there exists a subpath q of p of length at most 8δ which is not a geodesic. Such a subpath q must have label u_i or u_i^{-1} for some i. On the other hand, if $\ell(p) \leq 8\delta$, then since p is not a geodesic, $\mathbf{Lab}(p) = u_i$ for some i. In either case, the third condition is satisfied.

(2) $\implies$ (3): This is Exercise 3.17

In the interest of time, we will not prove that (3) $\implies$ (1). A proof can be found in Bridson–Haefliger III.H.2.9. ☐

In fact, if a group G satisfies a subquadratic isoperimetric inequality, then G satisfies a *linear* isoperimetric inequality and is therefore hyperbolic.

Historical note: Dehn proved Fuchsian groups admit Dehn presentations in 1912. Cannon extended this to all fundamental groups of closed, negatively curved manifolds in 1984. The equivalence of (1) and (2) is due to Gromov.

Corollary 3.19. *Let G be a finitely presented group which is not hyperbolic. Then the Dehn function of G is at least quadratic.*

It is worthwhile to note that there are other, non-hyperbolic groups with fast solutions to the word problem. For example, if G is a free abelian group, one can check if a word is trivial simply by checking that the exponent sum on each generator is zero.

There are several interesting results in the literature that relate solvability of the word problem to certain algebraic properties of the group. These are outside the scope of this course, but here is a sample of such a result:

Theorem 3.20 (Boone–Higman Theorem, 1977)**.** *A finitely presented group has solvable word problem if and only if it can be embedded in a simple subgroup of a finitely presented group.*

3.3 The conjugacy problem

The goal of this subsection is to prove the following theorem.

Theorem 3.21. *The conjugacy problem is solvable for hyperbolic groups.*

Before doing so, let's consider the more straightforward case of a free group $\mathbb{F}_n$. In a free group, two cyclically reduced words are conjugate if and only if they are cyclic permutations of each other. In light of this, here is an algorithm to determine if two words $u, v \in \mathbb{F}_n$ are conjugate. First, cyclically reduce u and v to form new words u' and v'. This can be done in $\|u\|$ and $\|v\|$ steps, respectively. Now, compare cyclic permutations of u' to v' to see if u' and v' are conjugate. This can be done in $\|u'\| \leq \|u\|$ steps. Since cyclic reduction corresponds to conjugation, u is conjugate to u' and v is conjugate to v'. Thus u and v are conjugate if and only if u' and v' are conjugate, which can be checked in $2\|u\| + \|v\|$ time. This is a linear time solution to the conjugacy problem, where here we mean linear in $\|u\| + \|v\|$ with constants independent of the choice of u and v.

Let $\langle S \mid R\rangle$ be a presentation for a hyperbolic group G, and let δ be such that G is δ–hyperbolic. The following lemma states that for any pair of conjugate elements in G, there is a conjugator whose length is bounded by a linear function of $\|u\|$ and $\|v\|$ (with constants independent of the choice of u and v).

Lemma 3.22. *If $u, v \in F(S)$ are conjugate in G, then there exists a conjugator $x \in F(S)$ (i.e., a word so that $xux^{-1} = v$) satisfying $\|x\| \leq \|u\| + \|v\| + |S^{4\delta}| + 1$, where $S^{4\delta}$ is the set of words in $F(S)$ of length at most 4δ.*

Proof. Suppose towards a contradiction that x is a conjugator of u and v of minimal length and that $\|x\| > \|u\| + \|v\| + |S^{4\delta}| + 1$. Consider the geodesic rectangle in $\Gamma(G, S)$ with sides labeled by $xux^{-1}v^{-1}$. We will call the sides labeled by x to be the top and bottom of the rectangle, while the sides labeled by u and v are the left and right sides, respectively. Label the vertices on the sides labeled by x $p_1, \dots, p_k$ and $q_1, \dots, q_k$, so that $k = \|x\|$ and p_i is on the top side while q_i is on the bottom side.

Since geodesic quadrilaterals in a δ–hyperbolic space are 2δ–thin, for each p_i with $\|u\| \leq i \leq k - \|v\|$, there is a point q_j on the bottom such that $d(p_i, q_j) \leq 2\delta$. We will show that in fact $d(p_i, q_i) \leq 4\delta$. Suppose p_i divides the top side of the rectangle into paths $\sigma_1\sigma_2$ and q_j divides the bottom of the rectangle into paths $\tau_1\tau_2$. Then if $d(q_j, q_i) > 2\delta$, then either the path $\sigma_1[p_i, q_j]\tau_2$ or the path $\tau_1[q_jp_i]\sigma_2$ is a path from p to q whose length is shorter than the path labeled by x. Let x' be the label of this shorter path. By cutting the quadrilateral along the path labeled by x' and gluing the two sides labeled by x, we see that x' conjugates u to v. But $\|x'\| < \|x\|$, which contradicts our choice of x. Therefore $d(q_j, q_i) \leq 2\delta$, and so $d(p_i, q_i) \leq 4\delta$.

Thus for each $\|u\| \leq i \leq k - \|v\|$, we have $d(p_i, q_i) \leq 4\delta$. Since $k > \|u\| + \|v\| + |S^{4\delta}| + 1$, there are at least $|S^{4\delta} + 1|$ paths of length at most 4δ connecting points p_i on the top of the quadrilateral to points q_i on the bottom of the quadrilateral. Since there are fewer than $|S^{4\delta} + 1|$ words of length 4δ in G, there must be two such paths, say $[p_i, q_i]$ and $[p_j, q_j]$ with $i \neq j$ which have the same label in G.

By construction, p_i and p_j and q_i and q_j subdivide the two sides labeled by x into $x_1x_2x_3$ (where the subsegments have the same label on both the top and bottom side). Cut the diagram along $[p_i, q_i]$ and $[p_j, q_j]$, remove the center quadrilateral, and glue $[p_i, q_i]$ to $[p_j, q_j]$. This results in a new quadrilateral whose top and bottom sides both have the label x_1x_3. Therefore x_1x_3 conjugates u to v and $\|x_1x_3\| < \|x\|$. This contradicts our choice of x.

Therefore, we conclude that $\|x\| \leq \|u\| + \|v\| + |S^{4\delta}| + 1$, as desired. $\square$

We can now use Lemma 3.22 to give a solution to the conjugacy problem in hyperbolic groups.

Proof of Theorem 3.21. Fix a Dehn presentation $\langle S \mid R\rangle$ for a δ–hyperbolic group G. Fix two elements $u, v \in F(S)$. We need to determine if there exists $g \in F(S)$ such that $gug^{-1} =_G v$. If such an element exists, then there exists a (possibly different) element $g' \in F(S)$ with $\|g'\| \leq \|u\| + \|v\| + |S^{4\delta}| + 1 = K$ by Lemma 3.22. There are only finitely many words in $F(S)$ whose length is at most K, and so we can check them one by one and see if $g'ug'^{-1}$ and v represent the same word in G, using the fact that the word problem is solvable in hyperbolic groups. If the two words are not equal for any such word g', then u and v are not conjugate. Otherwise, they are, and the algorithm also computes the conjugating element. $\square$

The algorithm described above yields an exponential time solution to the conjugacy problem (exponential in $\|u\| + \|v\|$, with constants independent of the choice of u and v), as there are exponentially many words g' for which one needs to check the appropriate relation (and the word problem is solvable in linear time). This is not the only possible algorithm one can write to solve the conjugacy problem in hyperbolic groups. Epstein and Holt show in [**?**, **?**] that it is actually solvable in *linear* time, just like the word problem.

3.4 The Isomorphism Problem

We will not spend time on the isomorphism problem in this class. It was only recently shown that the isomorphism problem is solvable for hyperbolic groups with torsion. Dahmani and Guirardel proved this in 2010. The case of torsion-free hyperbolic groups was proved by Sela in 1995. In fact, Sela only published the result for "rigid" torsion-free hyperbolic groups, but had a proof for the more general case. Dahmani and Groves then simplified Sela's original approach and published a proof for all torsion-free hyperbolic groups (and toral relatively hyperbolic groups, which we may have time to define at the end of the course) [**?**].

Theorem 3.23 ([**?**, **?**, **?**])**.** *The isomorphism problem is solvable for hyperbolic groups.*

The isomorphism is not solvable for all finitely presented groups. In 1955 and 1958, Adyan and Rabin used the groups constructed by Novikov and Boone (as examples of groups with unsolvable word problem) to prove that there does not exist an algorithm to determine if an arbitrary finite presentation defines the trivial group.

However, the isomorphism problem is also unsolvable for several "nice" classes of groups: free-by-free groups [**?**], (free abelian)-by-free groups [**?**], and the class of solvable groups of derived length 3 [**?**]. Moreover, the isomorphism problem is still open for generalized Baumslag-Solitar groups, Coxeter groups, and one-relator groups.

4 The Rips complex and finiteness properties

Recall that hyperbolic groups are finitely presented.

4.1 Rips complex

In this section, given a hyperbolic group G, we will construct a contractible simplicial complex on which G acts properly and cocompactly, called the Rips complex.

Definition 4.1. Given a metric space X and a constant $d \geq 0$, the *Rips complex* K_d is the simplicial complex that has a vertex for every point in X and an n–simplex $\{x_0, x_1, \ldots, x_{n-1}\}$ whenever $d_X(x_i, x_j) \leq d$ for all i, j.

Theorem 4.2. *For any geodesic δ–hyperbolic metric space X and any $d \geq 4\delta$, K_d is contractible.*

Proof. Fix δ so that X is δ–hyperbolic. Let $d \geq 4\delta$. We will show that the Rips complex $K = K_d$ satisfies the three properties.

Since K_d is CW complex, it is contractible if and only if $\pi_n(K_d) = 0$ for all n. In particular, we must show that any continuous map $\mathbb{S}^n \to K_d$ for $n \geq 2$ has image which is homotopic to a point. Since the image of any such map must lie in a finite subcomplex of K_d, it suffices to show that any such subcomplex is contractible. Let L be a finite subcomplex of K_d.

Fix a vertex $x_0 \in X$ (which is also a vertex of K_d, by definition). Let $v \in L$ be the point so that $d_X(x_0, v)$ is maximal. There are two cases to consider. In both cases we will show that L is homotopic to a subcomplex contained in a single simplex of K_d, which will prove that L is contractible.

Case 1. If $d_X(x_0, v) \leq d/2$, then every pair of points $y, z \in L$ satisfy $d_X(y, z) \leq d$. Thus L is a simplex (by definition of K_d), and so is contractible.

Case 2. Suppose $d_X(x_0, v) > d/2$, and let $[x_0, v]$ be a geodesic in X from x_0 to v. Let y be a point on this geodesic with $d_X(y, v) = d/2$. We will show that any simplex σ of L containing v is a face of a larger simplex σ' formed by joining y to σ.

Suppose $u \in L$ satisfies $d_X(u, v) \leq d$ (equivalently, let u be an arbitrary vertex of a simplex σ containing v). We will show that $d_X(u, y) \leq d$, which will imply that σ is a face of a larger simplex containing y. Consider the triangle $[x_0, v] \cup [v, u] \cup [x_0, u]$. Then there is some $w \in [x_0, u] \cup [u, v]$ such that $d_X(w, y) \leq \delta$. Suppose first that $w \in [x_0, u]$. Then

$$d_X(x_0, w) + \delta + d_X(y, v) \geq d_X(x_0, v) \geq d_X(x_0, u) = d(x_0, w) + d_X(w, u),$$

where the second inequality follows from our assumption that $v \in L$ maximizes the distance (in X) from x_0 to points in L. In particular,

$$d_X(w, u) - \delta \leq d_X(y, v) = d/2.$$

Thus

$$d_X(y, u) \leq \delta + d_X(w, u) \leq \delta + d/2 - \delta = d/2,$$

as desired.

Now suppose that $w \in [u, v]$. Then

$$d_X(y, v) \leq d_X(y, w) + d_X(w, v),$$

and so

$$d_X(w, v) \leq d/2 - \delta.$$

Thus

$$d_X(u, w) = d_X(u, v) - d_X(v, w) \leq d - d/2 + \delta \leq d/2 + \delta.$$

Finally, we see that

$$d_X(u, y) \leq d/2 + \delta + \delta \leq d,$$

where the last inequality follows since $d \geq 4\delta$.

Therefore, if σ is a simplex of L containing v, then $\sigma \cup \{y\}$ is also a simplex of K_d. Let $\sigma' = (\sigma \setminus \{v\}) \cup \{y\}$ be the simplex formed by replacing v with y. There is a natural affine homotopy which takes σ to σ' through $\sigma \cup \{y\}$. Thus we get a homotopy $L \to L'$, where L' is formed by replacing v with y in every simplex of L containing v. Since $d_X(x_0, y) < d_X(x_0, v)$, applying this procedure finitely many times (using that L is a finite subcomplex) shows that L is homotopic to a subcomplex contained in a single simplex, as in Case 1. Therefore L is contractible. $\square$

The above theorem holds if, more generally, we let $Y \subseteq X$ be a coarsely dense subset of a δ–hyperbolic space X (that is, a subspace such that there exists a constant k so that for each $x \in X$, there is some $y \in Y$ with $d_X(x, y) \leq k$), and let $K_d = K_d(Y)$ be the Rips complex associated to Y. The above proof holds with only minor modifications; we leave the details as an exercise.

Theorem 4.3. *Let G be a hyperbolic group. There exists a simplicial complex K and an action G on K such that:*

1. *G acts properly on K;*

2. *K/G is compact;*

3. *K is contractible.*

4. *K is locally finite and finite dimensional.*

Proof. Fix a finite presentation $G = \langle S \mid R \rangle$, let X be the vertex set of $\Gamma(G, S)$. Then X is a 1–coarsely dense subset of a geodesic δ–hyperbolic space. Letting $K = K_d$ for any $d \geq 4\delta$, (the generalization of) Theorem 4.2 will show that (3) holds. For (4), note that there are only finitely many vertices in any ball of radius d in $\Gamma(G, S)$.

The natural action of G on the vertex set of $\Gamma(G, S)$ naturally extends to a simplicial action of G on K with compact quotient. This action is free and transitive on the vertex set of K, and therefore the stabilizer of every simplex is finite. □

Note that if G has torsion, then it is possible that there are simplices which are stabilized setwise but not pointwise, in which case the quotient will not be a simplicial complex.

Also, notice that the dimension of K_d is one less than the cardinality of the largest set of vertices in $\Gamma(G, S)$ of diameter d; this is bounded by $|S|^d$ (assuming S is symmetric). Finally, the 1–skeleton of K_d is the Cayley graph of G with respect to the generating set $\mathcal{B}$, which is the set of all non-trivial vertices of $\Gamma(G, S)$ in a ball of radius d about the identity.

We now show how to use the Rips complex of a hyperbolic group to deduce various homological and finiteness properties of the group.

An *Eilenberg-MacLane* space is a CW complex whose fundamental group is isomorphic to a given group and which has a contractible universal cover. Such a space is denoted $K(G, 1)$ (where here, 1 refers to the condition on π_1). Given a group, a $K(G, 1)$ is unique up to homotopy equivalence. Every group G admits a $K(G, 1)$, constructed as follows.

Fix a (not necessarily finite) presentation $G = \langle S \mid R \rangle$ for G. Form a rose by taking a wedge of $\mathbb{S}^1$'s, one for each element of S. Now, for each $r \in R$, glue a 2–cell onto the rose along the path specified by r, starting at the single vertex. This results in a space X' with $\pi_1(X') \simeq G$. The universal cover $\tilde{X}'$ of X' is called the *Cayley complex* of G. However, $\tilde{X}'$ may not be contractible (i.e., there may be some $n \geq 1$ such that $\pi_n(\tilde{X}') \neq 0$). To deal with this, modify X' by adding n–cells for $n \geq 3$ (which does not affect the fundamental group) to ensure that $\tilde{X}$ is contractible. Then X is a $K(G, 1)$.

A couple asides.

1. There is also the more general notion of an Eilenberg-MacLane space $K(G, n)$, which is a CW complex such that the only nontrivial homotopy group of the universal cover is π_n, and we have $\pi_n(\tilde{X}) \simeq G$. We won't need this notion for this class.

2. Some sources only require a $K(G, 1)$ to be a path connected space (not a CW complex). Since any space can be approximated up to weak equivalence by a CW complex, this is not such a different notion. (Recall that two spaces are weakly equivalent if, roughly, there are bijections between all homotopy groups.)

If Y is a $K(G, 1)$ for a group G, then the homology and cohomology groups of G can be defined to be $H_n(G) = H_n(Y)$ and $H^n(G) = H^n(Y)$, respectively (here we assume all coefficients are in $\mathbb{Z}$).

Definition 4.4. A group G *is of type* F_n for $n \in \mathbb{N} \cup \{\infty\}$ if G admits a $K(G, 1)$ with finitely many cells in each dimension $\leq n$. G is *of type F* if it admits a finite $K(G, 1)$.

We have the following implications:

$$F_1 \Longleftarrow F_2 \Longleftarrow \cdots \Longleftarrow F_n \Longleftarrow F_\infty \Longleftarrow F.$$

It can be seen from the above construction of a $K(G,1)$ that a group G is finitely generated if and only if it is of type F_1, and G is finitely presented if and only if it if of type F_2.

The *geometric dimension* of a group, which we denote by $gd(G)$, is the minimal d such that G has a d–dimensional $K(G,1)$. Note that if $gd(G) = n$, then $H_k(G) = 0$ for all $k > n$. It is a (non-obvious) fact that if G is a non-trivial finite group, then there exists an arbitrarily large integer n such that $H_n(G) \neq 0$ (I believe this fact is originally due to Swan); in other words, there is no N such that $H_n(G) = 0$ for all $n \geq N$. So if G is a group with torsion, then there is some finite cyclic group H that is a subgroup of G. In particular, if X is a $K(G,1)$, there is a cover X_H of X such that $\pi_1(X_H) = H$. Since X and X_H have the same universal cover, it follows that X_H is a $K(H,1)$, and hence there is no N such that $H_k(H) = 0$ for all $k \geq N$. However, if X is finite dimensional, then so is X_H (as X_H is a cover of X). But this is a contradiction. This shows that if G is not torsion-free then G is not type F and $gd(G) = \infty$. However, it is still possible for a group with torsion to be type F_∞.

Corollary 4.5. *Let G be a hyperbolic group and K be the complex from Theorem 4.3. If G is torsion-free, then K/G is a finite $K(G,1)$. In particular, all torsion-free hyperbolic groups are type F.*

Proof. Recall that the stabilizer of any simplex of K is finite. If G is torsion-free, then the stabilizer of any simplex of K is trivial, and thus the action of G on K is free. In particular, it is a covering space action. Thus, under this assumption, the quotient K/G is a $K(G,1)$. Finally a CW complex is finite if and only if it is compact, so in fact K/G is a finite $K(G,1)$. □

Even if a hyperbolic group has torsion, the action on the Rips complex is still sufficient to give powerful results about its (co)homological and finiteness properties. We won't discuss these in detail; see Bekali and Kapovich's notes.

Corollary 4.6. *Let G be a hyperbolic group. Then:*

1. *If G is virtually torsion-free, then G has finite virtual cohomological dimension.*
2. *The group G is of type F_∞ and FP_∞.*
3. *There is no n_0 such that $H^n(G,\mathbb{Q}) = 0$ for all $n \geq n_0$.*
4. *$H_*(G,\mathbb{Q})$ and $H^*(G,\mathbb{Q})$ are finite dimensional.*

5 Subgroups of hyperbolic groups

It is natural to ask is whether all subgroups of hyperbolic groups are hyperbolic.

Exercises 5.1. *Find the mistake in the following proof.*

Claim 5.2. If G is hyperbolic and $H \leq G$, then H is hyperbolic.

Proof. Let S be a finite generating set for G and consider the metric d_S on H. Equivalently, consider H with the metric induced by considering it as a subspace of $\Gamma(G,S)$ (in other words, for any two points $h_1, h_2 \in H$, let $d_S(h_1,h_2)$ be the distance from h_1 to h_2 in $\Gamma(G,S)$. As $\Gamma(G,S)$ is hyperbolic, every subspace with the induced metric is hyperbolic: this follows from the Gromov product definition of hyperbolicity. Additionally, it is straightforward to check that the action of H on this subspace of $\Gamma(G,S)$ is proper and cobounded. Thus by the Milnor–Schwartz lemma, H is finitely generated and a Cayley graph of H is quasi-isometric to (H,d_S). Therefore H is hyperbolic. □

In fact, not all subgroups of hyperbolic groups are hyperbolic, simply because free groups contain infinitely generated subgroups, which cannot be hyperbolic. So one could ask, instead, whether all *finitely generated* subgroups of a hyperbolic group are hyperbolic. The answer to this question is still no: it is possible to construct such examples of finitely generated subgroups of hyperbolic groups which are not finitely presented (and thus are not hyperbolic) using small cancellation theory; we'll do this at the end of the semester if we have time. There special cases where this is true, though. For example, the answer is yes in free groups, because all subgroups of a free group are free (and all finitely generated free groups are hyperbolic). The answer is also yes for surface groups (i.e., the fundamental group of a compact orientable surface of genus at least 2); all finitely generated subgroups of surface groups are either free (if the subgroup is infinite index) or surface groups (if the subgroup if finite index).

In fact, even if you restrict to asking only about all finitely presented subgroups of hyperbolic groups, the answer is still no. Noel Brady constructs a finitely presented subgroup of a hyperbolic group (which is normal with quotient $\mathbb{Z}$) that is not hyperbolic. The construction uses techniques from CAT(0) cube complexes. However, the subgroup Brady constructs is not of type F_3. What if we require the subgroups to have higher finiteness properties?

The following question is open in general.

Question 5.3. *If G is hyperbolic and $H \leq G$ is of type F_3, is H hyperbolic?*

The question is, in fact, open if 3 is replaced with any $n \geq 3$. Is *any* finiteness property sufficient?

Question 5.4. *If G is hyperbolic and $H \leq G$ is of type F, is H hyperbolic?*

5.1 Cyclic subgroups

Our first result states that given an infinite order element g in a hyperbolic group G, the image of $\langle g \rangle$ in a Cayley graph of G is a quasigeodesic. This is not generally the case. For example, if G is the Baumslag-Solitar group $BS(1,2) = \langle a, t \mid tat^{-1} = a^2 \rangle$, then we have $t^k a t^{-k} = a^{2^k}$, and so $\langle a \rangle$ is not a quasigeodeisc. (Exercise: fill in the details of this argument!) In fact, this shows that $BS(1,2)$ cannot be a subgroup of a hyperbolic group. (Exercise: why?)

Theorem 5.5. *Let G be a hyperbolic group with Cayley graph X, and let g be an infinite order element of G. Then the image of $\langle g \rangle$ in X is a quasigeodesic.*

Proof. Fix any $R \geq 0$. Fix a presentation $\langle S \mid R \rangle$ of G, and let $X = \Gamma(G, S)$. Let δ be the hyperbolicity constant of X, and let g be an infinite order element of G.

First, there exists some $K \in \mathbb{N}$ such that $d(1, g^K) \geq 8R + 2\delta$. This is because there are only finitely many vertices of X in a ball of radius $8R + 2\delta$ about the identity and g is infinite order.

Consider a geodesic $[1, g^K]$ from 1 to g^K in X. Let m be the midpoint of $[1, g^K]$ (by which we mean, a vertex at distance at most 1/2 from the true midpoint). Let $p \in B_R(1)$ and $q \in B_R(g^K)$ and let m_1 be the midpoint of $[p, q]$.

We will prove the result using the following three claims.

Claim 1. Let I is the subsegment of $[1, g^K]$ of length $2R$ centered at m. Then m_1 is contained in the 2δ–neighborhood of I.

Claim 2. Let $N/2$ be the number of vertices in a ball of radius 2δ about the identity in X. There is a number $P(R) \leq NR$ such that $g^{P(R)} \notin B_R(1)$.

Claim 3. For all L, $|g^{NL}| \geq |L|$.

Assuming these three claims hold, we will now prove the result. Consider the path γ which is the image of $\langle g\rangle$ in X, that is, γ is formed by concatenating the images of the geodesic $[1, g]$ under $\langle g\rangle$, so that γ is the concatenation $\cdots \cup [g^{-1}, 1] \cup [1, g] \cup [g, g^2] \cup [g^2, g^3] \cup \cdots$. We will show that γ is a quasigodesic.

Let x, y be any two points on γ. Then there exist a, b such that $\ell(\gamma|_{[x,g^{aN}]}) \leq N|g|$ and $\ell(\gamma|_{[g^{bN},y]} \leq N|g|$. Then we have

$$\ell(\gamma|_{[x,y]} \leq 2N|g| + |b-a|N|g|. \tag{6}$$

Moreover, by Claim 3, we have

$$d(x, y) \geq d(g^{aN}, g^{bN}) - 2N|g| = d(g^{(b-a)N}, 1) - 2N|g| \geq |b-a| - 2N|g|.$$

Solving for $|b-a|$ and substituting into (6), we obtain

$$\ell(\gamma|_{[x,y]}) \leq N|g|d(x,y) + 2N^2|g|^2 + 2N|g|.$$

This proves the result.

It remains to prove the three claims.

Proof of Claim 1. Let m_2 be the midpoint of $[p, g^K]$. First, since $|d(p, g^K) - d(p, q)| \leq R$, we have $|d(p, m_1) - d(p, m_2)| \leq R/2$. A similar argument yields that $|d(m, g^K) - d(m_2, g^K)| \leq R/2$. Moreover, by our choice of K, we see that none of the midpoints m, m_1, m_2 are contained in a ball of radius $R+\delta$ around g^K or a ball of radius $R+\delta$ around 1.

Consider the triangle with vertices p, g^K, q. Then all three internal points of this triangle are at distance at most $R+\delta$ from g^K (see Section 2.1 for the definition of an internal point). The thin (not slim!) triangles definition of hyperbolicity applied to the point m_1 shows that there is a point $m_1' \in [p, g^k]$ such that the following hold:

(1) $d(m_1, p) = d(m_1', p)$ (in other words, m_1 and m_1' map to the same point in the comparison tripod);

(2) $d(m_1, m_1') \leq \delta$;

(3) $d(m_1', m_2) \leq R/2$ (this follows from (1) and the fact that $|d(p, m_1) - d(p, m_2)| \leq R/2$); and

(4) $d(m_1', 1) > R + \delta$.

The final item (4) holds, because our choice of K ensures that $d(p, q) > 7R+2\delta$, and so $d(p, m_2) > 7/2R+\delta$. Then (3) ensures that $d(p, m_1') > 3R + \delta$, and so $d(1, m_1') > 3R + \delta - R > R + \delta$.

Now, by applying the thin triangles definition of hyperbolicity to the triangle with vertices $1, p, g^K$ to the points m_1' and m_2 yields points $m_1'', m_2' \in [1, g^K]$ such that the following hold:

(a) $d(m_2', g^K) = d(m_2, g^K)$ and $d(m_1'', g^K) = d(m_1', g^K)$;

(b) $d(m_2', m_2) \leq \delta$ and $d(m_1'', m_1') \leq \delta$;

(c) $d(m_1'', m_2') = d(m_1', m_2)$ (this follows from (a) and (b)); and

(d) $d(m_2', m) < R/2$.

The final item (d) holds from (a) and the fact since m_2 and m are both midpoints, we have $|d(m, g^K) - d(m_2, g^K)| \leq R/2$. Note that in order to be able to apply thin triangles to the point m_1' and obtain a point m_1'' with these properties, we need (4) from the definition of m_1' to hold, since this implies that m_1' lies closer to g^K than the internal point on $[p, g^K]$ does. This then implies that m_1'' lies on $[1, g^K]$, not on $[1, p]$.

Combining all of this, we see that

$$d(m_1, m_1'') \leq d(m_1, m_1') + d(m_1', m_1'') \leq \delta + \delta = 2\delta,$$

and

$$d(m, m_1'') \leq d(m, m_2') + d(m_2', m_1'') \leq R/2 + d(m_1', m_2) \leq R/2 + R/2 = R.$$

where the second inequality follows from (d) and (c), while the third inequality follows from (3). Thus we see that $m_1'' \in I$, by the definition of I, as desired. $\square$

Proof of Claim 2. The 2δ–neighborhood of I contains at most NR vertices. Consider the translates of $[1, g^K]$ by $1, g, \ldots, g^{NR}$. Then the midpoints of each of these translates are distinct (else a power of g would fix a point in X, which would imply that g is elliptic), and there are at most $1+NR$ of them. If for all $0 \leq i \leq NR$, $g^i \in B_R(1)$, then $g^{K+i} \in B_R(g^K)$, and so Claim 1 implies that all of these distinct midpoints lie in the 2δ–neighborhood of I. However, this is a contradicts the number of distinct points in the 2δ–neighborhood of I. Therefore, there must exist some $0 \leq P(R) \leq NR$ such that $g^{P(R)} \notin B_R(1)$, completing the proof of Claim 2. $\square$

Note that Claim 2 implies that $|g^{P(R)}| = d(1, g^{P(R)}) > R$.

Proof of Claim 3. Suppose the statement of the claim does not hold. Then there exists some L_0 and ε so that $|g^{NL_0}| \leq |L_0| - \varepsilon$. Then for all $s \geq NL_0$, we may write $s = nNL_0 + L_1$, where $0 \leq L_1 < NL_0$ and $n \in \mathbb{N}$. Then, if for any n such that $n\varepsilon > |g^{L_1}|$, we have

$$\begin{aligned}
|g^s| &\leq |g^{nL_0R}| + |g^{L_1}| \\
&\leq n(L_0 - \varepsilon) + |g^{L_1}| \\
&< nL_0 \\
&< s/N.
\end{aligned}$$

In particular, this holds for all sufficiently large s. Choose a value of R so that $P(R) > L_0$. Then by Claim 2, we have $|g^{P(R)}| > R$. But the above arguments shows that $|g^{P(R)}| \leq P(R)/N \leq R$, which is a contradiction. $\square$

$\square$

This theorem implies, in particular, that if $g \in G$ is an infinite order element, then there exists a constant $L \geq 0$ (depending on g and the hyperbolicity constant δ) such that the following holds for any $i, j \in \mathbb{N}$. If x is a point on the geodesic $[g^i, g^j]$, then there is some $k \in \mathbb{N}$ such that $d(x, g^k) \leq L$. (Exercise: fill in the details.)

We now turn our attention to the centralizers of infinite order elements. Given an element $g \in G$, the *centralizer* of g in G, denoted $C_G(g)$ (or just $C(g)$ if no confusion is possible) is the set of all $h \in G$ such that h commutes with g.

Theorem 5.6. *Let G be a hyperbolic group and suppose $g \in G$ has infinite order. Then the centralizer $C(g)$ of g is virtually cyclic.*

Proof. Fix a finite generating set S for G, and let δ by the hyperbolicity constant of the Cayley graph $\Gamma(G, S)$. Since Theorem 5.5 shows that the image of $\langle g \rangle$ in $\Gamma(G, S)$ is a quasi-geodesic, there is some $M \geq 0$ (depending on g and δ) such that the geodesic $[1, g^n]$ is contained in the M–neighborhood of $\{1, g, g^2, \ldots, g^n\}$. Let $h \in C(g)$, and choose m so that $d(1, g^m) > 2d(1, h) + 2\delta$. Consider the rectangle $[1, g^m] \cup [g^m, hg^m] \cup [hg^m, h] \cup [h, 1]$:

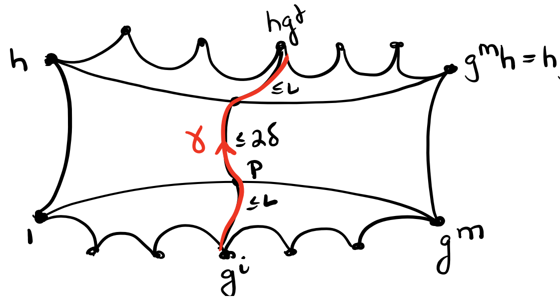

By our choice of m, there exists some $p \in [1, g^m]$ such that $d(p, [h, hg^m]) \leq 2\delta$. Therefore there exist $1 \leq i, j \leq m$ such that $d(g^i, hg^j) \leq 2\delta + 2M$. In particular, there is a path γ of length at most $2\delta + 2M$ from g^i to hg^j. Let $\mathbf{Lab}(\gamma) = u \in G$, so that $|u| \leq 2\delta + 2M$. Then $hg^j = g^i u$, and so $u = g^{j-i}h$, which in turn implies u is an element of the coset $\langle g \rangle h$. That is, every coset of $\langle g \rangle$ in $C(g)$ has a representative of length at most $2\delta + 2M$. Since there are only finitely many words of length at most $2\delta + 2M$, this implies that $|C(g) : \langle g \rangle| < \infty$. In particular, $C(g)$ is virtually cyclic. $\square$

The following two corollaries are immediate from Theorem 5.6.

Corollary 5.7. *Let G be a hyperbolic group and H an abelian subgroup containing an infinite order element. Then H is of the form $\mathbb{Z} \oplus A$ where A is a finite abelian group.*

Corollary 5.8. *The group $\mathbb{Z} \oplus \mathbb{Z}$ is not a subgroup of any hyperbolic group.*

We say $\mathbb{Z} \oplus \mathbb{Z}$ is a *poison subgroup* for hyperbolicity. We saw above that $BS(1, 2)$ is also a poison subgroup for hyperbolicity.

Corollary 5.9. *Abelian subgroups of hyperbolic groups are either finite or virtually cyclic.*

5.2 Quasi-convex subgroups

Definition 5.10. A subset Y of a metric space X is *σ-quasiconvex* if every geodesic in X between points of Y is contained in the σ–neighborhood of Y. In other words, if $y_1, y_2 \in Y$ and $[y_1, y_2]$ is a geodesic in X, then for every point $z \in [y_1, y_2]$ there is a point $w \in Y$ such that $d_X(z, w) \leq \sigma$.

A subgroup H of a group G is *quasiconvex* with respect to a finite generating set S of G if there exists $\sigma \geq 0$ such that H is a quasiconvex subset of the Cayley graph $\Gamma(G, S)$.

Example 5.11. Finite and finite-index subgroups of a finitely generated group are quasiconvex.

Exercises 5.12. *In general, quasiconvexity of a subgroup H depends on the choice of generating set for the group G. Give an example of a subgroup $H \leq \mathbb{Z}^2$ and two generating sets S, T of $\mathbb{Z}^2$ so that H is quasiconvex in G with respect to S but not with respect to T.*

Proposition 5.13. *If G is hyperbolic and $H \leq G$ is quasiconvex with respect to a finite generating set S of G, then H is quasiconvex with respect to every finite generating set of G.*

Proof. Let S, T be two finite generating sets of a hyperbolic groups G. Then there is a (λ, c)–quasi-isometry $\varphi \colon \Gamma(G, T) \to \Gamma(G, S)$ for some constants λ, c (depending only on S, T) which is the identity on all vertices of $\Gamma(G, T)$. Suppose $H \leq G$ is σ–quasiconvex with respect to S. We need to show that given any two elements of H, a geodesic in $\Gamma(G, T)$ joining them stays within a uniform distance of H. Without loss of generality, we may assume that one of the two elements in the identity element.

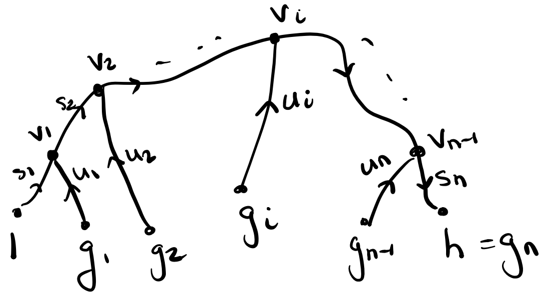

Let $h \in H$, and suppose γ is a geodesic from 1 to h in $\Gamma(G, T)$. Since φ is a (λ, c)–quasi-isometry, $\varphi(\gamma)$ is a (λ, c)–quasi-geodesic in $\Gamma(G, S)$ from 1 to h. Since $\Gamma(G, S)$ is hyperbolic, there is some constant M such that $[1, h]$ is contained in the M–neighborhood of $\varphi(\gamma)$. Moreover, since H is a σ–quasiconvex subgroup of G with respect to S, the geodesic $[1, h]$ is contained in the σ–neighborhood of H.

Let $x \in \gamma$ be an arbitrary point. Then $\varphi(x)$ lies on $\varphi(\gamma)$, and so by the previous paragraph, there is some $h' \in H$ such that $d_S(\varphi(x), h') \leq \sigma + M$. Again using the fact that φ is a (λ, c)–quasi-isometry, we see that

$$d_T(x, h') \leq \lambda d_S(\varphi(x), \varphi(h')) + c = \lambda d_S(\varphi(x), h') + c \leq \lambda(\sigma + M) + c.$$

Therefore γ is contained in the $(\lambda(\sigma + M) + c)$–neighborhood of H in $\Gamma(G, T)$, which shows that H is $(\lambda(\sigma + M) + c)$–quasiconvex with respect to T, as desired. $\square$

Proposition 5.13 implies that for a *hyperbolic* group G, we can refer to quasiconvex subgroups without specifying the associated finite generating set of G.

Examples 5.14. 1. If $g \in G$ is infinite order and G is hyperbolic, then $\langle g \rangle$ is quasiconvex by Theorem 5.5.

2. If $K < H < G$, K is finite index in H, and G is hyperbolic, then K is quasiconvex in G if and only if H is quasiconvex in G. Thus virtually cyclic subgroups of hyperbolic groups are quasiconvex, since they are finite extensions of infinite cyclic subgroups. In particular, if $g \in G$ is infinite order, then $C(g)$ is quasiconvex in G.

Proposition 5.15. *Let G be a hyperbolic group and $H \leq G$ a subgroup. Then H is quasiconvex if and only if H is finitely generated and quasi-isometrically embedded in G.*

Proof. We will prove each direction separately.

($\Longleftarrow$): Suppose H has a finite generating set T, let S be a finite generating set for G, and suppose the inclusion $\iota\colon \Gamma(H, T) \hookrightarrow \Gamma(G, S)$ induced by the inclusion $H \hookrightarrow G$ is a (λ, c)–quasi-isometric embedding. Let δ be the hyperbolicity constant of $\Gamma(G, S)$. Fix any $h \in H$, let γ be a geodesic from 1 to h in $\Gamma(H, T)$, and let α be a geodesic from 1 to h in $\Gamma(G, S)$. Then the image $\iota(\gamma)$ of γ in $\Gamma(G, S)$ is a (λ, c)–quasi-geodesic, and so there is a constant M (depending on δ, λ, and c) such that α is contained in the M–neighborhood of the vertex set of $\iota(\gamma)$. Thus given any $x \in \alpha$, there is some vertex $h' \in \iota(\gamma)$ such that $d_S(x, h') \leq M$. But since ι is the identity on vertices, $h' \in H$, and so $d_S(x, H) \leq M$. Since x was arbitrary, this implies that H is M–quasiconvex with respect to S, as desired.

($\Longrightarrow$): Fix a finite generating set S for G, and suppose H is σ–quasiconvex with respect to S. Fix any $h \in H$, and let γ be a geodesic in $\Gamma(G, S)$ from 1 to h. Label the vertices of γ by $v_0 = 1, v_1, v_2, \ldots, v_n = h$. Each edge of γ is labeled by a generator $s_i \in S$, so that $v_i = s_1 \ldots s_i$ for all $1 \leq i \leq n$. See Figure 5.2.

Since H is σ–quasiconvex with respect to S, for each $1 \leq i \leq n - 1$, there is some $g_i \in H$ such that $d_S(v_i, g_i) \leq \sigma$. Let u_i be the label of a geodesic from g_i to v_i, so that $d_S(1, u_i) \leq \sigma$. Let $g_0 = 1$ and $g_n = h$.

The for each i we have $u_i s_{i+1} u_{i+1}^{-1} = g_i^{-1} g_{i+1}$, which is an element of H since $g_i, g_{i+1} \in H$. Note that the length of the path labeled by $u_i s_{i+1} u_{i+1}^{-1}$ is at most $2\sigma + 1$.

We first find a finite generating set for H. Let T be the set of all elements of H which are in a ball of radius $2\sigma + 1$ about the identity in $\Gamma(G, S)$. This is a finite set. Moreover, by construction, $g_i^{-1} g_{i+1} \in T$ for each $1 \leq i \leq n-1$. Since

$$h = g_n = g_1(g_1^{-1} g_2) \cdots (g_{n-1}^{-1} g_n) = (g_1^{-1} g_2) \cdots (g_{n-1}^{-1} g_n), \tag{7}$$

we see that $h \in \langle T \rangle$. Since h was an arbitrary element of H, this shows that T generates H.

Moreover, (7) shows that $d_T(1, h) \leq n = d_S(1, h)$. Additionally, each $t \in T$ satisfies $d_S(1, t) \leq 2\sigma + 1$ by construction, and so

$$d_S(1, h) \leq (2\sigma + 1) d_T(1, h).$$

These two inequalities show that the inclusion $H \hookrightarrow G$ is a $(2\sigma + 1)$–quasi-isometric embedding. This completes the proof of the forward direction. $\square$

The following corollary follows immediately from Proposition 5.13 by applying Corollary 2.15.

Corollary 5.16. *Quasiconvex subgroups of hyperbolic groups are hyperbolic.*

Proposition 5.17. *Let G be a hyperbolic group and $A, B \leq G$ quasiconvex subgroups. Then $A \cap B$ is quasiconvex in G.*

Proof. Fix a finite generating set S for G. Suppose H_1, H_2 are σ–quasiconvex. Consider a point $h \in H_1 \cap H_2$, and let p be a geodesic from 1 to h. We will show that given an arbitrary vertex $v \in p$, there is a path from v to $H_1 \cap H_2$ of length at most $|B_\sigma(1)|^2$, where $B_\sigma(1)$ is the ball of radius σ about the identity in $\Gamma(G, S)$. This will show that $H_1 \cap H_2$ is $|B_\sigma(1)|^2$–quasiconvex in G.

Fix any $v \in p$, and let Σ be the set of all paths from v to $H_1 \cap H_2$ such that every vertex on q is distance at most 1 from both H_1 and H_2. The set Σ is non-empty, because the subpath of p starting at v and ending at h satisfies this property. Let q be the shortest path in Σ, and suppose $\mathbf{Lab}(q) = w$. Conflating the vertex v with its label and thinking of v as a word in g, let $vw = s_1 \dots s_n$ for some $s_i \in S$. Then by assumption $vw \in H_1 \cap H_2$. For each vertex $u \in q$ (which we also think of as an element of G), let $u_i \in H_i$ such that $d(u, u_i) \leq \sigma$. Now suppose $u, w \in q$ are the labels of vertices that appear in that order along q, so that $\ell(q|_{[v,u]}) < \ell(q|_{[v,w]})$. Let r_u, r_w be the prefixes of $\mathbf{Lab}(q)$ corresponding to the vertices u, w. In particular, we have that $u = vr_u$ and $w = vr_w$, and there exists r'_u and r'_w with $vr_u r'_u = vr_w r'_w$. Moreover, we must have that r'_w is a strictly shorter word than r'_u, by our assumption on the placement of u and w.

Let $x_i^u = u^{-1} u_i$ and $x_i^w = w^{-1} w_i$, so that x_i^u and x_i^w are the labels of paths from the vertex u to the vertex u_i and from w to w_i, respectively, for $i = 1, 2$. Then $vr_u x_i^u \in H_i$ and $vr_w x_i^w \in H_i$. If $u^{-1} u_i = w^{-1} w_i$, then $x_i^u = x_i^w$, and so $vr_u x_i^w \in H_i$.

Now, we show that the path with label $r_u r'_w$ is in Σ. It suffices to show that $vr_u r'_w \in H_1 \cap H_2$. Let the label of u be $s_1 \dots s_k$ and let the label of w be $s_1 \dots s_p$.

This in fact shows that the path with label $r_u r'_w$ is an element of Σ. However, this path is strictly shorter than q, which is a contradiction.

Therefore, for $u, w \in q$, the pairs (u_1, u_2) and (w_1, w_2) are distinct. Thus the number of distinct vertices on q is at most the number of distinct elements in $B_\sigma(1) \times B_\sigma(1)$. We conclude that $\ell(q) \leq |B_\sigma(1)|^2$, as desired, concluding the proof. $\square$

The proof of the following proposition is very similar to the proof of Theorem 5.6.

Proposition 5.18. *Let G be a hyperbolic group, and $H \leq G$ an infinite quasiconvex subgroup. If H is normal in G, then $|G : H| < \infty$.*

Proof. Let S be a finite generating set of G, suppose $\Gamma(G,S)$ is δ–hyperbolic, and suppose H is σ–quasiconvex in G. Let cH be a coset of H in G. We will show that this coset has a representative whose length is bounded by $2\sigma+2\delta$, which will prove the result.

Let $u \in H$ be such that $|u| \geq 2\delta+2|c|$ (which we can do because H is infinite). Consider the rectangle with vertices $1, c, u, uc$. Then by the thinness of quadrilaterals in $\Gamma(G,S)$, there are points $p \in [1,u]$ and $q \in [c, uc]$ with $d(p,q) \leq 2\delta$. Since H is σ–quasiconvex, there are vertices $h_1, h_2 \in H$ such that $d(p,h_1) \leq \sigma$ and $d(q, ch_2) \leq \sigma$. Thus there is a word w with $|w| \leq 2\delta + 2\sigma$ such that $h_1 w = ch_2$. Thus $h_1 w = (ch_2 c^{-1})c = h'c$ for some $h \in H$, since H is a normal subgroup of G. In particular, $w = h_1^{-1}h'c \in Hc$, and so each coset of H has a representative of length at most $2\sigma+2\delta$, which completes the proof. ☐

5.3 Elementary subgroups

A group G is *elementary* if it is virtually cyclic. Note that finite groups are all elementary, as $\{1\}$ is a cyclic subgroup of finite index. Every such group is a finite extension of $\mathbb{Z}$ or $D_\infty = \mathbb{Z}/2\mathbb{Z} * \mathbb{Z}/2\mathbb{Z} = \mathrm{Isom}(\mathbb{Z}) \simeq \mathbb{Z} \rtimes \mathbb{Z}/2\mathbb{Z}$. When thinking of this decomposition of the isometry group of $\mathbb{Z}$, you should think of the $\mathbb{Z}$ factor as acting on the line by translation, and the $\mathbb{Z}/2\mathbb{Z}$ factor as acting by switching ∞ and $-\infty$.

The following results are proven using standard group theoretic techniques, and we will not give details here.

Lemma 5.19. *Let E be an elementary group.*

1. *If E is torsion-free, then E is cyclic.*
2. *E contains normal subgroups $T \leq E^+ \leq E$, such that $|E : E^+| \leq 2$, T is finite, and $E^+/T \simeq \mathbb{Z}$. Moreover, if $E \neq E^+$, then $E/T \simeq D_\infty$.*

For the second property, if $E = D_\infty$, then $E^+ \simeq \mathbb{Z}$ and is the subgroup that fixes ∞ and $-\infty$. On the other hand, if $E = \mathbb{Z}$, then $E = E^+$.

We next define an elementary subgroup containing an infinite order element of a hyperbolic group and show that it satisfies nice properties. We will need the following fact, which we will prove when we define boundaries (soon!).

Lemma 5.20. *Let G be a hyperbolic group. If $g \in G$ has infinite order and $x^{-1}g^k x = g^\ell$ for some $x \in G$ and some $k, \ell \in \mathbb{Z}$, then $k = \pm\ell$.*

Definition 5.21. Let G be a hyperbolic group and $g \in G$ an infinite order element. Define

$$E(g) = \{x \in G \mid x^{-1}g^n x = g^{\pm n} \text{ for some } n = n(x) \in \mathbb{Z}\}.$$

The proofs of the following results are left as an exercise.

Lemma 5.22. *Let G be a hyperbolic group and $g \in G$ an infinite order element.*

1. *$E(g)$ is a subgroup of G.*
2. *$E^+(g) = \{x \in G \mid x^{-1}g^n x = g^n$ for some $n = n(x) \in \mathbb{Z}\}$ is a subgroup of order 2.*
3. *$E(g) = E(g^k)$ for any $k \in \mathbb{Z} \setminus \{0\}$.*
4. *$E(y^{-1}gy) = y^{-1}E(g)y$ for any $y \in G$.*
5. *$\langle g \rangle \leq C(g) \leq E^+(g) \leq E(g)$.*

Theorem 5.23. *Let G be a hyperbolic group and $g \in G$ an infinite order element. Then $|E(g) : \langle g \rangle| < \infty$, and $E(g)$ is the unique maximal virtually cyclic subgroup of G containing g.*

Proof. Let G be a hyperbolic group, fix a finite generating set S for G, and let $g \in G$ be infinite order. Then there exist constants λ, c (depending on the hyperbolicity constant of $\Gamma(G,S)$ and g) so that $\{\dots, g^{-2}, g^{-1}, 1, g, g^2, \dots\}$ is a (λ, c)–quasigeodesic in $\Gamma(G,S)$. Let K be a constant so that all (λ, c)–quasigeodesic quadrilaterals in $\Gamma(G,S)$ are K–slim.

For the first statement of the theorem, it suffices to show that $|E^+(g) : \langle g \rangle| < \infty$. Let $x \in E^+(g)$, so that $x^{-1}g^n x = g^n$. Let m be large enough mulitple of $2n$ so that $d(1, g^{m/2}) \geq 2K + d(1,x)$. Consider the (λ, c)–quasigeodesic quadrilateral consisting of geodesics $[1,x]$ and $[g^m, xg^m]$ and quasigeodesics $\{1, g, g^2, \dots, g^m\}$ and $\{x, xg, xg^2, \dots, xg^m\}$.

By our choice of m, there is a geodesic from the vertex $g^{m/2}$ (which is a power of g^n) on the bottom of the quadrilateral to the vertex $xg^{m/2}$ on the top of the quadrilateral that is labeled by x. Moreover, since this quadrilateral is K–thin, there is some $\ell \in \{0, \dots, m\}$ so that $d(xg^{m/2}, g^\ell) \leq K$.

In particular, $g^{-\ell}xg^{m/2} = g^{-k}x$ for some k, and $g^{-k}x \in \langle g \rangle x$. Additionally, we see that $K \geq d(xg^{m/2}, g^\ell) = d(g^{-\ell}xg^{m/2}, 1) = d(g^{-kx}, 1)$. Therefore $g^{-k}x$ is a representative of the coset $\langle g \rangle x$ whose length is uniformly bounded by K. Therefore there are only finitely many cosets of $\langle g \rangle$ in $E^+(g)$, and thus in $E(g)$.

For the second statement, note that if E is any elementary group containing g, then $\langle g^n \rangle$ is normal in E for some $n \in \mathbb{N}$ (because $\langle g \rangle$ is a subgroup of E of finite index, and thus contains a finite-index subgroup which is normal in E). This implies that $E \leq E(g)$. $\square$

We will see some geometric properties of $E(g)$ after we define the boundary of a hyperbolic group in the next section.

6 Boundaries and isometries of hyperbolic spaces

Before defining the boundary of a hyperbolic group, we begin with a few preliminary definitions. First, we will be regularly discussing when a sequence of geoesics converges. More precisely, by this we mean the following. A sequence of functions $f_i \colon X \to Y$ *converges uniformly on compact sets* to a function f if for every compact $K \subseteq X$ and for all $\varepsilon > 0$, there exists N such that for all $i \geq N$ and all $x \in K$, $d_Y(f_i(x), f(x)) \leq \varepsilon$.

Second, we will mostly work in the context of proper metric spaces (although much of the theory works in non-proper metric spaces, although things are often significantly more complicated). One of the main reasons to restrict attention to proper spaces is that we are then able to use the following useful theorem.

Theorem 6.1 (Arzela-Ascoli Theorem)**.** *If X is separable and Y is compact, then every equicontinuous sequence of maps $X \to Y$ has a subsequence which converges uniformly on compact sets to a continuous map $X \to Y$.*

The main application of this theorem in our context is the following corollary.

Corollary 6.2. *Let $\gamma_i \colon [0,\infty) \to X$ be a sequence of geodesic rays with $\gamma_i(0) = \gamma_j(0)$ for all i. Then there a subsequence which converges to a geodesic ray with the same basepoint.*

Proof. For each n, $\gamma_i([0,n])$ is contained in a ball B_n of radius n about $\gamma_i(0)$. Since X is proper, B_n is compact, and thus there a subsequence of $(\gamma_i([0,n])$ which converges. First consider the convergent subsequence for $n = 1$, and then pass to a further convergent subsequence for $n = 2$. Continuing in this way produces a subsequence of (γ_i) which converges to a continuous function γ. It is straightforward (and a good exercise!) to show that γ is a geodesic. $\square$

6.1 The boundary of a hyperbolic space

We will first define several boundaries (as sets) of a hyperbolic space, and then we will show that there are bijections between any pair. After doing so, we will put topologies on each of the boundaries and show that the bijections above induce homeomorphisms of the various boundaries. Thus any of the following definitions of boundaries of a hyperbolic space can be called "the boundary". This boundary we introduce is the *Gromov boundary* or *visual boundary* of a hyperbolic space.

Throughout this section, X will be a proper, geodesic hyperbolic space with a basepoint x_0, unless otherwise noted.

We will introduce the boundary of a hyperbolic space in three different ways: as equivalence classes of geodesics, as equivalence classes of quasi-geodesics, and as equivalences classes of sequences of points. We begin with the geodesic (and quasi-geodesic) viewpoint.

Let $\gamma_1 \colon [0, \infty)$ and $\gamma_2 \colon [0, \infty)$ be (quasi-)geodesic rays in X. We say that γ_1 and γ_2 are *asymptotic*, and write $\gamma_1 \sim \gamma_2$, if $d_{Haus}(\gamma_1, \gamma_2) < \infty$. Equivalently, there exists a $K \geq 0$ such that for all $t \in [0, \infty)$, $d(\gamma_1(t), \gamma_2(t)) \leq K$. The relation $\sim$ is an equivalence relation on the set of (quasi-)geodesic rays.

Definition 6.3 (Geodesic boundary)**.** Define the *relative geodesic boundary of X (with resepect to x_0)* as

$$\partial^g_{x_0} X = \{\gamma \subset X \mid \gamma \colon [0, \infty) \text{ a geodesic ray based at } x_0\}/\sim .$$

Define the *geodesic boundary of X* as

$$\partial^g X = \{\gamma \subset X \mid \gamma \colon [0, \infty) \text{ a geodesic ray}\}/\sim .$$

We make similar definitions using quasigeodesics rather than geodesics.

Definition 6.4. Define the *quasigeodesic boundary of X* as

$$\partial^q X = \{\gamma \subset X \mid \gamma \colon [0, \infty) \text{ a geodesic ray}\}/\sim .$$

One downside of the above definitions is that thy requires the space X to be (quasi)geodesic. There is a definition of the boundary which holds in non-geodesic spaces, involving sequences of points.

Definition 6.5. A sequence $(x_n)_{n\geq 1}$ in X *converges to infinity* if

$$\liminf_{i,j\to\infty} (x_i \mid x_j)_{x_0} = \infty.$$

This definition does not depend on the choice of basepoint. Two sequences $(x_n)_{n\geq 1}$ and $(y_n)_{n\geq 1}$ that converge to infinity are *equivalent* if

$$\liminf_{i,j\to\infty} (x_i \mid y_j)_{x_0} = \infty.$$

If this is the case, we write $(x_n) \sim (y_n)$.

As above, the relation $\sim$ is an equivalence relation on the set of sequences converging to infinity.

Definition 6.6 (Sequential boundary)**.** Let X be a (not necessarily geodesic) hyperbolic metric space. Define the *sequntial boundary of X* to be

$$\partial_s X = \{(x_n)_{n\geq 1} \mid (x_n) \text{ converges to } \infty \text{ in } X\}/\sim .$$