AN INTRODUCTION TO RELATIVELY HYPERBOLIC GROUPS

A thesis

submitted by

Murphy A. Page

in partial fulfillment of the requirements

for the degree of

Master of Science

in

Mathematics

TUFTS UNIVERSITY

August 2017

Adviser: Genevieve Walsh

Abstract

We look at a few of the various ways in which relatively hyperbolic groups are defined. It has been shown that, if a little care taken with hypotheses, all definitions are equivalent; here we examine the relationship between the definitions using the coned-off Cayley graph and fine δ-hyperbolic graphs. We also go over the definition using geometrically finite convergence groups. Finally, we use the free group F_2 as an example of a relatively hyperbolic group to explore the different definitions.

To my bicycle, Bruce, for bringing me everywhere I needed to be and never bringing me down.

Acknowledgements

I am eternally grateful to the professors that helped get me through this graduate school experience in (more or less) one piece. Thanks to the graduate director Kim Ruane, for always making everything go; to Moon Duchin, for being a stellar mentor; and to my advisor Genevieve Walsh, for convincing me to make good decisions for myself and for keeping me going through the last stretch of this.

Contents

List of Figures

An Introduction to Relatively Hyperbolic Groups

Chapter 1

Introduction

The study of relatively hyperbolic groups is relatively new in the field of geometric group theory but has developed extensively over the past twenty years. There is an abundance of definitions for the notion of "relatively hyperbolic," which, although more or less equivalent, can feel rather different from one another. As a result, the literature on the topic is somewhat labythinthine. This thesis is intended to be an approachable introduction to relatively hyperbolic groups, reviewing some of the various definitions and looking at simple, specific examples.

1.1 History and Significance

Hyperbolic groups were introduced and explored in depth by Gromov in his seminal 1987 paper [7]. In it Gromov also describes the notion of relative hyperbolicity as a generalization of hyperbolicity; this was specifically motivated by the fundamental groups of finite volume, cusped manifolds of negative curvature. Unlike hyperbolic groups, however, relatively hyperbolic groups were left largely untouched for about a decade after this publication.

A precise timeline is difficult to establish, but 1997/98 saw a surge in papers on the topic. Two of the main pioneers here are Bowditch, who developed a theory grounded in Gromov's original proposal in [2] (first draft in '97, not published until '12), and Farb, who proposed a new definiton based on the construction of a "coned-off Cayley graph" in [5]. There now exist around six definitions of relatively hyperbolic groups which have been shown to be equivalent (see [8], in which Hruska demonstrates the equivalence of definitions without requiring the groups to be finitely generated). However, this thesis will focus on only a few of these formulations.

The subject of relatively hyperbolic groups continues to be fruitful, and much work has been done translating concepts from hyperbolic groups to their more general, "relative" versions. For example, Dahmani constructs a relative Rips complex in [3], and there are many papers looking at boundaries of relatively hyperbolic groups. More recently, acylindrically hyperbolic groups have been a popular topic in which relatively hyperbolic groups naturally arise as examples. Specifically, it is known that non-virtually cyclic relatively hyperbolic groups are acylindrically hyperbolic [9].

1.2 Summary

In Chapter 2, we provide a brief overview of hyperbolic groups, going over relevant definitions, properties, etc. Note that we assume the reader is familiar with fundamental ideas from geometric group theory and has a little experience with hyperbolic spaces and groups.

Chapter 3 goes over various definitons of relatively hyperbolic groups. First we look Farb's construction using the coned-off Cayley graph (as an aside, this is the author's favorite definition). We then prove that when this graph is δ-hyperbolic, the bounded coset penetration property is equivalent to fineness, which brings us to a second definition. To close the chapter, we present geometrically finite convergence groups as a dynamical formulation for relatively hyperbolic groups.

Finally, in Chapter 4 we look at F_2 as a relatively hyperbolic group, both as the fundamental group of a punctured torus with a cusp and using the Coned-off Cayley graph.

Chapter 2

Overview of Hyperbolic Groups

In this chapter, we quickly review some fundamental and relevant information about δ-hyperbolic spaces and hyperbolic groups. Sources for this chapter are [1] and [6].

2.1 Preliminaries

Let X be a metric space with distance $d : X \times X \to [0, \infty)$. A path in X is a continuous function from an interval to X, $p : [0, t] \to X$. A geodesic is a path whose length realizes the distance between endpoints. Throughout this chapter we will let X be a proper geodesic metric space, i.e. there exists a (possibly non-unique) geodesic between any two points of X and every closed ball is compact.

Let $\lambda \geq 1, c \geq 0$ be constants. A (λ, c)-quasigeodesic between $x_0, x_1 \in X$ is a path with $p(0) = x_0$, $p(t) = x_1$ and

$$\frac{1}{\lambda} d(x_0, x_1) - c \leq d(x_0, x_1) \leq \lambda d(x_0, x_1) + c.$$

Let $(X, d_X), (Y, d_Y)$ be proper geodesic metric spaces. An isometry from X to Y is a map $f : X \to Y$ that preserves distances, i.e. for all $x, y \in X$ we have

$$d_Y(f(x), f(y)) = d_X(x, y).$$

Given constants $\lambda \geq 1$ and $c \geq 0$, a (λ, c)-quasiisometry between X and Y is a map $g :\to Y$ such that

- for all $x, y \in X$ we have $\frac{1}{\lambda} d_X(x, y) - c \leq d_Y(g(x), g(y)) \leq \lambda d_X(x, y) + c$, and
- for all $y \in Y$ there exists $x \in X$ such that $d_Y(g(x), y) \leq c$.

The second condition here is known as *coarse surjectivity*.

2.2 δ-hyperbolic Spaces

There are various ways of defining δ-hyperbolic spaces which make prescise the heuristic notion that these spaces are "approximately trees."

Definition 2.2.1 A geodesic metric space X is δ-hyperbolic (sometimes called "Gromov hyperbolic") if there is a constant $\delta > 0$ such that, given any geodesic triangle in X, one side of the triangle is contained in the union of the δ-neighborhoods of the other two.

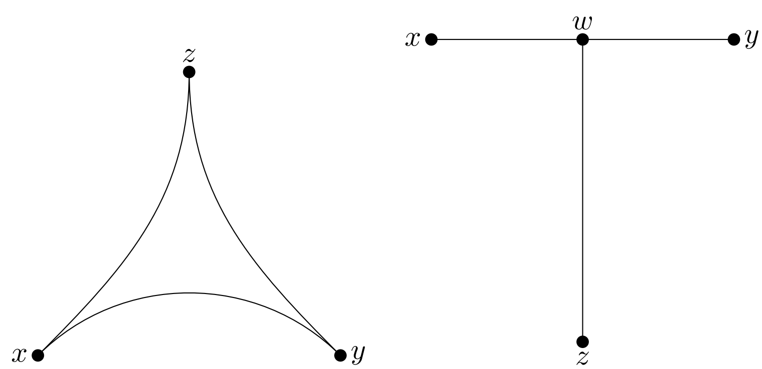

(a) δ-thin triangle (b) 0-thin triangle

The second way we'll look at defining δ-hyperbolic spaces uses the *Gromov product* of two points.

Definition 2.2.2 Fix a basepoint $z \in X$. The *Gromov product* based at z of two points $x, y \in X$ is

$$(x|y)_z := \frac{1}{2}\left(d(x,z) + d(y,z) - d(x,y)\right).$$

Keeping our heuristic in mind, given a triangle (or tripod) in a tree, the Gromov product tells you the length of one of the legs; for example, in figure 2.1b we see that $(x|y)_z = d(w,z)$.

Definition 2.2.3 A geodesic metric space X is δ-hyperbolic if there is some $\delta > 0$ such that for all $x, y, z, w \in X$ we have

$$(x|z)_w \geq \min((x|y)_w, (y,z)_w) - \delta.$$

Note that if for some fixed basepoint $w_0 \in X$, the above inequality holds for all $x, y, z \in X$, then the space X is δ-hyperbolic.

The following properties of δ-hyperbolic spaces will be useful for us in the next chapter:

Lemma 2.2.1 *Let X be δ-hyperbolic and p, p' be distinct geodesics in X beginning and ending at the same points. Then p and p' lie within a δ-neighborhood of each other.*

Proof: Two such geodesics form a bi-gon, which is geodesic triangle with one degenerate side. By the thin-triangles condition for δ-hyperbolic spaces, the bi-gon cannot be fatter than δ. □

Extending this idea to quasi-geodesics requires an important theorem known as the Morse lemma:

Theorem 2.2.1 *Given $\lambda \geq 1, c \geq 0$ there is a constant $r > 0$ depending on δ, λ, c such that any (λ, c)-quasigeodesic in X lies within an r neighborhood of some geodesic in X.*

By combining this theorem with lemma 2.2.1, we achieve the result that will be used later.

Lemma 2.2.2 *Let p, p' be (λ, c)-quasigeodesics beginning and ending at the same points. Then there is some constant $R > 0$ such that p and p' lie within an R-neighborhood of each other.*

Another important result, particularly to the study of hyperbolic groups, is that δ-hyperbolicity is a quasiisometry invariant:

Theorem 2.2.2 *Let X be a δ-hyperbolic space, and suppose there is a quasiisometry $f : X \to Y$ where Y is some geodesic metric space. Then there exists $\delta' > 0$ such that Y is a δ'-hyperbolic space.*

2.2.1 The Boundary of a Hyperbolic Space

Fix a basepoint $w \in X$. We can define the *Gromov boundary* of a δ-hyperbolic space using either geodesic rays or sequences that tend to infinity. Two rays $r_1 : [0, \infty) \to X, r_2 : [0, \infty) \to X$ originating at w are equivalent if they remain within some bounded distance of each other, and the boundary $\partial^r X$ is the set of equivalence classes of rays. A sequence $\{x_i\}$ in X tends to infinity if $\liminf_{i,j\to\infty}(x_i|x_j)_w = \infty$. Two such sequences $\{x_i\}, \{y_j\}$ are equivalent if $\liminf_{i,j\to\infty}(x_i, y_j)_w = \infty$. The boundary $\partial^s X$ is the set of equivalence classes of sequences tending to infinity.

The following proposition combines a number of well-known facts about these boundaries, the proofs of which can all be found in [1].

Proposition 2.2.1 *The definitions of $\partial_r X$ and $\partial_s X$ are equivalent, allowing us to simply refer to ∂X. Additionally,*

- *each definition is independent of the choice of basepoint,*
- *there are topologies for $\partial_r X$ and $\partial_s X$ which are homeomorphic, and*
- *∂X is compact and metrizable.*

A relevant example of the boundary of a hyperbolic space is the hyperbolic plane $\mathbb{H}^2$, whose boundary is the circle S^1. This is easily seen by looking at the Poincare disc model and thinking about $\partial^r X$: no two rays originating from the center of the disc will ever remain within bounded distance from each other, so none will be equivalent to another.

2.3 Hyperbolic Groups

Let G be a group acting on a space X. The action is *by isometries* if every $g \in G$ acts as an isometry of X. It is a *cocompact* action if the quotient space X/G is compact. When there are no accumulation points of the action, it is called *properly discontinuous*; that is, for all $K \subset X$ compact, there are only finitely many $g \in G$ such that $gK \cap K \neq \varnothing$.

Definition 2.3.1 The group G acts on X *geometrically* if the action is properly discontinuous, cocompact, and by isometries.

Definition 2.3.2 A group G is *hyperbolic* if there exists some δ-hyperbolic space on which G acts geometrically.

A useful fact is that any group acts geometrically on its Cayley graph; thus, if a group G has a presentation such that the corresponding Cayley graph is δ-hyperbolic, G is hyperbolic. Indeed, it is known that the Cayley graphs for any presentations of G are quasiisometric, so by theorem 2.2.2, if any Cayley graph is δ-hyperbolic, every Cayey graph is.

The free group on two generators is a wonderful example of a hyperbolic group, as we can see its hyperbolicity in various ways. Its Cayley graph is a tree and therefore 0-hyperbolic (see figure 2.1).

Now F_2 is also the fundamental group of a punctured torus, so we can see its hyperbolicity through a geometric action on a closed, convex subset S of the hyperbolic plane $\mathbb{H}^2$ (which is quasi-isometric to a tree). By choosing the appropriate Moebius transformations as the two generators for the group, specifically ensuring that their commutator is a hyperbolic transformation, we can obtain a compact punctured torus as the quotient of the subset S by F_2. We can similiary see that the fundamental group of the two-holed torus is hyperbolic; figure 2.2 shows a fundamental domain for the action on $\mathbb{H}^2$.

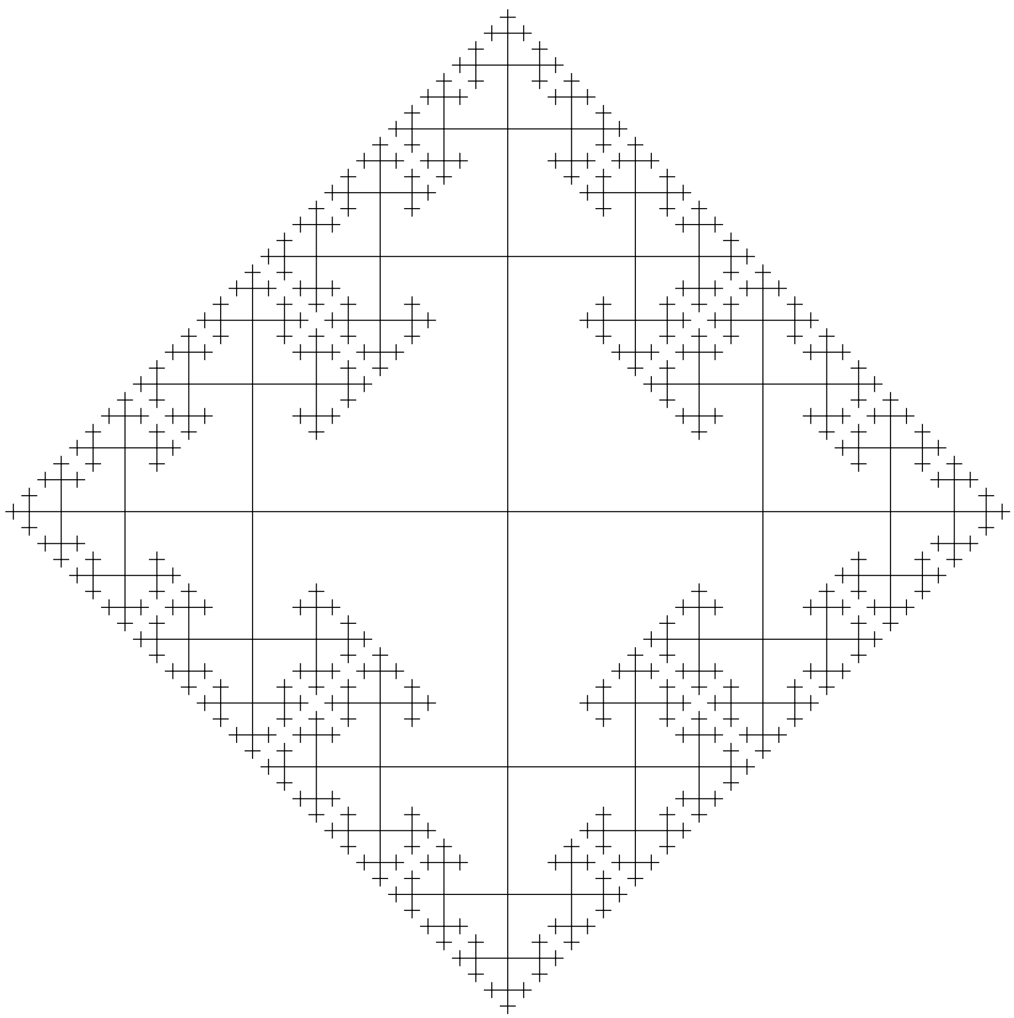

Figure 2.1: The Cayley graph for the free group on two generators

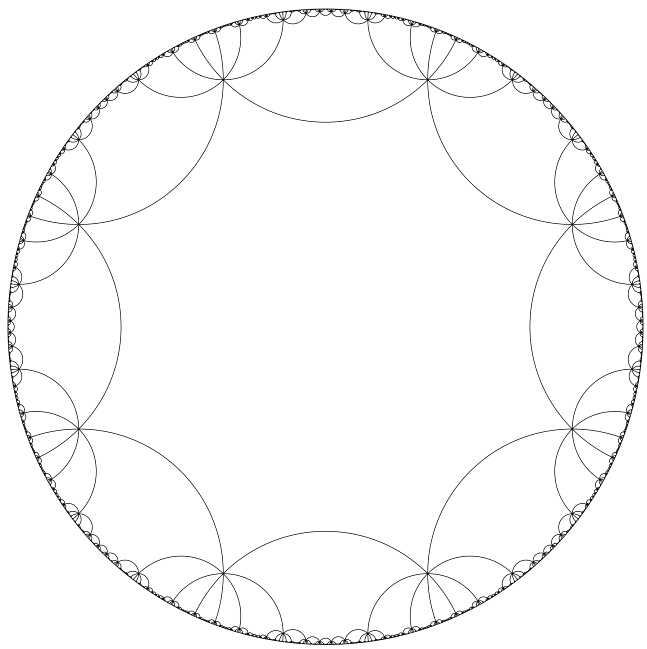

Figure 2.2: The two-holed torus arises as a quotient of $\mathbb{H}^2$ by "folding" the hyperbolic octagon. Image made with Curt McMullen's LIM (http://www.math.harvard.edu/~ctm/math101/www/programs/index.html

Chapter 3

Relatively Hyperbolic Groups

We will now look at some of the ways to define relatively hyperbolic groups. A group is hyperbolic when it acts geometrically on some δ-hyperbolic space, so there are various ways in which a group can *fail* to be hyperbolic: it could act geometrically only on spaces which are not δ-hyperbolic (such as $\mathbb{Z} \oplus \mathbb{Z}$) or it could fail to act geometrically on a δ-hyperbolic space in some way (such as the fundamental group of a cusped space, where the action fails to be cocompact). The hope of relatively hyperbolicity is to find the groups who miss the hyperbolic mark one way or another, but not by too much. When we think of a cusped space, for example, there is a sense that the cusp (which prevents hyperbolicity) can be easily controlled. This chapter will make precise the notion of being "not too far off" from hyperbolic.

3.1 The Coned Off Cayley Graph

In this section we will let $G = \langle S|R\rangle$ be a group with and Γ be its Cayley graph. For simplicity, we will conflate words in G and their corresponding paths in Γ, and we will use the word metric d_S to measure distances in G and Γ.

Definition 3.1.1 For a $H < G$, we construct the *coned off Cayley graph* $\widehat{\Gamma}(G,H)$ from Γ by adding a cone vertex $v(gH)$ for each distinct coset gH and an edge from a vertex $g_0 \in \Gamma$ to $v(gH)$ if $g_0 \in gH$. To make $\widehat{\Gamma}(G,H)$ a metric space, we declare that each new edge to a cone vertex has length $1/2$, and when there is no ambiguity around the subgroup, we will simply use the notation $\widehat{\Gamma}$.

(Note that we may replace H with a finite collection of subgroups $\{H_i\}$ with no complications, but for simplicity of notation in this section, we will only be considering one subgroup.)

We associate a path w in Γ to a path $\hat{w}$ in $\widehat{\Gamma}$ as follows: For every maximal subpath

$z = z_0 z_1 \cdots z_{n-1} z_n$ of w such that $n > 0$ and every vertex of z belongs to the same coset gH, replace $z_1 \cdots z_{n-1}$ with $v(gH)$ using the edges of length 1/2, $(z_0, v(gH))$ and $(v(gH), z_n)$. In the case that $n = 1$, replace the edge from the Cayley graph (z_0, z_1) with the two edges of length 1/2 passing through $v(gH)$.

A path w whose associated path $\hat{w}$ contains such a subpath $\hat{z} = z_0 \cdot v(gH) \cdot z_n$ is said to *penetrate* the coset gH with entering vertex z_0 and exiting vertex z_n. A path w is *without backtracking* if it never penetrates the same coset more than once. We call a path w a *relative (quasi)geodesic* if $\hat{w}$ is a (quasi)geodesic in $\widehat{\Gamma}$.

Definition 3.1.2 The coned off Cayley graph $\widehat{\Gamma}$ has the *Bounded Coset Penetration* (BCP) property if for every $\lambda \geq 1$ there is a constant $c(\lambda) > 0$ such that whenever w and w' are $(\lambda, 0)$-quasigeodesics without backtracking such that their initial endpoints are $w_- = w'_-$ and their terminal endpoints are w_+, w'_+ with $d_S(w_+, w'_+) \leq 1$, then the following hold:

1. If w penetrates a coset gH but w' does not, then the entering and exiting vertices of w in gH are at most c far from each other.

2. If w and w' both penetrate a coset gH, then their entering vertices are at most c far from each other, and similarly their exiting vertices are at most c far.

We are now ready to give the definiton of relative hyperbolicity attributed to Farb.

Definition 3.1.3 A group G is *hyperbolic relative to the subgroup H* if the coned off Cayley graph $\widehat{\Gamma}(G, H)$ is δ-hyperbolic and has the BCP property. If $\widehat{\Gamma}$ is δ-hyperbolic without the BCP property, we say that G is *weakly relatively hyperbolic.*

In the original formulation, weakly relatively hyperbolic was called relatively hyperbolic while the added condition of the BCP property made a group strongly relatively hyperbolic. However, the coned off Cayley graph can be δ-hyperbolic even in situations that feel *very* far from hyperbolic, and eventually the terminology was reformulated. One example where we see this is $\mathbb{Z} \oplus \mathbb{Z} = \langle a, b \,|\, [a, b] \rangle$, which is weakly hyperbolic relative to $\mathbb{Z} =< b >$. The coned off Cayley graph is pictured in Figure

??, and it is simple to verify that it is indeed δ-hyperbolic. To see that it does not satisfy the BCP property, consider the paths b^n and ab^n for increasingly large n.

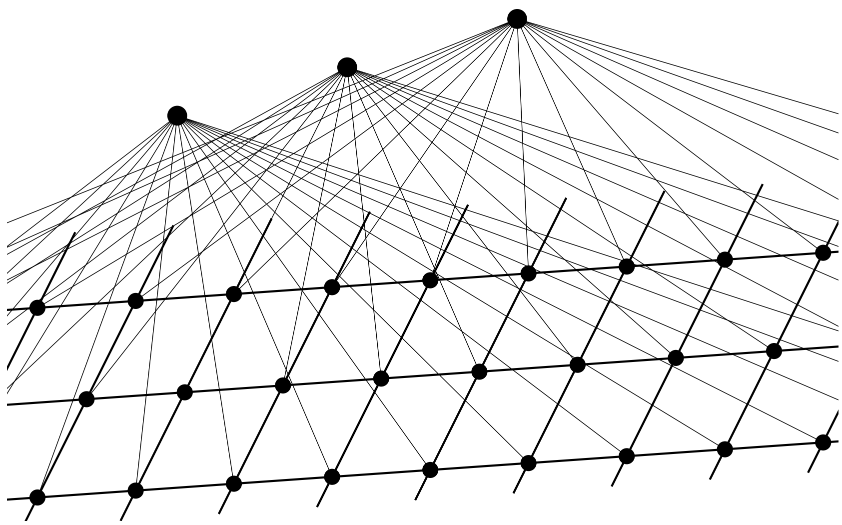

Figure 3.1: The coned off Cayley graph of $\mathbb{Z} \oplus \mathbb{Z}$ relative to $\mathbb{Z}$.

3.2 Fine Hyperbolic Graphs

When $\hat{\Gamma}$ is δ-hyperbolic, an equivalent property to the BCP property that we may look for instead is *fineness*.

Definition 3.2.1 A graph is *fine* if each of its edges is contained in only finitely many circuits of length n for every $n > 1$.

This property in turn is equivalent to the graph being *angularly locally finite*:

Definition 3.2.2 Given two edges $e_1 = (v, v_1), e_2 = (v, v_2)$ sharing one vertex, the angle θ between e_1 and e_2 is the length of the shortest path between v_1 and v_2 that does not include v. We define θ as being infinite if no such path exists.

A graph is *angularly locally finite* if for all angles θ, given an edge e, the set of edges which make an angle of θ with e is finite.

The following lemmas, along with the second half of the proof of the main proposition, come from the appendix of [4].

Lemma 3.2.1 *A graph is fine if and only if it is angularly locally finite.*

Proof: Suppose that a graph G is not angularly locally finite. Then there is some $\theta > 0$ and some edge e in G such that there are infinitely many edges making angles of θ with e. Then e is contained in infinitely many circuits of length $\theta + 2$. Also note that for two edges e_1, e_2 making angles of θ with e, the resulting circuit for e_1 does not contain e_2 and vice versa. This is by the definiton of a circuit as a simple loop and because e, e_1, and e_2 all have exactly one vertex in common. Thus there are infinitely many distinct such circuits, so G is not fine.

On the other hand, suppose that G is angularly locally finite. Let $n > 0$ and e be an edge of G. Any circuit of length n that contains e gives an edge e' adjacent to e making an angle of $n - 2$ with it. Thus there are only finitely many possibilities, so G is fine. □

Note that in the definitions of fine and angularly locally finite, we may replace "loops of length n" (or "angles of θ") with "loops of lenth $< n$" (angles $< \theta$).

Lemma 3.2.2 *If the coned-off Cayley graph $\widehat{\Gamma}$ satisfies the first criterion of BCP then, up to a change in constant, it satisfies the second criterion as well.*

Proof: Suppose the first criterion holds for $\lambda \geq 1$ with constant r. Let w_1, w_2 be two $(\lambda, 0)$-quasigeodesic paths without backtracking in $\widehat{\Gamma}$ hat penetrate the coset gH. Take w_1' to be the path obtained by cutting off w_1 after its entering vertex of gH. This is also a $(\lambda, 0)$-quasigeodesic.

Now take w_2' to be the path obtained first by cutting off w_2 after the vertex $v(gH)$ and then adding the edge from $v(gH)$ to the last vertex of w_1'. This is a $(\lambda', 0)$-quasigeodesic for some $\lambda' \geq \lambda$. So w_1' must also be a $(\lambda', 0)$-quasigeodesic. Thus, w_1' and w_2' meet the conditions of the first criterion, with w_2' penetrating gH and w_1' not.

So there is some constant $c(\lambda')$ such that the entering and exiting vertices of w_2' are at most c far from each other. But the entering vertex of w_2' is the entering vertex of w_2, and the exiting vertex of w_2' is the entering vertex of w_1, so we have shown that the entering vertices of w_1 and w_2 are at most c far from each other. Since λ' depends only on λ and is independent of choice of paths (we would use the exact same construction) then the constant c holds across all paths that meet the conditions of the second criterion. □

Proposition 3.2.1 *A δ-hyperbolic coned off Cayley graph has the BCP property if and only if it is fine.*

Proof: ($\Rightarrow$ by contrapositive) Suppose that $\widehat{\Gamma}$ is not fine. Then there is some edge e contained in infinitely many circuits of length n for some n. Since the only infinite-valence vertices are the cone vertices, then infinitely many of these circuits must penetrate some coset gH. From these circuits we can choose finite subpaths to obtain infinitely many paths w_i each of which have e as their first edge (thus all beginning at the same point) and penetrate the coset gH. Because they come from circuits, each path will be without backtracking.

We must show that for every $c > 0$ we can find two paths w_i, w_j such that (WLOG) their entering vertices are more than c far away from each other. Suppose not, i.e. the entering vertices of the w_i are all within some bounded distance from one another in the Cayley graph Γ. Since there must be infinitely many distinct entering vertices to give infinitely many paths penetrating gH, this leads to a bounded neighborhood in Γ containing infinitely many vertices, contradicting finite generation. Thus, $\widehat{\Gamma}$ cannot of the BCP property.

($\Leftarrow$ By contradiction) Suppose that $\widehat{\Gamma}$ is fine but does not have the BCP property and let $\lambda \geq 1$. Then for all $c > 0$ there are $(\lambda, 0)$-quasigeodesic paths w_1 and w_2 in $\widehat{\Gamma}$ such that

1. the paths start at the same point,
2. the paths end at most distance one from each other,
3. w_1 penetrates a coset gH with entering and exiting vertices further than c apart in Γ, and
4. w_2 does not penetrate the coset.

Since $\widehat{\Gamma}$ is δ-hyperbolic, by lemma 2.2.2, there is some constant $R(\delta, \lambda)$ such that any two paths satisfying 1 and 2 above stay at most R far from each other.

Now for any $\theta > 0$ there is a constant $r(\theta)$ such that the following holds: when $g, g' \in H$ and $d_S(g, g') > r$, the angle between the edges $(v(H), g)$ and $(v(H), g')$ in $\widehat{\Gamma}$ is greater than θ. Since for a fixed $g \in H$ there are infinitely many $g' \in H$ that will be further than r from it, if this were not true it would lead to infinitely many edges $(v(H), g')$ adjacent to the edge $(v(H), g)$ making an angle of less than θ with it, contradicting lemma 3.2.1.

So now we fix θ to be $64 \cdot \lambda \cdot R(\delta, \lambda) + 1$ and choose w_1 and w_2 that satisfy the four conditions above with constant $c = r(\theta)$. Since w_1 penetrates gH for more that r, the path passes through the vertex $v(gH)$ (while w_2 does not) and the angle of the adjacent edges at this vertex is greater than θ. Our goal is to contradict that this angle is $> \theta$.

First we parametrize the two paths by arc length to have $w_1 : [t_{-1}, t_1] \to \widehat{\Gamma}$ such that $w_1(0) = v(gH)$ and $w_2 : [0, t_2] \to \widehat{\Gamma}$. Let $t = \min(t_1, 10 \cdot \lambda \cdot R)$, so $w_1(t)$ occurs after the exiting vertex of w_1 from gH, and let $\sigma_1 = w_1|_{[0,t]}$.

Now choose t' so that $d(w_1(t), w_2(t'))$ (the distance in $\widehat{\Gamma}$) is minimized, and thus less than R. Let σ_2 be a geodesic from $w_1(t)$ to $w_2(t')$. Let $t'' = \min(t_2, t' + 20 \cdot \lambda \cdot R)$ and let $\sigma_3 = w_2|_{[t', t'']}$. Next, choose t''' so that $d(w_2(t''), w_1(t'''))$ is minimized (again, less than R), and let σ_4 be a geodesic from $w_2(t'')$ to $w_1(t''')$. And lastly, let $\sigma_5 = w_1|_{[t''', 0]}$ so that concatenating the σ_i in order gives us a loop. (See figure 3.2 for reference.)

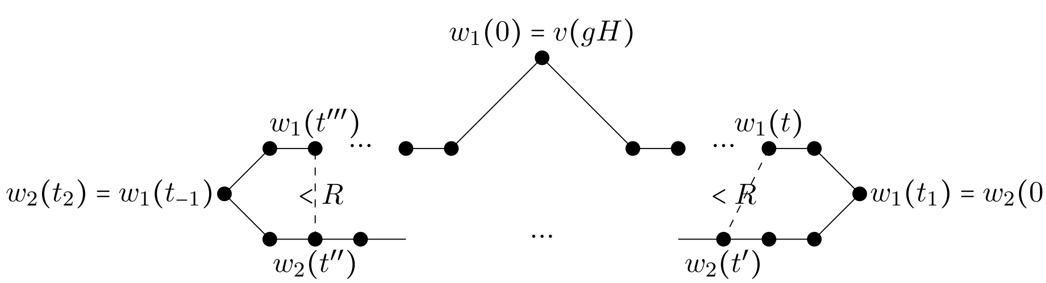

Figure 3.2: Two paths in $\widehat{\Gamma}$.

Since σ_3 lies entirely in w_2 it does not contain the vertex $v(gH)$. We can see that σ_2 also does not contain $v(gH)$: first suppose $t = 10 \cdot \lambda \cdot R$. Then $w_1(t)$ is at a distance of at least $10 \cdot R$ from $v(gH)$, and since σ_2 is shorter than R, it cannot contain $v(gH)$. Now suppose $t = t_1$. Then since the end points of w_1 and w_2 are at most one away from each other, σ_2 has length one and can be chosen as an edge from the Cayley graph Γ. Thus is does not contain $v(gH)$.

Now we want to see that σ_4 does not contain $v(gH)$ and that $w_1(t''')$ in fact occurs before (or at) the entering vertex of w_1 in gH. First suppose that $t'' = t' + 20 \cdot \lambda \cdot R$. Then the distance between $w_2(t')$ and $w_2(t'')$ is at least $20 \cdot R$. But if $w_1(t''')$ lies on the λ-quasigeodesic σ_1, we could construct a path starting at $w_2(t'')$ to $w_1(t''')$ then $w_1(t)$ and finally $w_2(t')$ which is shorter than $2 \cdot 10R$. So $w_2(t''')$ must occur before $v(gH)$. If $t'' = t_2$, then it works out similarly to the second case above.

Thus we have a loop that passes through v exactly once, so the angle at v, which is the distance between the entering and exiting vertices, both of which are included in the loop, is less than the length of the loop. This length in turn is less than $2(10 \cdot \lambda \cdot R + R + 20 \cdot \lambda \cdot R + R) \leq 2(10 \cdot \lambda \cdot R + \lambda \cdot R + 20 \cdot \lambda \cdot R + \lambda \cdot R) \leq 64 \cdot \lambda \cdot R = \theta$. Thus we have arrived at a contradiction to the angle θ being large. □

This proposition brings us around to another definition of relative hyperbolicity:

Definition 3.2.3 Let G be a group acting on a fine δ-hyperbolic graph with finite edge stabilizers and finitely many orbits of edges. Then if $\mathbb{P}$ is a set of representatives of conjugacy classes of infinite-valence vertex stabilizers, G is hyperbolic relative to $\mathbb{P}$.

So by its construction, when the coned-off Cayley graph is δ-hyperbolic and has the BCP property, then it meets the conditions of this definition.

3.3 Geometrically Finite Convergence Groups

This section presents the dynamical formulation of a relatively hyperbolic group from [2], which nicely generalizes the action of a Kleinian group on its limit set.

Definition 3.3.1 The action of a group G on a compact, metrizable space M with more than two points is a (non-elementary) *convergence group action* if the action of G on the space of distinct triples of points of M is properly discontinuous. A convergence action is *uniform* if the action of the distinct space of triples is also co-compact.

We may have elementary convergence group actions on spaces with two or fewer points as well: if M is empty, G must be finite, if M has one point, G must be countable, and if M has two points, G must be virtually cyclic.

A subgroup $P < G$ is a *parabolic subgroup* if it is infinite and contains no loxodromic element (i.e. an infinite-order element which fixes exactly two poitns). The unique fixed point of a parabolic subgroup P is called a *parabolic point*, and a parabolic point is *bounded* if $\text{Stab}_G(p)$ acts cocompactly on $M - \{p\}$. A point $\alpha \in M$ is a *conical limit point* if there is a sequence of group elements $\{g_i\}$ and distinct points β_0, β_1 such that $g_i\alpha \to \beta_0$, and for all $\gamma \in M - \{\alpha\}$, $g_i\gamma \to \beta_1$.

Definition 3.3.2 If $\mathbb{P}$ is a set of representatives of the conjugacy classes of maximal parabolic subgroups in G, then the action of the pair $(G, \mathbb{P})$ on M is *geometrically finite* if it is a convergence action and every point of M is either a conical limit point or a bounded parabolic point.

It is a theorem of Bowditch that G acts as a uniform convergence action on M if and only if every point is a conical limit point, i.e. there are no parabolic points. This is relevant to another theorem of his havin to do with hyperbolic groups, which leads directly into our next characterization of relative hyperbolicity.

Theorem 3.3.1 *If G acts as a uniform convergence group on a compact, metrizable space M, then G is hyperbolic and M is homeomorphic to ∂G.*

From this, we have the following result of Yaman in [11], which is a natural generalization of hyperbolicity to relative hyperbolicity.

Definition 3.3.3 If G acts as a geometrically finite convergence group on a compact, metrizable space M, then G is hyperbolic relative to $\mathbb{P}$, the parabolic stabilizers.

Now it is also known that if a group acts properly discontinuously on a proper δ-hyperbolic space X, the induced action on ∂X is a convergence group action [10]. Since ∂X is compact and metrizable (as stated in proposition 2.2.1), this brings us to the final definition of a relatively hyperbolic group:

Definition 3.3.4 Let G be a group that acts properly discontinuously on a proper δ-hyperbolic space X such that the induced convergence action on ∂X is geometrically finite. Then if $\mathbb{P}$ is a set of representatives of conjugacy classes of the maximal parabolic subgroups, $(G, \mathbb{P})$ is relatively hyperbolic.

This second definition is essentially a more specific version of the first, requiring M to be the boundary of a proper δ-hyperbolic space, but it is nice as it gives us the convergence action requirement for free.

Chapter 4

Examples of Relatively Hyperbolic Groups

In this chapter, we will take a look at the free group as a relatively hyperbolic group in two different ways. Rather than focusing on rigorously proving relative hyperbolicity, we aim more to develop some intuition around the ideas presented in the previous chapter.

4.1 F_2 as a Geometrically Finite Convergence Group

First we will look at F_2 as the fundamental group of a cusped, punctured torus. This arises as the quotient of $\mathbb{H}^2$ by the properly discontinuous action of $F_2 = \langle a, b\rangle$, where a and b are hyperbolic Moebius transformations (i.e. isometries that each fix exactly two boundary points) that have been chosen so their commutator $[a, b]$ is parabolic.

The large shaded square in figure 4.1 is a fundamental domain for the action; the generator a takes the left side to the right, while b takes the bottom side to the top. So when the square gets folded up, its four corners become the puncture, and we see this is a cusp by noting that these corners lay on the boundary at infinity. The commutator corresponds to the loop of the puncture, and so the commutator subgroup $\langle [a, b]\rangle$ corresponds to the cusp. We will see that F_2 is hyperbolic relative to $\mathbb{P}$, the set of conjugates of the commutator subgroup, a maximal parabolic subgroup, by observing that the action on $\partial\mathbb{H}^2 = S^1$ is geometrically finite.

The parabolic points are all the translates ot the four corners of the fundamental domain, which appear to be accumulation points in the figure (and in fact, the set of these points is dense in S^1). Heuristically, if we look at the checker-board-style tiling of $\mathbb{H}'^2$ by the fundamental domain, the parabolic points on the boundary are

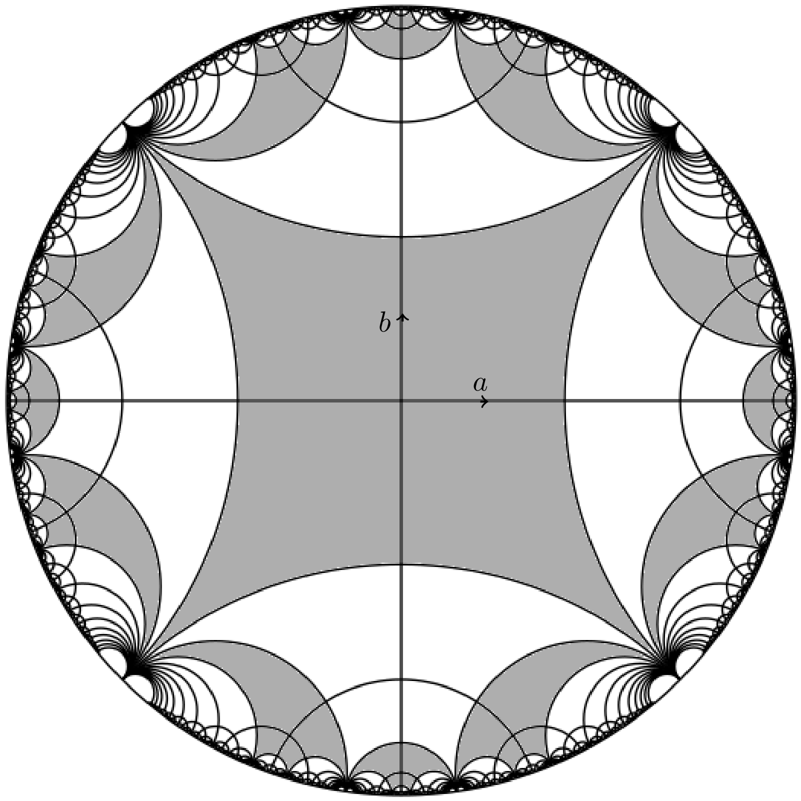

Figure 4.1: The Cayley graph of F_2 (see 2.1) overlayed on $\mathbb{H}^2$, tiled with a fundamental domain for the action of F_2 on the plane. Image generated with Curt McMullen's LIM (http://www.math.harvard.edu/~ctm/math101/www/programs/index.html

those that come from rays that eventually "stay in the same color." Note that for any of these points p, the quotient $S^1 - \{p\}$ by $Stab_p(G)$ is S^1, and since this is compact, each parabolic point is bounded (think about the bottom left corner of the fundamental domain, which is stabilized by $\langle [a,b] \rangle$).

Now what's left to see is that every other point on the boundary is a conical limit point. Going back to our heuristic, these other points can be viewed as rays that eventually pass through alternating colors in the checker board tiling. In our figure, we can also see these points as the ones corresponding to the boundary of the overlayed Cayley graph of F_2. We won't examine this part in detail, but note that the "north-south dynamics" of the hyperbolic isometries (e.g. the generators) show us that the endpoints of the hyperbolic axes of these isometries are conical limit points.

4.2 The Coned-off Cayley Graph of $(F_2, \langle a \rangle)$

We'll now look at $\langle a, b \rangle$ as being hyperbolic relative to the cyclic subgroup generated by a (it would work the same if instead we chose b). We construct the coned-off Cayley graph from figure 2.1 first by adding a vertex for every coset of $\langle a \rangle$ and then connecting these to every element in the cosets. This is easily visualized by placing a vertex "above" each of the horizonal axes in the graph, which correspond to group elements eventually ending in repeating a's and connecting it to each vertex of that axis.

So we need to check to see that the newly constructed graph has the bounded coset penetration property or that it is fine. In this case, the only circuits in the graph are the ones we created with the cone vertices, and the symmetry of the construction makes it very simple to check for fineness. In fact none of the original edges from the Cayley graph are contained in more than one circuit of a given length. Each of the new edges from the cone vertices is in only two circuits of a given length.

If we want to think about the bounded coset penetration property here, it is useful to compare this coned-off Cayley graph with that of $(\mathbb{Z} \oplus \mathbb{Z}, \mathbb{Z})$, shown in figure 3.1 which does not have the BCP property. With $\mathbb{Z} \oplus \mathbb{Z}$, we're "shortcutting" across flat parts in one direction when we cone off the Cayley graph. But there was flatness all over the graph, so while coning off created a δ-hyperbolic graph, the absence of the BCP property tells us that the original graph was too flat to begin with. In contrast, the flat copies of $\mathbb{Z}$ that we are coning off in the Cayley graph of F_2 are nicely separated from each other. This prevents us from having path that begin and end together with one path penetrating a coset arbirarily deeply.

Bibliography

[1] J. M. Alonso, T. Brady, D. Cooper, V. Ferlini, M. Lustig, M. Mihalik, M. Shapiro, and H. Short, *Notes on word hyperbolic groups*, in Group theory from a geometrical viewpoint (Trieste, 1990), World Sci. Publ., River Edge, NJ, 1991, pp. 3–63. Edited by Short.

[2] B. H. Bowditch, *Relatively hyperbolic groups*, Internat. J. Algebra Comput., 22 (2012), pp. 1250016, 66.

[3] François Dahmani, *Classifying spaces and boundaries for relatively hyperbolic groups*, vol. 86, 2003.

[4] ——, *Les groupes relativement hyperboliques et leurs bords*, 2003/13 (2003), pp. viii+87. Thèse, l'Université Louis Pasteur (Strasbourg I), Strasbourg, 2003.

[5] B. Farb, *Relatively hyperbolic groups*, Geom. Funct. Anal., 8 (1998), pp. 810–840.

[6] Étienne Ghys and Pierre de la Harpe, *Espaces métriques hyperboliques*, in Sur les groupes hyperboliques d'après Mikhael Gromov (Bern, 1988), vol. 83 of Progr. Math., Birkhäuser Boston, Boston, MA, 1990, pp. 27–45.

[7] M. Gromov, *Hyperbolic groups*, in Essays in group theory, vol. 8 of Math. Sci. Res. Inst. Publ., Springer, New York, 1987, pp. 75–263.

[8] G. Christopher Hruska, *Relative hyperbolicity and relative quasiconvexity for countable groups*, Algebr. Geom. Topol., 10 (2010), pp. 1807–1856.

[9] D. Osin, *Acylindrically hyperbolic groups*, Trans. Amer. Math. Soc., 368 (2016), pp. 851–888.

[10] Pekka Tukia, *Convergence groups and Gromov's metric hyperbolic spaces*, New Zealand J. Math., 23 (1994), pp. 157–187.

[11] Asli Yaman, *A topological characterisation of relatively hyperbolic groups*, J. Reine Angew. Math., 566 (2004), pp. 41–89.