Relatively Hyperbolic Groups

In memory of Yetta Farb

Benson Farb*
University of Chicago

Revised Version, May 14, 1998

1 Introduction

The theory of word-hyperbolic groups has been a central topic in geometric group theory. General references include [Gr1],[Ca],[GH], [CDP]. In this paper we introduce a theory of relatively hyperbolic groups.

Examples of groups which are not word-hyperbolic (or even automatic), and indeed which lie outside many of the techniques in geometric group theory, are fundamental groups of complex (resp. quaternionic, Cayley) hyperbolic manifolds with cusps. The reason these groups differ so much from the real hyperbolic case, where Epstein ([E et al.]) has proved that such groups are (bi)automatic, is that in the real hyperbolic case the geometry of the group turns out to be *nonpositively curved.* This is due essentially to the fact that the cusp groups are abelian. The fundamental groups of complex (resp. quaternionic, Cayley) hyperbolic manifolds with cusps do *not* exhibit nonpositively curved geometry. Instead they combine a nontrivial mix of both negatively curved and nilpotent geometry.

The techniques introduced in this paper are meant to provide some machinery for dealing with groups which exhibit more than one type of geometric behavior. Our methods pick out and exploit aspects of negative curvature in a group Γ when Γ itself is not a word-hyperbolic group. We emphasize that we do *not* simply consider coset graphs, which is the naive approach (and doesn't work for the motivating examples).

*Supported in part by an NSF postdoctoral fellowship and NSF grant DMS 9704640.

1.1 Definitions

Here is a rough idea of the definitions involved; for details see §3. Although the general theory applies to groups which are hyperbolic relative to a finite set of subgroups, we begin with the case of only one subgroup.

Let H be a finitely-generated subgroup of a finitely-generated group G. For any Cayley graph of Γ of G we form a quotient graph $\widehat{\Gamma}$ of Γ by identifying, for each $g \in G$, all vertices of Γ corresponding to elements lying in the left coset gH.

Definition (Relatively hyperbolic group): The group G is *hyperbolic relative to* H if $\widehat{\Gamma}$ is a negatively curved metric space.

We also isolate a property that can be satisfied by a pair (G, H), where G is hyperbolic relative to H: the Bounded Coset Penetration property (or BCP property for short). The BCP property strengthens the notion of malnormal subgroup, and will allow us to conclude information about G from corresponding information about H (see, e.g., Theorem 3.7 and Theorem 3.8).

1.2 Negatively curved manifolds with cusps

Eberlein, Gromov, and Margulis greatly clarified the structure of the ends of noncompact, complete, finite-volume Riemannian manifolds M^n with (pinched) negative sectional curvatures $-b^2 \leq K(M^n) \leq -a^2 < 0$ ([Eb],[BGS]). Their work shows that $\Gamma = \pi_1(M^n)$ is finitely presented. We take their work as a starting point for our analysis of the geometric and combinatorial structure of the group $\Gamma = \pi_1(M^n)$. Such groups Γ are the motivating examples for the ideas behind relatively hyperbolic groups.

For simplicity we assume that M^n has one cusp (see §5 for the general statement). Associated to the fundamental group Γ of such a manifold M^n is its *cusp subgroup* H, which is the unique (up to conjugacy) maximal parabolic subgroup of Γ. The group H is a finitely-generated nilpotent group.

Theorem 4.11 *Let $\Gamma = \pi_1(M^n)$ be as above, and let H be the cusp subgroup of Γ. Then Γ is hyperbolic relative to H and the pair (Γ, H) has the Bounded Coset Penetration property.*

Applying some of the general properties of relatively hyperbolic pairs (Γ, H) which satisfy the BCP property, we obtain the following two results.

Theorem 4.14 [Fast solution to word problem] *There is a curve-shortening algorithm which solves the word problem for Γ in time $O(n \log n)$.*

Theorem 4.12 [Dehn functions] *Γ satisfies precisely the same isoperimetric inequality as its cusp subgroup H. That is, the Dehn functions for Γ and H are equal.*

Theorem 4.12 was claimed by Gromov in Section 5.6 of [Gr2].

Our approach to understanding Γ involves studying the geometry of "Electric Space" - a space $\widehat{X}$ canonically associated to M^n which models the geometry of the coned-off Cayley graph $\widehat{\Gamma}$ (see §4). "First-order" geometric properties of $\widehat{\Gamma}$ (e.g., what geodesics look like) translate into "second-order" geometric properties of $\widetilde{M}$ (e.g., visual sizes of sets such as horospheres in $\widetilde{M}$).

1.3 Connections with other work

In [Gr1, Gr2] Gromov proposes a different definition of hyperbolicity relative to a subgroup. We did not see how to use that definition, so our approach is different. The two definitions are compared in [Sz]. An additional (and central) idea in our approach is the Bounded Coset Penetration property (see Section 3.3).

We believe that the ideas of relative hyperbolicity and the Bounded Coset Penetration property are applicable to many other examples. Indeed, since [Fa3], relative hyperbolicty and the Bounded Coset Penetration Property have been applied to analyzing certain Coxeter groups ([Kr]), non-cocompact actions on hyperbolic metric spaces and relative hyperbolization of polyhedra ([Sz]), mapping class groups [MM1, MM2], semiconjugacies between Kleinian group actions [Kla], and metabolicity [Ge3].

1.4 Contents and Acknowledgements

Section §3 contains the basic properties and examples of relatively hyperbolic groups. In §4 we provide some applications to the motivating examples of fundamental groups of negatively curved manifolds with cusps. In §5 we discuss briefly the definition of hyperbolicity relative to a finite set of subgroups.

The material in this paper was part of my PhD thesis [Fa3]. I am grateful to my advisor Bill Thurston and to John Stallings. I also thank Noel Brady, Martin Bridson, Mahan Mitra, and the many others who helped me. Finally, I thank the referee for many useful comments and corrections.

2 Background on hyperbolic groups

In this section we briefly review some facts on hyperbolic groups which we will need. For more details, background, and motivation, we refer the reader to [Gr1],[GH], [CDP].

To a group G with finite generating set X, one associates the *Cayley graph* $\Gamma = \Gamma(G, X)$, which is a directed graph whose vertex set consists of elements of G, with a directed edge labelled x going from g to $g \cdot x$ for each $g \in G, x \in X$. We make Γ into a metric space by assigning each edge length 1, and by defining the distance between two points to be the length of the shortest path between them.

Recall that a path-metric space X is *hyperbolic* if there exists some $\delta > 0$ so that the δ-neighborhood of any two sides of a geodesic triangle in X contains the third side. A finitely generated group is called a *(word) hyperbolic group* if its Cayley graph is a hyperbolic metric space. The property of being a hyperbolic path-metric space is invariant under quasi-isometry.

A (K, C)-*quasi-isometry* between metric spaces is a map $f : X \to Y$ such that, for some constants $K, C, C' > 0$:

1. $\frac{1}{K} d_X(x_1, x_2) - C \leq d_Y(f(x_1), f(x_2)) \leq K d_X(x_1, x_2) + C$ for all $x_1, x_2 \in X$.

2. The C'-neighborhood of $f(X)$ is all of Y.

A map satisfying (1) but not necessarily (2) is called a *quasi-isometric embedding* of X into Y. A path $\alpha : [0, p] \longrightarrow X$ is a K-*quasi-geodesic* if it is a K-quasi-isometric embedding of $[0, p]$ into X.

Let Γ be a finitely generated group with generating set X and relator set R. Let $\mathcal{A} = X \sqcup X^{-1}$, and let $\mathcal{A}^*$ denote the free monoid on $\mathcal{A}$. For a word $w \in \mathcal{A}^*$, we denote by $\bar{w}$ the image of w under the natural map $\mathcal{A}^* \longrightarrow \Gamma$. For $g \in \Gamma$, we let $||g|| = d(1, g)$, where $d = d_\Gamma$ denotes the word metric in Γ (with respect to X). A word $w \in \mathcal{A}^*$ is identified with an eventually constant $w : [0, \infty) \longrightarrow \Gamma$.

Let $\ell(w)$ denote the length of w in $\mathcal{A}^*$. Given $w \in \mathcal{A}^*$ with $\overline{w} = 1$, we can write

$$w = \prod_{j=1}^{N} z_j R_j z_j^{-1}, \quad z_j \in F(X), R_j \in R \cup R^{-1} \tag{1}$$

with the equality in $F(X)$, the free group on X. A function f is an *isoperimetric function* for Γ if for every $w \in F(X)$ with $\overline{w} = 1$, the word w can be written as in 1 with $N \leq f(\ell(w))$. The smallest isoperimetric function for a group G is called the *Dehn function* for the group G.

Isoperimetric functions give a measure of the complexity of the word problem for groups. It is shown in [Ge2] that a finitely presented group G has solvable word problem iff there exists a recursive isoperimetric function for G iff the Dehn function for G is recursive.

3 Relatively Hyperbolic Groups

3.1 Definition and Examples

Let G be a finitely generated group, and let H be a finitely generated subgroup of G. We begin with the Cayley graph Γ of G, and we form a new graph $\widehat{\Gamma}$ as follows : for each left coset gH of H in G, add a vertex $v(gH)$ to Γ, and add an edge $e(gh)$ of length $1/2$ from each element gh of gH to the vertex $v(gH)$. We call this new graph the *coned-off Cayley graph* of G with respect to H, and denote it by $\widehat{\Gamma} = \widehat{\Gamma}(H)$. We give this graph the path metric. Note that $\widehat{\Gamma}$ is not a proper metric space (i.e. closed balls are not always compact).

Remarks:

1. In general the graph $\widehat{\Gamma}$ is very different from the quotient graph $H \setminus \Gamma$ obtained by quotienting out by the left action of H on Γ. This is due to the difference between left and right cosets. When $H \lhd G$, the graphs $\widehat{\Gamma}$ and $H \backslash \Gamma$ are quasi-isometric.

2. It is easily seen that $\widehat{\Gamma}$ is quasi-isometric to the graph obtained from Γ by identifying each left coset to a point.

Example: Let $\Gamma = F(a, b)$ be the fundamental group of a punctured torus, and let H be its cusp subgroup, which is the cyclic subgroup $H =$

$< aba^{-1}b^{-1} >$. The coned-off Cayley graph $\widehat{\Gamma}$ of Γ relative to H is quasi-isometric to the 1-skeleton of the Farey tesselation. It is easy to check that $\widehat{\Gamma}$ is a hyperbolic metric space.

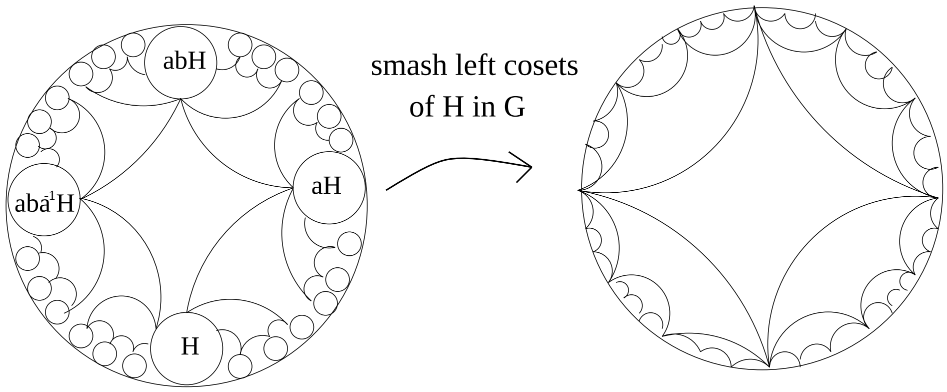

Figure 1. A finite part of the picture of what happens when each coset of $H = \langle aba^{-1}b^{-1} \rangle$ in $G = F(a, b)$ is identified to a point. The result is the quadralateral version of the Farey graph, which is a negatively curved metric space.

Definition (Relatively hyperbolic group): The group G is *hyperbolic relative to* H if the coned-off Cayley graph $\widehat{\Gamma}$ of G with respect to H is a negatively curved metric space.

Note: There is a useful notion of hyperbolicity of a group relative to a *finite set* of subgroups - see §5. This arises in many examples (e.g. manifolds with several cusps, small cancellation theory over free products, etc.).

Examples:

1. Let $\Gamma = \pi_1(M^n)$, where M^n is a complete, finite volume Riemannian manifold with (pinched) negative sectional curvatures and one cusp, and let H be the fundamental group H of the cusp, which is virtually a nilpotent group. Then Γ is hyperbolic relative to H. These examples are discussed at length in Section 4.

2. (Gromov) In [Gr1, Gr2] Gromov proposes the following definition of relative hyperbolicity. Let G be a group acting properly discontinuously by isometries on a δ-hyperbolic metric space Y, such that the quotient X/G is quasi-isometric to $[0,\infty)$. Let H denote the stabilizer subgroup of the endpoint on ∂X of a lift of this ray to X. Then G is said to be hyperbolic relative to H in the Gromov definition. Note that this example includes the previous example.

   A. Szczepanski has shown ([Sz]) that, if G is hyperbolic relative to H in the Gromov defintion, then G is hyperbolic relative to H in our definition, but not conversely.

3. If H is normal in Γ, then Γ is hyperbolic relative to H if and only if Γ/H is a word-hyperbolic group. This follows from the fact that, when $H \lhd \Gamma$, the graph $\widehat{\Gamma}$ is quasi-isometric to the coset graph of Γ/H, which in turn is just the Cayley graph of the word-hyperbolic group Γ/H. Note that the property of being a hyperbolic metric space is a quasi-isometry invariant ([GH]).

4. Gersten has shown ([Ge1]) that a word-hyperbolic group is hyperbolic relative to any quasi-convex subgroup. For example, a surface group is hyperbolic relative to every (finitely generated) subgroup.

5. If H is any finitely generated group, then $N * H$ is hyperbolic relative to H if and only if N is word-hyperbolic (see Proposition 3.4).

6. Let $B_{1,2} =< a, b : aba^{-1} = b^2 >$ be the $(1,2)$-Baumslag-Solitar group, and let $H =< b >$ be the cyclic subgroup generated by b. Then $\widehat{\Gamma}$ is quasi-isometric to the infinite tree of (uniform) valence three, so that G is hyperbolic relative to H.

Example (M. Mitra): Let M^3 be an irreducible, non-Haken 3-manifold finitely covered by a surface bundle over the circle. So there is a non-embedded, immersed surface in M^3 which lifts to a fiber S in a finite surface-bundle cover $p : N^3 \to M^3$, where without loss of generality N^3 is a normal cover (this insures that all lifts of $p(S)$ are embedded). Let $\Gamma_1 = \pi_1(N^3), \Gamma_2 = \pi_1(M^3), H = \pi_1(S)$, so $H \lhd \Gamma_1 \lhd \Gamma_2$ and Γ_1 has finite index in Γ_2. Let $\widehat{\Gamma}_i, i = 1, 2$ be the coned-off Cayley graph of Γ_i with respect to H. Clearly $\widehat{\Gamma}_1$ is quasi-isometric with the real line (cf. Example 3 above). We claim that $\widehat{\Gamma}_2$ is finite.

Consider intersecting copies S and S' of the lifts of the immersed surface to N^3. Denote by t the generator of the circle in the fundamental group $\Gamma_1 \approx \pi_1(S) \times_\phi \mathbf{Z}$ of the bundle. There exists $k > 0$ so that $st^k \in \pi_1(S')$ for some $s \in \pi_1(S)$. It follows that the diameter of $\widehat{\Gamma}_2$ is finite.

If M^3 is an irreducible, non-Haken 3-manifold with surface subgroup $\pi_1(S)$, is $\pi_1(M^3)$ hyperbolic relative to $\pi_1(S)$? The answer to this question is "yes" for hyperbolic 3-manifolds M^3 by work of Thurston-Bonahon. Note that in this case, $\pi_1(M^3)$ is hyperbolic relative to any surface subgroup S, and the usual trichotomy is captured by the boundary $\partial(\widehat{\Gamma})$ of the negatively curved metric space $\widehat{\Gamma}$:

1. S virtually fibers. $\partial(\widehat{\Gamma})$ is empty.
2. S is a true fiber. $\partial(\widehat{\Gamma})$ has 2 points.
3. S is geometrically finite, i.e. quasi-convex. $\partial(\widehat{\Gamma})$ contains infinitely many points; cf. Example 4 on page 7.

Example (Masur-Minsky): Let $\Gamma_{g,r}$ denote the mapping class group of the surface $\Sigma_{g,r}$ of genus g with r punctures. For any simple closed curve C in $\Sigma_{g,r}$, let H_C denote the subgroup of elements of $\Gamma_{g,r}$ fixing C. One of the main results of [MM1] is that $\Gamma_{g,r}$ is hyperbolic relative to the subgroup H. In [MM2] this fact is used to help give a fast solution to the conjugacy problem for $\Gamma_{g,r}$.

3.2 Invariance and Closure Properties

In this section we show that relative hyperbolicity is independent of choices for generating sets for G and H. This will follow from the more general invariance under quasi-homomorphisms quasi-preserving H. We also explore certain closure properties of the class of relatively hyperbolic groups.

The notion of quasi-homomorphism was introduced by Brooks and Gromov ([Gr1]). A map $\phi : G_1 \longrightarrow G_2$ between groups is called a *quasi-homomorphism* if $\phi(1) = 1$, and if there exists a constant $E > 0$ so that $d_{G_2}(\phi(gh), \phi(g)\phi(h)) \leq E$ for all $g, h \in G_1$. We use the notation $hd(A, B)$ to denote the Hausdorff distance between A and B.

Proposition 3.1 *Let $f : G_1 \longrightarrow G_2, f' : G_2 \longrightarrow G_1$ be a pair of quasi-isometries, both of which are quasi-homomorphisms of groups. For $i = 1, 2$,* *let H_i be a finitely generated subgroup of G_i, let $\widehat{\Gamma}_i$ be the coned-off Cayley graph of G_i with respect to H_i, and assume that there exists $D > 0$ so that $hd(f(H_1), H_2) \leq D$. Then $\widehat{\Gamma}_1$ is quasi-isometric to $\widehat{\Gamma}_2$; in particular, $\widehat{\Gamma}_1$ is negatively curved if and only if $\widehat{\Gamma}_2$ is negatively curved.*

Proof: It clearly suffices to show that G_1 is quasi-isometric to G_2 in the $\widehat{\Gamma}_i$ metrics. Let f, f' be K-quasi-isometries. Note that

$$\begin{aligned} hd(f(gH_1), f(g)H_2) &\leq hd(f(gH_1), f(g)f(H_1)) + hd(f(g)f(H_1), f(g)H_2) \\ &\leq E + hd(f(H_1), H_2) \\ &\leq E + D \end{aligned}$$

Let $x, y \in G_1 \subset \widehat{\Gamma}_1$ be given, and let α be a $\widehat{\Gamma}$-geodesic from x to y. There are two types of subsegments of α: those which lie in $G_1 \subset \widehat{\Gamma}_1$, and those consisting of an edge-path of length 1 which begins in at some vertex $p \in gH_1$, goes to the vertex $v(gH_1)$, and ends in some vertex $q \in gH_1$. Segments of the first type are clearly stretched by a factor of at most K. For segments of the second type note that, since $hd(f(gH_1), f(g)H_2) \leq D + E$, there are vertices $p' \in f(g)H_2, q' \in f(g)H_2$ so that $d_{G_2}(p, p'), d_{G_2}(q, q') \leq D + E$. Since p' lies in the same coset of H_2 as q', there is an edge-path of length 1 in $\widehat{\Gamma}_2$ from p' to q'. Hence there is a path of length at most $2(D + E) + 1$ in $\widehat{\Gamma}_2$ from $f(p)$ to $f(q)$.

From these two observations it follows that

$$d_{\widehat{\Gamma}_2}(f(x), f(y)) \leq (K + 2(D + E) + 1) \cdot d_{\widehat{\Gamma}_1}(x, y) \quad \text{for all } x, y \in G_1.$$

Since $d_{G_2}(f' \circ f(x), x)$ is bounded (independently of x), it is easy to show that $hd(f'(gH_2), f'(g)H_1)$ is finite. By the same argument as above,

$$d_{\widehat{\Gamma}_1}(f'(x), f'(y)) \leq (K + 2(D + E) + 1) \cdot d_{\widehat{\Gamma}_2}(x, y) \quad \text{for all } x, y \in G_2$$

hence $\widehat{\Gamma}_1$ is quasi-isometric to $\widehat{\Gamma}_2$. Since both $\widehat{\Gamma}_1$ and $\widehat{\Gamma}_2$ are geodesic metric spaces, one is negatively curved if and only if the other is (see [GH], and note that the proof does not require the spaces in question to be proper). □

Proposition 3.1 has two immediate corollaries.

Corollary 3.2 *The property of a group G being hyperbolic relative to a subgroup H is independent of the choice of (finite) generating sets for both G and H.*

Corollary 3.3 *Suppose that $N < H < G$ are finitely generated groups with $[H : N] < \infty$. Then G is hyperbolic relative to H if and only if G is hyperbolic relative to N.*

Finally, we list two "combination theorems" for relatively hyperbolic groups. We leave their straightforward proofs as an exercise.

Proposition 3.4 *Let H be any finitely generated group. Then $H * N$ is hyperbolic relative to N if and only if H is a word-hyperbolic group.*

Proposition 3.5 *Suppose that H is a subgroup of the groups G_1 and G_2, and that each G_i is hyperbolic relative to H. Then the free product of G_1 and G_2 amalgamated over H is hyperbolic relative to H.*

3.3 Bounded Coset Penetration

For a fixed generating set S of Γ, choose a fixed set of words $y_i \in F(S)$ representing generators of the subgroup H. A path w in Γ gives a path $\widehat{w}$ in $\widehat{\Gamma}$ as follows: search through w, reading from left to right, for the subwords y_i (choose only the left-most word in case of overlaps). For each maximal substring z of y_i's, say z goes from g to $g \cdot \overline{z}$ in Γ, replace the path given by z with one edge from the vertex g to the cone point $v(gH)$, and an edge from $v(gH)$ to the vertex $g \cdot \overline{z}$. Do this for each substring of y_i's. We will call such paths z *coset subwords*, for we think of z as "moving along a coset". Note that this depends on the choice of generators for the subgroup H. We denote this replacement which takes a path in Γ and gives a path in $\widehat{\Gamma}$ by

$$\Gamma \to \widehat{\Gamma}$$

$$w \mapsto \widehat{w}$$

Note that this map is clearly a surjection. If $\widehat{w}$ passes through some cone point $v(gH)$, we say that w (or $\widehat{w}$) *penetrates* the coset gH.

Remark: In the special case when the generating set for Γ contains the generating set for H, the substrings z above are simply the maximal subwords of generators of H.

Definition (Relative (quasi)geodesic): If $\widehat{w}$ is a geodesic in $\widehat{\Gamma}$, we call w a *relative geodesic* in Γ. If $\widehat{w}$ is a P-quasi-geodesic in $\widehat{\Gamma}$, we call w a *relative P-quasi-geodesic* in Γ. A path w in Γ (or $\widehat{w}$ in $\widehat{\Gamma}$) is said to be a path *without* *backtracking* if, for every coset gH which $\widehat{w}$ penetrates, $\widehat{w}$ never returns to gH after leaving gH.

Definition (Bounded coset penetration): The pair (Γ, H) is said to satisfy the *bounded coset penetration property* (or BCP property for short) if, for every $P \geq 1$, there is a constant $c = c(P) > 0$ so that if u and v are relative P-quasigeodesics without backtracking with $d_\Gamma(\overline{u}, \overline{v}) \leq 1$, then the following conditions hold :

1. If u penetrates a coset gH but v does not penetrate gH, then u travels a Γ-distance of at most c in gH.

2. If both u and v penetrate a coset gH, then the vertices in Γ at which u and v first enter gH lie a Γ-distance of at most c from each other; similarly for the vertices $u(t_0)$ and $v(s_0)$ at which u and v last exit gH (i.e. $u(t) \notin gH$ and $v(s) \notin gH$ for all $t > t_0, s > s_0$).

Examples:

1. Suppose Γ is torsion free. Then the Bounded Coset Penetration property implies that the subgroup H is *malnormal*; that is, $gHg^{-1} \cap H = \{1\}$ for all $g \in \Gamma \backslash H$. For otherwise $gHg^{-1} \cap H$ would be infinite for some g, and then it is not difficult to construct a pair of relative geodesics with the same endpoint, one of which moves far in H, while the other does not penetrate H but moves far in gH, violating the BCP property.

2. If $\Gamma = \pi_1(M)$ is the fundamental group of a complete, finite volume, (pinched) negatively curved Riemannian manifold M with a cusp, and if $H < \Gamma$ is the cusp subgroup, then (Γ, H) has the Bounded Coset Penetration property (Theorem 4.11).

3. The free abelian group of rank 2: $\Gamma =< a, b : ab = ba >$ is hyperbolic relative to the subgroup $H =< a >$, but (Γ, H) does not have the Bounded Coset Penetration property: the paths b^3a^n and b^4a^n are relative geodesics (of $\widehat{\Gamma}$-length 4 and 5, respectively) ending a distance 1 apart in Γ, but clearly violate condition (1) of the BCP property when n is large enough.

4. (D. Allcock) As noted in Example 6 on page 7, the group $B_{1,2} = < a, b : aba^{-1} = b^2 >$ is hyperbolic relative to the infinite cyclic subgroup $H = < b >$ generated by b. It is easy to check that $(B_{1,2}, H)$ does not satisfy the Bounded Coset Penetration property. The cyclic subgroup $H = < a >$ is an example of a subgroup which is malnormal in $B_{1,2}$, but the pair $(B_{1,2}, H)$ does not have the BCP property.

3.4 A Curve-Shortening Algorithm

Suppose that Γ is hyperbolic relative to the subgroup H, and that (Γ, H) has the Bounded Coset Penetration property. In this section we describe a curve-shortening algorithm for paths in Γ in the spirit of J. Cannon's "Dehn's algorithm" for hyperbolic groups [Ca].

Recall that a path w in a path metric space X is called a *k-local geodesic* if every subpath of w of length at most k is a geodesic. We need the following fact about k-local geodesics:

Lemma 3.6 ([Ca, Sh]) *Let X be a δ-hyperbolic metric space, and let $k = 4\delta$. If u is a k-local geodesic in X, and if v is a geodesic connecting the endpoints of u, then u lies in a 3δ-neighborhood of v.*

We now describe a curve-shortening process for closed loops in Γ; this will provide a solution to the word problem for Γ. As always, we are assuming that H is a finitely generated subgroup, and that H has solvable word problem.

The pair (Γ, H) provides us with two constants: the geodesic metric space $\widehat{\Gamma}$ is a δ-hyperbolic metric space, and let $c = c(2)$ be the constant given for 2-quasi-geodesics in the definition of the Bounded Coset Penetration property. Let $k = 4\delta$. Now suppose we are given a closed loop $w \in F(X)$, where X is the given generating set for Γ. Let Y be a finite set of words in $F(X)$ representing the given set of generators for the subgroup H. We form the following finite sets:

$$\mathcal{L}_1 = \{z \in F(Y) : \ell_Y(z) \leq c\}$$

$$\mathcal{L}_2 = \{w \in F(X) : \ell_X(w) \leq (8\delta - 1) \cdot c\}$$

We first run through w from left to right and find each coset subword. Recall that coset subwords are chosen to be maximal. We then perform the following procedure:

Coset subword reduction: For each coset subword z, considered as a word in Y^*, check to see if $\overline{z} = \overline{\alpha}$ for some $\alpha \in \mathcal{L}_1$; this can be done since H has solvable word problem. If this is the case, replace the coset subword z of w by the word α. Of course, we are really replacing z by the word α with each letter of α itself being written out as a word in $F(X)$.

Now, either $\widehat{w}$ is or is not a k-local geodesic.

If $\widehat{w}$ is not a k-local geodesic, then there exists a subsegment $\widehat{u}$ of $\widehat{w}$ of length $\ell_{\widehat{\Gamma}}(\widehat{u}) \leq k$ which is not a geodesic, and so that any subsegment of $\widehat{u}$ is a geodesic. Note that this implies that $\widehat{u}$ is a path without backtracking, and also that $\widehat{u}$ is a 2-quasi-geodesic. Let $\widehat{v}$ be a geodesic between the endpoints of $\widehat{u}$ with the property that the length of the shortest path v in Γ projecting to $\widehat{v}$ (under the natural map $\Gamma \to \widehat{\Gamma}$) is minimal.

If $\widehat{u}$ and $\widehat{v}$ penetrate some coset gH, say at $\widehat{u}(t)$ and $\widehat{v}(s)$, then it must be that $s = t$ since $\widehat{v}$ and any subsegment of $\widehat{u}$ are geodesics. But then $\widehat{v}([t+1,\infty])$ would be shorter than $\widehat{u}([t+1,\infty])$, contradicting the fact that every subsegment of $\widehat{u}$ is geodesic. Hence $\widehat{u}$ and $\widehat{v}$ penetrate distinct cosets. By the Bounded Coset Penetration property for 2-quasi-geodesics, both $\widehat{u}$ and $\widehat{v}$ travel a Γ-distance of at most c in any coset. By the coset subword reduction procedure and by the minimality of v, every coset subword of $\widehat{u}$ (and $\widehat{v}$) has Γ-length at most c. Hence

$$\begin{aligned} \ell_\Gamma(vu^{-1}) &\leq c \cdot \ell_{\widehat{\Gamma}}(\widehat{v}\widehat{u}^{-1}) \\ &\leq c \cdot (k + (k-1)) \\ &= c \cdot (8\delta - 1) \end{aligned}$$

so that $vu^{-1} \in \mathcal{L}_2$. So, after performing the coset subword reduction procedure, we need only search through w for more than half of a word in $\mathcal{L}_2$: if we find it, replace it by the other (shorter!) half. This gives a new word w' with $\overline{w'} = \overline{w}$ with $\ell_{\widehat{\Gamma}}(\widehat{w'}) < \ell_{\widehat{\Gamma}}(\widehat{w})$.

There is the following possibility (caused by backtracking): it may happen that a subword z_1uz_2 of w is shortened to z_1z_2, where each z_i is a coset subword (with $\ell_Y(z_i)$ possibly greater than c). If this ever happens, we apply coset subword reduction to z_1z_2.

As an illustration, suppose we begin with $w = z_1uz_2$ where $\overline{w} = 1$, each z_i is very long, and u is a loop in Γ leaving H. Then the first time we perform the coset subword reduction procedure, neither of the z_i will be shortened. Our algorithm eventually shortens u to the empty word, and we are left with a long loop $z_1 z_2$ in H. Applying the coset subword reduction procedure then reduces $z_1 z_2$ to the empty word.

By continuing the above process (re-applying the coset subword reduction procedure when necessary), we can shorten (the path in $\widehat{\Gamma}$ corresponding to) the given word until we are left with a word w with $\widehat{w}$ a k-local geodesic in $\widehat{\Gamma}$. By Lemma 3.6, $\widehat{w}$ stays in a 3δ-neighborhood of any geodesic, which in this case (if $\overline{w} = 1$) can be taken to be the null path beginning and ending at $1 \in \Gamma$. But then $\ell_{\widehat{\Gamma}}(\widehat{w}) \leq 3\delta$, since if $\widehat{w}$ had an initial subpath of length $3\delta + 1$, this subpath would be a geodesic (since $3\delta + 1 \leq 4\delta = k$) and hence would leave the 3δ-neighborhood of the identity. But since $\ell_{\widehat{\Gamma}} \leq 3\delta < k$, $\widehat{w}$ is a geodesic; that is, w is the empty word.

The above gives a simple algorithm in the spirit of Dehn's algorithm to solve the word problem in Γ; namely, given a word w, we perform the coset subword reduction procedure, then continue to shorten w as much as possible. The above discussion shows that we are left with the empty word if and only if $\overline{w} = 1$ in Γ. Doing these procedures in an efficient way, we may deduce the following:

Theorem 3.7 *Suppose Γ is hyperbolic relative to H, the pair (Γ, H) has the Bounded Coset Penetration property, and H has word problem solvable in time $O(f(n))$. Then the curve-shortening algorithm gives an $O(f(n)\log n)$-time solution to the word problem for Γ.*

Proof: By the above discussion, the only thing to check is the time bound. Let $w \in F(X)$ with $\ell_X(w) = n$ be given. Clearly we can identify all coset subwords of w (with respect to a given generating set for H) in time $O(n)$.

Recall that, after we perform the coset subword reduction procedure to the initially input word, we only repeat this procedure when we see that two coset subwords come together. The maximal possible number of times this procedure needs to be repeated occurs for an input word w with $\ell_X(w) = n$ when w has approximately $\log_2(n)$ coset subwords of length $\sim n/\log_2(n)$, each separated by small loops leaving H. After deleting these trivial loops, the coset subword combine into half as many coset subwords, each of twice the length - suppose these had actually been separated by small loops leaving H. Again after these small loops are killed, these coset subwords combine. Repeating this procedure, we see that the number of steps to reduce w to the empty word takes at most (approximately):

$$n + f(n) + 2f(n/2) + 4f(n/4) + \cdots + \log_2(n) f(n/\log_2(n)) \preceq O(f(n)\log n)$$

steps. $\square$

By counting the number of small loops needed to reduce a big loop to the empty word in the curve-shortening algorithm above, one immediately obtains the following:

Theorem 3.8 *Suppose Γ is hyperbolic relative to H and that the pair (Γ, H) has the Bounded Coset Penetration property. Then any isoperimetric function for H is an isoperimetric function for Γ.*

Gromov showed ([Gr1]) that word-hyperbolic groups are characterized by the property that they have a linear isoperimetric function (for details, see [Sh]). It follows from Theorem 3.8 that if, in addition to the hypotheses of the theorem, the subgroup H is a word-hyperbolic group, then Γ is a word-hyperbolic group.

4 Applications to Negatively Curved Manifolds with a Cusp

In this section we apply the theory of relatively hyperbolic groups to analyze fundamental groups $\Gamma = \pi_1(M^n)$ of (noncompact) complete, finite-volume Riemannian manifolds with pinched negative sectional curvatures. The most important examples are when M^n admits a complex (resp. quaternionic, Cayley) hyperbolic metric, in which case the curvatures are pinched between -4 and -1. In these cases Γ is not automatic, or even combable ([Fa2]). For simplicity we assume that M^n has only one cusp, with cusp group H. All of our results hold in the case of several cusps (see §5).

4.1 Negatively Curved Manifolds

Pinched Hadamard Manifolds

Let H be a Hadamard manifold; that is, H is a connected, simply connected, complete Riemannian manifold with nonpositive sectional curvatures. We will be most interested in the case when H is the universal cover

$H = \widetilde{M}$ of a complete, finite-volume negatively curved Riemannian manifold M, and where M (hence H) has pinched negative sectional curvatures: $-b^2 \leq K(M) \leq -a^2 < 0$. When the latter condition is satisfied, H is called a *pinched Hadamard manifold*. See [BGS, Kl] as general references for Hadamard manifolds.

We begin by recalling precisely how orthogonal projection onto a geodesic in a Hadamard manifold changes lengths of paths.

Proposition 4.1 ([Kl], Prop. 3.9.11) *Let H be a Hadamard manifold with $-b^2 \leq K(H) \leq -a^2 \leq 0$. Let $\gamma(t)$ be a geodesic in H, and let $\beta(t), 0 \leq t \leq r$ be any curve from $p = \beta(0)$ to $q = \beta(r)$ with $d_H(p,\gamma) = d_H(q,\gamma) = K$ and $d_H(\beta(t),\gamma) \geq K$ for all $t \in [0,r]$. Let p' (resp. q') denote the foot of the perpindicular from p (resp. q) to γ. Then $\ell(\beta) \geq d_H(p',q')\cosh(aK)$; hence*

$$d_H(p',q') \leq \ell(\beta) \cdot \cosh^{-1}(aK) \leq \ell(\beta) \cdot e^{-aK}.$$

In order to analyze negatively curved manifolds with cusps, it is first necessary to understand the geometry of horospheres in pinched Hadamard manifolds. So suppose $x \in H$, z is a point at infinity, and γ is the geodesic ray from x to z. Then the *horosphere* through x with center z is defined to be the limit as $t \to \infty$ of the sphere of radius t in H with center $\gamma(t)$. More formally, horospheres are the level surfaces of the Busemann function $F = \lim F_t$, where F_t is defined by $F_t(p) = d_H(p,\gamma(t)) - t$. For a horosphere S in H, we denote by d_S the induced path metric on S; that is, $d_S(x,y)$ is the infimum of the lengths of all paths in S from x to y. We will need the following fact about projections onto horospheres:

Proposition 4.2 ([HI], Thm. 4.9) *Let H be a Hadamard manifold with $-b^2 \leq K(H) \leq -a^2 < 0$. If γ is a geodesic tangent to a horosphere S, and if p and q are the projections of $\gamma(\pm\infty)$ onto S, then*

$$\frac{2}{b} \leq d_S(p,q) \leq \frac{2}{a}$$

We will also need the following:

Proposition 4.3 (Projections of horospheres on horospheres are bounded) *Let H be a Hadamard manifold with $-b^2 \leq K(H) \leq -a^2 < 0$. Let S and S' be nonintersecting horospheres based at distinct points of ∂H. Then the S-diameter of the projection $\pi_S(S')$ is at most $4/a$.*

Proof: Suppose S' is based at $x \in \partial H$. By Proposition 4.2, any geodesic segment γ emanating from x and intersecting S in at most 1 point has the property that the S-diameter of $\pi_S(\gamma)$ is at most $2/a$. But every point of S' intersects such a γ, and all the sets $\pi_S(\gamma)$ have a common point, namely $\pi_S(x)$. The proposition clearly follows. $\square$

Finally, we note that, since $d_H(S,*)$ is a convex function ([BGS]), a geodesic intersects a horosphere S in 0,1, or 2 points.

Finite-volume, negatively curved manifolds

Our analysis of complete, finite-volume Riemannian manifolds M with pinched negative sectional curvatures relies on information about the structure of the ends (cusps) of M. Such results are due to Eberlein, Gromov and Margulis. In this subsection we recall, following Eberlein ([Eb]), the basic picture of what the cusps of M look like. We refer the reader to [Eb] and [BGS] for details and proofs.

Now $\widetilde{M}$ is a (pinched) Hadamard manifold with a properly discontinuous action of $\Gamma = \pi_1(M)$ by isometries. Now M has finitely many ends, and each end is a parabolic, Riemannian collared end. It follows that Γ is finitely presented. An end E is *parabolic* if there exists a divergent geodesic ray $\gamma : [0,\infty) \to M$ that converges to E and can be expressed as $\pi \circ \tilde{\gamma}$, where $\pi : \widetilde{M} \to M$ denotes the covering projection and the geodesic ray $\tilde{\gamma}$ determines a point at infinity in $\widetilde{M}$ which is the fixed point of a parabolic isometry of Γ. An end is *Riemannian collared* if there exists a neighborhood U of E, a compact C^2 codimension 1 submanifold N of M and a C^1 diffeomorphism $F : N \times (0,\infty) \to U$ such that the curves $t \to F(n,t), n \in N$ are unit speed distance minimizing geodesics of M that intersect each hypersurface $F(N \times \{s\})$ orthogonally.

N is the projection of a precisely invariant horosphere S in $\widetilde{M}$ at a point at infinity x fixed by some parabolic isometry in Γ, and U is the projection of the corresponding open horoball in $\widetilde{M}$. The stabilizer subgroup Γ_x is a maximal virtually nilpotent subgroup of Γ (indeed every such nilpotent subgroup is of this form), and Γ_x is finitely generated. The compact manifold

N is diffeomorphic to S/Γ_x. We call the group $\pi_1(N) \approx \Gamma_x$ the *cusp subgroup* of Γ.

Finite-volume quotients of the (real) rank one symmetric spaces provide the most explicit examples of negatively curved manifolds. Gromov-Thurston ([GT]) have constructed examples, in dimension 4 and higher, of manifolds M (compact and noncompact) with metrics of pinched negative curvature which are not quotients of any rank one symmetric space. It still seems to be open whether or not the fundamental group of a negatively curved manifold is linear. We will work out some combinatorial properties of these groups.

4.2 The Electric Space Associated to a Cusped Manifold

Our combinatorial study of the fundamental group $\Gamma = \pi_1(M)$ of a negatively curved manifold with one cusp takes place in the coned-off Cayley graph $\widehat{\Gamma}$ of Γ with respect to the cusp subgroup H (see Section 2). It is much easier to do geometry in a "thicker" space which is quasi-isometric to $\widehat{\Gamma}$, and where some Riemannian geometry is available. We now introduce such a space.

Recall that Γ acts freely by isometries on the pinched Hadamard manifold $\widetilde{M}$ with one orbit of parabolic fixed points. Choose a Γ-invariant set of disjoint horoballs centered on the parabolic fixed points (see [BGS]). These horoballs can be thought of as lifts of the cusps of M. Let X be the space formed by deleting the interiors of all of these horoballs, and give X the path metric. This makes each component of the boundary of X a totally geodesic horosphere. Now Γ acts freely and cocompactly by isometries on X. Choosing a basepoint $x \in X$ lying on a horosphere, the natural map $\gamma \mapsto \gamma \cdot x$ gives a quasi-isometry of Γ with X; each distinct coset gH sits on its own horosphere.

Definition (Electric space): To the group $\Gamma = \pi_1(M)$ we have already associated a space X which is quasi-isometric to Γ. The *electric space* [1] $\widehat{X}$ is the quotient of X obtained by identifying points which lie in the same horospherical boundary component of X. As a quotient $\widehat{X}$ has a path pseudometric $d_{\widehat{X}}$ induced from the path metric d_X. Another way to describe the pseudometric $d_{\widehat{X}}$ is as follows: First let

[1] The term *electric space* was suggested to me by Bill Thurston, since geodesics in this space behave like lightning bolts shooting between metal plates (the horospheres).

$$d_Y(x,y) = \begin{cases} 0 & \text{if } x,y \in S \text{ for some horosphere boundary component of } X \\ \\ d_X(x,y) & \text{otherwise} \end{cases}$$

Then $d_{\widehat{X}}(x,y)$ is defined to be the infimum of $\sum d_Y(x_i, x_{i+1})$ over all sequences of points $x = x_1, x_2, \ldots, x_n = y$. Hence the pseudo-metric $d_{\widehat{X}}$ can be thought of as a pseudo-metric on X, where the distance between two points is the length of the shortest path between them, but path-length along a horosphere boundary component of X is measured as zero length. Locally $d_{\widehat{X}}$ agrees with $d_{\widetilde{M}}$ outside these horospheres.

When drawing a path γ in $\widehat{X}$, called an *electric path*, we will for simplicity always draw any subsegment of γ lying on a horosphere S as a geodesic on S, even though it is only the endpoints and not the path that matters. The *electric length* of an electric path γ, denoted by $\ell_{\widehat{X}}(\gamma)$, is just the sum of the X-lengths of the subpaths of γ lying outside every horosphere.

An *electric geodesic* between $x, y \in \widehat{X}$ is a path γ in $\widehat{X}$ from x to y such that $\ell_{\widehat{X}}(\gamma)$ is minimal. It is not hard to see that for an electric geodesic γ between x and y, the subpaths of γ lying outside every horosphere qualitatively consist of the following: the shortest X-path from x to some horosphere, followed by a union of paths which are shortest X-paths between two horospheres, followed by the shortest X-path from some final horosphere to y. By looking at the first variation of a path, it is clear that each of these subpaths intersects the horospheres at right angles (however we will never use this fact). An *electric P-quasi-geodesic* is simply a P-quasi-geodesic in the (pseudo-)path-metric space $\widehat{X}$.

We have seen that the cocompact, properly discontinuous action by isometries of Γ on the neutered space X gives a (P,ϵ) quasi-isometry $f : \Gamma \to X$ for some $P, \epsilon > 0$. The quasi-isometry f can be used to define a quasi-isometry $\widehat{f} : \widehat{\Gamma} \longrightarrow \widehat{X}$ by simply defining $\widehat{f}(v) = f(v)$ for all $v \in \Gamma$ (we define $\widehat{g}$ similarly). Recall that quasi-isometries need only be defined on nets, and Γ is clearly a net in $\widehat{\Gamma}$. Hence we have the following commutative diagram:

$$\begin{array}{ccc} \Gamma & \xrightarrow{f} & X \\ \downarrow & & \downarrow \\ \widehat{\Gamma} & \xrightarrow{\widehat{f}} & \widehat{X} \end{array}$$

Noting that paths in $\widehat{\Gamma}$ correspond via f to paths in $\widehat{X}$, it is not difficult to verify that $\widehat{f}$ is a $(2RP, \epsilon+1)$-quasi-isometry, where $R > 0$ is such that the $\widetilde{M}$-distance between any two horospheres of the Γ-equivariant collection of chosen horospheres is at least R. The fact that $\widehat{\Gamma}$ is quasi-isometric to $\widehat{X}$ allows us to prove (quasi-)geometric facts about $\widehat{\Gamma}$ by first proving them in $\widehat{X}$, which is done by doing geometry in $\widetilde{M}$ (see Section 4.3). But (quasi-)geometric properties of $\widehat{\Gamma}$ give us combinatorial properties of $\widehat{\Gamma}$, which in turn give combinatorial properties of the group Γ.

4.3 Electric Geometry

In this section we study the behavior of quasi-geodesics in the electric space $\widehat{X}$. This gives corresponding properties for (quasi)geodesics in $\widehat{\Gamma}$. Throughout this section we assume that any two of the deleted horoballs in $\widehat{X}$ are a distance (in the $\widetilde{M}$ metric) of at least 1 from each other.

We will prove the following key facts about quasi-geodesics in the Electric Space $\widehat{X}$:

- Electric quasi-geodesics electrically track $\widetilde{M}$-geodesics (Lemma 4.5).
- Electric Space is a δ-hyperbolic pseudo-metric space (Proposition 4.6).
- The way that electric quasi-geodesics without backtracking penetrate horospheres is close to the way that their tracking $\widetilde{M}$-geodesics penetrate those horospheres (Lemmas 4.8 and 4.9).
- Electric quasi-geodesics satisfy an analog of the Bounded Coset Penetration property, where 'coset' is replaced by 'horosphere' (Proposition 4.10).

We begin with a simple property of horospheres in a pinched Hadamard manifold Y.

Definition (Visual size of a horosphere): Let S be a horosphere in Y, and let γ be a geodesic in Y not intersecting S. Let T be the set of points $s \in S$ so that there exists some t for which $\overline{\gamma(t)s} \cap S = \{s\}$. Then the visual size of S (with respect to γ) is defined to be the diameter of T in the metric d_S. The *visual size* of the horosphere S is defined to be the supremum of the visual size of S with respect to γ, where the supremum is taken over all geodesics γ which do not intersect S.

Lemma 4.4 (Horospheres are visually bounded) *Horospheres in a pinched Hadamard manifold Y have (uniformly) bounded visual size.*

Proof: Let S be a horosphere in Y determining a horoball T, let π_S be orthogonal projection onto S, and let γ be a geodesic not intersecting S. First off, Proposition 4.2 implies that $diam(\pi_S(\gamma)) \leq 2/a$.

Now consider any geodesic from a point $x \in \gamma$ to some point $y \in S$. Consider the triangle in Y with vertices $x, y, \pi_S(X)$. Triangles in Y are uniformly thin: there exists $\delta = \delta(a, b)$ so that the δ-neighborhood of the union of two sides of any triangle in Y contains the third side. This follows easily from the corresponding fact in hyperbolic space, using comparison triangles (see [GH]). Since $\overline{xy}$ and $\overline{x\pi_S(x)}$ lie in the complement of T, the geodesic $\overline{y\pi_S(x)}$ lies within a Hausdorff distance δ of S. On the other hand, this Hausdorff distance is pinched between the corresponding distances in the spaces of constant curvatures $-a^2$ and $-b^2$ ([HI], Lemma 4.2). In particular it approaches ∞ as $d_S(y, \pi_S(x)) \to \infty$. Hence $d_S(y, \pi_S(x)) \leq C$ for some constant C, independent of x, β, γ. Hence the visual size of S is bounded by $2/a + 2C$. $\square$

The following lemma is a generalization of the fact that a quasi-geodesic in a pinched Hadamard manifold lies in a uniformly bounded nieghborhood of a geodesic. This fact was first noticed by Morse in the 1920's, and rediscovered/exploited later in the proof of Mostow Rigidity. Lemma 4.5 says that, even if you are allowed to move without cost in horospheres, it is *still* more efficient to travel close to a geodesic (as long as there is a uniform bound on the distance between any two of the horospheres).

Lemma 4.5 (Electric quasi-geodesics electrically track $\widetilde{M}$-geodesics) *Let $P > 0$ be given. Then there exist constants $K = K(P), L = L(P) > 0$, with the following property: let β be any electric P-quasi-geodesic from x to y, and let γ be the $\widetilde{M}$-geodesic from x to y. Then any subsegment of β which lies outside $Nbhd_{\widehat{X}}(\gamma, K)$ must have electric length at most L. In particular, any electric P-quasigeodesic from x to y stays completely inside $Nbhd_{\widehat{X}}(\gamma, K + L/2)$.*

Proof: We assume that β does not lie in an electric unit neighborhood of a horosphere, since this case is trivial.

Choose K so that $K \geq \frac{1}{a}\log(10P)$. Suppose that β' is a subsegment of β lying completely outside $Nbhd_{\widehat{X}}(\gamma, K)$; say $z = \beta'(0)$ satisfies $d_{\widehat{X}}(z, \gamma) = K$

and and let w be the last point of β' with $d_{\widehat{X}}(w,\gamma) = K$. Note that since β' lies outside $Nbhd_{\widehat{X}}(\gamma, K)$, in particular any horoball which β' penetrates must also lie outside this neighborhood, so that any such horoball must not intersect γ.

Recall that we have chosen the horoballs so that each horoball is at least a distance 1 away from any other horoball. Also recall that $\widetilde{M}$-lengths of paths lying outside of $Nbhd_{\widetilde{M}}(\gamma, K)$ (hence outside $Nbhd_{\widehat{X}}(\gamma, K)$) decrease by a factor of at least e^{-aK} under orthogonal projection onto γ ([Kl], Prop. 3.9.11). We know from Lemma 4.4 that horospheres have visual size at most D for some constant $D > 0$.

Let z' (resp. w') denote the image of z (resp. w) under orthogonal projection onto γ. Now β (hence β') is an electric P-quasi-geodesic, so

$$\begin{aligned} \ell_{\widehat{X}}(\beta') &\leq P \cdot [d_{\widehat{X}}(z,z') + d_{\widehat{X}}(z',w') + d_{\widehat{X}}(w',w)] \\ &\leq P \cdot [d_{\widetilde{M}}(z,z') + d_{\widetilde{M}}(z',w') + d_{\widetilde{M}}(w',w)] \\ &\leq P \cdot [K + \ell_{\widehat{X}}(\beta') \cdot e^{-aK} + D \cdot (\text{ \# of horospheres}) \cdot e^{-aK} + K] \quad (*) \end{aligned}$$

where "(# of horospheres)" means the number of horospheres penetrated by β'. The third inequality holds since orthogonal projection onto γ decreases lengths of paths outside $Nbhd_{\widetilde{M}}(\gamma, K)$ by a factor of at least e^{-aK}, and since the visual size of a horoball is at most D by Lemma 4.4, where D depends only on $\widetilde{M}$. Now horoballs are separated by an $\widetilde{M}$-distance of at least 1, so

$$1 \cdot (\text{\# of horospheres } + 1) \leq \ell_{\widehat{X}}(\beta')$$

and so

$$(\text{\# of horospheres }) \leq \ell_{\widehat{X}}(\beta') - 1.$$

Plugging back into the inequality $(*)$ above gives:

$$\ell_{\widehat{X}}(\beta') \leq P \cdot [2K + \ell_{\widehat{X}}(\beta') \cdot e^{-aK} + D \cdot (\ell_{\widehat{X}}(\beta') - 1) \cdot e^{-aK}]$$

and so

$$\ell_{\widehat{X}}(\beta') \cdot [1 - (D+1)Pe^{-aK}] \leq 2KP - DPe^{-aK}.$$

Choosing K large enough, we can make $[1 - (D+1)Pe^{-aK}] \geq \frac{1}{2} > 0$, so then :

$$\ell_{\widehat{X}}(\beta') \leq \frac{2KP - DPe^{-aK}}{1/2} \leq 4PK$$

Now let $L = 4PK$, and we are done. $\square$

Note: It is easy to see that Lemma 4.5 is not true if the horoballs used to construct the Electric Space are allowed to be arbitrarily close to each other.

Lemma 4.5 is a useful tool for determining the behavior of geodesics and quasi-geodesics in the electric space. To begin with, we have the following proposition, which is central to our approach to analyzing the structure of fundamental groups of (pinched) negatively curved manifolds with cusps.

Proposition 4.6 (Electric Space is δ-hyperbolic) *The electric space $\widehat{X}$ is a δ-hyperbolic pseudometric space for some $\delta > 0$.*

Proof: Suppose we are given an electric triangle $\Delta(x, y, z)$. By 'triangle' we mean the union of the 3 electric geodesics $\overline{xy}, \overline{yz}$, and $\overline{xz}$. Now consider the $\widetilde{M}$-geodesics between each pair of vertices. The resulting triangle in $\widetilde{M}$ is δ''-thin (measured in the $\widetilde{M}$ metric) for some $\delta'' > 0$.

Let $K = K(1), L = L(1)$ be the constants given by Lemma 4.5. Now suppose we are given a point p on $\overline{xy}$. Then there is a point p' on the $\widetilde{M}$-geodesic from x to y so that $d_{\widehat{X}}(p, p') \leq K + L/2$, and there is a point q' on, say, the $\widetilde{M}$-geodesic from x to z with

$$d_{\widehat{X}}(p', q') \leq d_H(p', q') \leq \delta''$$

(see Figure 2).

Now Lemma 4.5 implies that there is a point q on $\overline{xz}$ so that $d_{\widehat{X}}(q, q') \leq K + L/2$; without changing we may assume that q' does not lie on a horosphere. It follows from these observations that

$$d_{\widehat{X}}(p, q) \leq (K + L/2) + \delta'' + (K + L/2).$$

Hence triangles in $\widehat{X}$ are (electrically) $(2K + L + \delta'')$-thin. $\square$

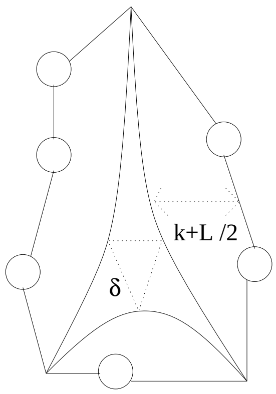

Figure 2. Why electric triangles are (uniformly) electrically thin.

Penetration Properties of Electric Quasi-Geodesics

We've just seen that electric quasi-geodesics electrically track $\widetilde{M}$-geodesics. The rest of this section is devoted to understanding more precise information: namely, how the sequence of horospheres which an electric quasi-geodesic β penetrates compares to the the sequence penetrated by the $\widetilde{M}$-geodesic α which β (electrically) tracks. Furthermore, if both β and α penetrate a certain horosphere, do they have to enter and exit close-by to each other?

Before answering these questions we need :

Lemma 4.7 (Projections of electric quasi-geodesics onto horospheres) *Suppose β is an electric P-quasi-geodesic which does not penetrate S. Then there exists a constant $D = D(P)$ so that the projection of β onto S has S-length at most $D \cdot \ell_{\widehat{X}}(\beta)$.*

Proof: Recall that we are assuming that any two deleted horoballs are an $\widetilde{M}$-distance of at least 1 from each other. It follows that the the projection of any other horosphere onto the horosphere S has S-length bounded by some $D' > 0$. Furthermore, the subsegments of β connecting two horospheres are P-quasi-geodesics in the $\widetilde{M}$-metric, so they stay a bounded distance from an $\widetilde{M}$-geodesic. As $\widetilde{M}$-geodesics have projections of length at most $2/a$ (by Proposition 4.2), these subsegments of β have projections of S-length at most D'' for some $D'' > 0$. Hence the S-length of the projection of β onto S is bounded by

$$(\# \text{ of horospheres}+1)\cdot D' + (\# \text{ of horospheres}+1)\cdot D'' \leq \ell_{\widehat{X}}(\beta)\cdot D' + \ell_{\widehat{X}}(\beta)\cdot D''$$

where the inequality comes from the fact that the number of horospheres that β penetrates is at most $\ell_{\widehat{X}}(\beta) - 1$ (this fact was noted in the proof of Lemma 4.5). Choosing $D = D' + D''$ completes the proof. $\square$

Definition: An electric quasi-geodesic β is said to be a *quasi-geodesic without backtracking* if β is a quasi-geodesic in $\widehat{X}$, and if, for each horosphere S in $\widehat{X}$ which β penetrates, β never returns to S after leaving S.

In order to simplify statements concerning how paths penetrate horospheres, we will only consider paths without backtracking. When faced with a situation where paths may backtrack, we will be careful to reduce the problem to this simpler case. Note that the image under $\widehat{f} : \widehat{\Gamma} \to \widehat{X}$ of a geodesic in $\widehat{\Gamma}$ is a quasi-geodesic without backtracking in $\widehat{X}$.

Lemma 4.8 (Comparing penetration patterns I) *Let β be an electric P-quasi-geodesic (without backtracking) from x to y, and let γ be the $\widetilde{M}$-geodesic from x to y. Then there exists some constant $D = D(P)$ so that if precisely one of $\{\beta, \gamma\}$ penetrates S, then the S-distance from the point of entry of this path into S to its exit point is at most D.*

Proof: First suppose γ intersects S and β does not intersect S. Let x' and y' be the points of intersection of γ with S (the case $x' = y'$ is trivial); we must bound $d_S(x', y')$. Denote by π_S the orthogonal projection onto S, and denote by S' the complement of the horoball corresponding to S. Note that π_S is length decreasing on paths in S' (exponentially so, in fact).

Choose $K = K(P)$ and $L = L(P)$ as in Lemma 4.5, and $E = E(P)$ as in Lemma 4.7.

We claim that the S-length of the orthogonal projection $\pi_S(\beta)$ is bounded by a constant depending only on P. To prove this, first note that by Lemma 4.2 and Lemma 4.7, we have that each component of $\pi_S(Nbhd_{\widehat{X}}(\gamma, K))$ has S-diameter at most $E + 4/a$. Hence each component of

$$\pi_S(\beta \cap (Nbhd_{\widehat{X}}(\gamma, K) \cap S'))$$

has S-diameter at most $E+4/a$. By Lemma 4.5, $\beta \cap (Nbhd_{\widehat{X}}(\gamma, K))^c \cap S')$ has electric length at most L. So by Lemma 4.7,

$$\ell_S(\pi_S(\beta \cap (Nbhd_{\widehat{X}}(\gamma, K))^c \cap S'))) \leq EL$$

These two projection estimates immediately imply (since β is connected) that

$$\ell_S(\pi_S(\beta)) \leq (E+4/a) + EL + (E+4/a) \leq L'$$

for $L' = (L+2)E + 8/a$.

Now Lemma 4.2 shows that $d_S(\pi_S(x), x') \leq 2/a$ and $d_S(\pi_S(y), y') \leq 2/a$. Since $\pi_S(\beta)$ is a path from $\pi_S(x)$ to $\pi_S(y)$, the lemma now follows setting $D = 2/a + L' + 2/a$.

The proof of the case when β intersects S and γ does not intersect S is similar. □

Lemma 4.9 (Comparing penetration patterns II) *Let β be an electric P-quasi-geodesic (without backtracking) from x to y, and let γ be the $\widetilde{M}$-geodesic from x to y. Then there exists a constant $D = D(P)$ so that if both β and γ penetrate some horosphere S, the point of entry of β into S is an S-distance of at most D from the point of entry of γ into S; similarly for the exit points.*

Proof: Let $\gamma(s_0)$ denote the first point of entry of γ into S, and let $\beta(t_2)$ denote the first point of entry of β into S; we need to bound $d_S(\beta(t_2), \gamma(s_0))$.

Choose $K = K(P)$ as in Lemma 4.5, and let t_1 be the largest element of $\{t \leq t_2 : \beta(t) \in Nbhd_{\widehat{X}}(\gamma, K)\}$.

By Lemma 4.5, $\beta([t_1, t_2])$ has electric length at most $L = L(K)$ for some constant L; hence $\pi(\beta([t_1, t_2]))$ has S-length at most $L' = L'(P)$ for some constant L', by Lemma 4.7. Also note that Lemma 4.2, together with the fact that π is distance decreasing outside the horoball corresponding to S, that $d_S(\pi(\beta(t_1)), \gamma(s_0)) \leq K + 2/a$. Hence we have $d_S(\beta(t_2), \gamma(s_0)) \leq L' + (K + 2/a)$, as desired.

The same argument works for the points where β and γ exit S. □

We are now ready to prove the analogue of the Bounded Coset Penetration property for electric quasi-geodesics. From this it will actually follow that the pair (Γ, H) has the Bounded Coset Penetration property (Theorem 4.11).

Proposition 4.10 (Bounded Horosphere Penetration) *Let $P, R > 0$ be given. Suppose u and v are electric P-quasi-geodesics (without backtracking) which begin at the same point and end an $\widetilde{M}$-distance of at most R from each other. Then there exists a constant $C = C(P, R)$ so that, for any horosphere S :*

1. *If u enters S but v does not enter S, then u travels an S-distance of at most C.*
2. *If both u and v enter S, then the point at which u enters (exits) S is an S-distance of at most C from the point at which v enters (resp. exits) S.*

Proof: Let γ_u and γ_v be the $\widetilde{M}$-geodesics between the endpoints of u and v, respectively. Note that since $\widetilde{M}$ is a negatively curved path-metric space, γ_u and γ_v are so-called *k-fellow travellers* in $\widetilde{M}$; that is, there exists a constant $k > 0$, depending only R and on the pinching constants for the curvature of $\widetilde{M}$, so that

$$d_{\widetilde{M}}(\gamma_u(t), \gamma_v(t)) \leq k \quad \text{for all } t \quad (**).$$

First suppose that u enters S but v doesn't enter S. By Lemma 4.8, there is a constant $D = D(P)$ so that the point of entry (if any) of γ_v into S is within an S-distance of at most D from the point at which γ_v exits S. By (**) above, there exists a constant $D' = D'(P)$ so that the point (if any) of entry of γ_u into S is within an S-distance of at most D' of its exit point. Hence, by either Lemma 4.8 or Lemma 4.9 (depending on whether γ_u penetrates S or not), u travels a bounded distance in S.

Now suppose that both electric P-quasi-geodesics u and v enter S. If one of γ_u or γ_v does not enter S (say γ_u doesn't enter), then by $(**)$ the entry point of γ_v into S (if there is one) is a bounded distance from its exit point. It follows from Lemma 4.8 that u travels a bounded amount in S, and it follows from either Lemma 4.8 or Lemma 4.9 (depending on whether or not γ_v penetrates S) that v also travels a bounded distance in S.

Now if both γ_u and γ_v enter S, it follows from $(**)$ and two applications of Lemma 4.9 that the entry (exit) point of u into S is a bounded S-distance from the entry (resp. exit) point of v into S, and we are done.

□

4.4 Hyperbolicity of $\pi_1(M)$ Relative to the Cusp Subgroup

As always, M will denote a complete Riemannian manifold of finite-volume, pinched negative sectional curvatures, and with a single cusp. $\Gamma = \pi_1(M)$ will denote the fundamental group of M, and $H < \Gamma$ will denote the cusp subgroup. With the geometry of the electric space $\widehat{X}$ in hand, we are ready to prove the main theorem of this section:

Theorem 4.11 *The group Γ is hyperbolic relative to H, and the pair (Γ, H) has the Bounded Coset Penetration property.*

Proof: Recall (Section 4.2) that there is a quasi-isometry $\widehat{f} : \widehat{\Gamma} \to \widehat{X}$ with the following property : there is a correspondence between left cosets gH of $\Gamma \subset \widehat{\Gamma}$ and horospheres $gS \subset X$, where X is the neutered space. Since we have the commutative diagram

$$\begin{array}{ccc} \Gamma & \xrightarrow{f} & X \\ \downarrow & & \downarrow \\ \widehat{\Gamma} & \xrightarrow{\widehat{f}} & \widehat{X} \end{array}$$

it follows that if a $\widehat{\Gamma}$-geodesic u enters a coset gH at $x \in \Gamma$ and exits gH at $y \in \Gamma$, then the electric quasi-geodesic $\widehat{f}(u)$ enters the horosphere S corresponding to gH at $\widehat{f}(x) = f(x)$, and exits S at $\widehat{f}(y) = f(y)$. In particular, since u is a geoedesic and so only enters and exits any coset at most once, it follows that $\widehat{f}(u)$ is an electric quasi-geodesic without backtracking. More generally, the image under $\widehat{f}$ of a quasi-geodesic without backtracking in $\widehat{\Gamma}$ is a quasi-geodesic without backtracking in $\widehat{X}$.

From these observations it easily follows from Proposition 4.6 (Electric Space is hyperbolic) that $\widehat{\Gamma}$ is a δ-hyperbolic metric space for some $\delta > 0$, and it follows from Proposition 4.10 (Bounded horosphere penetration) the pair (Γ, H) has the Bounded Coset Penetration Property.

$\square$

Applications to isoperimetric functions for Γ

Let M^n denote a complete, finite-volume, (pinched) negatively curved Riemannian manifold with a single cusp. Let $\Gamma = \pi_1(M^n)$, and let H denote the cusp subgroup.

Recall that the *Dehn function* for a group Γ is a minimal isoperimetric function for Γ. Gromov claims in Section 5.6 of [Gr2] that the Dehn function for Γ equals the Dehn function for H.

Theorem 4.12 *The Dehn function for $\Gamma = \pi_1(M^n)$ equals the Dehn function for the cusp subgroup H.*

Proof: It follows immediately from Theorem 4.11 and Theorem 3.8 that any isoperimetric function for the cusp subgroup is an isoperimetric function for Γ.

For the converse direction, first note that the Dehn function for Γ is equivalent to that of the neutered space associated to Γ (by a slight variation of [Bu]). But there is a length-decreasing (hence area-decreasing) retraction of the neutered space X associated to Γ onto a single horosphere $S \subset X$. Hence if $\gamma \subset S$ is a loop, the minimal area disk D with $\partial D = \gamma$ must also lie in S. Hence the Dehn function of the neutered space is at least as big as that of the horosphere S, which is equivalent to that of the cusp subgroup. $\square$

If $M^{2n}, n > 2$ is a finite volume complex-hyperbolic n-manifold with cusps (here $2n$ is the real dimension), the cusp subgroup of $\Gamma = \pi_1(M)$ is commensurable with the Heisenberg group

$$H_{2n-1} =< x_1, y_1, \ldots, x_{n-1}, y_{n-1}, z : [x_i, y_i] = z, z \text{ central} >$$

Since the Dehn functions for these groups have been determined in [E et al.] for $n = 2$ and [Al] for $n > 2$, we may draw the following consequence from Theorem 4.12.

Corollary 4.13 (Dehn functions for complex hyperbolic manifolds) *Let Γ be the fundamental group of a complete, finite volume, complex-hyperbolic n-manifold M^{2n}, where $n \geq 2$. Then the Dehn function for Γ is linear if M is compact, and otherwise is cubic if $n = 2$, and quadratic if $n > 2$.*

We believe the technique of [Al] can be generalized to other two-step nilpotent groups, so that Theorem 4.12 should also imply an analogue of Corollary 4.13 for Quaternionic and Cayley hyperbolic manifolds.

4.5 A Fast Solution to the Word Problem for $\Gamma = \pi_1(M)$

Let M be a complete Riemannian manifold of finite-volume, pinched negative sectional curvatures, and with a single cusp. In this section we use the fact that γ is hyperbolic relative to its cusp subgroup H to give a fast solution to the word problem for Γ. Since H is (virtually) nilpotent, it is first necessary to discuss normal forms in nilpotent groups.

Nilpotent groups

Let N be a nilpotent group with an ordered set $\{a_1, \dots a_m\}$ of generators so that a_i occurs further down in the lower central series than a_{i+1} for each i. There is a normal form $a_1^{s_1} a_2^{s_2} \cdots a_m^{s_m}$ for elements of N. We keep track of this normal form by a counter $(s_1, \dots, s_m)$, with each s_i written in binary. Given a word w in normal form, the normal form of $w \cdot a_i$ is determined by a simple polynomial equation which tells you how the numbers $(s_1, \dots, s_m)$ are to be updated; this polynomial is a polynomial in m variables, of degree depending only on the nilpotency class of N. For example, for the 3-dimensional integral Heisenberg group $N = < a, b, c : [a, b] = c, [a, c] = [b, c] = 1 >$, the updates look like:

$$(i, j, k) \xmapsto{a} (i + 1, j, k + j)$$

$$(i, j, k) \xmapsto{b} (i, j + 1, k)$$

$$(i, j, k) \xmapsto{c} (i, j, k + 1)$$

Now suppose we are given a word w of length n, and we want to put w into normal form. We start out with $(0, 0, 0)$, read in one letter at a time, and simply update the counter. How many steps are needed in each update? If we choose a model of computation where multiplication of arbitrary numbers takes constant time, then each update takes constant time. Since we are using a Turing machine as our model of computation, updating counters with large entries takes time. Using crude estimates on multiplying numbers in binary, it is not hard to see that any individual update takes $O((\log n)^p)$ for some integer $p > 0$. Jin-yi Cai (a complexity theorist) has informed me that, in terms of the fastest asymptotic estimates of Strassen-Schonhage, each update can be done in time

$$s(n) = O(\log n \cdot \log \log n \cdot \log \log \log n).$$

Let $T(n)$ denote the time needed to put a word of length n into normal form. By using the "divide and conquer" method, we find that

$$T(n) = 2T(n/2) + s(n).$$

It is now easy to show by induction that this gives a time bound of, say

$$(\text{constant}) \cdot n - (\log n)^2 = O(n).$$

To check whether or not a given word represents the identity element of N, one simply puts the word into normal form, and checks whether or not the counter reads $(0, \ldots, 0)$. Hence the word problem for N is solvable in linear (in word length) time.

The word problem for Γ

The curve-shortening algorithm for relatively hyperbolic groups provides a fast solution to the word problem for fundamental groups of finite-volume, negatively curved manifolds.

Theorem 4.14 *Let $\Gamma = \pi_1(M^n)$, where M^n is a complete, finite-volume (pinched) negatively curved Riemannian manifold with one cusp. Then there is an $O(n \log n)$-time solution to the word problem in Γ.*

Proof: The Theorem follows immediately from Theorem 4.11, Theorem 3.7, and the fact that there is a linear time algorithm to solve the word problem for nilpotent groups. □

Theorem 4.14 also holds in the case when M^n has more than one cusp (see §5). This immediately implies the following special, though interesting consequence:

Corollary 4.15 *There is an $O(n \log n)$-time algorithm to solve the word problem for hyperbolic knot and link complements.*

Corollary 4.15 is useful because most link complements are hyperbolic; i.e., they have the structure of a complete, finite-volume, hyperbolic 3-manifold, with a cusp for each component of the link. This is the fastest algorithm we know of for hyperbolic link groups.

5 Groups Hyperbolic Relative to a Finite Set of Subgroups

In this section we define what it means for a group to be hyperbolic relative to a finite set of subgroups. This occurs in many examples, and is meant to make the methods described above more widely applicable. An important example to keep in mind is the fundamental group of a complete, (non-compact) finite volume, (pinched) negatively curved Riemannian manifold relative to its (finite set of nonconjugate) cusp subgroups. The theory for more than one subgroup is also needed, for example, for the inclusion of small cancellation theory over free products (with amalgamation) into the framework of relatively hyperbolic groups.

Let G be a finitely generated group with Cayley graph Γ, and let $\{H_1, \ldots, H_r\}$ be a finite set of finitely generated subgroups of Γ. We form a new graph $\widehat{\Gamma} = \widehat{\Gamma}(\{H_1, \ldots, H_r\})$ as follows: for each coset $gH_i, 1 \leq i \leq r$ of H_i in G, add a vertex $v(gH_i)$ to Γ, and add an edge $e(gh_i)$ of length $1/2$ from each element gh_i of gH_i to the vertex $v(gH_i)$. We call this the *coned-off Cayley graph* of G with respect to $\{H_1, \ldots, H_r\}$.

Definition: The group G is *hyperbolic relative to* $\{H_1, \ldots, H_r\}$ if the coned-off Cayley graph $\widehat{\Gamma}$ of G with respect to $\{H_1, \ldots, H_r\}$ is a negatively curved metric space.

It is clear how to define the Bounded Coset Penetration property in this context, as the BCP property was defined in terms of paths in $\widehat{\Gamma}$ penetrating a coset gH. Now we make the definition for each coset $gH_1, \ldots, gH_r$. With these alterations in mind, the properties of relatively hyperbolic groups easily extend to the case of a finite set of subgroups.

Note that this more general concept applies to negatively curved manifolds with several cusps. As an illustrative example, suppose that Γ is the fundamental group of a finite-volume hyperbolic 3-manifold with 2 cusps; let H_1 and H_2 be the two cusp subgroups. We form the space X consisting of hyperbolic 3-space with an equivariant union of horoballs deleted (with the path metric); here there are two types of horoballs, which we shall call H_1-horoballs and H_2-horoballs. By adjusting the sizes of these horoballs, we may assume that any H_1-horoball touches any H_2-horoball in at most 1 point. Choosing a point $x \in X$ which lies at one of these intersection point, we obtain the quasi-isometry of Γ with X given by $g \mapsto g \cdot x$.

Now the left cosets of H_i in Γ are sitting on the copies of the boundaries of H_i-horoballs. As each $g \in \Gamma$ is contained in a left H_1-coset and a left H_2-coset, each vertex $g \cdot x$ of the Cayley graph is contained in precisely two horospheres - an H_1-horosphere and an H_2-horosphere. So touching each H_1-horoball are infinitely many H_2-horoballs, and touching each H_2-horoball are infinitely many H_1-horoballs. This pattern of horoballs is illustrated for dimension 2 in Figure 3.

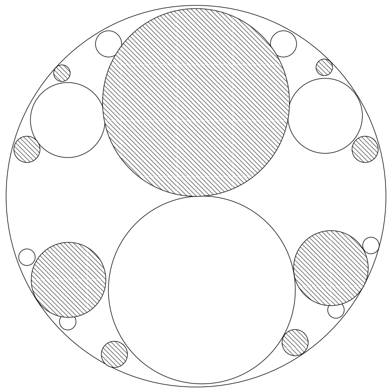

Figure 3. This figure is a schematic of the pattern of horoballs which must be deleted from the hyperbolic plane in forming the space X which is quasi-isometric to the fundamental group of a hyperbolic 2-manifold with two cusps.

Now shrink the horoballs so that any two of them are a hyperbolic distance of at least one apart; this gives a new space X' quasi-isometric to X. For each vertex of the Cayley graph Γ, think of it as splitting into two (one copy following an H_1-horosphere, one following an H_2-horosphere) under this shrinking process, with an edge of finite length connecting the two copies of the vertex. This gives a graph Γ' which is quasi-isometric to Γ. From these observations it is clear that the coned-off Cayley graph $\widehat{\Gamma} = \widehat{\Gamma}(H_1, H_2)$ is quasi-isometric to the electric space $\widehat{X'}$.

All of the results on electric geometry (Sections 4.2 and 4.3) then go through for the case of several cusps. For completeness we state the following.

Theorem 5.1 (Multiple cusps) *Let $\Gamma = \pi_1(M^n)$, where M^n is a complete, (noncompact) finite-volume Riemannian manifold with (pinched) neg-*

ative sectional curvatures. Let $\{H_1, \ldots, H_r\}$ denote the cusp subgroups of Γ. Then the following are true:

1. *Γ is hyperbolic relative to the set $\{H_1, \ldots, H_r\}$ of cusp subgroups, and the pair $(\Gamma, \{H_1, \ldots, H_r\})$ has the Bounded Coset Penetration property.*

2. *There is a curve-shortening algorithm which solves the word problem for Γ in time $O(n \log n)$.*

3. *Γ satisfies precisely the same isoperimetric inequality as any of its cusp subgroups.*

References

[Al] D. Allcock, An isoperimetric inequality for the Heisenberg groups, *Geom. and Func. Anal.*, to appear.

[Alo] J. Alonso, Inégalités isopérimétriques et quasi-isometries, *C. R. Acad. Sci. Paris Serie 1* **311** (1990), p.761-4.

[BGS] W. Ballmann, M. Gromov and V. Schroeder, *Manifolds of Nonpositive Curvature*, Progress in Mathematics, Vol. 61, Birkaüser, 1985.

[Bo] F. Bonahon, Bouts des variétés hyperboliques de dimension 3, *Annals of Math.* 124 (1986), p.71-158..

[Bu] J. Burillo, Equivalence of geometric and combinatorial Dehn functions, preprint, Univ. of Utah, 1995.

[Ca] J. Cannon, The Combinatorial Structure of Cocompact Discrete Hyperbolic Groups, *Geometriae Dedicata* 16 (1984), p. 123-148.

[CDP] M. Coornaert, T. Delzant, A. Papadopoulos, *Notes sur les groupes hyperboliques de Gromov*, Springer Lecture Notes, vol. 1441.

[Eb] P. Eberlein, Lattices in spaces of nonpositive Curvature, *Annals of Math.* 111 (1980), p.435-476.

[E et al.] D.B.A. Epstein, J. Cannon, D.F. Holt, S. Levy, M.S. Patterson and W.P. Thurston, *Word Processing in Groups*, Jones and Bartlett, 1992.

[Fa1] B. Farb, Automatic groups: a guided tour, *L'Enseignement Math.*, Vol.38 (1992), p.291-313.

[Fa2] B. Farb, Combing lattices in semisimple Lie groups, in *Groups Korea 1994*, ed. by A.C. Kim and D.L. Johnson, Walter De Gruyter and Co., 1995.

[Fa3] B. Farb, Relatively hyperbolic and automatic groups with applications to negatively curved manifolds, Princeton University thesis, May 1994.

[Ge1] S.M. Gersten, Subgroups of word-hyperbolic groups in dimension 2, *J. London Math. Soc.* (2) 54 (1996), no. 2, p. 261-283.

[Ge2] S. Gersten, Isoperimetric and isodiametric functions of finite presentations, in : *Geometric Group Theory, Volume 1*, G. Niblo and M. Roller, eds., LMS Lecture Notes 182, Cambridge Univ. Press, 1993.

[Ge3] S. Gersten, Metabolicity of surface groups and amalgams, August 1997 preprint.

[GH] E. Ghys and P. de la Harpe, editors, *Sur les Groupes Hyperboliques d'aprés Mikhael Gromov*, Progress in Mathematics, Vol. 83, Birkaüser, 1990.

[Gr1] M. Gromov, Hyperbolic Groups, *Essays in Group Theory*, Springer Verlag, MSRI series Vol. 8, S.M. Gersten, ed., 1987, pp. 75-263.

[Gr2] M. Gromov, Asymptotic Invariants of Infinite Groups, in *Geometric Group Theory*, G. Niblo and M. Roller, eds., LMS Lecture Notes 182, Cambridge Univ. Press, 1993.

[Gr3] M. Gromov, Infinite Groups as Geometric Objects, Plenary Address, *Proceedings of the International Congress of Mathematicians*, Warsaw, 1983.

[GT] M. Gromov and W.P. Thurston, Pinching Constants for Hyperbolic Manifolds, *Invent. Math.* 89, p.1-12 (1987).

[HI] E. Heintze and H-C. Im Hof, Geometry of Horospheres, *Jour. Diff. Geom.* 12, p. 481-491 (1977).

[HU] J.E. Hopcroft and J.D. Ullman, *Introduction to Automata Theory, Languages and Computation*, Addison-Wesley 1979.

[Kl] W. Klingenberg, *Riemannian Geometry*, de Gruyter, Berlin, 1982.

[Kla] E. Klarreich, Semiconjugacy between Kleinian group actions on the Riemann sphere, 1997 preprint.

[Kr] D. Krammer, An extension of the Moussong complex, 1995 preprint.

[Mi] J. Milnor, A Note on Curvature and Fundamental Group, *Jour. Diff. Geom.* 2, p.1-7.

[MM1] H. Masur and Y. Minsky, Geometry of the complex of curves I: Hyperbolicity, May 1996 preprint.

[MM2] H. Masur and Y. Minsky, Geometry of the complex of curves II: Hierarchical structure, 1998 preprint.

[Sh] H. Short (ed.), Notes on Negatively Curved Groups, in *Group Theory from a Geometrical Viewpoint*, E. Ghys, A. Haefliger, A. Verjovsky, eds., Proc. ICTP Trieste, World Scientific, 1992.

[Sz] A. Szczepanski, Relatively hyperbolic groups, *Michigan Math. Jour.*, to appear.

[Th] W.P. Thurston, *The Geometry and Topology of Three-Manifolds*, Lecture Notes, Princeton University.

Benson Farb
Math Department, University of Chicago
5734 University Ave.
Chicago, Illinois 60637
E-mail: farb@math.uchicago.edu