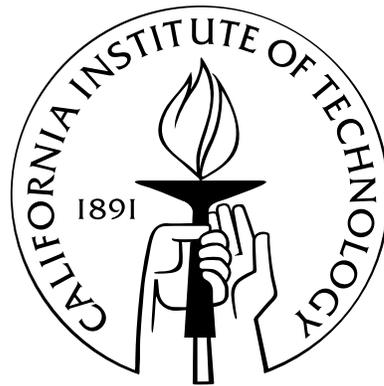


Quasiconvex Subgroups and Nets in Hyperbolic Groups

Thesis by
Thomas Patrick Mack

In Partial Fulfillment of the Requirements
for the Degree of
Doctor of Philosophy



California Institute of Technology
Pasadena, California

2006
(Defended May 12, 2006)

For my grandparents

Acknowledgements

I would like to thank my adviser, Danny Calegari, for his advice and guidance; Nathan Dunfield and Dinakar Ramakrishnan for their helpfulness in preparing this work; Michael Aschbacher and Hee Oh for graciously agreeing to serve on my Thesis Advisory Committee; and my family and friends for their patience and support during my research.

Abstract

Consider a hyperbolic group G and a quasiconvex subgroup $H \subset G$ with $[G : H] = \infty$. We construct a set-theoretic section $s : G/H \rightarrow G$ of the quotient map (of sets) $G \rightarrow G/H$ such that $s(G/H)$ is a net in G ; that is, any element of G is a bounded distance from $s(G/H)$. This section arises naturally as a set of points minimizing word-length in each fixed coset gH . The left action of G on G/H induces an action on $s(G/H)$, which we use to prove that H contains no infinite subgroups normal in G .

Contents

Acknowledgements	iv
Abstract	v
1 Introduction	1
2 The Geometry of Hyperbolic Spaces	3
3 Quasiconvexity	16
4 Finite Automata	22
5 Nets in Groups	32
6 Nets in Hyperbolic Groups	36
Bibliography	45

Chapter 1

Introduction

Let G be a group with finite generating set Σ . The Cayley graph $\Gamma = C(G, \Sigma)$ is defined to be the graph with vertex set G and edges connecting those $g, g' \in G$ with $g = g'\sigma$ for some $\sigma \in \Sigma$. (We assume that Σ is closed under inversion, so that this relation is symmetric.) A hyperbolic group is one for which Γ has the large-scale structure of a tree. That is, geodesics are “almost” unique, in the sense that there exists a constant $C > 0$ such that any two geodesics $\gamma(t), \gamma'(t)$ between the same points satisfy $d(\gamma(t), \gamma'(t)) < C$ for all t .

Let H be a finitely generated subgroup of G , and choose a finite generating set Σ' for it. Assume without loss of generality that $\Sigma' \subset \Sigma$. Then the Cayley graph $\Gamma' = C(H, \Sigma')$ naturally embeds in Γ , and so there are two natural metrics on $\Gamma' \subset \Gamma$: the path-length metric d' considering chains that remain in Γ' for all time, and the path-length metric d considering all chains in Γ . A quasiconvex subgroup is one for which these two metrics differ by no more than a constant multiplicative factor.

The purpose of this paper is to prove the following theorem:

Theorem. *Let G be a hyperbolic group, and let $H \subset G$ be a quasiconvex subgroup. If $[G : H] = \infty$, then there exists a (set-theoretic) section $s : G/H \rightarrow H$ of the quotient map $G \rightarrow G/H$ such that $s(G/H)$ is a net in G ; that is, $\sup_{g \in G} d(g, s(G/H))$ is finite.*

The argument depends on showing that for suitable s , there exists a finite automaton recognizing the language L of points in $s(G/H)$. By a geometric argument, any point in Γ lies within a bounded distance of the prefix closure \overline{L} of L . Since L is a regular language, it follows that any point in \overline{L} is a bounded distance from a point in L , proving the theorem.

The first three preliminary sections of the paper summarize general results in hyperbolic geometry that are used in the subsequent sections. Section 2 is a broad overview of hyperbolic topological spaces and hyperbolic groups, including many examples of such objects. Section 3 covers the properties of quasiconvex subgroups and several methods of their construction. General references for these two sections include [7], which discusses hyperbolic spaces and groups in detail; [2], which covers the large-scale geometry of general metric spaces; and [8], which outlines many results in combinatorial

group theory used throughout the paper. Section 4 is an overview of finite automata and regular languages. This machinery is useful not only for its direct use in the proof of the main theorem, but also because arbitrary hyperbolic groups have an automatic structure; the set of geodesics in Γ can be recognized by a finite automaton. The general material on finite automata can be found in [5], and [3] contains specific applications to hyperbolic groups. Section 5 is a brief summary of the problem of finding nets as in the main theorem for arbitrary groups, including a few examples and counter-examples. Section 6 contains the proof of the main theorem and a result about normal subgroups embedded in quasiconvex subgroups that follows from it.

Chapter 2

The Geometry of Hyperbolic Spaces

In this section, we recall the basic definitions of hyperbolic spaces and machinery of coarse geometry. Most of the material on hyperbolic geometry below is contained in [7] or [2], and the general machinery of combinatorial group theory is covered in [8]. There are many equivalent definitions of hyperbolic metric spaces, but we will mainly consider them as spaces in which geodesic triangles are *thin*: any side is contained in a bounded neighborhood of the other two sides, with the size of that bound independent of the particular triangle chosen. It is clear that Euclidean space does not have this property above dimension 1. On the other hand, it does hold in hyperbolic space \mathbb{H}^n . A group is hyperbolic if it has the large-scale geometry of a hyperbolic space when considered as a metric space. In order to make this more precise, we first make two definitions.

Definition. Let (X, d) be a metric space. A *geodesic* on X is a map f from an interval $[0, N]$ or $[0, \infty)$ in \mathbb{R} to X such that $d(ft, ft') = |t' - t|$ for all t, t' . If the former case, we call f a *geodesic segment* of length N ; otherwise, we call f a *geodesic ray*. For fixed $K > 0$ and $\epsilon \geq 0$, a (K, ϵ) -*quasigeodesic segment* (resp. (K, ϵ) -*quasigeodesic ray*) is a map $f : [0, N] \rightarrow X$ (resp. $f : [0, \infty) \rightarrow X$) such that $K^{-1}|t' - t| \leq d(ft, ft') \leq K|t' - t| + \epsilon$ for all t, t' . The space X is *geodesic* if there exists a geodesic between any two of its points.

A geodesic is thus an isometry from an interval of $\mathbb{R}^{\geq 0}$ into X . We denote a geodesic $f : [0, N] \rightarrow X$ with endpoints $f(0) = p$ and $f(N) = p'$ by $[pp']$. Such a path is not necessarily unique; consider geodesics joining antipodal points on S^n , for example.

Definition. A geodesic metric space X is *hyperbolic* if it satisfies the following two equivalent properties:

- (i) There exists a constant $\delta \geq 0$ such that for any geodesic triangle in Γ with vertices p, q, r , the side $[pq]$ lies in $\overline{N_\delta([pr] \cup [qr])}$, where $\overline{N_\epsilon(S)} = \{x \in X : d(x, p) \leq \epsilon \text{ for some } p \in S\}$ denotes the closed ϵ -neighborhood around S .

(ii) There exists a constant $\delta \geq 0$ such that

$$(p.q)_x \geq \min \{(p.r)_x, (q.r)_x\} - \delta \quad (2.1)$$

for all $p, q, r, x \in \Gamma$, where $(\cdot \cdot)$ denotes the *Gromov product*

$$(p.q)_x = \frac{1}{2} (d(p, x) + d(q, x) - d(p, q)).$$

To motivate the definition of the Gromov product, consider a tree X . Fix vertices $p, q, r \in X$. Since X is a tree, there exists a unique simple path between any two of its nodes. Any simple path in X is therefore a geodesic; in particular, geodesics in X are unique. Thus X satisfies condition (i) in the definition above for any $\delta > 0$. Consider the geodesics $[rp]$ and $[rq]$, and write $[rp] = (r = x_0, \dots, x_n = p)$ and $[rq] = (r = y_0, \dots, y_m = q)$ for some $x_i, y_i \in X$. Let k denote the largest index such that $x_i = y_i$. Suppose that $x_i = y_i$ for some $l > k$. Assuming without loss of generality that l is the smallest such index, the chains (x_k, \dots, x_l) and (y_k, \dots, y_l) are distinct simple paths from $x_k = y_k$ to $x_l = y_l$. Thus $x_i \neq y_i$ for all $i > k$. The chain $(p = x_n, x_{n-1}, \dots, x_k, y_{k+1}, \dots, y_{m-1}, y_m = q)$ is then a simple path from p to q , and is thus a geodesic. Thus

$$\begin{aligned} (p.q)_r &= \frac{1}{2} (d(p, r) + d(q, r) - d(p, q)) \\ &= \frac{1}{2} (|(x_0, \dots, x_n)| + |(x_0, \dots, x_m)| - |(x_n, x_{n-1}, \dots, x_k, y_{k+1}, \dots, y_{m-1}, y_m)|) \\ &= k. \end{aligned} \quad (2.2)$$

Thus $(p, q)_r$ measures the length of time for which geodesics $[rp]$ and $[rq]$ coincide. It follows immediately from (2.2) or this geometric interpretation that $(p.q)_x \geq \min \{(p.r)_x, (q.r)_x\}$ for all $p, q, r, x \in X$.

If X satisfies condition (i) or (ii) in the preceding definition for some value of δ , then it also satisfies it for all $\delta' \geq \delta$. Thus call X δ_0 -hyperbolic if it satisfies conditions (i) and (ii) above with $\delta = \delta_0$; X is then hyperbolic iff it is δ -hyperbolic for some δ . One of the most important basic results in hyperbolic geometry is the Morse Lemma, which states that quasigeodesics stay uniformly close to geodesics. In fact, this property characterizes hyperbolic spaces; any geodesic space for which the conclusion of the Morse Lemma holds is necessarily hyperbolic [2, 8.4]. Using this criterion, it is easily verified that hyperbolic space \mathbb{H}^n is indeed a hyperbolic metric space for all n .

Lemma 2.1 (Morse Lemma, [2, 8.4.20]). *Let (X, d) be a δ -hyperbolic geodesic space. For any $K > 0$ and $\epsilon \geq 0$, there exists a constant $C > 0$, depending only on δ, K , and ϵ , such that any (K, ϵ) -quasigeodesic segment with endpoints $p, q \in X$ lies in the C -neighborhood of a geodesic $[pq]$.*

Since we are interested in only the coarse geometry of metric spaces, it is useful to consider a

class of maps slightly broader than isometries. For any metric spaces (X, d_X) and (Y, d_Y) , a function $f : X \rightarrow Y$ (not necessarily continuous) is a (C, ϵ) -quasi-isometric embedding for $C > 0, \epsilon \geq 0$ if the inequalities

$$C^{-1}d_Y(fx, fx') \leq d_X(x, x') \leq Cd_Y(fx, fx') + \epsilon$$

hold for all $x, x' \in X$. If the precise values of C or ϵ are irrelevant, call f a C -quasi-isometric embedding or simply a quasi-isometric embedding. A quasi-isometry $X \rightarrow Y$ is a quasi-isometric embedding $f : X \rightarrow Y$ such that $\overline{N_r(f(X))} = Y$ for some r . Two spaces X, Y are quasi-isometric if there exists a quasi-isometry $X \rightarrow Y$. The equivalence relation of quasi-isometry is the large-scale analogue of isometry. As such, it is important to note that hyperbolicity is a quasi-isometry invariant.

Corollary 2.2. *Let X, X' be quasi-isometric length spaces. If X is hyperbolic, then so is X' .*

Proof. The corollary follows immediately from the Morse Lemma. \square

Since geodesic triangles in a hyperbolic space are thin, geodesics to nearby points remain close for long periods of time. For trees, it was shown above that two geodesics $[rp], [rq]$ coincide for a time equal to the Gromov product $(p.q)_r$. A similar result holds in an arbitrary hyperbolic space.

Lemma 2.3. *Let (X, d) be a δ -hyperbolic geodesic space. Suppose r, r' are geodesic segments in X that have endpoints q, q' , respectively, and the same initial point p . For all $0 \leq t \leq (q.q')_p$, we have $d(r_t, r'_t) \leq 4\delta$. In particular, we have $d(r_t, r'_t) \leq 4\delta$ for all time t if $q = q'$.*

Proof ([2, 8.4.2]). Fix a geodesic $[qq']$. For any $x \in [qq']$ and $y \in X$, we have

$$\begin{aligned} (y.q)_x + (y.q')_x &= d(y, x) + \frac{1}{2}(d(q, x) - d(y, q) + d(q', x) - d(y, q')) \\ &= d(y, x) + \frac{1}{2}(d(q, q') - d(y, q) - d(y, q')) \\ &= d(y, x) - (q.q')_y. \end{aligned}$$

It follows that $d(y, [qq']) \geq (q.q')_y$. For any t , we have

$$\begin{aligned} (q.q')_{r_t} &= \frac{1}{2}(d(q, r_t) + d(q', r_t) - d(q, q')) \\ &\geq \frac{1}{2}(d(q, p) + d(q', p) - d(q, q')) - d(p, r_t) \\ &= (q.q')_p - t. \end{aligned}$$

Thus if $t < (q.q')_p - \delta$, then $d(r_t, [qq']) \geq (q.q')_{r_t} > \delta$. Since X is δ -hyperbolic, the point $r_t \in [pq]$ must therefore lie in $\overline{N_\delta(r')}$. Hence $d(r_t, r'_t) \leq 2\delta$ for $t < (q.q')_p - \delta$. The lemma follows. \square

Corollary 2.4. *Let (X, d) be a δ -hyperbolic geodesic space, and let r and r' be geodesic rays from a point p . If $d(r_t, r'_t)$ is bounded as $t \rightarrow \infty$, then $d(r_t, r'_t) \leq 4\delta$ for all time t .*

Proof. Set $C = \sup_t d(r_t, r'_t)$. For any $t \geq 0$,

$$\begin{aligned} (r(t+C), r'(t+C)) &= \frac{1}{2} (|r(t+C)| + |r'(t+C)| - d(r(t+C), r'(t+C))) \\ &= t + C - \frac{1}{2} d(r(t+C), r'(t+C)) \\ &\geq t. \end{aligned}$$

Thus $d(r_t, r'_t) \leq 4\delta$ for all time $t \geq 0$ by Lemma 2.3. \square

Corollary 2.5. *Let X be a δ -hyperbolic space. For any $p, q, q' \in X$, let $\alpha_p(q, q')$ denote the minimum value of t for which there exist geodesic segments $\gamma = [pq]$ and $\gamma' = [pq']$ such that $d(\gamma_t, \gamma'_t) \geq 4\delta$; if all such geodesics remain a distance less than 4δ apart for all time, set $\alpha_p(q, q') = \min \{d(p, q), d(p, q')\}$. Then $(q, q')_p \leq \alpha_p(q, q') \leq (q, q')_p + 2\delta$.*

Proof. Let $p, q, q' \in X$, and choose geodesic segments $\gamma = [pq]$ and $\gamma' = [pq']$. To simplify notation, set $t = \alpha_p(q, q')$. By the definition of α , we have $d(\gamma_t, \gamma'_t) \leq 4\delta$. Hence

$$\begin{aligned} 2(q, q')_p &= d(p, q) + d(p, q') - d(q, q') \\ &= d(p, \gamma_t) + d(\gamma_t, q) + d(p, \gamma'_t) + d(\gamma'_t, q) - d(q, q') \\ &= 2t + d(\gamma_t, q) + d(\gamma'_t, q) - d(q, q') \\ &\geq 2t - 4\delta + d(\gamma_t, q) + d(\gamma'_t, q) - d(q, q') \\ &\geq 2t - 4\delta. \end{aligned}$$

Thus $t \leq (q, q')_p + 2\delta$. The opposite inequality follows immediately from Lemma 2.3. \square

In order to apply this geometric machinery to a group G , we consider its canonical action on a certain space, its Cayley graph. Hyperbolic groups are ones for which this associated space is hyperbolic in the sense defined above. Since G acts freely on this graph by isometries, hyperbolicity of the Cayley graph imposes important constraints on the algebraic structure of G . To simplify notation, we make the following convention throughout this paper:

Convention. All generating sets are assumed to be finite and closed under inversion.

Definition. Let G be a group with generating set Σ . Define the Cayley graph $\Gamma = C(G, \Sigma)$ to be the graph with vertex set G and edges connecting those vertices g, g' with $g' = gc$ for some $c \in \Sigma$.

We often implicitly identify G with the set of vertices in $\Gamma = C(G, \Sigma)$, or even with Γ itself. The free group $F(\Sigma)$ with basis Σ admits a homomorphism onto G sending a word $[c_1] \cdots [c_n]$ to the

corresponding element $c_1 \cdots c_n$ of G . Denote this evaluation map by either $w \rightarrow \pi(w)$ or $w \rightarrow \bar{w}$. To simplify notation, we often write \bar{w} or $\pi(w)$ simply as w when the intended meaning is clear. For any $g \in G$, the *length* of G with respect to Σ is defined to be

$$|g|_\Sigma = \min \{|w| : \alpha \in \pi^{-1}(w)\},$$

where $|w|$ is the usual word length in $F(\Sigma)$; explicitly, $|1| = 0$ and

$$|\sigma_1^{n_1} \cdots \sigma_r^{n_r}| = |n_1| + \cdots + |n_r|$$

for any nontrivial reduced word in $F(\Sigma)$. Metrize Γ by setting $d(x, y) = |x^{-1}y|_\Sigma$ for all vertices $x, y \in G$, and extend this metric over the edges to make Γ into a geodesic length space. In particular, the length of a chain (x_0, \dots, x_n) in the graph Γ is simply n .

Definition. A group G with generating set Σ is (*word- or Gromov-*)*hyperbolic* if the Cayley graph $C(G, \Sigma)$ is a hyperbolic geodesic space.

Any two generating sets Σ, Σ' of G satisfy $C^{-1}|x|_{\Sigma'} \leq |x|_\Sigma \leq C|x|_{\Sigma'}$, where

$$C \geq \max \{|\sigma|_{\Sigma'} : \sigma \in \Sigma\}, \max \{|\sigma|_\Sigma : \sigma \in \Sigma'\}.$$

Hence the identity map $C(G, \Sigma) \rightarrow C(G, \Sigma')$ is a quasi-isometry. By Corollary 2.2, word-hyperbolicity is therefore independent of the choice of generating set Σ . Since the Cayley graph of G is independent of the choice of generating set up to quasi-isometry, we write $C(G)$ for $C(G, \Sigma)$ when only its coarse geometry is important.

It was shown above that any tree is 0-hyperbolic. Arbitrary hyperbolic spaces are therefore ones that have, in a rough sense, the large-scale geometry of a tree. For any free group G with basis $\{x_1, \dots, x_n\}$, the Cayley graph $C(G, \{x_1^{\pm 1}, \dots, x_n^{\pm 1}\})$ is a tree of degree $2n$. Conversely, any Cayley graph $\Gamma = C(G, \Sigma)$ of an arbitrary group G defines a presentation $G = \langle \Sigma \mid R \rangle$, where R is the set of cycles in Γ with basepoint 1. Thus if Γ is a tree, then G is the free group generated by Σ . Hyperbolic groups therefore represent, in the same sense as above, groups with large-scale structure similar to that of a free group. To illustrate this correspondence, we list several examples of hyperbolic groups below.

- (i) Any finite group G has a Cayley graph quasi-isometric to a single point, which is clearly 0-hyperbolic. Hence G is also hyperbolic.
- (ii) Let G be a discrete, cocompact subgroup of the group $\text{Isom}^+(\mathbb{H}^n) = \text{SO}^+(n, 1)$ of orientation-preserving isometries of \mathbb{H}^n . If G acts freely on \mathbb{H}^n , then the map $g \rightarrow x_0.g$ is a quasi-isometry $C(G) \rightarrow \mathbb{H}^n$ for any basepoint $x_0 \in \mathbb{H}^n$ [3]. Hence G is hyperbolic. The quotient \mathbb{H}^n/G is

also a compact hyperbolic manifold, and its fundamental group is isomorphic to G . In low dimensions, we have $SO^+(1,1) = \mathbb{R}$, $SO^+(2,1) = PSL_2(\mathbb{R})$, and $SO^+(3,1) = PSL_2(\mathbb{C})$.

- (iii) Conversely, let M^n be a closed, hyperbolic manifold. The fundamental group $G = \pi_1 M$ acts on the universal cover $\tilde{M} = \mathbb{H}^n$ by isometries. Fix a basepoint $x_0 \in \mathbb{H}^n$. The map $g \rightarrow x_0.g$ is then a quasi-isometry $C(G) \rightarrow \mathbb{H}^n$ by the same proof as in (ii). Hence G is hyperbolic.
- (iv) More generally, consider a group G acting on a proper geodesic metric space X by isometries. If the action is proper and cocompact, then G is quasi-isometric to X [8, IV.23]. Thus if X is a hyperbolic space, then G is a hyperbolic group.
- (v) As indicated above, free groups of finite rank are hyperbolic.
- (vi) Under certain conditions, discussed in Propositions 3.9 and 3.10 below, amalgamated free products and HNN-extensions of hyperbolic groups are also hyperbolic. (For the definitions of these two constructions, see Section 3 below.)
- (vii) Let G be a hyperbolic group, and let $H \subset G$. Suppose $[G : H]$ is finite. Choose a generating set Σ for H , then extend it to a generating set Σ' of G . Then the inclusion $C(H, \Sigma) \rightarrow C(G, \Sigma')$ is a quasi-isometry, so H is hyperbolic. Conversely, the same quasi-isometry shows that a group G is hyperbolic if it contains a hyperbolic subgroup of finite index. For example, the commutator subgroup of $SL_2(\mathbb{Z})$ has index 12 and is free of rank 2 [1, IX.6], so $SL_2(\mathbb{Z})$ is hyperbolic.
- (viii) As a special case of the preceding construction, extensions of hyperbolic groups by finite groups are also hyperbolic. An *elementary* hyperbolic group G is one that is a finite extension of a cyclic group. Thus G is either finite or an extension of \mathbb{Z} by a finite group.
- (ix) Let N be a closed surface of genus $g \geq 2$. Consider two foliations $\mathcal{F}^s, \mathcal{F}^u$ on N that are transverse except at a common set of singularities x_1, \dots, x_n . Assume each x_i is a p_i -pronged saddle singularity for some $p_i \geq 3$, as shown in Figure 2.1. Let μ^s and μ^u be transverse measures on \mathcal{F}^s and \mathcal{F}^u , respectively. An orientation-preserving homeomorphism $\varphi : N \rightarrow N$ is *pseudo-Anosov* [4, §6] if there exists a constant $\lambda > 1$ such that

$$\varphi^* \mu^s(\alpha) = \lambda \mu^s(\alpha); \quad \varphi^* \mu^u(\beta) = \lambda^{-1} \mu^u(\beta)$$

for all arcs α transverse to \mathcal{F}^s and β transverse to \mathcal{F}^u . Given such a map, define M to be its mapping torus. Thus $M = N \times [0, 1] / \sim$, where $(x, 0) \sim (\varphi(x), 1)$ for all $x \in N$. By Thurston's Hyperbolization Theorem [13], M is a hyperbolic manifold. Hence $\pi_1 M$ is a hyperbolic group. Since N and M are hyperbolic and thus have contractible universal covers, all $\pi_1 N$ and $\pi_1 M$

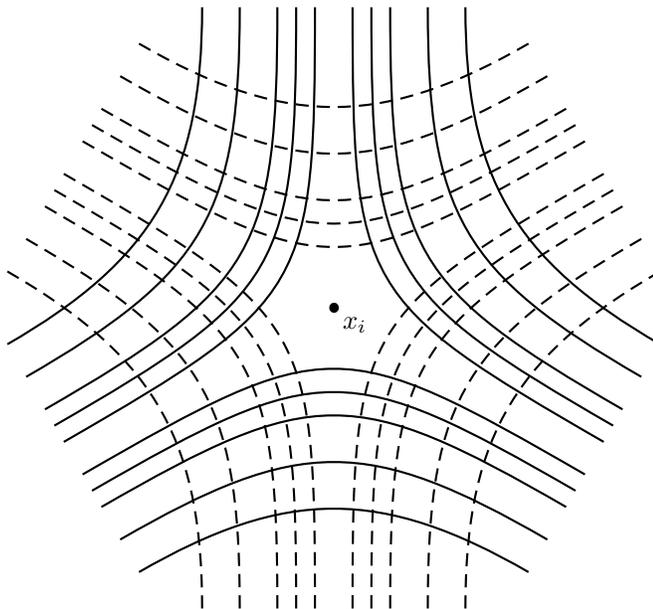


Figure 2.1: A p -pronged saddle singularity at x_i for $p = 3$ (cf. Figure 5 in [14]). The foliations \mathcal{F}^s and \mathcal{F}^u are shown with dashed and solid lines, respectively.

vanish for $i \geq 2$. The long exact sequence corresponding to the bundle $N \rightarrow M \rightarrow S^1$ gives an exact sequence

$$1 \longrightarrow \pi_1 N \longrightarrow \pi_1 M \longrightarrow \mathbb{Z} \longrightarrow 1, \quad (2.3)$$

where a generator t of \mathbb{Z} acts on $\pi_1 N$ by $t^{-1}xt = \varphi_*(x)$. This sequence splits, giving an isomorphism $\pi_1 M = \pi_1 N \rtimes \mathbb{Z}$. Thus if N has genus $g \geq 2$, then $\pi_1 M$ has a presentation

$$\pi_1 M = \left\langle x_1, y_1, \dots, x_g, y_g, t \mid [x_1, y_1] \cdots [x_g, y_g] = 1, t^{-1}x_i t = \varphi_*(x_i), t^{-1}y_i t = \varphi_*(y_i) \right\rangle. \quad (2.4)$$

To simplify notation, we adopt the following convention:

Convention. Throughout this paper, G denotes a hyperbolic group with some fixed generating set Σ and Cayley graph $\Gamma = C(G, \Sigma)$. Denote the metric on Γ by $d(\cdot, \cdot)$, and write $|g|$ for $|g|_\Sigma = d(g, 1)$.

Since the Cayley graph Γ of a hyperbolic group is a hyperbolic geodesic space with only finitely many geodesics between any two points, we can extend the Morse Lemma 2.1 to geodesic rays on Γ .

Lemma 2.6. *Let G be a δ -hyperbolic group, and set $\Gamma = C(G, \Sigma)$ for some generating set Σ of G .*

For any $K > 0$ and $\epsilon \geq 0$, there exists a constant $C > 0$, depending only on δ, K , and ϵ , such that any (K, ϵ) -quasigeodesic ray from a point p lies within a distance C of a geodesic ray from p .

Proof. Let $r : [0, \infty) \rightarrow \Gamma$ be a (K, ϵ) -quasigeodesic ray from p . Let r^t denote the geodesic segment $r|_{[0, t]}$ for each $t \geq 0$. By the Morse Lemma, there exists some constant $C > 0$ and geodesic segments $\gamma^t = [pr(t)]$ for all integers $t > 0$ such that $d(r^t(s), \gamma^t(s)) < C$ for all time $0 \leq s \leq t$. Define a sequence $x_0, x_1, \dots \in \Gamma$ inductively as follows. Set $x_0 = 1$. For each $n > 0$, choose some $x_n \in \Gamma$ such that $d(p, x_n) = n$ and there exist infinitely many t such that $x_0, \dots, x_{n-1}, x_n \in \gamma^t$; since the set $\{g \in G : d(p, g) = n\}$ is finite, such a point x_n must exist. For any $i, j \geq 0$, there exists some $t > i, j$ such that $x_i, x_j \in \gamma^t$. Then $d(x_i, x_j) = d(\gamma^t(i), \gamma^t(j)) = |i - j|$. Hence the set $\{x_i\}$ defines a geodesic ray from p . Furthermore, we have $d(x_i, r) = d(\gamma^t(i), r) \leq d(\gamma^t, r^t(i)) < C$. Thus r is the required geodesic. \square

Any isometry of \mathbb{H}^n extends across the sphere at infinity $\partial\mathbb{H}^n$. Thus any group of isometries acts on $\partial\mathbb{H}^n = S^{n-1}$ by homeomorphisms. Similarly, an arbitrary hyperbolic group acts on another analogous topological space, its hyperbolic boundary.

Definition. Let X be a hyperbolic space, and fix $x_0 \in X$. Let $\Omega(X)$ denote the set of sequences of points in X with $(p_i \cdot p_j)_{x_0} \rightarrow \infty$ as $i, j \rightarrow \infty$. Write $(p_i) \sim (p'_i)$ if $\liminf_{i, j \rightarrow \infty} (p_i \cdot p'_j)_{x_0} = \infty$. By (2.1), \sim is an equivalence relation on $\Omega(X)$. The set of equivalence classes $\Omega(X)/\sim$ is the *hyperbolic boundary* ∂X . For a hyperbolic group G , set $\Omega(G) = \Omega(\Gamma)$ and $\partial G = \partial\Gamma$.

Definition. Let X and Y be hyperbolic spaces. For any quasi-isometry $f : X \rightarrow Y$, let $f_\infty : \partial X \rightarrow \partial Y$ denote the map $f(p_i) = (fp_i)$.

It is clear that ∂G is independent of the choice of generating set Σ . We also have

$$|(p \cdot q)_x - (p \cdot q)_{x_0}| = \frac{1}{2} |d(p, x) - d(p, x_0) + d(q, x) - d(q, x_0)| \leq d(x, x_0) \quad (2.5)$$

for any $x, x_0 \in G$, so the definition is also independent of the choice of x_0 . To simplify notation, we take $x_0 = 1$ in defining ∂G and write $x \cdot x' = (x \cdot x')_1$. Let $\Omega_0(X)$ denote the set of geodesic rays from this fixed basepoint in X . For a group G , set $\Omega_0(G) = \Omega_0(\Gamma)$; by the Morse Lemma, $\Omega_0(G)$ is independent of the choice of Σ up to quasi-isometry. Each $r \in \Omega_0(X)$ satisfies $r_t \cdot r_{t'} = \frac{1}{2} (|t| + |t'| - |t - t'|) = \min\{t, t'\}$, so (r_t) lies in ∂X . The next lemma shows that every element of the boundary arises in this way.

Lemma 2.7. *For any δ -hyperbolic space X , the map $\Omega_0(X) \rightarrow \partial X$ given by $r \rightarrow (r_t)$ is surjective.*

Proof. Let (p_i) be an arbitrary sequence in ∂X . Since $p_i \cdot p_j \rightarrow \infty$, there exists for each $t > 0$ some N_t such that $p_i \cdot p_j > t$ for $i, j \geq N_t$. Choose a geodesic $[1p_{N_t}]$ for each t , and let q_t denote the point a distance t from 1 along this path. It is clear from the definition of ∂X that the sequence

$(p'_i) = (p_{N_i})$ is equivalent to (p_i) . For each i and j , we have $p'_i \cdot p'_j > \min\{i, j\}$. Since q_j lies on the chosen geodesic $[1p'_j]$, the Gromov product $q_j \cdot p'_j = \frac{1}{2}(|q_j| + |p'_j| - d(q_j, p'_j)) = |q_j| = j$. Thus

$$p'_i \cdot q_j \geq \min\{p'_i \cdot p'_j, q_j \cdot p'_j\} - \delta \geq \min\{i, j\} - \delta.$$

Hence $\liminf p'_i \cdot q_j = \infty$, and so $(p'_i) \sim (q_i)$.

For any $t \leq t'$, any geodesics $[1p_{N_t}]$ and $[1p_{N_{t'}}]$ stay a distance of at most 4δ apart until time t by Lemma 2.3. Hence

$$t' - t = |q_{t'}| - |q_t| \leq d(q_t, q_{t'}) = t + t' - 2(t, t') \leq t' - t + 4\delta. \quad (2.6)$$

For each i , choose a geodesic s_i connecting q_i to q_{i+1} . Let s denote the infinite chain (s_0, s_1, \dots) . By (2.6), s is a quasigeodesic. Thus by Lemma 2.7, there exists a geodesic ray such that $d(r_t, s)$ is bounded. The geodesic $r'_t = r_0^{-1}r_t \in \Omega_0(X)$ satisfies

$$r'_t \cdot s_t = \frac{1}{2}(|r'_t| + |s_t| - d(r'_t, s_t)) \geq t - \frac{1}{2}d(r'_t, s_t)$$

for all time t . Thus $r'_t \cdot s_t \rightarrow \infty$ as $t \rightarrow \infty$, and so $(r'_t) \sim (s_t) \sim (q_t) \sim (p_t)$. \square

The map in Lemma 2.7 is not a bijection in general. Consider $G = \mathbb{Z} \times \mathbb{Z}_2$, for example. The chains

$$\gamma_k(t) = \begin{cases} (t, 0) & \text{if } t \leq k; \\ (k, 1) & \text{if } t = k + 1; \\ (t - 1, 1) & \text{if } t \geq k + 2 \end{cases}$$

are distinct geodesics in G for all positive k . All the sequences $(\gamma_k(t))$ are equivalent, however, since $d(\gamma_k(t), \gamma_{k'}(t))$ is bounded for each k, k' . On the other hand, $\mathbb{Z} \times \mathbb{Z}_2$ is quasi-isometric to its finite index subgroup \mathbb{Z} . The Cayley graph $C(\mathbb{Z}, \{\pm 1\})$ is a tree, and the only geodesic rays from 0 in it are the two paths $\gamma_t = t$ and $\gamma'_t = -t$. Since $\gamma_t \cdot \gamma'_t = \frac{1}{2}(t + t' - (t' + t)) = 0$ for all t and t' , the sequences (γ_t) and (γ'_t) in ∂G are inequivalent. The map $\Omega_0(G) \rightarrow \partial G$ is thus a bijection in this case. For a hyperbolic space, the map $\Omega_0(\mathbb{H}^n) \rightarrow \partial \mathbb{H}^n$ is also bijective. To prove this result, suppose $(p_i) \sim (p'_i)$ in $\Omega(\mathbb{H}^n)$. By Lemma 2.7, there exist geodesics $r, r' \in \Omega_0(\mathbb{H}^n)$ such that $(p_t) \sim (r_t)$ and $(p'_t) \sim (r'_t)$. Thus $r_t \cdot r'_t \rightarrow \infty$ as $t \rightarrow \infty$. Thus $d(r_t, r'_t) \leq 4\delta$ for all time t by Lemma 2.3. But it follows from a brief computation [12, 11.6.8] that no two distinct geodesics in \mathbb{H}^n stay a bounded distance from each other for all time, so $r = r'$.

The Gromov product can be extended across the boundary ∂X , and the resulting function defines

a metric on ∂X . Fix some basepoint $p_0 \in X$. For any $x, y \in \partial X$, define

$$(x.y)_{p_0} = \sup_{\substack{(x_i) \sim x \\ (y_j) \sim y}} \liminf_{i,j \rightarrow \infty} (x_i.y_j)_{p_0}. \quad (2.7)$$

For any sequences $(x_i) \sim (x'_i)$ in $\Omega(G)$, we have

$$\begin{aligned} \liminf_{i,j \rightarrow \infty} (x'_i.y_j)_{p_0} &\geq \liminf_{i,j \rightarrow \infty} \min \{(x'_i.x_j)_{p_0}, (x_i.y_j)_{p_0}\} - \delta \\ &= \liminf_{i,j \rightarrow \infty} (x_i.y_j)_{p_0} - \delta, \end{aligned}$$

since $\liminf (x_i.p'_j)_{p_0} = \infty$. Applying the same argument to p_i gives

$$\liminf_{i,j \rightarrow \infty} (p_i.q_j)_{p_0} - \delta \leq \liminf_{i,j \rightarrow \infty} (p'_i.q_j)_{p_0} \leq \liminf_{i,j \rightarrow \infty} (p_i.q_j)_{p_0} + \delta.$$

In particular, the supremum in (2.7) is finite. Give ∂X the topology generated by the basis of closed sets $\overline{N_r(x)} = \{y \in \partial X : (x.y)_{p_0} \geq r\}$ for all $x \in \partial X$ and $r \geq 0$. Note that this topology is independent of the choice of the basepoint by (2.5). The following proposition describes the structure of ∂X as a metrizable space. Since we do not directly use any of the topological properties of ∂X below, we omit its proof.

Proposition 2.8. *Let X be a hyperbolic space.*

- (i) *The topology defined above for ∂X is metrizable.*
- (ii) *If X is proper, then ∂X is compact.*
- (iii) *For any hyperbolic space Y and any quasi-isometry $f : X \rightarrow Y$, the map $f_\infty : \partial X \rightarrow \partial Y$ is continuous.*

Proof. See Section 1.8 in [7]. □

In hyperbolic space \mathbb{H}^n , it can be shown [12, 11.6] that the boundary $\partial\mathbb{H}^n$ defined above is homeomorphic to the usual sphere at infinity S^{n-1} . Furthermore, the point in $\partial\mathbb{H}^n$ corresponding to a sequence $(p_i) \in \Omega(\mathbb{H}^n)$ is its limit in this sphere [8, V.58]. By considering the geometric action of an arbitrary hyperbolic group G on its boundary ∂G , many algebraic properties of G can be determined. The following three useful propositions are examples of the results this technique produces.

Proposition 2.9 ([8, V.58]). *Let $g \in G$. If g is not torsion, then the centralizer $C_G(g)$ of g is virtually cyclic. (That is, $C_G(g)$ is a finite extension of a cyclic group.)*

Proposition 2.10 ([8, V.58]). *If G is not elementary, then it contains a nonabelian free group.*

In Example (iii) above, we showed that any closed, hyperbolic manifold has word-hyperbolic fundamental group. This result does not hold without the assumption of compactness. Consider the manifold M obtained by removing a tubular neighborhood of the figure-8 knot in \mathbb{R}^3 . This space can be constructed by gluing two regular ideal tetrahedra together, and therefore has the structure of a hyperbolic manifold of finite volume [11, 4.4.2]. The manifold M is the total space of a bundle over S^1 with fiber T' , where T' denotes the torus T^2 with a point removed. The long exact sequence in homology corresponding to the bundle $T' \rightarrow M \rightarrow S^1$ gives an exact sequence

$$1 \longrightarrow \mathcal{F}_2 \longrightarrow \pi_1 M \longrightarrow \mathbb{Z} \longrightarrow 1,$$

where $\mathcal{F}_2 = \langle a, b \rangle$ is the free group of rank 2. Considering the monodromy of this bundle gives a presentation

$$\pi_1 M = \langle a, b, t \mid t^{-1}at = aba, t^{-1}bt = ba \rangle; \quad (2.8)$$

that is, $\pi_1 M$ is the HNN-extension of \mathcal{F}_2 by the homomorphism $a \rightarrow aba, b \rightarrow ba$ [11, 4.4.3]. Thus

$$t^{-1}[a^{-1}, b^{-1}]t = t^{-1}aba^{-1}b^{-1}t = (aba)(ba)(a^{-1}b^{-1}a^{-1})(a^{-1}b^{-1}) = aba^{-1}b^{-1} = [a^{-1}, b^{-1}].$$

It is clear from the presentation (2.8) that $\langle t \rangle$ and $\langle [a, b] \rangle$ have trivial intersection, so $\langle t, [a^{-1}, b^{-1}] \rangle = \mathbb{Z}^2$. But no hyperbolic group can contain a subgroup isomorphic to \mathbb{Z}^2 by Proposition (2.9), so $\pi_1 M$ is not word-hyperbolic.

In addition to its boundary, an arbitrary hyperbolic space also acts on its corresponding Rips complex, another topological object associated to it. This space is a finite-dimensional CW-complex, and its quotient under the action by G is compact. Although we do not use it directly, we define the Rips complex and sketch a few of its most useful applications below.

Definition. For any fixed $d \geq 0$, define the *Rips complex* $P_d(G)$ of G to be the simplicial complex consisting of all subsets $S \subset G$ such that $d(x, y) \leq d$ for all $x, y \in S$.

Since G acts on Γ by isometries, it also acts on $P_d(G)$ by $gS = \{gs : s \in S\}$ for any $S \in P_d(G)$. The significance of the Rips complex to hyperbolic groups is a consequence of the proposition below.

Proposition 2.11. *For $d \geq 4\delta + 1$, the Rips complex $P_d(G)$ is finite-dimensional, cocompact, and contractible.*

Proof. For any fixed $g_0 \in G$, there are at most $\#\overline{N_d(g_0)} = \#\overline{N_d(1)}$ points $g \in G$ with $d(g, g_0) \leq d$. Thus $\dim P_d(G) \leq \#\overline{N_d(1)}$. For any $S \in P_d(G)$ and $s \in S$, the set $sS^{-1} \subset G$ contains 1. Hence each element of the quotient $G \backslash P_d(G)$ has a representative containing 1. Any element of $G \backslash P_d(G)$

therefore has a representative in the finite set $\overline{N_d(1)}$. Thus $P_d(G)$ is cocompact under the action of G . For the proof of the contractibility of $P_d(G)$, see [7, 1.7]. \square

The contractibility of $P_d(X)$ for sufficiently large d actually holds for any hyperbolic space X , not just Cayley graphs of hyperbolic groups. We only use the Rips complex in the context of the two results below, however, which are specific to hyperbolic groups. The first follows from the fact that $P_d(G)$ is simply-connected, and the second can be proved by considering the action of G on $P_d(G)$.

Proposition 2.12 ([8, V.56]). *The hyperbolic group G has a finite presentation.*

Proposition 2.13 ([8, V.56]). *There are only finitely many conjugacy classes of G that consist of torsion.*

As a final application of the Rips complex, we prove that hyperbolic groups have finite cohomology over \mathbb{Q} . The proof below is adapted from the derivation in [1, VII] of a spectral sequence computing the homology $H_*(G \backslash X)$ for a free G -complex X in terms of the homology of G with coefficients in the G -module $H_*(X)$. The basic argument is given in [1, VII.7], and the preliminary details are covered in [1, VII.5].

Proposition 2.14. *The \mathbb{Q} -module $H^*(G, \mathbb{Q})$ (where G acts trivially on \mathbb{Q}) is finitely generated. In particular, G has finite cohomological dimension over \mathbb{Q} .*

Proof. Set $X = P_{4\delta+1}(G)$. Let $C_* = C_*(X) \otimes \mathbb{Q}$ denote its cellular chain complex over \mathbb{Q} , considered as a module over $\mathbb{Z}G$. For any projective resolution F_* of G , we have a spectral sequence [1, VII.5.3]

$$E_{pq}^1 = H_q(G, C_p) = H_q(F_* \otimes_{\mathbb{Z}G} C_p) \Rightarrow H_{p+q}(F_* \otimes_{\mathbb{Z}G} C_*).$$

Let X_p denote the set of p -cells of X , and let Σ_p denote the set of equivalence classes of X_p modulo the action of G . Since $G \backslash X$ is a finite complex, Σ_p is finite for each p . Consider the action of G on each $\sigma \in \Sigma_p$. Since Σ_p is finite, there exists a subgroup $G' \subset G$ of finite index such that each $g \in G'$ acts on all $\sigma \in \Sigma_p$ by orientation-preserving maps. The index $[G : G']$ is invertible in \mathbb{Q} , so $H^*(G, \mathbb{Q})$ embeds in $H^*(G', \mathbb{Q})$ [1, III.10.4]. It is therefore sufficient to prove the result for G' . Thus assume without loss of generality that all $g \in G$ preserve the orientation of each cell of Σ_p . As an abelian group, the cellular chain complex C_p satisfies

$$C_p = \bigoplus_{\sigma \in X_p} \mathbb{Q}.$$

Thus we have an isomorphism of $\mathbb{Z}G$ -modules

$$C_p = \bigoplus_{\sigma \in \Sigma_p} \text{Ind}_{G_\sigma}^G \mathbb{Q}.$$

For fixed p , we therefore have

$$\begin{aligned} H_*(G, C_p) &= \bigoplus_{\sigma \in \Sigma_p} H_q(G, \text{Ind}_{G_\sigma}^G \mathbb{Q}) \\ &= \bigoplus_{\sigma \in \Sigma_p} H_q(G_\sigma, \mathbb{Q}) \\ &= \bigoplus_{\sigma \in \Sigma_p} H_q(G_\sigma) \otimes \mathbb{Q} \end{aligned} \tag{2.9}$$

by Shapiro's lemma [1, III.6.2]. Let $g \in G_\sigma$ for some p -cell σ of X . The group G_σ permutes the $p+1$ vertices of σ , so $|G_\sigma| \leq (p+1)!|H|$, where H is the stabilizer of some fixed vertex $\{x_0, \dots, x_p\}$ of σ . Each $h \in H$ satisfies $\{hx_0, \dots, hx_p\} = \{x_0, \dots, x_p\}$. It follows that either H is trivial or $\{x_0, \dots, x_p\} = \{1, g, \dots, g^{p-1}\}$ for some $g \in G$ of order p . In either case, H is cyclic of order dividing p . Thus G_σ is finite for all σ . Each group $H_p(G_\sigma)$ is hence annihilated by $|G_\sigma|$ for $p > 0$ [1, III.10.1]. Thus $H_p(G_\sigma) \otimes \mathbb{Q} = 0$ for all such p . By (2.9), we therefore have

$$H_*(G, C_p) = \bigoplus_{\sigma \in \Sigma_p} H_0(G_\sigma) \otimes \mathbb{Q} = \mathbb{Q}^{|\Sigma_p|}. \tag{2.10}$$

for all p . We have a spectral sequence [1, VII.7.2]

$$K_{pq}^2 = H_q(G, H_q(X, \mathbb{Q})) = H_q(F_p \otimes_{\mathbb{Z}G} H_q(C_*)) = H_q(F_p \otimes_{\mathbb{Z}G} C_*) \Rightarrow H_{p+q}(F_* \otimes_{\mathbb{Z}G} C_*). \tag{2.11}$$

By Lemma 2.11, the space X is contractible. Thus the sequence (2.11) degenerates to an isomorphism $H_p(G, \mathbb{Q}) = H_{p+q}(F_* \otimes_{\mathbb{Z}G} C_*)$. Hence

$$E_{pq}^1 = H_q(G, C_p) \Rightarrow H_{p+q}(G, \mathbb{Q}).$$

Since each E_{pq}^1 is a finite \mathbb{Q} -vector space with $E_{pq}^1 = 0$ for $q > \dim X$, each $H_q(G, \mathbb{Q})$ is finite-dimensional with $H_q(G, \mathbb{Q}) = 0$ for sufficiently large q . The required statement follows from the universal coefficient theorem for cohomology. \square

Chapter 3

Quasiconvexity

In a hyperbolic space, quasigeodesics remain uniformly close to geodesics by the Morse Lemma. Quasiconvex subspaces, those in which geodesic segments starting in the subspace remain in a bounded neighborhood of it for all time, therefore share many of the properties of convex subspaces. In the case of hyperbolic groups, many constructions hence carry over to quasiconvex subgroups. In particular, any quasiconvex subgroup of a hyperbolic group is also hyperbolic by the Morse Lemma. Recall that we continue to use the convention that G denotes a δ -hyperbolic group with generating set Σ and Cayley graph $\Gamma = C(G, \Sigma)$.

Definition. Let X be a hyperbolic space, and let $Y \subset X$. For any constant $K \geq 0$, the subspace Y is K -*quasiconvex* if any geodesic $[pq]$ with $p, q \in Y$ lies in $\overline{N_K(Y)}$. Call Y *quasiconvex* if it is K -quasiconvex for some K . For a hyperbolic group G with generating set Σ , a subgroup $H \subset G$ is *quasiconvex* if the set of vertices of the Cayley graph $\Gamma(G, \Sigma)$ corresponding to elements of H is quasiconvex.

Since geodesics in G between points in H remain a bounded distance from H for all time, the word-length of elements of H over a generating set of H should approximate word-length over a generating set of the larger group G . Geodesics in G between the same points remain a bounded distance apart for all time, so the same property holds for geodesics in H ; that is, H is hyperbolic. In order to make this argument rigorous, we require the following lemma.

Lemma 3.1 ([8, IV.49]). *Let $H \subset G$. If H is K -quasiconvex in G , then H is finitely generated.*

Proof. For each $x \in \Gamma$, choose a point $f(x) \in H$ minimizing $d(x, fx)$. Fix an arbitrary point $x \in H$, and choose a shortest path $[1x]$ in Γ . Write $[1x]$ as a chain $(1 = x_0, \dots, x_n = x)$ for some $x_i \in G$. Since H is K -quasigeodesic,

$$d(fx_{i+1}, fx_i) \leq d(x_i, fx_i) + d(x_{i+1}, fx_{i+1}) + d(x_{i+1}, x_i) = 2K + 1$$

for all i . Thus

$$x = fv_n(fv_{n-1})^{-1}fv_{n-1}(fv_{n-2})^{-1}\cdots(fv_2)^{-1}fv_2(fv_1)^{-1}fv_1$$

with each $fv_{i+1}(fv_i)^{-1} \in H \cap \overline{N_{2K+1}(1)}$. The finite set $H \cap \overline{N_{2K+1}(1)}$ therefore generates H . \square

Any quasiconvex subgroup $H \subset G$ therefore has a (finite) generating set Σ' . Assume without loss of generality that $\Sigma' \subset \Sigma$. Then the Cayley graph $\Gamma' = C(H, \Sigma')$ embeds in $C(G, \Sigma)$. By the Morse Lemma, this inclusion is a quasi-isometric embedding. The following lemma shows that this condition is not only necessary for quasiconvexity, but also sufficient.

Lemma 3.2. *Let G be a hyperbolic group, and let $H \subset G$ be a finitely generated subgroup. Choose a finite generating set Σ' for H , and assume without loss of generality that $\Sigma \supset \Sigma'$. Set $\Gamma' = C(H, \Sigma')$, and let d' denote the metric on Γ' . Then H is quasiconvex in G iff the inclusion map $(\Gamma', d') \rightarrow (\Gamma, d)$ is a quasi-isometric embedding.*

Proof. Suppose H is K -quasiconvex in G . For each $x \in \Gamma$, choose a point $f(x) \in \Gamma'$ minimizing $d(x, f(x))$. Any geodesic $[yy']$ in Γ with endpoints in Γ' lies in $\overline{N_K(\Gamma')}$, so any two points $p, p' \in [yy']$ satisfy $d(fp, fp') \leq d(p, p') + 2K$. Write $[yy']$ as a minimum-length chain $(y = x_0, \dots, x_n = y')$. Then $d(fx_i, fx_{i+1}) \leq 2K + 1$ for all i , so $d'(fx_i, fx_{i+1}) \leq K'$, where $K' = \max\{|g|_{\Sigma'} / |g|_{\Sigma} : |g|_{\Sigma} \leq 2K + 1\}$. Thus $f[yy']$ is a K' -quasigeodesic in Γ' . Hence $d'(y, y') \leq K'd(y, y')$ for all $y, y' \in \Gamma'$. Clearly $d'(y, y') \leq d(y, y')$, so $(\Gamma', d') \hookrightarrow (\Gamma, d)$ is a quasi-isometric embedding. The converse follows immediately from the Morse Lemma. \square

Corollary 3.3. *Let G be a hyperbolic subgroup, and let $H \subset G$. Then the hyperbolicity of a subgroup $H \subset G$ is independent of the choice of generating set for Σ .*

This corollary does not hold without the hyperbolicity assumption. For example, the subgroup $H = \langle(1, 1)\rangle$ of \mathbb{Z}^2 is clearly quasiconvex with respect to the generating set $\{\pm(1, 0), \pm(1, 1)\}$. The chain

$$\gamma_t = \begin{cases} (0, t) & \text{if } t = 0, \dots, n; \\ (t - n, n) & \text{if } t = n + 1, \dots, 2n, \end{cases}$$

however, is a shortest path with respect to the generating set $\{\pm(1, 0), \pm(0, 1)\}$, but $\max_t d(\gamma_t, H)$ is unbounded as $n \rightarrow \infty$ [8, IV.49].

Corollary 3.4. *Let G be a hyperbolic group, and let $H \subset G$. If H is quasiconvex, then it is a hyperbolic group.*

Proof. Suppose H is K -quasiconvex in G . By Lemma 3.1, H is finitely generated. Choose a generating set Σ' for H , and set $\Gamma = C(H, \Sigma')$. Let p_1, p_2, p_3 be arbitrary distinct points in Γ' , and let $\gamma_1 = [p_2 p_3], \gamma_2 = [p_3 p_1], \gamma_3 = [p_1 p_2]$ be geodesics in Γ . For each $x \in \Gamma$, choose a point $f(x) \in \Gamma'$ minimizing $d(f(x), x)$. It was shown in the proof of Lemma 3.2 that each $f\gamma_i$ is a $(1, 2K)$ -quasigeodesic in Γ' . Let C' denote the constant in the Morse Lemma corresponding to $C = 1$ and $\epsilon = 2K$ (in the notation of Lemma 2.1). Then there exist geodesics γ'_i in Γ' that have the same endpoints of γ_i and satisfy $d(\gamma'_i(t), \gamma_i(t)) < C'$ for all time t . Since G is hyperbolic, we have $\gamma_1 \subset \overline{N_\delta(\gamma_2 \cup \gamma_3)}$. It follows that $\gamma_1 \subset \overline{N_{\delta+C'}(\gamma'_2 \cup \gamma'_3)}$, and so H is $(\delta + C')$ -hyperbolic. \square

Lemma 3.1 can be used to prove that a given subgroup is not quasiconvex. In the opposite direction, the following lemmas give constructions for creating quasiconvex subgroups.

Lemma 3.5 ([8, IV.49]). *Let G be either a free group of finite rank or the fundamental group of a closed, hyperbolic surface. Then any finitely generated subgroup $H \subset G$ is quasiconvex.*

Lemma 3.6 ([8, IV.49]). *Let $H, K \subset G$. If H and K are both quasiconvex in G , so is $H \cap K$.*

Lemma 3.7. *Let G be a δ -hyperbolic group, and let $\varphi \in \text{Aut}(G)$. If φ has finite order, then the group $G^\varphi = \{g \in G : \varphi(g) = g\}$ is quasiconvex in G .*

Proof. Choose a finite generating set for G closed under φ . Fix $g \in G$, and let $g_1 \cdots g_n$ be a geodesic path to g . Then $\varphi(g_1) \cdots \varphi(g_n)$ is also a geodesic path to g , so $x_k = g_1 \cdots g_k$ satisfies $d(x_k, \varphi(x_k)) \leq 4\delta$ for $k = 1, \dots, n$. For each $c \in \overline{N_{4\delta}(1)}$, fix some h_c with $h_c^{-1}\varphi(h_c) = c$ if such an element exists. For any $x, y \in G$, we have $x^{-1}\varphi(x) = y^{-1}\varphi(y)$ iff $x^{-1}y \in G^\varphi$. It follows that for each k , we have $x_k \in G^\varphi h_c$ for some h_c . Hence G^φ is K -quasiconvex for $K = \max |h_c|$. \square

Thus, for example, the centralizer $C_G(x)$ of any torsion element of G is a quasiconvex subgroup. For elements of infinite order, the following lemma shows that $\langle x \rangle$ itself is quasiconvex. By Proposition 2.9, $C_G(x)$ is also quasiconvex in this case.

Lemma 3.8 ([7, 8.1.D]). *Any cyclic subgroup of G is quasiconvex.*

This result does not hold in arbitrary groups; in fact, it is useful as a criterion to determine whether a group is hyperbolic. For example, consider the Heisenberg group

$$H = \left\{ \begin{pmatrix} 1 & a & b \\ 0 & 1 & c \\ 0 & 0 & 1 \end{pmatrix} : a, b, c \in \mathbb{Z} \right\} \subset SL_3(\mathbb{Z}).$$

Set

$$Z = \begin{pmatrix} 1 & 0 & 1 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}.$$

For all integers n ,

$$\begin{aligned} \left[\begin{pmatrix} 1 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}^n, \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 1 \\ 0 & 0 & 1 \end{pmatrix}^n \right] &= \left[\begin{pmatrix} 1 & n & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & n \\ 0 & 0 & 1 \end{pmatrix} \right] \\ &= \begin{pmatrix} 1 & n & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & n \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & -n & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & -n \\ 0 & 0 & 1 \end{pmatrix} \\ &= \begin{pmatrix} 1 & 0 & n^2 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \\ &= Z^{n^2}. \end{aligned} \tag{3.1}$$

Thus over some generating set of H , the inequality $|Z^{n^2}| \leq 4n$ holds for all $n > 0$. The subgroup $\langle Z \rangle \subset H$ is therefore not quasiconvex.

For any positive integers p and q , the Baumslag-Solitar group $B(p, q)$ is defined by the presentation $\langle a, b \mid a^{-1}b^pa = b^q \rangle$; it is the HNN-extension of \mathbb{Z} by the isomorphism $p\mathbb{Z} \rightarrow q\mathbb{Z}$ given by $n \rightarrow (q/p)n$. Suppose $p < q$. For any n , we have

$$a^{-n}b^{p^n}a^n = a^{-(n-1)}(a^{-1}b^pa)^{p^{n-1}}a^{n-1} = a^{-(n-1)}b^{p^{n-1}q}a^{n-1} = \left(a^{-(n-1)}b^{p^{n-1}}a^{n-1}\right)^q.$$

It follows that $a^{-n}b^{p^n}a^n = b^{q^n}$ for all n . Thus over some generating set of $B(p, q)$, we have $|b^{q^n}| \leq 2n + p^n = o(q^n)$ as $n \rightarrow \infty$, since $p < q$. The subgroup $\langle b \rangle \subset B(p, q)$ is therefore not quasiconvex. Furthermore, no hyperbolic group can contain a subgroup isomorphic to $B(p, q)$ for $p < q$. Clearly $B(p, q) = B(q, p)$, and $B(p, p)$ contains a subgroup $\langle a, b^p \rangle$ isomorphic to \mathbb{Z}^2 . By Proposition 2.9, hyperbolic groups also do not contain a subgroup isomorphic to $B(p, q)$ for any $p, q > 0$.

Quasiconvexity is also useful in constructing hyperbolic groups. Let X and X' be groups with (not necessarily finite) presentations $X = \langle S \mid R \rangle$ and $X' = \langle S' \mid R' \rangle$. For any subgroups $A \subset X$ and $A' \subset X'$ and isomorphism $\alpha : X \rightarrow X'$, define the *amalgamated free product* $X *_\alpha X'$ by the presentation $\langle S, S' \mid R, R', a = \alpha(a) \text{ for all } a \in A \rangle$. When there is a canonical choice of α (for example, X and X' are subgroups of another group Y , $A = A'$, and α is the map resulting from

the inclusion of X and X' in Y), the product $X *_\alpha X'$ is often denoted by $X *_A X'$ or $X *_A' X'$. For $X = X'$, define the *HNN-extension* $X *_\alpha = \langle S, t \mid R, t^{-1}st = \alpha(s) \text{ for all } s \in S \rangle$. As the name implies, this group is an extension of X ; that is, the canonical map $X \rightarrow X *_\alpha$ is injective. Hence the isomorphism $\alpha : A \rightarrow A'$ in X extends to an inner automorphism of $X *_\alpha$. Both of these constructions are independent of the particular choices of presentations for X and X' [10, IV.2].

Arbitrary amalgamated free products and HNN-extensions of hyperbolic groups are not necessarily hyperbolic; for example, the Baumslag-Solitar group $B(p, q)$ is an HNN-extension of the elementary hyperbolic group \mathbb{Z} . With a few additional assumptions, however, these groups remain hyperbolic in the quasiconvex case. Following [9], call a subgroup $H \subset G$ *conjugate separated* if $H \cap H^g$ is finite for all $g \notin H$, where $H^g = g^{-1}Hg$.

Proposition 3.9 ([9]). *Let G be a hyperbolic group, and let $\alpha : H \rightarrow K$ be an isomorphism between two subgroups $H, K \subset G$. Suppose that $H \cap K^g$ is finite for all $g \in G$ and that H is conjugate separated in G . If H and K are quasiconvex in G , then the HNN-extension $G *_\alpha$ is a hyperbolic group.*

Proposition 3.10 ([9]). *Let G and G' be hyperbolic groups, and let $\alpha : H \rightarrow H'$ be an isomorphism between subgroups $H \subset G$ and $H' \subset G'$. Suppose H is conjugate separated in G . If H and H' are quasiconvex in G and G' , respectively, then the amalgamated free product $G *_\alpha G'$ is a hyperbolic group.*

As a final example of quasiconvexity, consider the group $G = \pi_1 M$ in Thurston's construction (2.4). The cyclic subgroup $H = \langle t \rangle \subset G$ is quasiconvex by Lemma 3.8, but there is also a direct proof in this case. Take the generating set $\Sigma = \{x_i^{\pm 1}, y_i^{\pm 1}, t^{\pm 1}\}$ for G . Let f denote the quotient map $G \rightarrow \mathbb{Z}$ in (2.3). Then for any word $w = z_1 t_1^{n_1} z_2 t_2^{n_2} \cdots z_k t_k^{n_k}$ representing t^n , we have

$$|w| \geq |n_1| + \cdots + |n_k| \geq |n_1 + \cdots + n_k| = |f(w)| = |n|.$$

Thus $|t^n| = n$ for all n . The group $H = \langle t \rangle$ is hence actually convex in this case, not just quasiconvex: every geodesic with endpoints in H lies entirely within H . More generally, suppose that $G = N \rtimes_\alpha H$ for some $N \triangleleft G$ and $\alpha : H \rightarrow \text{Aut}(N)$. Assume H is finitely generated. Choose a generating set Σ' for H , and extend it to a generating set Σ of G . Let $i : H \rightarrow G$ denote the inclusion map, and let $\pi : G \rightarrow G/N = H$ denote the quotient map. Then $(\pi i)h = h$ for all $h \in H$. Thus $|h|_{\Sigma'} = |\pi i h|_{\Sigma'} \leq C |ih|_{\Sigma}$, where $C = \max\{|\pi \sigma|_{\Sigma'} : \sigma \in \Sigma\}$. By the definition of Σ , we have $|h|_{\Sigma'} \geq |ih|_{\Sigma}$. Thus the inclusion $i : H \rightarrow G$ is a C -quasi-isometric embedding. By Lemma 3.2, H is quasiconvex in G .

Even in the hyperbolic case, arbitrary subgroups are not necessarily quasiconvex. Consider instead the subgroup $H = \pi_1 N$ of $G = \pi_1 M$ for N, M as in Example (ix) of Section 2. Then

$G = H \rtimes \mathbb{Z}$, where a generator of \mathbb{Z} acts on H by φ_* for some pseudo-Anosov map $\varphi : N \rightarrow N$. Although both H and G are hyperbolic, H is not a quasiconvex subgroup of G [8, IV.49].

Chapter 4

Finite Automata

By Lemma 3.8, geodesic segments in the Cayley graph of a hyperbolic group remain close for all time. As a result, these groups are particularly amenable to analysis by combinatorial group theory. In this section, we recall some of the basic definitions and lemmas concerning finite automata, which are used both in deriving an important result of Cannon in this section and in proving the main result of this paper. Most of the definitions and results below are taken from [5], which also covers the material outlined here in more detail.

Let Σ be a finite set. Define Σ^* to be the set of sequences (x_1, \dots, x_n) with all $x_i \in \Sigma$; we include the sequence $()$ of length 0 in Σ^* , called the *null string* ϵ . A *language* over Σ is a subset of Σ^* . For any $w = (w_1, \dots, w_n)$ and $w' = (w'_1, \dots, w'_m)$ in Σ^* , set $ww' = (w_1, \dots, w_n, w'_1, \dots, w'_m)$. In particular, $\epsilon w = w\epsilon = w$ for all words w .

Definition. Let L, L' be languages over a fixed alphabet Σ .

- (i) The *concatenation* LL' is the language consisting of words ww' with $w \in L$ and $w' \in L'$.
- (ii) The *union* $L \cup L'$ is the union of the sets $L, L' \subset \Sigma^*$.
- (iii) The *star closure* L^* consists of ϵ and all words $w_1 w_2 \cdots w_n$ for $w_i \in L$.

The language L is *regular* if it can be formed from the languages $\{\sigma\}$ for $\sigma \in \Sigma$ and $\{\epsilon\}$ by the operations (i), (ii), (iii) above.

Definition. A *deterministic finite automaton* M is a quintuple (S, Σ, μ, Y, s_0) satisfying the following properties:

- (i) The *state set* S is a finite set.
- (ii) The *alphabet* Σ is a finite set.
- (iii) The *transition function* μ is a map $S \times \Sigma \rightarrow S$.
- (iv) The *set of accept states* Y is a subset of S .

(v) The *start state* s_0 is an element of S .

Write sx for $\mu(s, x)$. Extend μ to Σ^* inductively by setting $s(x_1 \cdots x_n) = (s(x_1 \cdots x_{n-1}))x_n$. For any automaton $M = (S, \Sigma, \mu, Y, s_0)$, define the *language recognized by M* to be the subset $L(M) \subset \Sigma^*$ of all words w with $s_0w \in Y$. We can represent a finite automaton (S, Σ, μ, Y, s_0) visually as shown in the graph in Figure 4.1. Each circle in the graph corresponds to a state in S . For each state $s \in S$ and each $x \in \Sigma$, we draw an arrow from s to the state $\mu(s, x)$ and label it with x . The accept states are drawn with a double border. The start state is the target of the bold arrow. In this representation, the state corresponding to $s(w_1 \cdots w_n)$ is the endpoint of the path from s labelled w_1, \dots, w_n . For simplicity, we often omit arrows that can never be used by words accepted by the automaton. In other words, the automaton fails to recognize a word if no appropriate arrow to follow exists at any point.

Conceptually, a deterministic finite automaton represents a machine that inputs a word in Σ^* and outputs either yes (corresponding to an accept state) or no (corresponding to a reject state; that is, a state not in Y), depending on whether the input has some desired property. The automaton in Figure 4.1, for example, accepts exactly those inputs that end in either 01 or 10. The set of all words $w \in \Sigma^*$ with $s_0w \in Y$ is the *language $L(M)$ accepted by M* .

It is often advantageous to relax the definition above slightly as follows.

Definition. A *non-deterministic finite automaton M* is a quintuple (S, Σ, μ, Y, s_0) satisfying the following properties:

- (i) The *state set* S is a finite set.
- (ii) The *alphabet* Σ is a finite set.
- (iii') The *transition function* μ is a map $S \times (\Sigma \cup \{\epsilon\}) \rightarrow 2^S$.
- (iv) The *set of accept states* Y is a subset of S .
- (v') The *set of start state* S_0 is a subset of S .

As in the deterministic case, extend the map μ to Σ^* inductively by setting

$$s(x_1 \cdots x_n) = (s(x_1 \cdots x_{n-1}))x_n = \bigcup_{s' \in s(x_1 \cdots x_{n-1})} s'x_n,$$

where $sx = \mu(s, x)$ for $s \in S$ and $x \in \Sigma$. Figure 4.2 shows an example of a non-deterministic finite automaton, using the same conventions as in the deterministic case. This automaton recognizes the same language as the deterministic one in Figure 4.1. The states corresponding to $s(w_1, \dots, w_n)$ in this representation are all possible endpoints of paths labelled w_1, \dots, w_n . For a non-deterministic finite automaton $M = (S, \Sigma, \mu, Y, s_0)$, we also define the language $L(M)$ recognized by M to be the

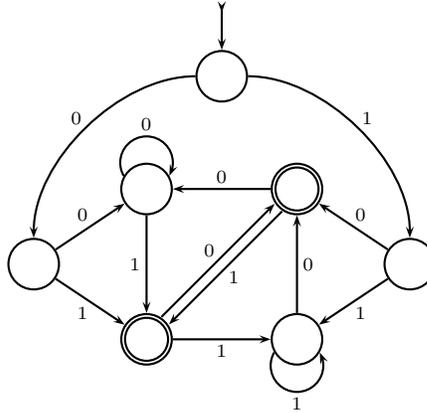


Figure 4.1: A deterministic finite automaton.

set of words $w \in \Sigma^*$ with $Y \cap s_0 w$ non-empty for some $s_0 \in S_0$. We can relax the definition of a finite automaton even further, as follows.

Definition. A *generalized finite automaton* is a quintuple (S, Σ, μ, Y, s_0) satisfying the following properties:

- (i) The *state set* S is a finite set.
- (ii) The *alphabet* Σ is a finite set.
- (iii'') The *transition function* μ is a map $S \times \Sigma^* \rightarrow 2^S$ such that $M(s, s') = \{w \in \Sigma^* : s' \in \mu(s, w)\}$ is a regular language over Σ for each fixed $s, s' \in S$.
- (iv) The *set of accept states* Y is a subset of S .
- (v'') The *set of start states* S_0 is a subset of S .

We can represent a generalized finite automaton by a directed graph in the same manner as for deterministic and non-deterministic automata; the only difference is that edges are labelled with regular languages instead of elements of $\Sigma \cup \{\epsilon\}$. An example of such a graph is shown in Figure 4.3; it recognizes the same language as the automata in Figures 4.1 and 4.2. Define $sw \in 2^S$ for all words $w \in \Sigma^*$ and $s \in S$ as in the non-deterministic case. Under this representation, sw corresponds to all possible states that can be taken by following paths labelled with regular languages containing u_1, \dots, u_n for any words u_i with $w = u_1 \cdots u_n$. Define the language $L(M)$ recognized by M to be the set of words $w \in \Sigma^*$ with $Y \cap s_0 w$ non-empty for some $s_0 \in S_0$.

Theorem 4.1 (Kleene, Rabin, and Scott, [5, 1.2.7]). *Let L be a language over an alphabet Σ . The following conditions on L are equivalent:*

- (i) L is regular.

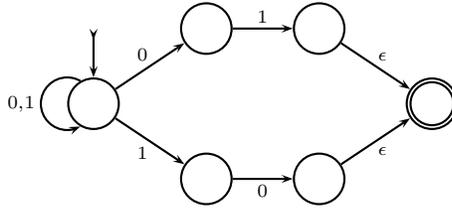


Figure 4.2: A non-deterministic finite automaton.

- (ii) *There exists a deterministic finite automaton M recognizing L .*
- (iii) *There exists a non-deterministic finite automaton M recognizing L .*
- (iv) *There exists a generalized finite automaton M recognizing L .*

Not every language is regular. For example, suppose the language $L = \{x_1^n x_2^n : n \geq 0\}$ over $\Sigma = \{x_1, x_2\}$ is recognized by a deterministic finite automaton M . For each $n \neq m$, the word $x_1^n x_2^n$ is an accept state and $x_1^m x_2^n$ a reject state. It follows that $s_0 x_1^m$ and $s_0 x_1^n$ are distinct states of M for distinct n and m , where s_0 denotes the start state of M . But M has only finitely many states, so L cannot be regular.

Definition. Let $\Sigma_1, \dots, \Sigma_n$ be alphabets. Choose a *padding character* $\$ \notin \bigcup \Sigma_i$. Let $w_i \in \Sigma_i^*$ for $i = 1, \dots, n$, and set $N = \max |w_i|$. Define the *padded extension* $(w_1, \dots, w_n)^\$$ to be the pair $(\tilde{w}_1, \dots, \tilde{w}_n)$, where \tilde{w}_i is formed by appending $N - |w_i|$ copies of $\$$ to w_i . A *padded string* over $\Sigma = \Sigma_1 \times \dots \times \Sigma_n$ is a string of the form $(w_1, \dots, w_n)^\$$ for some $w_i \in \Sigma_i$, and a *language* over Σ is a set of padded strings over Σ .

For example, over the alphabet $\{0, 1\} \times \{0, 1\}$, we have $(10, 1000)^\$ = (10\$, 1000)$. To simplify notation, we write the latter string as simply $(10, 1000)$. For any languages L over Σ and L' over Σ' , define $L \times L'$ to be the language $\{(u, v)^\$: u \in L, v \in L'\}$ over $\Sigma \times \Sigma'$. All of the operations defined above extend to languages over product alphabets in the obvious way; for example, the union of two languages L, L' over such an alphabet is $\{w : w^\$ \in L \cup L'\}^\$$. The important lemma below guarantees that many such operations preserve regularity.

Lemma 4.2. *Let L, L' be regular languages over an arbitrary alphabet Σ , and let L'' be a regular alphabet over Σ^2 . Then the following languages are also regular:*

- (i) L^* .
- (ii) LL' .
- (iii) $L \cup L'$

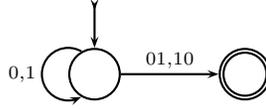


Figure 4.3: A generalized finite automaton.

- (iv) $\neg L = \{u \in \Sigma^* : u \notin L\}$
- (v) $L \cap L'$
- (vi) $\{u \in \Sigma^* \mid \exists v \in L : uv \in L'\}$
- (vii) $\bar{L} = \{u \in \Sigma^* \mid \exists v \in \Sigma^* : uv \in L'\}$.
- (viii) $L \times L'$.
- (ix) $\{x \in L \mid \exists x' \in \Sigma^* : (x, x') \in L''\}$.
- (x) $\{x \in L \mid \exists x' \in L' : (x, x') \in L''\}$.
- (xi) $\{x \in L \mid \forall x' \in \Sigma^* : (x, x') \in L''\}$.
- (xii) $\{x \in L \mid \forall x' \in L' : (x, x') \in L''\}$.

Proof ([5, 1.2–1.4]). Choose finite automata

$$M = (S, \Sigma, \mu, Y, s_0) \quad M' = (S', \Sigma, \mu', Y', s'_0) \quad M'' = (S'', \Sigma^2, \mu'', Y'', s''_0)$$

recognizing L , L' , and L'' , respectively. We construct finite automata recognizing each of the twelve languages in the lemma below. Although the regularity of the first three languages follows immediately from the definition of a regular language, showing that the set of languages recognized by finite automata is closed under star closure, concatenation, and union is an important step in the proof of Theorem 4.1.

- (i) The required automaton recognizing L^* is formed by adding an ϵ -move from each accept state of M to the initial state s_0 .
- (ii) Construct an automaton N by attaching an ϵ -move from each state of M to the start state s'_0 of M' . More explicitly, set $N = (M \amalg M', \Sigma, \nu', Y', s_0)$, where

$$\nu(s, x) = \begin{cases} \mu(s, x) \cup \{s'_0\} & \text{if } s \in S; \\ \mu'(s, x) & \text{if } s \in S'. \end{cases}$$

It is clear that $L(N) = L(M)L(M')$.

(iii) Construct a non-deterministic automaton N by adding a state t_0 to $M \cup M'$ and attaching arrows labelled with ϵ from t_0 to s_0 and to s'_0 . Designate t_0 as the initial state of N . More precisely, set $N = (S \amalg S' \amalg \{t_0\}, \Sigma, \nu, Y \cup Y', t_0)$, where

$$\nu(s, x) = \begin{cases} \mu(s, x) & \text{if } x \in S; \\ \mu'(s, x) & \text{if } x \in S'; \\ \{s_0, s'_0\} & \text{if } x = t_0. \end{cases}$$

Then $tw = s_0w \cup s'_0w$ for all $w \in \Sigma^*$, so $L(N) = L(M) \cup L(M')$.

(iv) The automaton $(S, \Sigma, \mu, S - Y, s_0)$ recognizes exactly those words w with $s_0w \notin Y$; that is, the language $\neg L$.

(v) The regularity of $L \cap L' = \neg(\neg L \cup \neg L')$ follows immediately from parts (iii) and (iv).

(vi) Let Z denote the set of states $s \in S'$ such that $sw \in Y$ for some $w \in L$. Then $N = (M, \Sigma, \mu, Z, s_0)$ is the required automaton.

(vii) Take $L = \Sigma^*$ in part (vi).

(viii) Set $N = (S \times S', \Sigma^2, \nu, Y \times Y', (s_0, s'_0))$, where

$$\nu((s, s'), (x, x')) = \begin{cases} (\mu(s, x), \mu(s', x')) & \text{if } x, x' \neq \$; \\ (s, \mu(s', x')) & \text{if } x = \$; \\ (\mu(s, x), s') & \text{if } x' = \$. \end{cases}$$

Then $(s, s')(w, w') = (sw, s'w)$, so N recognizes the language $L \times L'$.

(ix) Let N denote the automaton formed by replacing every arrow label (x, x') in M'' with x , producing a non-deterministic finite automaton over Σ . Explicitly, $N = (S'', \Sigma, Y'', \nu, s''_0)$, where

$$\nu(s, x) = \bigcup_{x' \in \Sigma} \nu(s, (x, x')).$$

Clearly N recognizes the given language.

(x) The language $\Sigma^* \times L'$ is regular by part (viii), so

$$\{x \in L \mid \exists x' \in L' : (x, x') \in L''\} = \{x \in L \mid \exists x' \in \Sigma^* : (x, x') \in L''\} \cap (\Sigma^* \times L')$$

is regular by part (ix).

(xi) The regularity of

$$\{x \in L \mid \forall x' \in \Sigma^* : (x, x') \in L''\} = \neg \{x \in L \mid \exists x' \in \Sigma^* : (x, x') \in \neg L''\}$$

follows from parts (iv) and (ix).

(xii) The language

$$\{x \in L \mid \forall x' \in L' : (x, x') \in L''\} = \{x \in L \mid \forall x' \in \Sigma^* : (x, x') \in L''\} \cap (\Sigma^* \times L')$$

is regular by the same argument as in part (x). \square

Note that for any group X with (finite) generating set Σ , the evaluation map $w \rightarrow \bar{w}$ from $F(\Sigma)$ to G extends to a map $\Sigma^* \rightarrow G$ by considering each word in Σ as an element of $F(\Sigma)$. This identification is not unique because of inverses, but the resulting map is still well-defined. We also denote the evaluation map on Σ^* by either $w \rightarrow \pi(w)$ or $w \rightarrow \bar{w}$.

The connection between hyperbolic groups and finite automata is the property of *automation*. In an automatic group X , there exists a finite automaton determining whether two words w, w' in a certain set of representatives of elements of X satisfy $\pi(w') = \pi(w)g_0$ for any fixed $g_0 \in G$. In particular, the word problem for G is solvable. That is, there exists an algorithm determining whether a given word w (restricted to a certain subset of Σ^*) represents the identity in X . As a result, many problems concerning hyperbolic groups are amenable to the techniques of geometric group theory.

Definition. Let X be an arbitrary finitely generated group with generating set Σ . Call X *automatic* if there exist finite automata W over Σ and M_x over Σ^2 for each $x \in \Sigma \cup \{\epsilon\}$ that satisfy the following two properties:

- (i) For each $g \in X$, there exists some $w \in L(W)$ with $\bar{w} = g$.
- (ii) The language recognized by M_x is $L(M_x) = \{(u, v) \in L(W) \times L(W) : \bar{u}x = \bar{v}\}$.

The automaton W is called a *word acceptor*, and M_x is called a *multiplier automaton* for $x \neq \epsilon$ and an *equality recognizer* for $x = \epsilon$. The condition of having an automatic structure is called *automation*.

Lemma 4.3 ([5, 2.4.1]). *Automation is independent of the choice of generating set Σ .*

To prove that hyperbolic groups are automatic, we first introduce Cannon's idea of *cone types* [3].

Definition. For any $x, y \in G$, write $x \leq y$ if there exists a geodesic segment $[1y]$ passing through x . Define an equivalence relation on G by setting $x \sim x'$ if the inequality $x \leq y$ holds exactly when $x' \leq (x'x^{-1})y$. The quotient $\mathcal{C}(G) = G/\sim$ is the *set of cone types* of G , and the image of $x \in G$ in $\mathcal{C}(G)$ is the *cone type* of x , denoted by $C(x)$.

To simplify notation, write $u \leq v$ for words $u, v \in \Sigma^*$ if $\bar{u} \leq \bar{v}$. For any $C \in \mathcal{C}(G)$, write $C \leq x$ if $g \leq gx$ for some (and hence every) $g \in G$ with $C(G) = C$. For any fixed $x \in G$, the relation $g \leq gx$ clearly depends only on the cone type of g . Thus write $C(g)x$ for $C(gx)$ if $C(g) \leq x$.

Theorem 4.4 (Cannon). *For any hyperbolic group G , the set of cone types $\mathcal{C}(G)$ is finite.*

Proof. See [3] for a geometric proof or [5, 3.2] for a more combinatorial one. \square

Definition. Define the *language of geodesics* $\Lambda(G)$ to be the set of all words $w \in \Sigma^*$ with $|w| = |\bar{w}|_\Sigma$. (Note that Λ also depends on the choice of Σ .)

Thus elements of $\Lambda(G)$ correspond to geodesics in Γ from 1 to any other point. These geodesics will form the language $L(W)$ in the automatic structure for G . It is not true in general, even for hyperbolic groups, that any language can be chosen for $L(W)$. There exists no finite automaton recognizing the language $\{x_1^n x_2^n : n \geq 0\}$, for example, so \mathbb{Z} has no automatic structure with $L(W) = 0 \cup \{-1, 1\}^*$.

The automation of any hyperbolic group G will follow from the three corollaries of Theorem 4.4 below.

Corollary 4.5. *The language $\Lambda(G)$ over Σ is regular.*

Proof ([5, 3.2]). Define an automaton M as follows. For the set of states S of M , take $\mathcal{C}(G)$ with one additional state r . The transition function of M is given by

$$\mu(C, x) = \begin{cases} Cx & \text{if } C \leq x; \\ r & \text{otherwise} \end{cases}$$

on states $C \in \mathcal{C}(G)$, and $\mu(r, x) = r$ for all x . The automaton $M = (S, \Sigma, \mu, \mathcal{C}(G), C(1))$ then accepts exactly those words $w \in \Sigma^*$ with $1 \leq \pi(w_1) \leq \pi(w_1 w_2) \leq \dots \leq \pi(w)$. These words are precisely the elements of $\Lambda(G)$. \square

Corollary 4.6. *For any fixed $g \in G$, the language $L = \{w \in \Lambda(G) : w \leq wg\}$ is regular.*

Proof. Choose a word $u \in \Lambda(G)$ with $\bar{u} = g$. For any $w \in \Lambda(G)$, we have $w \in L$ iff $wu \in \Lambda(G)$. Hence

$$L = \Lambda(G) \cap \{w \in \Sigma^* : wu \in \Lambda(G)\}.$$

The language $\Lambda(G)$ is regular by Corollary 4.5. By parts (v) and (vi) (with $L' = \{w\}$) of Lemma 4.2, L is therefore also regular. \square

Corollary 4.7. *For any fixed cone type $C \in \mathcal{C}(G)$, the language $L = \{w \in \Lambda(G) : C(\bar{w}) = C\}$ is regular over Σ .*

Proof. Since $\mathcal{C}(G)$ is finite, there exist s_1, \dots, s_n and s'_1, \dots, s'_m in Σ^* such that C is the only cone type for which the inequality $C \leq s$ holds for $s = s_1, \dots, s_n$ but fails for $s = s'_1, \dots, s'_m$. Thus

$$L = \bigcap_{i=1}^n \{w \in \Lambda(G) : w \leq s_i\} \cap \bigcap_{i=1}^m \neg \{w \in \Lambda(G) : w \leq s'_i\}.$$

By Lemma 4.2 and Corollary 4.6, L is therefore regular. \square

Proposition 4.8. *Every δ -hyperbolic group G is automatic.*

Proof ([5, 2.3.4]). By Corollary (4.5), $\Lambda(G)$ is regular. Choose a deterministic finite automaton M representing $\Lambda(G)$. Let S denote the set of states of M , let $Y \subset S$ denote the subset of accept states, and let s_0 denote the initial state. For any $x \in \Sigma^*$, define a deterministic finite automaton M_x as follows. For the set of states of M_x , add a special state r to the product $S \times S \times X$, where $X = \overline{N_{4\delta+4|x|}(1)}$. Define Σ^2 to be the alphabet of M_x . The transition function μ of M_x is given by $\mu(r, (y_1, y_2)) = r$ for any y_1, y_2 and

$$\mu((s_1, s_2, g), (y_1, y_2)) = \begin{cases} (s_1 y_1, s_2 y_2, y_1^{-1} g y_2) & \text{if } y_1^{-1} g y_2 \in X; \\ r & \text{otherwise} \end{cases}$$

for any other state. The accept states of M_x are all states of the form (s_1, s_2, x) with $s_1, s_2 \in Y$, and the initial state of M_x is $(s_0, s_0, 1)$. For any (padded) $(y, y') = (y_1 \cdots y_n, y'_1 \cdots y'_m)$ in $\Lambda(G) \times \Lambda(G)$, we have $(s_0, s_0)(y, y') = r$ if the geodesic chains $\pi(y_1 \cdots y_t)$ and $\pi(y'_1 \cdots y'_m)$ remain a distance at most $4\delta + 4|x|$ apart for all time t . Otherwise, the state $(s_0, s_0)(y, y')$ is $(s_0 y, s'_0 y', y^{-1} y')$. Hence the language

$$L'_x = (\Lambda(G) \times \Lambda(G)) \cap L(M_x)$$

consists of all pairs (y, y') of geodesics in $\Lambda(G)$ such that $y^{-1} y' = x$ and y, y' stay a distance of at most $4\delta + 4|x|$ apart for all time. By Lemma 2.3, any two geodesics y, y' with $y^{-1} y' = x$ stay a distance at most 4δ apart until time

$$y \cdot y' = \frac{1}{2} (|y| + |y'| - d(y, y')) = \frac{1}{2} (|y| + |yx| - |x|) \geq |y| - |x|.$$

It follows that y, y' stay at most $4\delta + 4|x|$ apart for all time. Hence

$$L'_x = \{(y, y') \in \Lambda(G) \times \Lambda(G) : \pi(y') = \pi(yx)\}.$$

Thus G is automatic with word acceptor $\Lambda(G)$, equality recognizer M_ϵ , and multiplier automata M_x for each $x \in \Sigma$. \square

The converse of Proposition 4.8 is false. For example, any finitely generated abelian group is automatic [5, 4.2.4]. The group \mathbb{Z}^n , however, does not satisfy Proposition 2.9 for $n > 1$, and is therefore not hyperbolic.

The automation of G also has the following more geometric consequence.

Proposition 4.9. *If G is infinite, then there exists a constant $C > 0$ such that for every $g \in G$, some geodesic ray $r \in \Omega_0(G)$ passes through $\overline{N_C(g)}$.*

Proof. We first claim that $\Omega_0(G)$ is non-empty. By Corollary 4.5, there exists a deterministic finite automaton M with $L(M) = \Lambda(G)$. Suppose that no word $w \in \Lambda(G)$ revisits the same state twice. Then for each word $w = w_1 \cdots w_{|w|} \in \Lambda(G)$, the states $s_0(w_1 \cdots w_k)$ are distinct for $k = 0, \dots, |w|$, where s_0 is the initial state of M . It follows that $|w|$ is bounded by the number of states C of M . Since the evaluation map $\Lambda(G) \rightarrow G$ is surjective, G must be finite, contradicting the hypothesis of the lemma. Thus there exists a geodesic word $w \in \Lambda(G)$ such that $s_0(w_1 \cdots w_n) = s_0(w_1 \cdots w_m)$ for some $n < m$. By definition, the prefix-closure

$$\overline{\Lambda(G)} = \{x \in \Sigma^* \mid \exists y \in \Sigma^* : xy \in \Lambda(G)\}$$

is $\Lambda(G)$ itself. Hence the state $s = s_0(w_1 \cdots w_n)$ is an accept state. Hence

$$s_0 w_1 \cdots w_n (w_{n+1} \cdots w_m)^k = s$$

for all k . The word $w_1 \cdots w_n (w_{n+1} \cdots w_m)^k$ thus lies in $\Lambda(G)$ for all k , and so gives the geodesic ray required by the claim.

Choose a geodesic $w \in \Lambda(G)$ with $\overline{w} = g$, and write $w = w_1 \cdots w_n$ with each $w_i \in \Sigma$. If the states $s_0(w_1 \cdots w_k)$ are all distinct for $k = 0, \dots, n$, then $k \leq C$. Hence $|g| \leq C$. In this case, every geodesic in $\Omega_0(G)$ intersects $\overline{N_C(g)}$. Otherwise, there exists some $n < m$ such that

$$s_0(w_1 \cdots w_n) = s_0(w_1 \cdots w_m) \tag{4.1}$$

As in the proof of the preceding claim, the words $w_1 \cdots w_n (w_{n+1} \cdots w_m)^k$ for $k \geq 0$ define a geodesic ray $r \in \Omega_0(G)$. Assuming without loss of generality that n is the smallest index satisfying (4.1), $|n| < C$. Thus $d(g, r) \leq n < C$, as required. \square

Chapter 5

Nets in Groups

Let M be a metric space (not necessarily hyperbolic). For any $C > 0$, a subspace $M' \subset M$ is a C -net if $M \subset \overline{N_C(M')}$. Call M' a net in M if M' is a C -net for some C ; or, equivalently, if $d(p, M')$ is bounded for all $p \in M$. Similarly, for a finitely generated group X , a subgroup $X' \subset X$ is a net if it is a net in the metric space $C(X, \Sigma)$ for some finite generating set Σ . Note that we do not require that X be hyperbolic or that X' be finitely generated. Since the Cayley graph $C(X, \Sigma)$ is independent of Σ up to quasi-isometry, the condition of being a net in X is independent of the particular choice of generating set Σ for X . Consider the problem of finding pairs (X, X') with $X' \subset X$ that satisfy the following property:

There exists a section $s : X/X' \rightarrow X$ such that $s(X/X')$ is a net in G . (*)

In (*), X/X' denotes the space of right cosets of X' in X ; we do not require $X' \triangleleft X$. In particular, the desired map $s : X/X' \rightarrow X$ is only a map of sets, not a group homomorphism or a continuous map. The goal of this paper is to prove that the pair (G, H) satisfies (*) if G is hyperbolic and $H \subset G$ is a quasiconvex subgroup of infinite index. In order to motivate this result, we provide a few examples and counterexamples of pairs satisfying (*) in this section.

Lemma 5.1. *Let $1 \rightarrow N \rightarrow E \xrightarrow{\pi} Q \rightarrow 1$ be an exact sequence of groups. Suppose E is finitely generated. Then the pair (E, N) satisfies (*) iff N is finite.*

Proof. It is clear that (E, N) satisfies (*) if N is finite. Thus assume N is infinite. Fix generating sets Σ and $\Sigma' = \pi(\Sigma)$ for E and Q , respectively. Suppose instead that there exists a section $s : Q \rightarrow E$ of π such that $s(Q)$ is a C -net (with respect to Σ) in $C(E, \Sigma)$ for some constant C . Clearly $|\pi g|_{\Sigma'} \leq |g|_{\Sigma}$ for all g . Thus for any $g, g' \in E$, we have

$$d(g, g') = |g^{-1}g'|_{\Sigma} \geq |(\pi g)^{-1}(\pi g')|_{\Sigma'} = d(\pi g, \pi g'). \quad (5.1)$$

Fix $q_0 \in Q$. Since N is infinite, there exists some $g \in E$ with $g \in s(q_0)N$ and $d(g, s(q_0)) > C$ for

all $q \in Q$ with $d(q, q_0) \leq C$. By (5.1), $d(g, s(q)) > C$ for all $q \in Q$. The lemma follows from this contradiction. \square

Thus pairs (E, N) with $N \triangleleft E$ cannot satisfy $(*)$ because the preimages $\pi^{-1}(q)$ remain uniformly separated: $d(\pi^{-1}q, \pi^{-1}q') \geq d(q, q')$ for all $q, q' \in Q$. Hence to construct a section $s : X/X' \rightarrow X$ of the required type, we need to consider subgroups X' for which the cosets gX' exhibit more complicated behavior. It was proved in Lemma 5.1 that for any exact sequence $1 \rightarrow N \rightarrow E \rightarrow Q \rightarrow 1$, the pair (E, N) does not satisfy $(*)$. If this sequence splits, then we can embed Q in E and consider the pair (E, Q) ; that is, we consider (E, Q) for E a semidirect product $N \rtimes Q$. In order to analyze this problem, it is useful to consider the projection $f : E \rightarrow s(Q)$ instead of s itself. Note that f moves points a bounded distance iff $s(Q)$ is a net in E . The following lemma makes this observation more precise.

Lemma 5.2. *Let $X = N \rtimes X'$ for some action of X' on N , and let $\pi : X \rightarrow X/N = X'$ denote the quotient map. Suppose X is finitely generated. Let d and d' denote the metrics on $C(X, \Sigma)$ and $C(X', \Sigma')$, respectively, for some fixed generating sets Σ of X and $\Sigma' = \pi(\Sigma)$ of X' . Then (X, X') satisfies $(*)$ iff there exists a function $\varphi : X \rightarrow N$ and a constant C that satisfy the following two properties:*

- (i) *If $\varphi(\alpha) = \varphi(\alpha')$ for $\alpha \neq \alpha'$ in X , then $d'(\pi\alpha, \pi\alpha') \leq C$.*
- (ii) *The distance $d(\alpha, (\varphi\alpha, \pi\alpha)) \leq C$ for all $\alpha \in X$.*

Proof. Suppose first that (X, X') satisfies $(*)$. Then there exists a section $s : N \rightarrow X$ such that $S = s(N)$ is a $\frac{1}{2}C$ -net in the Cayley graph $C(G)$ for some C . For each $\alpha \in X$, choose some $f(\alpha) \in S$ with $d(\alpha, f(\alpha)) \leq C/2$. Write $f(\alpha) = (\varphi(\alpha), \psi(\alpha))$ for some functions $\varphi : X \rightarrow X'$ and $\psi : X \rightarrow N$. Since s is a section, we have $\psi(\alpha) = \psi(\alpha')$ for some $\alpha, \alpha' \in X$ iff $\varphi(\alpha) = \varphi(\alpha')$. But

$$d'(\psi\alpha, \pi\alpha) = d(\pi f\alpha, \pi\alpha) \leq d(f\alpha, \alpha) \leq C/2,$$

so φ satisfies property (i).

Since $\Sigma' = \pi(\Sigma)$, we have $|\pi\alpha|_{\Sigma'} \leq |\alpha|_{\Sigma}$ for all Σ . Assume without loss of generality that $(1, \sigma') \in \Sigma$ for all $\sigma' \in \Sigma'$. Then $|(1, x)|_{\Sigma} \leq |x|_{\Sigma'}$ for all $x \in X$. It follows that $|(1, x)|_{\Sigma} = |x|_{\Sigma'}$ for all $x \in X$. We therefore have

$$\begin{aligned} d((n, x), (n, x')) &= |(x^{-1}.n^{-1}, x^{-1})(n, x')|_{\Sigma} \\ &= |(1, x^{-1}x')|_{\Sigma} \\ &= |x^{-1}x'|_{\Sigma'} \\ &= d'(x, x') \end{aligned} \tag{5.2}$$

for all $n \in N$ and $x \in X'$. Hence for all $\alpha \in X$,

$$\begin{aligned}
d(\alpha, (\varphi\alpha, \pi\alpha)) &\leq d(\alpha, (\varphi\alpha, \psi\alpha)) + d_{\Sigma}((\varphi\alpha, \psi\alpha), (\varphi\alpha, \pi\alpha)) \\
&= d'(\psi\alpha, \pi\alpha) + C/2 \\
&= d'(\pi f\alpha, \pi\alpha) + C/2 \\
&\leq d(f\alpha, \alpha) + C/2 \\
&\leq C.
\end{aligned}$$

The conclusion of the lemma therefore holds.

Conversely, suppose that such a function φ exists. Consider the set $S = \{(\varphi(n, x), x) : (n, x) \in X\}$. By property (ii), S is a C -net in $\Gamma = C(G, \Sigma)$. For any points $(n, x), (n, x') \in X$ for fixed $n \in N$, we have $d'(x, x') \leq C$ by (i). Hence $d((n, x), (n, x')) \leq C$ by (5.2). Let S' be a subset of S such that $S' \cap (n \times X)$ contains at most one point for each $n \in N$. Then S' is a $2C$ -net in Γ . Choose an arbitrary (set-theoretic) section $s : N \rightarrow X$ such that $S' \cap (n \times X) = \{s(n)\}$ whenever this intersection is nonempty. Then $s(N) \supset S'$ is also a $2C$ -net in Γ , so (X, X') satisfies (*). \square

Using Lemma 5.2, we now describe a method for using sections to satisfying (*) to construct such sections over larger groups.

Lemma 5.3. *Let Q be a finitely generated group, and let Q act on groups N and N' . Suppose $E = N \rtimes Q$ and $E' = N' \rtimes Q$ are finitely generated. Set $X = (N \times N') \rtimes Q$, where Q acts on $N \times N'$ via the diagonal map. If (E, Q) and (E', Q) satisfy (*), then (X, Q) also satisfies it.*

Proof. Choose finite generating sets Σ and Σ' for E and E' , respectively. Assume without loss of generality that Σ contains $(1, q)$ for any $(n, q) \in \Sigma$, and similarly for Σ' . Let $(n \times n', q) \in X$, and let $(n, q) = (n_1, q_1) \cdots (n_r, q_r)$ and $(n', q) = (n'_1, q'_1) \cdots (n'_s, q'_s)$ with all $(n_i, q_i) \in E$ and $(n'_i, q'_i) \in E'$. Then

$$(n \times n', q) = (1 \times n_1, q_1) \cdots (n_r, q_r)(1, q_r^{-1}) \cdots (1, q_1^{-1})(n'_1 \times 1, q'_1) \cdots (n'_s \times 1, q'_s).$$

Hence X is generated by

$$S = \{(n \times 1, q) : (n, q) \in \Sigma\} \cup \{(1 \times n, q) : (n, q) \in \Sigma'\};$$

furthermore, we have

$$|(n \times n', q)|_{\Sigma \times \Sigma'} \leq 2(n, q)_{\Sigma} + |(n', q)|_{\Sigma'}.$$

Equivalently,

$$|\alpha|_{\Sigma \times \Sigma'} \leq 2|\rho\alpha|_{\Sigma} + |\rho'\alpha|_{\Sigma'} \quad (5.3)$$

for all $\alpha \in X$, where ρ and ρ' denote the quotient maps $\rho : X \rightarrow X/N' = E$ and $\rho' : X \rightarrow X/N = E'$.

Suppose (E, Q) and (E', Q) satisfy $(*)$. Then there exist functions $\varphi : E \rightarrow N$ and $\varphi' : E' \rightarrow N'$ satisfying properties (i) and (ii) in Lemma 5.2 for some constant C . Consider the function $\psi : X \rightarrow N \times N'$ defined by $\psi = \varphi\rho \times \varphi'\rho'$. To simplify notation, denote the three quotient maps $E \rightarrow E/N = Q$, $E' \rightarrow E'/N' = Q$, and $X \rightarrow X/(N \times N') = Q$ by π . If $\psi(\alpha) = \psi(\beta)$, then property (i) forces

$$d_{\pi(S)}(\pi\alpha, \pi\beta) \leq d_{\pi(\Sigma)}(\pi\rho\alpha, \pi\rho\beta) \leq C. \quad (5.4)$$

By property (ii), any $\alpha \in X$ satisfies

$$\begin{aligned} d_S(\alpha, (\psi\alpha, \pi\alpha)) &\leq 2d_{\Sigma}(\rho\alpha, (\rho\psi\alpha, \pi\alpha)) + d_{\Sigma'}(\rho'\alpha, (\rho'\psi\alpha, \pi\alpha)) \\ &\leq 2d_{\Sigma}(\rho\alpha, (\varphi\rho\alpha, \pi\rho\alpha)) + d_{\Sigma'}(\rho'\alpha, (\varphi\rho'\alpha, \pi\rho'\alpha)) \\ &\leq 3C \end{aligned} \quad (5.5)$$

by (5.3). Combining (5.4) and (5.5) shows that (X, Q) satisfies $(*)$ by Lemma 5.2. \square

Chapter 6

Nets in Hyperbolic Groups

In this section, we prove that for any hyperbolic group G and any quasiconvex subgroup $H \subset G$ with $[G : H]$ infinite, the pair (G, H) satisfies $(*)$; that is, there exists a section $s : G/H \rightarrow G$ of the quotient map (of sets) $G \rightarrow G/H$ such that $s(G/H)$ forms a net in the Cayley graph of G . Fix a δ -hyperbolic group G , a generating set Σ of G , and a K -quasiconvex subgroup H . Let $\Gamma = C(G, \Sigma)$, and let $d(x, y) = |x^{-1}y|_\Sigma$ denote the metric on Γ . Set $\Lambda = \Lambda(G)$ and $\mathcal{C} = \mathcal{C}(G)$. To simplify notation, abbreviate $|\cdot|_\Sigma$ as $|\cdot|$.

For each $g \in G$, set $\sigma(g) = \min\{|gh| : h \in H\}$ and $S = \{g \in G : |g| = \sigma(g)\}$. We form the desired section simply by choosing one point in each coset $S \cap gH$. The diameter of the set $S \cap gH$ is uniformly bounded for all $g \in G$, so it suffices to prove that S itself is a net. The crucial step is doing so is showing that the set of geodesics to points in S is a regular language. The prefix closure \bar{S} of S is then a finite distance in Γ away from S itself. Thus if \bar{S} is a net in Γ , then so is S . As a subset of Γ , the set \bar{S} consists of all points that lie on geodesic rays from 1 that intersect S . We use the quasiconvexity of H to prove that any point in Γ is a bounded distance from such a geodesic if $[G : H] = \infty$, completing the proof. The first step in this argument is the following lemma, which provides a convenient bound or estimate for the distance between points on the same coset gH .

Lemma 6.1. *Let $H \subset G$ be a quasiconvex subgroup, and fix some $g \in G$. For all $x_1, x_2 \in gH$, we have $d(x_1, x_2) \leq C_1 + |x_1| + |x_2| - 2\sigma(g)$, where $C_1 > 0$ is a constant depending only on G and H .*

Proof. Choose geodesics $[1x_1]$, $[1x_2]$, and $[x_1x_2]$. Set $f(p) = d(p, [1x_1]) - d(p, [1x_2])$ for all $p \in [x_1x_2]$. For any two adjacent vertices p and p' , we have $|f(p) - f(p')| \leq 2$. Since $f(x_1) \leq 0$ and $f(x_2) \geq 0$, it follows that there exists some $p_0 \in [x_1x_2]$ with $|d(p_0, [1x_1]) - d(p_0, [1x_2])| \leq 2$. By the δ -hyperbolicity of Γ , we have $\min\{d(p, [1x_1]), d(p, [1x_2])\} \leq \delta$ for each $p \in [x_1x_2]$. Thus $d(p_0, [1x_1]), d(p_0, [1x_2]) \leq \delta + 2$; choose $x'_1 \in [1x_1]$ and $x'_2 \in [1x_2]$ realizing these inequalities. The quasiconvexity of H implies that there exists a point $p'_0 \in gH$ with $d(p_0, p'_0) < K$. See Figure 6.1 for an illustration of this

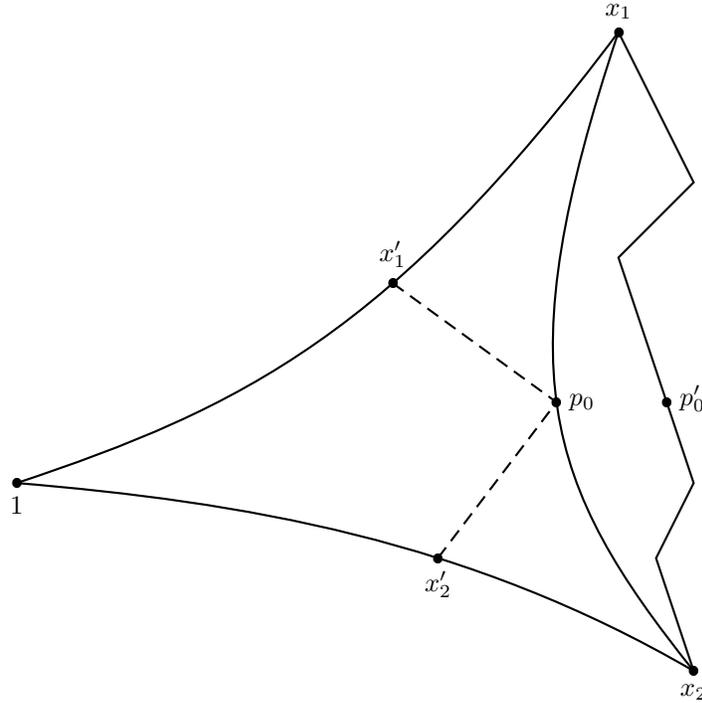


Figure 6.1: The construction in Lemma 6.1. The dotted lines each have length at most $\delta + 2$, and the quasigeodesic containing p'_0 stays in the K -neighborhood of the geodesic containing p_0 .

construction. We have

$$\sigma(g) \leq |p'_0| \leq |x'_i| + d(x'_i, p_0) + d(p_0, p'_0) \leq |x'_i| + \delta + K + 2 = |x_i| - d(x_i, x'_i) + \delta + K + 2.$$

Thus $d(x_i, x'_i) \leq |x_i| - \sigma(g) + \delta + K + 2$. We therefore have

$$\begin{aligned} d(x_1, x_2) &\leq d(x_1, x'_1) + d(x'_1, p_0) + d(p_0, x'_2) + d(x'_2, x_2) \\ &\leq 2\delta + 4 + d(x_1, x'_1) + d(x_2, x'_2) \\ &\leq 4\delta + 2K + 8 + |x_1| + |x_2| - 2\sigma(g), \end{aligned}$$

as required. □

For $x_1, x_2 \in S$, Lemma 6.1 forces $d(x_1, x_2) \leq C_1$. The intersections $S \cap gH$ for $g \in G$ thus have bounded diameter. In defining the section s by choosing one point in $S \cap gH$ for each coset gH , the particular choice of points is therefore irrelevant in the large-scale geometry of Γ . In particular, the condition that s be a net is independent of this choice.

Set $\langle g, g' \rangle = 2(g^{-1}.g') = |g| + |g'| - |gg'|$ for $g, g' \in G$. Following the conventions in previous

sections, we also write $\langle x, y \rangle$ for $\langle \bar{x}, \bar{y} \rangle$ with $x, y \in \Sigma^*$ to simplify notation. The set S consists precisely of those points $g \in G$ with $|g| \leq |gh|$ for all h . Thus

$$S = \{g \in G : \langle g, gh \rangle \leq |h| \text{ for all } h \in H\}.$$

In order to prove that S is regular, we first prove that the condition $\langle x, y \rangle = n$ defines a regular language $L_n \subset \Lambda^2$ for each n . The idea of the proof is to split each y with $(x, y) \in L_n$ into subwords $w, c, w' \in \Sigma^*$ such that $\langle x, y \rangle = \langle x, wcv' \rangle = \langle xy, c \rangle + \langle xwc, w' \rangle$ with $\langle x, wcv' \rangle, \langle xy, c \rangle < n$. By showing that the set of such words (y, w, c, w') is regular, we therefore conclude that L_n is regular by induction on n .

Lemma 6.2. *For any fixed $n \geq 0$, the language $L_n = \{(x, y) \in \Lambda^2 : \langle x, y \rangle = n\}$ over Σ^2 is regular.*

Proof. By Corollary 4.5, there exists a deterministic finite automaton M over Σ with $L(M) = \Lambda$. Let M' denote the obtained by replacing each arrow labelled $c \in \Sigma$ by arrows labelled (c, ϵ) and $(\$, c)$ (with the same source and target). The language recognized by M' consists of all $x, y \in \Sigma^*$ such that $xy \in \Lambda$. Thus $L_0 = L(M') \cap \Lambda^2$, and so is regular. Fix some positive integer n , and suppose that L_m is regular for all $m < n$. For any $x \in \Lambda$ and $c \in \Sigma$,

$$\langle x, c \rangle = |x| + 1 - |xc| = \begin{cases} 0 & \text{if } x \leq xc; \\ 2 & \text{if } x \geq xc; \\ 1 & \text{otherwise.} \end{cases}$$

Set $P_i(c) = \{x \in \Lambda \mid \langle x, c \rangle = i\}$ for $i = 0, 1, 2$ and $c \in \Sigma$. Let \mathcal{C}_c denote the set of cones C such that $x \leq xc$ if $C(x) = C$. Then

$$\begin{aligned} P_0(c) &= \{x \in \Lambda \mid C(x) \in \mathcal{C}_c\} \\ P_2(c) &= \{x \in \Lambda \mid \exists y \in \Lambda : \bar{xc} = \bar{y}, C(y) \in \mathcal{C}_{c^{-1}}\} \\ P_1(c) &= \Lambda \cap \neg(L_0 \cup L_2) \end{aligned}$$

Since \mathcal{C}_c is finite, all three languages P_i are regular by Lemma 4.2.

Let $R \subset \Lambda^4$ consist of all quadruples (x, y, z, w) with $y = y_1 \cdots y_n$, $z = xy_1 \cdots y_i$, and $w = y_{i+1} \cdots y_n$, where i is the largest index such that $xy_1 \cdots y_i \in \Lambda$. We claim that the language R is regular. Construct a finite automaton M over Σ^4 as follows. Let M have one state s_C for each cone type $C \in \mathcal{C}$, and add two states s_r and s_a . Designate s_a as the only accept state. For each state s_C , add arrows labelled $(u, \epsilon, u, \epsilon)$ and $(\$, u, u, \epsilon)$ from s_C to the state s_{Cu} for each $u \in \Sigma$ with $x \leq xu$ for $C(x) = C$. For any other $u \in \Sigma$, attach an arrow from s_C to s_r labelled $(\$, u, \epsilon, u)$. Add an arrow from s_r to itself labelled $(\$, u, \$, u)$ for each $u \in \Sigma$. Attach an arrow labelled $(\$, \$, \$, \$)$ from every

state in M to s_a . It is clear that for any geodesics $x, y, z, w \in \Lambda$, the resulting automaton accepts the quadruple (x, y, z, w) iff it lies in R . Thus $R = L(M') \cap \Lambda^4$. The claim follows.

Let R_c denote the set of quadruples $(x, y, z, w) \in R$ with $w_1 = c$, and let R'_c denote the set of sextuples $(x, y, z, w, z', w') \in \Lambda^6$ such that $(x, y, z, w) \in R_c$, $\overline{z'} = \overline{zc}$, and $w' = w_2 \cdots w_{|w|}$. Since R is regular, it is clear that R_c and R'_c are also regular for each $c \in \Sigma$. For any $x, y, z \in G$, we have

$$\begin{aligned} \langle x, y \rangle + \langle xy, z \rangle - \langle y, z \rangle &= (|x| + |y| - |xy|) + (|xy| + |z| - |xyz|) - (|y| + |z| - |yz|) \\ &= |x| + |yz| - |xyz| \\ &= \langle x, yz \rangle. \end{aligned} \tag{6.1}$$

Fix $x, y \in \Lambda$, and write $y = y_1 \cdots y_n$ with each $y_i \in \Sigma$. Let $k \geq 0$ denote the largest index such that $xy_1 \cdots y_k \in \Lambda$. By (6.1),

$$\begin{aligned} \langle x, y \rangle &= \langle xy_1 \cdots y_{k+1}, y_{k+2} \cdots y_n \rangle + \sum_{i=0}^k \langle xy_1 \cdots y_i, y_{i+1} \rangle \\ &= \langle xy_1 \cdots y_k, y_{k+1} \rangle + \langle xy_1 \cdots y_{k+1}, y_{k+2} \cdots y_n \rangle. \end{aligned}$$

Since $xy_1 \cdots y_{k+1} \notin \Lambda$, we have $\langle xy_1 \cdots y_k, y_{k+1} \rangle > 0$. Thus $\langle x, y \rangle = n$ iff $\langle xy_1 \cdots y_k, y_{k+1} \rangle = i$ and $\langle xy_1 \cdots y_{k+1}, y_{k+2} \cdots y_n \rangle = n - i$ for $i = 1$ or 2 . Hence L_n satisfies

$$\begin{aligned} L_n &= \bigcup_{\substack{c \in \Sigma \\ i=1,2}} \{(x, y) \in \Lambda^2 \mid \exists (z, w, z', w') \in \Lambda^4 : (x, y, z, w, z', w') \in R'_c, \langle z, c \rangle = i, \langle z', w' \rangle = n - i\} \\ &= \bigcup_{\substack{c \in \Sigma \\ i=1,2}} \{(x, y) \in \Lambda^2 \mid \exists (z, w, z', w') \in \Lambda^4 : (x, y, z, w, z', w') \in R'_c, z \in P_i(c), (z', w') \in L_{n-i}\}, \end{aligned}$$

and so is regular. The lemma follows by induction on n . \square

Lemma 6.3. *The language $L = \{x \in \Lambda : \overline{x} \in S\}$ is regular.*

Proof. Assume without loss of generality that the generating set Σ of G contains a generating set Σ' for the hyperbolic group H . By Corollary 4.5, the language $\Lambda' \subset \Sigma'^* \subset \Sigma^*$ of geodesics in $C(H, \Sigma')$ is regular. For any $g \in S$ and $h \in H$, we have

$$\langle g, gh \rangle = |g| + d(g, gh) - |gh| \leq |g| + (C_1 + |gh| - |g|) - |gh| = C_1$$

for some constant C_1 by Lemma 6.1. Hence

$$S = \{x \in \Lambda \mid \forall y \in \Lambda' : \langle x, y \rangle \leq |y|\} = \{x \in \Lambda \mid \forall y \in \Lambda' : \langle x, y \rangle \leq \min(|y|, C_1)\}.$$

Thus

$$S = \{x \in \Lambda \mid \forall y \in \Lambda' : \langle x, y \rangle \leq C_1\} \cap \bigcup_{r=0}^{C_1} \left(\bigcap_{y \in \Lambda' \cap \overline{N_r(1)}} \{x \in \Lambda : \langle x, y \rangle \leq r\} \right). \quad (6.2)$$

By Lemma 6.2, the language

$$\{(x, y) \in \Lambda \times \Lambda' : \langle x, y \rangle \leq n\} = (\Lambda \times \Lambda') \cap \bigcup_{i=0}^n \{(x, y) \in \Lambda^2 : \langle x, y \rangle = i\}$$

is regular for all n . Thus S is regular by (6.2) and Lemma 4.2. \square

We now prove the main theorem.

Theorem 6.4. *Let G be a hyperbolic group, and let $H \subset G$ be a quasiconvex subgroup. If $[G : H] = \infty$, then there exists a (set-theoretic) section $s : G/H \rightarrow H$ of the quotient map $G \rightarrow G/H$ such that $s(G/H)$ is a net in G .*

Proof. We first claim that there exists a constant C_2 such that for any $g \in G$, there exists a point $g' \in G$ and a geodesic ray r through g' such that $d(g, g') < C_2$ and $d(r(t), H)$ is unbounded as $t \rightarrow \infty$. By Proposition 4.9, there exists a point g' with $d(g, g') < C$ for some constant C (independent of g) and a geodesic ray $r \in \Omega_0(G)$ through g' . Suppose $r \subset \overline{N_l(H)}$ for some $l > 0$. Set $x = r(|g| + l)$, and choose some $h \in H$ such that $d(x, h) \leq l$. Then

$$x.h = \frac{1}{2} (|x| + |h| - d(x, h)) \geq |x| - d(x, h) \geq |g|.$$

Since H is K -quasiconvex, there exists a geodesic $[1h]$ in G lying in $\overline{N_K(H)}$. By Lemma 2.3, any two geodesics $[1g'], [1h]$ stay a distance at most 4δ apart until time $g'.h$. Hence

$$d(g, H) \leq d(g, [1h]) + K \leq d(g, g') + d(g', [1h]) + K \leq K + C + 4\delta. \quad (6.3)$$

Set $C' = K + C + 4\delta$. If $d(g, H) > C'$, contradicting (6.3), then the distance $d(r(t), H)$ must be unbounded; the claim then holds with $C_2 = C$. Suppose instead that $d(g, H) \leq C'$. Choose some $p \in \overline{N_{C'}(1)}$ such that $g \in Hp$. Since $[G : H]$ is infinite, there exists some $t \in G$ with $d(t, H) > C'$. Fix some such t minimizing $|t|$. The point $gp^{-1}t$ then satisfies $d(gp^{-1}t, H) \geq d(Ht, H) \geq d(t, H) > C'$ and $d(gp^{-1}t, g) \leq C_2$, where

$$C_2 = C' + \min \left\{ |t'| : t' \in G, t' \notin \overline{N_{C'}(H)} \right\}.$$

The claim therefore holds for all $g \in G$.

Fix $g \in G$. By the claim above, there exist $g', x \in G$ such that $d(g, g') < C_2$, $g' \leq x$, and $d(x, H) \geq |g| + \frac{1}{2}C_1$. Choose a point $x' \in S \cap xH$. The Gromov product $x.x'$ satisfies

$$x.x' = \frac{1}{2}(|x| + |x'| - d(x, x')) \geq |x'| - \frac{1}{2}C_1 \geq |g|$$

by Lemma 6.1. Any two geodesics $[1x]$ and $[1x']$ remain a distance no greater than 4δ apart until time $x.x'$, so $d(g, [1x']) \leq d(g, g') + d(g', [1x']) < C_2 + 4\delta$. Set $L = \{x \in \Lambda : \bar{x} \in S\}$. Then any $g \in G$ satisfies $d(g, \pi(\bar{L})) < C_2 + 4\delta$, where

$$\bar{L} = \{x \in \Sigma^* \mid \exists y \in \Sigma^* : xy \in L\}$$

denotes the prefix closure of L . By Lemmas 4.2 and 6.3, \bar{L} is regular. Let M be a deterministic finite automaton with $L(M) = L$. Since M has only finitely many states, any word in \bar{L} is within a bounded distance of a word in L ; explicitly, any $x \in \bar{L}$ satisfies $d(x, L) < C_3$, where C_3 is the number of states of M . Hence

$$d(g, L) \leq d(g, g') + d(g', L) < C_2 + C_3 + 4\delta$$

By Lemma 6.1, each coset g_0H contains at most $\#\overline{N_{C_1}(1)}$ elements of S . Choosing exactly one point in each intersection $S \cap g_0H$ produces a section $s : G/H \rightarrow G$ such that $d(p, s(G/H)) < \#\overline{N_{C_1}(1)} + C_2 + C_3 + 4\delta$ for all vertices $p \in \Gamma$. \square

Since the image $S = s(G/H)$ of the section $s : G/H \rightarrow G$ in Theorem 6.4 is a net, it is also a hyperbolic metric space. The left action of G on the right coset space G/H induces an action on S , given by $g.s(g') = s(gg')$. By considering the corresponding homeomorphisms of the boundary induced by this action, we prove two results about the intersection of conjugate subgroups of H below. Set $H^g = g^{-1}Hg$ for any $g \in G$. This conjugate depends only on the image of g in the left coset space $H \backslash G$. As such, we write $H^\gamma = \gamma^{-1}H\gamma = H^g$ for a coset $\gamma = Hg \in H \backslash G$. In [6], it is proved that any quasiconvex subgroup H of a hyperbolic group G has finite width; that is, $H^\gamma \cap H$ is finite for all but finitely many cosets $\gamma \in G/H$. Using a completely different method, the section $s : G/H \rightarrow G$ of Theorem 6.4, we prove a weaker version of this result. Specifically, we show in Proposition 6.7 below that quasiconvex subgroups of infinite index in G contain no infinite groups normal in G . We require the following elementary lemma:

Lemma 6.5. *If G is a finite extension of an infinite cyclic group, then any infinite cyclic subgroup $H \subset G$ has finite index.*

Proof. Choose an exact sequence $1 \rightarrow N \rightarrow G \xrightarrow{\pi} Q \rightarrow 1$ with N cyclic and Q finite. Then Q contains $\pi(H) = H/(H \cap N)$, so $[H : H \cap N]$ is finite. In particular, $H \cap N$ is non-trivial. Thus

$[N : H \cap N]$ is finite. The index $[G : H] \leq [G : H \cap N] = [G : N][N : H \cap N] = \#Q[N : H \cap N]$ is therefore also finite. \square

We also need the following result, which is interesting independently of its use in proving Proposition 6.7.

Proposition 6.6. *For any $g \in G$, let L_g denote the isometry $L_g(g') = gg'$ on the vertices of Γ . Extend L_g to a graph automorphism of Γ . Suppose G is not elementary. Then the homomorphism $g \rightarrow (L_g)_\infty$ has finite kernel.*

Proof. Let K denote the group of $g \in G$ with $(L_g)_\infty = \text{id}$, and fix $g \in K$. We first claim that there exists some constant $N = N(g)$ such that $[x^N, g] = 1$ for all $x \in G$. By Proposition 2.11, the supremum N_0 of the order of all torsion elements of G is finite. The claim therefore holds immediately for all torsion $x \in G$ with $N = N_0!$. Thus let $x \in G$ be an arbitrary element of infinite order. By Proposition 4.9, there exists a constant C , independent of g and x , such that $d(r, x) \leq C$ for some geodesic r . Since $(L_g)_\infty$ acts trivially on $(r_t) \in \partial G$, the distance $d(gr_t, r_t)$ is bounded. Both $[1g]$ and gr are geodesics, so the union $[1g] \cup gr$ is a $(1, 2|g|)$ -quasi-geodesic. Thus there exists some geodesic ray r' such that $d(gr_t, r'_t) \leq C'$ for some constant $C' = C'(g)$ by Lemma 2.6. The distance $d(r_t, r'_t) \leq d(r_t, gr_t) + d(gr_t, r'_t)$ is then bounded. But

$$r_t \cdot r'_t = \frac{1}{2} (|r_t| + |r'_t| - d(r_t, r'_t)) = t - d(r_t, r'_t),$$

so $d(r_t, r'_t) \leq 4\delta$ for all time t by Lemma 2.3. Choosing some t with $d(r_t, x) \leq C$, we therefore have

$$\begin{aligned} d(gx, x) &\leq d(gx, gr_t) + d(gr_t, r_t) + d(r_t, x) \\ &= 2d(x, r_t) + d(gr_t, r_t) \\ &\leq 2d(x, r_t) + d(gr_t, r'_t) + d(r_t, r'_t) \\ &\leq 2C + C' + 4\delta. \end{aligned}$$

Thus $|x^{-1}gx| \leq 2C + C' + 4\delta$ for all $x \in G$ of infinite order. Fix such an $x \in G$. Then $|x^{-n}gx^n| \leq 2C + C' + 4\delta$ for all $n > 0$, so there exist distinct n, m such that $x^{-n}gx^n = x^{-m}gx^m$ and $0 < n < m < K$, where $K = \#\overline{N_{2C+C'+4\delta}(1)}$. Hence any $x \in G$ of infinite order satisfies $[x^{K!}, g] = 1$. The claim therefore holds for arbitrary $x \in G$ with $N = N_0!K!$.

By Proposition 2.10, there exist $x_1, x_2 \in G$ such that $\langle x_1, x_2 \rangle$ is a free group of rank 2. The commutators $[x_1^N, g]$ and $[x_2^N, g]$ vanish by the claim above. The centralizers $C_G(x_1^N)$ and $C_G(x_2^N)$ are finite extensions of $\langle x_1^N \rangle$ and $\langle x_2^N \rangle$, respectively, by Proposition 2.9 and Lemma 6.5. Suppose $\langle g \rangle$ is infinite. By Lemma 6.5, both $\langle g \rangle$ and $\langle x_i^N \rangle$ have finite index in $C_G(x_i^N)$ for each $i = 1, 2$. It follows that $\langle g \rangle \cap \langle x_1^N \rangle \cap \langle x_2^N \rangle$ has finite index in $\langle g \rangle$. But $\langle x_1^N, x_2^N \rangle \subset \langle x_1, x_2 \rangle$ is free, so $\langle x_1^N \rangle \cap \langle x_2^N \rangle$

is trivial. It follows that g must have finite order.

Thus K consists entirely of torsion. Let $g_1, \dots, g_r \in K$ be distinct representatives of the conjugacy classes of torsion in G that intersect K . By the claim above, there exist constants $N(g_i)$ for each $\sigma \in \Sigma$ and $i = 1, \dots, r$ such that $[x^{N(g_i)}, g_i] = 1$ for all $x \in G$. Set $N = N(g_1) \cdots N(g_r)$, and let $H = \langle x^N : x \in G \rangle$. Then $H \triangleleft G$ and $[H, g_i] = 1$ for all g_i . Write x^y for the conjugate $y^{-1}xy$. Then for all $x \in G$ and $h \in H$,

$$(g_i^x)^h = h^{-1}g_i^x h = \left((h^{x^{-1}})^{-1} g_i (h^{x^{-1}}) \right)^x = g_i^x$$

for each g_i . Thus $[h, g_i^x] = 1$ for all $h \in H$ and $x \in G$. The commutator $[h, k]$ therefore vanishes for all $k \in K$. By Proposition 2.10, G contains an element $x \in G$ of infinite order. Then the centralizer $C_G(x^N)$ is a finite extension of \mathbb{Z} by Proposition 2.9. Hence $\langle x^N \rangle$ has finite index in $C_G(x^N)$ by Lemma 6.5. But $C_G(x^N)$ contains $\langle x^N \rangle \times K$, so

$$\#K = [\langle x^N \rangle \times K : \langle x^N \rangle] \leq [C_G(x^N) : \langle x^N \rangle] < \infty,$$

as required. □

Proposition 6.7. *Let G be a hyperbolic group, and let $H \subset G$ be a quasiconvex subgroup with $[G : H] = \infty$. Let $K \subset H$ with normal closure K^G in G . If K is infinite, then $[K^G : K] = \infty$. In particular, any subgroup of H normal in G is finite.*

Proof. The corollary is trivial if G is finite or a finite extension of \mathbb{Z} , so assume without loss of generality that G is non-elementary. Suppose instead that K^G/K is a finite set of order n . Each element of G acts on the coset space K^G/K by conjugation, giving a homomorphism $\rho : G \rightarrow S_n$. The kernel $G' = \ker \rho$ has finite index in G , so $\partial G' = \partial G$. Replacing G by the quasi-isometric group G' , we can therefore assume that $n = 1$; that is, $K \triangleleft G$.

Let $s : G/H \rightarrow G$ denote the section given by Theorem 6.4. Then G acts on $S = s(G/H) \subset G$ by $g.s(x) = s(gx)$. Since s is an injective function from the right coset space G/H to G , this action is well-defined. The stabilizer of any $s(x) \in S$ is $H^x \supset K$, so K acts trivially on S . Since S is a net in G , the inclusion $i : S \rightarrow G$ is a $(1, \epsilon)$ -quasi-isometry for some ϵ . Hence i induces a bijection (in fact, a homeomorphism) $i_\infty : \partial S \rightarrow \partial G$. Since $s(x)$ minimizes $|\cdot|$ in xH , we have

$$|s(x)| \leq |g^{-1}s(gx)| \leq |g^{-1}| + |s(gx)| = |g| + |s(gx)|.$$

Thus $|s(gx)| \geq |s(x)| - |g|$. By Lemma 6.1,

$$d(g.s(x), gs(x)) = d(s(gx), gs(x)) \leq C_1 + |sg(x)| + |gs(x)| - 2|s(gx)| \leq C_1 + 2|g| \quad (6.4)$$

for all x . Thus the diagram

$$\begin{array}{ccc}
 G & \xrightarrow{g \rightarrow (L'_g)_\infty} & \text{Homeo}(\partial S) \\
 & \searrow_{g \rightarrow (L_g)_\infty} & \uparrow (i_\infty)^* \\
 & & \text{Homeo}(\partial G)
 \end{array} \tag{6.5}$$

commutes, where $L_g(g') = gg'$ and $L'_g s(x) = g.s(x)$. We hence have $(L_g)_\infty = 1$ for all $g \in K$. By Proposition 6.6, K is finite. \square

Bibliography

- [1] K. Brown, *Cohomology of Groups*, Springer-Verlag, New York, 1982.
- [2] D. Burago, Y. Burago, and S. Ivanov, *A Course in Metric Geometry*, American Mathematical Society, Graduate Studies in Mathematics Vol. 33, Providence, Rhode Island, 2000.
- [3] J. W. Cannon, *The Combinatorial Structure of Cocompact Discrete Hyperbolic Groups*, *Geom. Dedicata* **16** (1984), 123-148.
- [4] A. Casson and S. Bleiler, *Automorphisms of Surfaces after Nielsen and Thurston*, Cambridge University Press, Cambridge, 1988.
- [5] D. B. A. Epstein et al., *Word Processing in Groups*, Jones and Bartlett Publishers, Boston, 1992.
- [6] R. Gitik et al., *Widths of Subgroups*, *Trans. Amer. Math. Soc.* **350** (1998), pp. 321–329.
- [7] M. Gromov, *Hyperbolic Groups*, published in “Essays in Group Theory,” ed. S. Gersten, Springer-Verlag, New York, 1987.
- [8] P. de la Harpe, *Topics in Geometric Group Theory*, University of Chicago Press, Chicago, 2000.
- [9] O. Kharlampovich and A. Myasnikov, *Hyperbolic Groups and Free Constructions*, *Trans. Amer. Math. Soc.* **350** (1998), pp. 573–613.
- [10] R. Lyndon and P. Schupp, *Combinatorial Group Theory*, Springer-Verlag, Berlin, 1977.
- [11] C. MacLachlan and A. Reid, *The Arithmetic of Hyperbolic 3-Manifolds*, Springer-Verlag, New York, 2003.
- [12] J. G. Ratcliffe, *Foundations of Hyperbolic Manifolds*, Springer-Verlag, New York, 1994.
- [13] W. Thurston, *Hyperbolic Structures on 3-Manifolds, II: Surface Groups and 3-Manifolds which Fiber over the Circle*, preprint.
- [14] ———, *On the Geometry and Dynamics of Diffeomorphisms of Surfaces*, *Bull. Amer. Math. Soc. (N.S.)* **19** (1988), pp. 417–431.