BEHAVIOR OF GEODESIC RAYS IN SPACES WITH GEOMETRIC GROUP ACTIONS

BY DANIEL STALEY

A dissertation submitted to the

Graduate School—New Brunswick

Rutgers, The State University of New Jersey

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

Graduate Program in Mathematics

Written under the direction of

Steve Ferry

and approved by

____

____

____

____

New Brunswick, New Jersey

May, 2010

ABSTRACT OF THE DISSERTATION

Behavior of Geodesic Rays in Spaces with Geometric Group Actions

by Daniel Staley

Dissertation Director: Steve Ferry

This dissertation studies certain groups by studying spaces on which they act geometrically. These spaces are studied by examining the behavior of geodesic rays in these spaces, which gives geometric data about the space that can translate into algebraic data about the group.

First, we investigate the amenability of Thompson's group F by studying the geometry of its Cayley graph. We apply the uniformly finite homology of Block and Weinberger to subsets of this graph. Many large subsets of the Cayley graph are shown to be nonamenable by exhibiting certain arrangements of geodesic rays which we call "tree-like quasi-covers".

We then examine CAT(0) boundaries. If a group acts geometrically on two CAT(0) spaces X and Y, then one obtains a G-equivariant quasi-isometry from X to Y. One may look at the image of a geodesic ray in X, and look at its closure in ∂Y. We show that this "boundary image" can have the homeomorphism type of any compact, connected subset of Euclidean space.

Acknowledgements

I would like to thank every member of my defense committee. Feng Luo, for countless discussions, seminars, and lectures that taught me so much about hyperbolic geometry and geometric group theory. Xiaochun Rong, who gave me a wonderful introduction to Gromov-Hausdorff limits which fueled my interest in CAT(0) boundaries. Kim Ruane, for the energized and excited discussions we've had that always seem to end in new results.

I'd like to thank my advisor, Steve Ferry, for constantly providing me with insights and avenues of research. He was always there, patiently pointing me in new and interesting directions, even when I stubbornly took my time finding an area that called to me.

I'd like to thank my family. Mom, Dad, and Kevin, thanks for all the phone calls and the support. My trips home kept me energized when I came back.

Thanks to Simon Thomas, for helpful discussions and a wealth of knowledge on hyperbolic groups.

Thanks to Nan Li, for office discussions on spherical geometry which helped me finish my proofs (particularly, the proof of Lemma 3.6.5 is entirely due to Nan).

Finally, thanks to my wife, Binnaz. Without her love and support, I would never have finished my degree.

Table of Contents

Chapter 1

Introduction

Geometric group theory studies groups by studying metric spaces on which they act. Given a group G acting on a metric space X, we say the action is *proper* if for each compact $K \subset X$, $gK \cap K = \emptyset$ for all but finitely many $g \in G$. We say the action is *cocompact* if X/G is compact. The action is *by isometries* if $d(x, y) = d(gx, gy)$ for all $g \in G$ and $x, y \in X$.

We are interested in actions of the following type:

Definition 1.0.1. An action of a group G on a metric space X is *geometric* if it is proper, cocompact, and by isometries.

If G is finitely generated, then G always acts geometrically on its *Cayley graph*. The Cayley graph is the graph constructed as follows: Given a generating set S of G, the vertex set is simply the elements of G. We then connect two vertices g and g' by an edge if $g' = gs$ for some $s \in S$. Since left multiplication by a fixed element does not change this relation, we have that the action of G on its Cayley graph via left multiplication is a geometric action.

Another important notion in geometric group theory is that of *quasi-isometry*:

Definition 1.0.2. Given two metric spaces X and Y, a function $f\colon X \to Y$ is a *quasi-isometric embedding* if there exist $\lambda \geq 1, B \geq 0$ such that for all $x_1, x_2 \in X$,

$$\frac{1}{\lambda} d(x_1, x_2) - B \leq d\big(f(x_1), f(x_2)\big) \leq \lambda d(x_1, x_2) + B.$$

If f further has the property that for all $y \in Y$ there is an $x \in X$ with $d(f(x), y) \leq B$, then f is a *quasi-isometry*.

It is a basic result in geometric group theory that if there is a quasi-isometry from X to Y, then there is a quasi-isometry from Y to X. Two metric spaces are called *quasi-isometric* if there exists a quasi-isometry between them. A fundamental theorem of geometric group theory is the following theorem of Milnor and Svarc:

Theorem 1.0.3. *Suppose G acts geometrically on a metric space X. Then X is quasi-isometric to any Cayley graph of G.*

This theorem allows us to study a wide variety of metric spaces with geometric group actions, and obtain results about the groups that act upon them.

Given $x, y \in X$, a *geodesic segment connecting x and y* is a map p from any real interval $[a, b]$ into X such that $p(a) = x, p(b) = y$, and $d\big(p(t), p(t')\big) = |t - t'|$ for any $t, t' \in [a, b]$. For ease of notation, we will occasionally fail to distinguish between a geodesic segment and its image.

Definition 1.0.4. A metric space X is a *geodesic metric space* if for any two points $x, y \in X$, there is a geodesic segment connecting x and y.

We include the edges (as copies of the unit interval) in the Cayley graph of a group G and give it the path metric, which makes it a geodesic metric space.

A *geodesic ray* in X is a path $p\colon [0, \infty) \to X$ such that $d\left(p(t), p(t')\right) = |t - t'|$ for all $t, t' \in [0, \infty)$. Geodesic rays give important information about the large-scale behavior of a metric space. Since large-scale behavior is typically all that is captured by a quasi-isometry, this makes geodesic rays a useful tool for studying properties which are invariant under quasi-isometry.

Using Theorem 1.0.3, we will study two open questions about groups by using the behavior of geodesic rays in geodesic metric spaces with geometric group actions.

In Chapter 2, we study Thompson's group F. This is a particular group with numerous interesting properties. For several decades, the question has been open as to whether F is *amenable*. We will define amenability in chapter 2. We do not answer the question, but we show that certain subgraphs of the Cayley Graph of F are not amenable. These graphs are "large" in the sense that their union is the entire Cayley graph.

In order to show that these subgraphs are not amenable, we use *uniformly finite homology*, which was developed by Block and Weinberger [4] but based on ideas originated by Følner in 1955 [14]. We construct "Ponzi schemes" on these graphs, by "flowing" along geodesic rays. By choosing an appropriate set of starting points for these rays, we find that the graphs are sufficiently "tree-like" to admit a Ponzi scheme.

In Chapter 3, we study CAT(0) metric spaces. CAT(0) spaces are geodesic metric spaces which satisfy a certain notion of nonpositive curvature. Groups which act geometrically on these spaces are called CAT(0) groups, and have been extensively studied. A CAT(0) space has a well-defined boundary consisting of geodesic rays emitting from a chosen basepoint. However, this boundary is not an invariant of the group which acts on it. Croke and Kleiner [9] constructed the first example of two CAT(0) spaces, acted upon by the same group, whose boundaries are not homeomorphic.

It is natural to then ask what sort of equivalence one has between boundaries of spaces acted upon by the same group. Bestvina [3] has shown that all such boundaries are *shape equivalent*, which we will define in Chapter 3. It is an open question as to whether these boundaries are *CE-equivalent*, another notion which we will define in chapter 3.

We do not answer this question, but we do construct spaces with erratic behavior of geodesic rays, which highlights the complexity of the situation. The construction demonstrates that quasi-geodesics can have pathological behavior, even if we impose the condition that they are images of true geodesics under G-equivariant maps.

Chapter 2

Thompson's Group F and Uniformly Finite Homology

2.1 Background

In 1965, Richard Thompson introduced his group F. F has many interesting properties: it is finitely presented, has exponential growth, and contains subgroups isomorphic to $F \times F$, yet it is "almost simple" in the sense that every proper quotient of F is a quotient of $\mathbb{Z} \times \mathbb{Z}$. The question as to whether F is amenable was first posed in 1979. F is, in a sense, "on the edge of amenability", as it is not elementary amenable but does not contain a free subgroup on two generators (see Brin and Squier [7]). If F is not amenable, it provides a finitely-presented counterexample to the Von Neumann conjecture. Very few such examples are known (Ol'shanskii and Sapir provided the first in 2000 [21]).

In 1955, Følner provided a geometric criterion for the amenability of a group based on the existence of subsets of the Cayley graph that satisfy a "small boundary" condition [14]. This criterion holds for semigroups as well (one may find a proof in Namioka [20]), and allows one to extend the definition of "amenable" to graphs of bounded degree. In 1992, Block and Weinberger extended the definition to a broad class of metric spaces [4]. They defined the *uniformly finite homology groups* $H_n^{uf}(M)$ of a metric space M and proved that M is amenable if and only if $H_0^{uf}(M) \neq 0$.

In this chapter we use the results of Block and Weinberger to define the notion of a *tree-like quasi-cover*, and show that an amenable graph cannot possess a tree-like quasi-cover. We then apply this result to subgraphs of the Cayley graph of Thompson's group F.

Thompson's group F can be described as the group with the following infinite presentation:

$$\langle x_0, x_1, x_2, x_3, ... \mid x_j x_i = x_{i+1} x_j \text{ for } i > j \rangle$$

The main result of this chapter is the following:

Theorem. *Let k, l be nonnegative integers, with $l > 0$. Let Γ_k^l be the subgraph of the Cayley graph of F consisting of vertices that can be expressed in the form*

$$a_1 ... a_m b_1 ... b_n,$$

where $m \leq k$, $a_1, ..., a_m \in \{x_0, x_1, x_2, ...\}$, and $b_1, ..., b_n \in \{x_0, ..., x_l\}$. Let the edge set of Γ_k^l include all edges in the Cayley Graph of F that connect such vertices. Then Γ_k^l is not amenable.

The case $k = 1, l = 1$ was proved by D Savchuk in [23].

A corollary of this theorem is that all finitely-generated submonoids of the positive monoid of F are not amenable. It follows that if F is amenable these sets have measure zero.

For ease of notation, we will sometimes identify a graph with its vertex set. If the graph is the Cayley Graph of a group, we will occasionally identify vertexes with group elements.

Many of the proofs in this chapter appear in [25].

2.2 Thompson's Group F

From the presentation of F given in the previous section, we see that $x_{i+1} = x_0 x_i x_0^{-1}$ for $i \geq 1$. Thus this group is finitely generated by $\{x_0, x_1\}$. F is finitely presented as well (see [1] or [15] for a proof). However, it is still useful to consider the infinite generating set $\{x_0, x_1, x_2, ...\}$. We have the following definition:

Definition 2.2.1. The *positive monoid* of F is the submonoid of F consisting of elements that can be expressed as words in $\{x_0, x_1, x_2, ...\}$ without using inverses.

Any element of F can be expressed as an element of the positive monoid multiplied by the inverse of such an element. Elements of F have a normal form which is such a product. The reader is referred to [1] or [15] for proofs.

In [1], the group F is studied using *two-way forest diagrams.* We will make extensive use of these diagrams when studying the positive monoid in Section 4. We describe the two-way forest diagrams of the positive monoid here, referring the reader to [1] for the proofs.

Definition 2.2.2. A *binary forest* is an infinite sequence of binary trees, such that all but finitely many of the trees are trivial (ie have a single node):

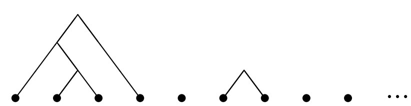

Each binary tree can be thought of as a collection of "carets". Removing the top caret of a nontrivial tree leaves two (possibly trivial) trees, which we call the *left child* and *right child* of the tree.

Definition 2.2.3. A *pointed forest* is a binary forest with a distinguished, or "pointed", tree:

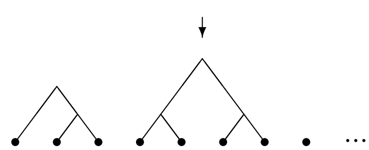

Henceforth, we will omit the ellipses and assume a forest diagram or pointed forest diagram has an infinite number of trivial trees continuing to the right.

Each element of the positive monoid of F can be identified with a pointed forest. The identity element is the pointed forest consisting only of trivial trees, with the the pointer on the leftmost tree.

Right multiplication by x_0 moves the pointer one tree to the right (Figure 2.1).

Figure 2.1: Multiplication by x_0

Right multiplication by x_1 adds a caret between the pointed tree and the tree immediately to its right, making a new tree whose left child is the pointed tree and whose right child is the tree to its right. This new tree becomes the pointed tree (Figure 2.2).

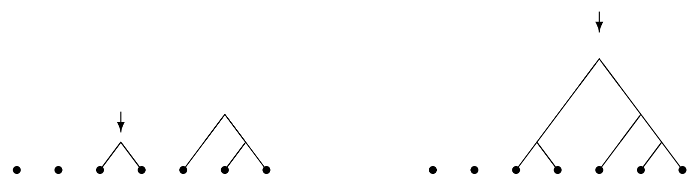

Figure 2.2: Multiplication by x_1

Since $x_i = x_0^{i-1}x_1x_0^{-(i-1)}$, we can see that right multiplication by x_i moves the pointer $i-1$ trees to the right, adds a caret, and then moves the pointer $i-1$ trees to the left again. This is equivalent to adding a caret between the trees $i-1$ and i steps to the right of the pointed tree (Figure 2.3).

Figure 2.3: Multiplication by x_3

Multiplication of pointed forests corresponds to "putting one on top of the other": If P and Q are pointed forests, then PQ is the forest obtained by using the trees of P as the nodes of Q with the pointed tree in P attaching to the leftmost node of Q (Figure 2.4).

The pointer is then placed above whatever tree was pointed in Q.

It has been a longstanding open question as to whether F is *amenable*:

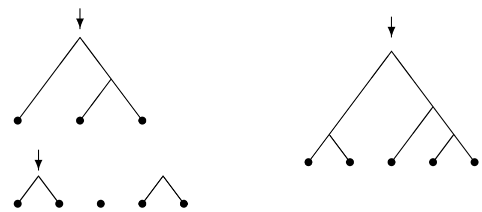

Figure 2.4: Multiplying the two pointed forests on the left yields the pointed forest on the right

Definition 2.2.4. A group G is called *amenable* if there exists a *right-invariant measure* on G–a function μ that assigns to each subset $A \subset G$ a value $0 \leq \mu(A) \leq 1$ such that:

1. $\mu(G) = 1$.

2. μ is finitely additive: If A and B are disjoint subsets of G, then $\mu(A) + \mu(B) = \mu(A \cup B)$.

3. μ is G-invariant: For any $g \in G$ and any $A \subset G$, $\mu(A) = \mu(Ag)$.

Many texts define amenability as the existence of a *left*-invariant measure, but the existence of left- and right- invariant measures is equivalent on a group. A useful result for determining amenability is *Følner's Criterion*, which uses the Cayley graph of G. Recall that the Cayley graph of G is the graph obtained by taking a generating set S and using G as the vertex set, connecting two vertices g and g' by an edge if $g' = gs$ for some $s \in S$.

Theorem (Følner's Criterion)**.** *A group G is amenable if and only if, for any $\epsilon > 0$, there exists a finite subset A of vertices in the Cayley graph of G such that*

$$\frac{\#\partial(A)}{\#A} < \epsilon,$$

where $\#A$ is the number of vertices in A, and $\#\partial(A)$ is the number of edges connecting vertices of A to vertices outside A.

Følner's criterion can be applied to any graph of finite valence. In particular, we say such a graph is amenable if Følner's criterion holds for that graph. This allows us to state the following proposition:

Proposition 2.2.5. *Let* Γ *be the subgraph of the Cayley graph of Thompson's group* F *(using the* x_0, x_1 *generating set) consisting of vertices in the positive monoid of* F *and all edges between such vertices. Then* Γ *is amenable if and only if* F *is amenable.*

For a proof see Savchuk [23]. Roughly, using the normal form alluded to above, any finite subset of F can be left-translated into the positive monoid. The result then follows from Følner's criterion.

We finish this section with a proposition about geodesic rays in the Cayley graph of F which we will need later:

Proposition 2.2.6. *Let* Γ *be as above, and let* v *be a vertex of* Γ*. Let* g *be either* x_0 *or* x_1*. Define* $c\colon [0, \infty) \to \Gamma$ *to be the path such that for* $n \in \mathbb{N} \cup \{0\}$*,* $c(n) = vg^n$*, and* c *maps each interval* $[n, n+1]$ *isometrically to the edge connecting* vg^n *to* vg^{n+1}*. Then* c *is a geodesic ray.*

Proof. It suffices to show that $d(v, vx_0^n) = d(v, vx_1^n) = n$ in Γ. Since left multiplication induces an isometry of the Cayley graph, it is sufficient to show that both x_0^n and x_1^n have distance n from the identity. This is evident from their pointed forest diagrams, as x_1^n has n carets and traversing an edge can add at most one caret, while x_0^n has the pointer on the $(n+1)^{\text{th}}$ tree, and traversing each edge can cause the pointer to be at most one additional tree to the right. □

2.3 Uniformly Finite Homology

This section describes the uniformly finite homology of Block and Weinberger defined in [4]. We will always be considering a graph Γ of bounded degree, though many of their results apply to a much broader class of metric spaces.

Definition 2.3.1. Let Γ be a graph of bounded degree with vertex set V. A *uniformly finite 1-chain with integer coefficients* on Γ is a formal infinite sum $\sum a_{x,y}(x, y)$, where the (x, y) are ordered pairs of vertices of Γ, $a_{x,y} \in \mathbb{Z}$, such that the following properties are satisfied:

1. There exists $K > 0$ such that $|a_{x,y}| < K$ for all vertices x and y.

2. There exists $R > 0$ such that $a_{x,y} = 0$ whenever $d(x, y) > R$.

Notice that condition (2) guarantees that for any fixed $x \in V$, the set of pairs (x, y) such that $a_{x,y} \neq 0$ is finite. This allows us to make the following definition:

Definition 2.3.2. A uniformly finite 1-chain is a *Ponzi scheme* if, for all $x \in \Gamma$, we have $\sum_{v\in\Gamma} a_{v,x} - \sum_{v\in\Gamma} a_{x,v} > 0$.

We now state the main result of [4] that we will use in this chapter:

Theorem 2.3.3. *Let Γ be a graph of bounded degree. A Ponzi scheme exists on Γ if and only if Γ is not amenable.*

We will use a rephrased version of this theorem for the case of our graphs:

Definition 2.3.4. Let Γ be a graph of bounded degree with vertex set V. A *Ponzi flow on* Γ will mean a function $\Phi \colon V \times V \to \mathbb{Z}$ with the following properties:

(i) $\Phi(v, w) = 0$ if there is no edge from v to w in Γ,.

(ii) $\Phi(v, w) = -\Phi(w, v)$ for all $v, w \in V$.

(iii) The function Φ is bounded.

(iv) For each $v \in V$, $\sum_{w\in V} \Phi(w, v) > 0$.

Note that the sum in condition *(iv)* is guaranteed to be finite by condition *(i)*. A Ponzi flow is almost exactly a Ponzi scheme in different language, with the exception that all "pairs" must be of distance 1. Our next proposition shows that this difference is unimportant:

Proposition 2.3.5. *Let Γ be a graph of bounded degree. There exists a Ponzi scheme on Γ if and only if there exists a Ponzi flow on Γ.*

Proof. The "if" direction is trivial: Given a Ponzi flow, we simply define our formal sum to be $\sum \Phi(x, y)(x, y)$. This will be a uniformly finite 1-chain with integer coefficients, as condition (1) is implied by *(iii)*, and condition (2) is implied by *(i)*. This 1-chain will be a Ponzi scheme by conditions *(ii)* and *(iv)*.

To see the "only if" direction, we start with a Ponzi scheme $\sum a_{x,y}(x,y)$, and alter it so that $a_{x,y} = 0$ if $d(x,y) > 1$. For each $a_{x,y}$ such that $d(x,y) = n > 1$, let $x = v_0, v_1, ..., v_{n-1}, v_n = y$ be a sequence of vertices forming a path of minimal length from x to y, and add $S_{x,y} = \sum_{i=0}^{n-1} a_{x,y}(v_i, v_{i+1}) - a_{x,y}(x,y)$ to the Ponzi scheme coefficientwise.

Each adjacent pair (v,w) has no more than $\binom{d^{R+1}}{2}$ pairs of vertices within distance R, where d is the bound on degree and R is the radial bound from condition 2). That means that each edge will be part of no more than $\binom{d^{R+1}}{2}$ of the paths we constructed above. Thus, the coefficientwise addition is a finite sum for each coefficient, and so we obtain a well-defined formal sum $\sum a'_{x,y}(x,y)$.

This formal sum will still satisfy the inequality of Definition 2.3.2, since each vertex in $S_{x,y}$ appears exactly twice, once contributing $a_{x,y}$ and once contributing $-a_{x,y}$, leaving the sum in Definition 2.3.2 unchanged. For (x,y) with $d(x,y) > 1$, the coefficients of (x,y) in the original Ponzi Scheme and in $S_{x,y}$ will cancel, leaving the new coefficient 0. Thus we have canceled the coefficients $a_{x,y}$ whose vertices are not adjacent while only affecting other coefficients with adjacent vertices.

We have $a'_{x,y} = 0$ if $d(x,y) > 1$, since these coefficients have been canceled. Furthermore, we have that each $a'_{x,y}$ is bounded by $K(\binom{d^{R+1}}{2} + 1)$, thus condition 1) of 2.3.1 still holds (with a different bound), so we have a uniformly finite 1-chain with $R = 1$. Since we did not change the sums in 2.3.2, $\sum a'_{x,y}(x,y)$ is still a Ponzi scheme. Now we simply define $\Phi(x,y) = a_{x,y} - a_{y,x}$, and condition (i) is now true by our construction, condition (ii) is clear, (iii) is implied by 1), and (iv) is implied by the inequality of 2.3.2, thus Φ is a Ponzi flow on Γ. $\square$

A quantitative treatment of Proposition 2.3.5 and of Ponzi flows can be found in Benjamini, Lyons, Perez and Schramm [2].

If a Ponzi flow exists on a Cayley graph, we then have that there can be no right-invariant measure on the group since the group cannot be amenable by Proposition 2.3.5 and Theorem 2.3.3. We supply here a direct proof that existence of a Ponzi flow implies no right-invariant measure exists, for any graph on which an appropriate notion of "right invariant" can be defined:

Definition 2.3.6. By a *labeled directed graph* we shall mean a directed graph Γ of bounded degree, each of whose edges are labeled by elements from a finite set $\{g_1, ..., g_n\}$, such that each vertex of Γ has at most one incoming edge and one outgoing edge with each label.

The motivating example is when Γ is a subgraph of a Cayley graph of a group generated by $\{g_1, ..., g_n\}$, but the argument here holds for all labeled directed graphs.

Definition 2.3.7. Let Γ be a labeled directed graph. Suppose S is a subset of vertices of Γ. For $1 \leq i \leq n$, we say S is *g_i-translatable* if each vertex in S has an outgoing edge labeled g_i. In such a case we denote by Sg_i the set of vertices of Γ with an incoming edge labeled g_i whose opposite vertex lies in S.

In the case where Γ is a subgraph of a Cayley graph, Sg_i is just the right-translate of the elements of S under the group multiplication. We will abuse notation slightly in the case of one-element sets, so that if v has an outgoing edge labeled by g_i, we will call the vertex on the other side of the edge vg_i.

For vertex sets S and S', we will say $S' = Sg_i^{-1}$ if $S'g_i = S$, and similarly for vertices v and v', we will say $v' = vg_i^{-1}$ if $v'g_i = v$.

Definition 2.3.8. Let μ be a finitely additive measure on the vertex set of Γ, such that $\mu(\Gamma) = 1$. Then we say μ is *right-invariant* if for each g_i and each g_i-translatable subset $S \subseteq \Gamma$, $\mu(S) = \mu(Sg_i)$.

If Γ is the full Cayley graph of a finitely generated group or semigroup, then this definition of right-invariant measure coincides with the standard one, since in this case every set of vertices is g_i-translatable for every i.

Theorem 2.3.9. *Suppose Γ is any labeled directed graph, and has a Ponzi flow Φ. Then there is no finitely additive, right-invariant measure on Γ.*

Proof. Let K be the bound on Φ. Consider the set of symbols $S = \{g_1, g_2, ..., g_n\} \cup \{g_1^{-1}, g_2^{-1}, ..., g_n^{-1}\} \cup \{h\}$. h has no meaning here except as a placeholder symbol.

For each vertex v of Γ, we define two functions $p_v, q_v \colon S \to \mathbb{N} \cup \{0\}$, representing the positive and negative parts of the flow Φ. For each $s \neq h$ in S, if vs does not exist or if $\Phi(v, vs) = 0$, then we set $p_v(s) = q_v(s) = 0$. Otherwise, if $\Phi(v, vs) > 0$, set $p_v(s) = \Phi(v, vs)$ and $q_v(s) = 0$, and if $\Phi(v, vs) < 0$, set $q_v(s) = -\Phi(v, vs)$ and $p_v(s) = 0$. This defines p_v and q_v on $S \setminus \{h\}$. Set $p_v(h)$ and $q_v(h)$ to be the values such that

$$\sum_{s \in S} p_v(s) = 2nK, \qquad \sum_{s \in S} q_v(s) = 2nK + 1.$$

Note that since p_v and q_v are defined on $2n$ elements besides h, and each of these takes a value at most K, we have that $p_v(h)$ and $q_v(h)$ are indeed nonnegative integers. Also note that since Φ is a Ponzi flow, by Definition 2.3.4 (ii) and (iv) we have

$$\sum_{g \neq h \in S} q_v(g) > \sum_{g \neq h \in S} p_v(g),$$

and thus

$$q_v(h) \leq p_v(h) \tag{2.1}$$

for all v.

For any function $p \colon S \to \mathbb{N} \cup \{0\}$ whose sum over all $s \in S$ is $2nK$, we define a set A_p of vertices of Γ. A_p consists of all vertices v for which $p_v = p$. Note that A_p may be empty. Since each vertex has a unique list associated to it, we have that the A_p are pairwise disjoint, and $\Gamma = \bigcup_p A_p$.

Similarly, for any function $q \colon S \to \mathbb{N} \cup \{0\}$ whose sum over all $s \in S$ is $2nK + 1$, we define Z_q to be the set of vertices v such that $q_v = q$. Again, the Z_q are pairwise disjoint, and $\Gamma = \bigcup_q Z_q$.

Suppose Γ has a right-invariant measure μ. Since there are a finitely many nonnegative functions p or q whose sum over S is $2nK$ or $2nK + 1$, we have

$$\sum_p \mu(A_p) = \mu\left(\bigcup_p A_p\right) = \mu(\Gamma) = 1$$

$$\sum_q \mu(Z_q) = \mu\left(\bigcup_q Z_q\right) = \mu(\Gamma) = 1$$

For any $g \in S$ and integer j with $1 \leq j \leq K$, define $B_{j,g}$ to be the set of vertices v such that $p_v(g) \geq j$. Let $\Lambda(j,g)$ be the set of all functions $p\colon S \to \mathbb{N} \cup \{0\}$ which sum to $2nK$ with $p(g) \geq j$. We then have that $B_{j,g} = \bigcup_{p \in \Lambda(j,g)} A_p$. Each $\Lambda(j,g)$ is finite, thus

$$\sum_{j,g} \mu(B_{j,g}) = \sum_{j,g} \mu \left(\bigcup_{p \in \Lambda(j,g)} A_p \right) = \sum_{j,g} \sum_{p \in \Lambda(j,g)} \mu(A_p) \tag{2.2}$$

Note that, for a fixed $g \in S$, a function p is contained in exactly $p(g)$ of the sets $\Lambda(j,g)$. This means that each function p is contained in precisely $2nK$ of the sets $\Lambda(j,g)$, and so each set A_p appears exactly $2nK$ times in the double sum on the right-hand side of Equation 2.2.

This allows us to explicitly calculate the sum of the measures of the $B_{j,g}$:

$$\sum_{g,j} \mu(B_{j,g}) = 2nK \sum_{p} \mu(A_p) = 2nK.$$

For $g \in S$ and $1 \leq j \leq K$, define $Y_{j,g}$ to be the set of vertices v such that $q_v(g) \geq j$. A similar argument as above shows that

$$\sum_{g,j} \mu(Y_{j,g}) = (2nK+1) \sum_{q} \mu(Z_q) = 2nK + 1.$$

For each g_i and $1 \leq j \leq K$, B_{j,g_i} consists of vertices with an outgoing edge labeled g_i, thus B_{j,g_i} is g_i-translatable. Indeed, we have that $Y_{j,g_i} = B_{j,g_i} g_i$. Similarly, we have that $Y_{j,g_i^{-1}} = B_{j,g_i^{-1}} g_i^{-1}$. By the right-invariance of μ, we have then have that $\mu(Y_{j,g}) = \mu(B_{j,g})$ for every $g \neq h$ in S. By Equation 2.1, we have that $Y_{j,h} \subset B_{j,h}$, and thus $\mu(Y_{j,h}) \leq \mu(B_{j,h})$. Thus,

$$2nK + 1 = \sum_{g,j} \mu(Y_{j,g}) \leq \sum_{g,j} \mu(B_{j,g}) = 2nK,$$

a contradiction. Thus, no right-invariant measure can exist on Γ.

$\square$

We remark that in [4], Ponzi Schemes are considered with real coefficients as well as with integer coefficients. The same results hold in the real setting, including Theorem

2.3.9, as long as we stipulate that for each Ponzi flow Φ there is some $\epsilon > 0$ such that either $\Phi(x, y) = 0$ or $|\Phi(x, y)| > \epsilon$ for all vertices $x, y \in \Gamma$.

Corollary 2.3.10. *Let Γ be any labeled directed graph. If Γ is amenable but contains a nonamenable subgraph P, then for any right-invariant measure μ on Γ, $\mu(P) = 0$.*

Proof. If $\mu(P) > 0$, then we can define a measure μ' on P by setting $\mu'(A) = \frac{\mu(A)}{\mu(P)}$ for $A \subset P$. Since $\mu(P)$ is constant, and any g_i-translatable subset of P is g_i-translatable in Γ, μ' will inherit the properties of right invariance and finite additivity from μ. We have $\mu'(P) = \frac{\mu(P)}{\mu(P)} = 1$, thus μ' is a finitely additive, right-invariant measure on P. But since P is nonamenable it has a Ponzi flow, and then 2.3.9 says no such μ' can exist, yielding a contradiction. □

We close this section by constructing another method for detecting amenability, based on uniformly finite homology.

Definition 2.3.11. Let Γ be a graph of bounded degree, and let A be a set of geodesic rays in Γ. We say A is a *tree-like quasi-cover* of Γ if:

1. There exists an $N > 0$ such that for every vertex v of Γ, no more than N of the rays in A contain v in their image.

2. There exists a $B > 0$ such that for every vertex v of Γ, there is a $c \in A$ such that $d(c(0), v) < B$.

Theorem 2.3.12. *Let Γ be a graph of bounded degree, and suppose Γ has a tree-like quasi-cover A. Then Γ is not amenable.*

Proof. For each $c \in A$, we define a function $\Psi_c\colon\ \Gamma \times \Gamma \to \mathbb{Z}$. Let d be the bound on the degree of Γ. For each $t \in \mathbb{R}^+$ with $c(t)$ a vertex, define $\Psi_c\big(c(t+1), c(t)\big) = B^d + 1$ and $\Psi\big(c(t), c(t+1)\big) = -(B^d + 1)$. Let $\Psi_c(x, y) = 0$ for all other pairs of vertices. Define

$$\Psi\colon\ \Gamma \times \Gamma \to \mathbb{Z} = \sum_{c \in A} \Psi_c$$

The sum on the right is well-defined, since condition (1) in Definition 2.3.11 guarantees that at most N of the rays traverse each edge, thus all but finitely many terms in the sum are 0 for each pair of vertices.

Call a vertex $v \in \Gamma$ *terminal* if $v = c(t)$ for some $c \in A, t \in [0, 1)$. For each vertex u which is not terminal, there is a terminal vertex v such that $d(u, v) < B + 1$, by Definition 2.3.11 (2). Choose a minimal-length path in Γ from v to u, calling the vertices $v = x_0, x_1, x_2, ..., x_{n-1}, x_n = u$. Define a function $\Phi_u\colon \Gamma \times \Gamma \to \mathbb{Z}$ by $\Phi_u(x_i, x_{i+1}) = 1$, $\Phi_u(x_{i+1}, x_i) = -1$ for $0 \leq i < n$, and $\Phi(x, y) = 0$ for all other vertex pairs.

If u is a terminal vertex, we define $\Phi_u(x, y) = 0$ for all $x, y \in \Gamma$. Now, define

$$\Phi\colon \Gamma \times \Gamma \to \mathbb{Z} = \Psi + \sum_{u \in \Gamma} \Phi_u.$$

Again, the sum on the right is well-defined, as each edge is part of no more than $\binom{(B+1)^d+1}{2}$ paths used in the Φ_u, and thus all but finitely many terms are 0 for each vertex pair. We claim that Φ is a Ponzi flow.

Properties (i) and (ii) in Definition 2.3.4 clearly hold, as Φ is a sum of functions with these properties. For each pair $x, y \in \Gamma$, we have that $\Psi(x, y)$ is bounded by $N(B^d + 1)$, and at most $\binom{(B+1)^d+1}{2}$ of the Φ_u have $\Phi_u(x, y) = \pm 1$, with the rest taking the value zero. Thus Φ is bounded and property (iii) holds.

For property (iv), note that for each $a \in \Gamma$,

$$\sum_{b \in \Gamma} \Phi(b, a) = \sum_{c \in A} \sum_{b \in \Gamma} \Psi_c(b, a) + \sum_{u \in \Gamma} \sum_{b \in \Gamma} \Phi_u(b, a)$$

If a is non-terminal, all the sub-sums in the first summation are zero, and all the sub-sums in the second summation are zero except when $a = u$, in which case $\sum_{b \in \Gamma} \Phi_u(b, a) = 1$. Thus property (iv) holds for these vertices.

If a is terminal, then the only non-zero sub-sums in the first summation occur when c is a geodesic for which a is terminal, in which case they are $B^d + 1$. The terms in the right sum can only be negative if u is within distance B of a, in which case they are -1. Since there are at most B^d such vertices, property (iv) still holds, and so Φ is a Ponzi Flow and Γ is not amenable.

□

2.4 Large Nonamenable Subgraphs of F

In this section we will prove Theorem 2.1.

We begin by characterizing the two-way forest diagrams of Γ_k^l. Given any binary tree T on n nodes, we define $s(T)$ to be the forest obtained by removing all the carets along the left edge of T (Figure 2.5).

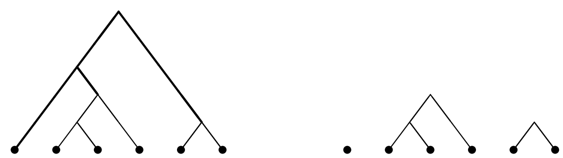

Figure 2.5: Applying s to a tree removes the left carets as shown

We extend the definition of s to apply to forests, as well as single trees, by applying s separately to each nontrivial tree in the forest. We will define the *complexity* of a tree or forest to be the minimum number of applications of s required to turn it into a forest of only trivial trees.

Note that applying s to a tree T gives a forest whose rightmost tree is the right child of T, and the remainder of the forest is s applied to the left child of T. This gives the following:

Proposition 2.4.1. *The complexity of a tree is the maximum of the complexity of its left child and one more than the complexity of its right child.*

We record here two basic properties of complexity and the function s:

Proposition 2.4.2. *Let T be a tree on n nodes, and let R be an n-tree forest. Denote by RT the tree obtained by attaching the roots of R to the nodes of T. If T has complexity j, then $s^j(RT)$ consists only of carets in R, ie every caret from T is removed by s^j. (Figure 2.6.)*

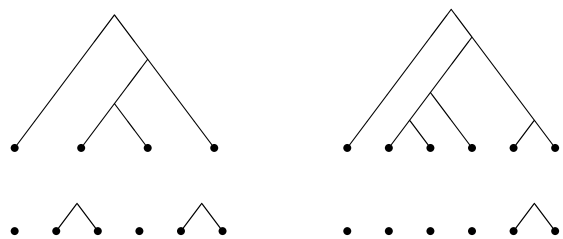

Figure 2.6: A 4-node tree T of complexity 2(top left) and a 4-tree forest R(bottom left) multiply to give RT (top right), and each caret of T is removed in $s^2(RT)$(bottom right)

Proof. This is easy to see, as we can determine whether a caret is removed by s^j by examining its relationship with those above it. Namely, when we examine the unique path from a caret to the root of the tree, it consists of moves to the right and moves to the left. An application of s removes all carets whose path consists only of moves to the right. Further, any caret's path to the root hits the left edge at some point, and consists only of moves to the right afterwards. After s is applied, the path is truncated, starting from the move that reaches the left edge (which is a move to the left). Thus each new path from a remaining caret to the root of its new tree is left with one less move to the left after applying s. So s^j removes all carets whose paths contain $j-1$ or fewer moves to the left. Since this property is unchanged in the carets of T whether or not it sits on R, the effect of s^j is the same on carets of T, ie it removes them all. $\square$

Proposition 2.4.3. *A pointed forest diagram consisting of a single nontrivial tree T of complexity j in the leftmost position, with the pointer on that tree, can be expressed as word in $x_1, ..., x_j$.*

Proof. We will proceed by induction on the number of carets in the tree T. Suppose the statement is true for all trees with n or fewer carets, and let T be a tree with $n+1$ carets and complexity j. Then the left child of T has no more than n carets and complexity no more than j by Proposition 2.4.1. Thus by the inductive hypothesis the left child can be constructed as a word w in $x_1, ..., x_j$. The right child of T has no more than n carets and complexity no more than $j-1$ by Proposition 2.4.1, thus can be constructed as a word v in $x_1, ..., x_{j-1}$. We can construct the desired pointed forest as $wx_0vx_0^{-1}x_1$, since this will construct the left child, move the pointer to the right, construct the right child, move the pointer back to the left child, and finally construct the top caret. However, since $x_i = x_0x_{i-1}x_0^{-1}$, by inserting $x_0^{-1}x_0$ between each letter of v we can rewrite $x_0vx_0^{-1}$ as a word in $x_2, ..., x_j$. Thus, the word $wx_0vx_0^{-1}x_1$ can be rewritten as a word in $x_1, ..., x_j$, and the proposition is proved.

□

For a positive integer j, we define a function ϕ_j from pointed forests to forests in the following way: Apply s^j to the pointed tree and every tree to its left. For each positive integer $q < j$, apply s^{j-q} to the tree that is q trees to the right of the pointed tree. That is, apply s^{j-1} to the tree to the immediate right of the pointed tree, s^{j-2} to next tree to the right, etc.

For the proof of the main theorem we will use the following lemma. Recall that Γ_k^l is the subgraph of the Cayley graph of F consisting of vertices that can be expressed in the form $a_1...a_mb_1...b_n$, with $m \leq k$, $a_1, ..., a_m \in \{x_0, x_1, x_2, ...\}$, and $b_1, ..., b_n \in \{x_0, ..., x_l\}$.

Lemma 2.4.4. *A pointed forest P lies in Γ_k^l if and only if $\phi_l(P)$ has k or fewer carets.*

Proof. Let $P \in \Gamma_k^l$. First suppose that $k = 0$. In this case P can be expressed as a word v in $\{x_0, ..., x_l\}$, and the proposition says it is annihilated by ϕ_l, ie $\phi_l(P)$ consists only of trivial trees. We proceed by induction on the length of v. If $\phi_l(v)$ consists only of trivial trees, then so does $\phi_l(vx_0)$, since $\phi_l(vx_0)$ is a subforest of $\phi_l(v)$ (each tree has s applied to it either the same number of times or one more time, since the pointer has simply moved one tree to the right).

For $0 < i \leq l$, multiplying by x_i adds a caret to the right of the tree $i-1$ trees from the pointer, combining it with the next tree to make a new tree. Since the left and right children of this new tree were $i-1$ and i trees to the right of the caret, respectively, by induction their respective complexities are no more than $(l-(i-1))$ and $(l-i)$. Thus by Proposition 2.4.1 the new tree has complexity no more than $(l-i+1)$. Since this new tree is $i-1$ trees to the right of the pointer, it is still annihilated by ϕ_l. The trees to the left of the new caret are unchanged, and the trees to the right of the caret have each been brought 1 tree closer to the pointer since two of the intervening trees have been merged. Thus ϕ_l applies s an additional time to each of these trees. This means that they will still be annihilated by ϕ_l, and so the new pointed forest is still turned into a trivial forest by ϕ_l.

The above argument shows that $\phi_l(v)$ is trivial if $v \in \Gamma_0^l$. Now let $P = wv$, where $w = a_1...a_m$ with $a_i \in \{x_0, x_1, ...\}$ and $m \leq k$ as in Theorem 2.1. Then wv attaches the trees of w to the nodes of v. Thus all the carets added in each tree of v are still removed by ϕ_l by Proposition 2.4.2, since s is applied the same number of times to each tree. Thus $\phi_l(wv)$ has at most the same number of carets as w, ie k or fewer. This proves the "only if" direction of the Lemma.

To prove the reverse direction, suppose that P is a pointed forest such that $\phi_l(P)$ has k or fewer carets. We can then create w as above to put these carets in place without moving the pointer (the generator x_i creates a caret on the i^{th} tree without moving the pointer).

Consider the element $w^{-1}P$. This is the pointed forest obtained by taking the trees in P that remain after applying ϕ_l, and replacing them with trivial trees (Figure 2.7).

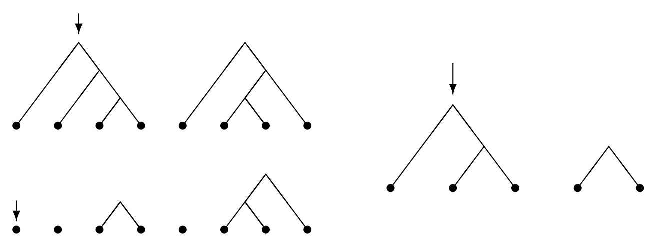

Figure 2.7: For $l = 2$, if P is the forest in the top left then w is $\phi_2(P)$ with the pointer on the first tree (bottom left), and $w^{-1}P$ is shown on the right.

The resulting pointed forest is then annihilated by ϕ_l by Proposition 2.4.2, and so each tree under or to the left of the pointer has complexity at most l. Thus, we may construct these trees as words in $x_1, ..., x_l$ using Proposition 2.4.3 and inserting x_0 between each word. This will result in building the first tree, moving the pointer to the right, building the next tree, etc. Further, the tree that is j trees to the right of the pointer has complexity at most $l - j$, and so Proposition 2.4.3 says we can construct it as $x_0^j u x_0^{-j}$, where u is a word in $x_1, ..., x_{l-j}$. As above, we then insert $x_0^{-j} x_0^j$ between each letter of u, which allows us to rewrite it as a word in $x_{j+1}, ..., x_l$. Repeating this for each j and appending these words in increasing order constructs all trees to the right of the pointer. This completes the construction of $w^{-1}P$ as a word in $x_0, ..., x_l$, which we will call v. Thus, $P = ww^{-1}P = wv$, and the proof is complete.

□

We are now ready to prove the main theorem, which will occupy the remainder of this section.

Proof of Theorem 2.1. Let $P \in \Gamma_k^l$. Note that applying ϕ_l to P affects at most l trees under or to the right of the pointer. Thus, by Lemma 2.4.4 there are at most $k + l$ nontrivial trees under or to the right of the pointer in P, otherwise, $\phi_l(P)$ would have more than k nontrivial trees and thus certainly have more than k carets.

For each $P \in \Gamma_k^l$, define a path $c_P \colon [0, \infty) \to \Gamma_k^l$ such that for $n \in \mathbb{N} \cup \{0\}$, $c_P(n) = Px_1^n$, and c_P maps the interval $[n, n+1]$ isometrically to the edge connecting Px_1^n and Px_1^{n+1}. Then by Proposition 2.2.6, c_P is a geodesic ray. Define

$$A = \left\{ c_P \mid P \in \Gamma_k^l, Px_1^{-1} \notin \Gamma_k^l \right\}$$

We claim A is a tree-like quasi-cover. The c_P have disjoint images, since if $Px_1^n = P'$, with $P \in \Gamma_k^l$, then $P'x_1^{-1} = Px_1^{n-1} \in \Gamma_k^l$, since Γ_k^l is clearly closed under right multiplication by x_1. Thus, each vertex of Γ_k^l lies in the image of at most one $c \in A$, and so property (1) in Definition 2.3.11 holds.

Any pointed forest P whose pointed tree is trivial has the property that Px_1^{-1} is not even in the positive monoid of F, and thus $c_P \in A$. But since every pointed forest in Γ_k^l has at most $k + l$ nontrivial trees under or to the right of the pointer, every pointed forest is within distance $k+l$ of a pointed forest with trivial pointed tree, which is $c_P(0)$ for some $c_P \in A$. Thus, condition (2) in Definition 2.3.11 holds as well.

Thus A is a tree-like quasi-cover of Γ_k^l, and so by Theorem 2.3.12, Γ_k^l is not amenable.

□

We close with some immediate corollaries:

Corollary 2.4.5. *If F is amenable, then for any right-invariant measure μ, $\mu(\Gamma_k^l) = 0$.*

Proof. By Theorem 2.1, Γ_k^l has a Ponzi flow, and thus by Proposition 2.3.10 it always has measure zero. □

Corollary 2.4.6. *If F is amenable, then for any right-invariant measure μ, and any finitely generated submonoid M of the positive monoid, $\mu(M) = 0$.*

Proof. Letting $p_1, ..., p_n$ be generators of M, express each as a word in the $x_0, x_1, x_2, ...$ generating set. Let L be the maximum index of the x_i used to express the p_j (or let $L = 1$, if this maximum is 0); then M is a subset of the monoid generated by $x_0, x_1, ..., x_L$. But this monoid is exactly Γ_0^L, which by the previous corollary has measure zero. Thus, $\mu(M) = 0$. □

Chapter 3

Boundary Behavior of CAT(0) Geodesics under G-equivariant Maps

3.1 Background

3.1.1 CAT(0) Groups

We begin this chapter with a brief overview of CAT(0) spaces and the groups that act upon them. The material in this subsection, as well as proofs of the theorems, can be found in Bridson-Haefliger [6].

Throughout this chapter we will use the convention $[a, b]$ to denote the standard real interval, unless a and b are already defined as points in a geodesic metric space. In this case $[a, b]$ will denote a geodesic segment connecting a and b.

Suppose X is a geodesic metric space. Given three points $a, b, c \in X$, a *geodesic triangle* $\triangle\, abc$ is the union of three geodesic segments $[a, b] \cup [b, c] \cup [a, c]$. Given such a geodesic triangle, we can find three *comparison points* $\overline{a}, \overline{b}, \overline{c}$ in the Euclidean plane $\mathbb{E}^2$, such that $d(a, b) = d(\overline{a}, \overline{b}), d(b, c) = d(\overline{b}, \overline{c})$, and $d(a, c) = d(\overline{a}, \overline{c})$. We call the triangle $\triangle\, \overline{a}\overline{b}\overline{c}$ the *comparison triangle* for $\triangle\, abc$.

Let x lie on one of the segments of $\triangle\, abc$, say $[a, b]$. Since the segments $[a, b]$ and $[\overline{a}, \overline{b}]$ are isometric, we can find a corresponding point $\overline{x}$ on the comparison triangle such that $d(x, a) = d(\overline{x}, \overline{a})$ and $d(x, b) = d(\overline{x}, \overline{b})$. We can find comparison points in the same way for any point on any segment of $\triangle\, abc$.

Definition 3.1.1. A geodesic triangle $\triangle\, abc$ satisfies the *CAT(0) inequality* if, for every pair of points $x, y \in \triangle\, abc$ with comparison points $\overline{x}, \overline{y} \in \triangle\, \overline{a}\overline{b}\overline{c}$, $d(x, y) \leq d(\overline{x}, \overline{y})$.

In a sense, a triangle satisfies the CAT(0) inequality if it is “no fatter” than its comparison triangle in Euclidean space.

Definition 3.1.2. A geodesic metric space X is a *CAT(0) metric space* if every geodesic triangle in X satisfies the CAT(0) inequality.

Recall that a metric space is *proper* if closed balls are compact. There has been much study of CAT(0) metric spaces which are not proper, but this chapter is only concerned with proper spaces. Henceforth, a CAT(0) space will be assumed to be proper and finite dimensional unless otherwise stated. We will also assume that our metric spaces are complete unless otherwise stated.

Examples of CAT(0) metric spaces include Euclidean space $\mathbb{E}^n$, hyperbolic space $\mathbb{H}^n$, and any simply connected Riemannian manifold with nonpositive sectional curvature. This definition also encompasses a wide class of spaces which are not manifolds. One way to construct such spaces is by using the following theorems, which can be found on page 167 of [6]:

Theorem 3.1.3. *Let X and Y be CAT(0) metric spaces. Then $X \times Y$ is a CAT(0) metric space.*

Theorem 3.1.4. *Let X and Y be CAT(0) metric spaces. Suppose $A \subset X$ and $B \subset Y$ are closed and convex, and A and B are isometric. Then the space $X \cup_{A,B} Y$, obtained by identifying A and B, is CAT(0).*

For example, metric trees, wedge sums of CAT(0) spaces, and spaces obtained by gluing together geodesics in proper CAT(0) spaces are all CAT(0), as are products of such spaces.

We note here some important properties of CAT(0) spaces:

Proposition 3.1.5. *CAT(0) metric spaces are unique geodesic spaces, and these unique geodesics vary continuously with their endpoints.*

Proof. Let c and c' be two geodesic segments connecting points x and y in a CAT(0) metric space. Consider the geodesic triangle consisting of c, c', and a degenerate third side consisting only of x. Then the comparison triangle is also degenerate, and so for $w \in c$, $w' \in c'$ with $d(x, w) = d(x, w')$, the CAT(0) inequality tells us that $d(w, w') = 0$. Thus, c and c' consist of the same points and are the same geodesic segment.

To see that geodesics vary continuously with their endpoints, let $\{x_i\}$ and $\{y_i\}$ be sequences of points converging to x and y, respectively. Consider the geodesic segments $[x_i, y_i]$. Let m be the midpoint of $[x, y]$, let m_i be the midpoint of $[x_i, y_i]$, and let m_i' be the midpoint of $[x, y_i]$. The CAT(0) inequality applied to the triangle $\triangle\, xy_iy$ and $\triangle\, xy_ix_i$ shows that both $d(m, m_i')$ and $d(m_i', m_i)$ approach zero, thus the m_i converge to m.

Since midpoints converge to midpoints, by expressing any point on the interval as the limit of a sequence of dyadic rationals and passing to a diagonal sequence we can see that the $[x_i, y_i]$ converge pointwise to $[x, y]$. $\square$

Proposition 3.1.6. *CAT(0) spaces are contractible.*

Proof. Let X be a CAT(0) space, and fix any point $p \in X$. For any $x \in X$ and $t \in \mathbb{R}$ with $0 \leq t < d(x, p)$, let x_t be the unique point on $[p, x]$ such that $d(x_t, p) = t$. Let $f\colon [0, 1) \to [0, \infty)$ be any order-preserving homeomorphism. Define a homotopy $h\colon X \times I \to X$ so that:

$$h(x,t) = \begin{cases} x & \text{if } t = 0 \text{ or } d(x,p) < f(1-t) \\ x_{f(1-t)} & \text{otherwise} \end{cases}$$

$h(x, 0)$ is then the identity on X, and $h(x, 1) = p$. Each point is simply moved uniformly along the geodesic segment connecting it to p, and since geodesic segments vary continuously with their endpoints by Proposition 3.1.5, this map is continuous. Thus, h is a strong deformation retraction from X to p. $\square$

We say a metric space has *nonpositive curvature* if it is locally CAT(0), that is, each point has a neighborhood which is CAT(0). The proof of the following theorem can be found in Chapter II.4 of [6]:

Theorem 3.1.7 (Cartan-Hadamard Theorem)**.** *The universal cover of a metric space of nonpositive curvature is a CAT(0) metric space.*

In particular, any metric space of nonpositive curvature is aspherical. This theorem gives us another method of constructing CAT(0) metric spaces; namely, taking the universal covers of spaces of nonpositive curvature.

There is a final basic property of CAT(0) spaces which we will use. Suppose X is a complete CAT(0) space, and $A \subset X$ is nonempty, closed and convex.

Proposition 3.1.8. *For every $p \in X$, there is a unique $a \in A$ that minimizes $d(a, p)$.*

Proof. The existence of such a point is guaranteed by the fact that A is closed and X is proper. To see uniqueness, suppose that for $a \neq a' \in A$, $d(a, p) = d(a', p)$. Consider the geodesic triangle $\triangle paa'$. Since A is convex, the segment $[a, a']$ is contained in A. For any point q in the interior of this segment, $d(q, p) < d(a, p) = d(a', p)$ by the CAT(0) inequality. Thus a and a' do not minimize the distance to p. $\square$

The following proposition is proved on page 177 of [6]:

Proposition 3.1.9. *Let $\rho_A \colon X \to A$ be the function which takes each $x \in X$ to the unique point on A of minimal distance to X. Then ρ_A is continuous and does not increase distances, ie $d(x, y) \geq d\big(\rho_A(x), \rho_A(y)\big)$.*

We will now define the boundary of a CAT(0) metric space.

Definition 3.1.10. Let c and c' be geodesic rays in X. We say c and c' are *asymptotic* if there exists $B \in \mathbb{R}$ such that $d\big(c(t), c'(t)\big) < B$ for all $t \geq 0$.

We define a relation $\sim$ on geodesic rays by $c \sim c'$ iff c and c' are asymptotic. It follows immediately from the definition and the axioms of a metric space that $\sim$ is an equivalence relation.

Definition 3.1.11. The *boundary* ∂X of a CAT(0) space X is the set of equivalence classes of geodesic rays in X under the equivalence relation $\sim$.

We will topologize this set shortly, but first we give an important proposition:

Proposition 3.1.12. *For any point $p \in X$ and any equivalence class $\xi \in \partial X$, there is a unique geodesic ray c in the class ξ such that $c(0) = p$.*

Proof. The proof of existence is slightly technical and we omit it (the reader is referred to page 261 of [6]). Essentially, given a geodesic ray $c' \in \xi$, the pointwise limit of the geodesic segments $[p, c'(t)]$ is a geodesic ray emitting from p and asymptotic to c'.

To see uniqueness, let c and c' be geodesic rays emitting from p, and suppose for some $T > 0$ we have $c(T) \neq c'(T)$. Let $x = c(T), y = c'(T)$, and consider the comparison triangle $\triangle\,\overline{xyp}$ in Euclidean space. For $t > T$, let $\overline{x}'$ be the point on $\overrightarrow{\overline{px}}$ with $d(\overline{x}', p) = t$, and let $\overline{y}'$ be the point on $\overrightarrow{\overline{py}}$ with $d(\overline{y}', p) = t$.

Then if $d\big(c(t), c'(t)\big) < d(\overline{x}', \overline{y}')$, the comparison triangle for $\triangle\, c(t)c'(t)p$ would have a strictly smaller angle at $\overline{p}$ than $\triangle\,\overline{x}'\overline{py}'$. The CAT(0) inequality would then imply that $d(x, y) < d(\overline{x}, \overline{y})$, which is impossible since the two distances were constructed to be equal. Thus, $c(t)$ and $c'(t)$ must grow apart at least as quickly as the Euclidean rays $\overrightarrow{\overline{px}}$ and $\overrightarrow{\overline{py}}$, and thus cannot be asymptotic. □

Proposition 3.1.12 implies that we can choose any basepoint $p \in X$, and points in ∂X correspond exactly to geodesic rays emitting from p.

We now define a topology on $X \cup \partial X$. Consider ∂X as geodesic rays emitting from a chosen basepoint p. For $c \in \partial X$ and $R, \epsilon > 0$, define

$$U(c, R, \epsilon) = \big\{c' \in \partial X \mid d\big(c(R), c'(R)\big) < \epsilon\big\} \cup$$

$$\big\{c'(t) \mid c' \in \partial X, d\big(c(R), c'(R)\big) < \epsilon, t > R\big\}$$

The basis for the topology on $X \cup \partial X$ will be a basis for the topology of X, together with the sets $U(c, R, \epsilon)$. This is called the *cone topology*, and its restriction to ∂X gives us a topology on the boundary. It turns out that the cone topology is independent of choice of p; the proof is somewhat technical and so we again omit it, referring readers to [6].

It is an unfortunate fact that, if a group G acts geometrically on a CAT(0) space X, then the boundary ∂X is not an invariant of G. In [9], Croke and Kleiner construct a group which acts geometrically on two CAT(0) spaces X and Y such that ∂X is not homeomorphic to ∂Y. This naturally leads to the question of what sort of equivalence must exist between ∂X and ∂Y in this situation.

3.1.2 Shape Equivalence and CE-Equivalence

We begin this subsection with a definition:

Definition 3.1.13. Two compact, metrizable, finite dimensional spaces X and Y are *shape equivalent* if they can be written as inverse limits of finite polyhedra:

$$X = \varprojlim \Big(X_1 \leftarrow X_2 \leftarrow X_3 \leftarrow \dots \Big)$$
$$Y = \varprojlim \Big(Y_1 \leftarrow Y_2 \leftarrow Y_3 \leftarrow \dots \Big)$$

and there exist maps $f_i \colon X_i \to Y_i$, $g_i \colon Y_i \to X_{i-1}$, such that the diagram

$$\begin{tikzcd}
X_1 \arrow[d, "f_1"'] & X_2 \arrow[l] \arrow[d, "f_2"'] & X_3 \arrow[l] \arrow[d, "f_3"'] & \dots \arrow[l] & X_i \arrow[l] \arrow[d, "f_i"'] & \dots \arrow[l] \\
Y_1 & Y_2 \arrow[l] \arrow[ul, "g_2"] & Y_3 \arrow[l] \arrow[ul, "g_3"] & \dots \arrow[l] \arrow[ul, "g_4"] & Y_i \arrow[l] \arrow[ul, "g_i"] & \dots \arrow[l] \arrow[ul, "g_{i+1}"]
\end{tikzcd}$$

commutes up to homotopy.

An overview of shape equivalence and shape theory can be found in [11] or [16]. The following theorem is due to Bestvina [3]:

Theorem 3.1.14. *Suppose a group G acts geometrically on two CAT(0) spaces X and Y. Then ∂X and ∂Y are shape equivalent.*

Definition 3.1.15. A map $f \colon X \to Y$ between compact, metrizable, finite dimensional spaces is *cell-like* if f is surjective, and for each $y \in Y$, $f^{-1}(y)$ has the shape of a point (ie, is shape equivalent to a point).

Definition 3.1.16. Two compact, finite dimensional, metrizable spaces X and Y are *CE-equivalent* if there exists a compact, finite dimensional, metrizable space Z and cell-like maps $f \colon Z \to X$ and $g \colon Z \to Y$.

A full definition and development of CE-equivalence can be found in Ferry [12]. Ferry proves that CE-equivalence is indeed an equivalence relation and is implied by homotopy equivalence. The reverse implication does not hold, since for example the topologist's sine curve can be written as an inverse limit of contractible spaces, and thus is shape equivalent to a point. It is proved in Sher [24] that shape equivalence is implied by CE-equivalence for finite-dimensional spaces. However, the reverse implication does not hold; in [13] Ferry describes two spaces which are shape equivalent but not CE-equivalent.

In [3], Bestvina asked whether, given a group G acting geometrically on CAT(0) spaces X and Y, one always has that ∂X and ∂Y are CE-equivalent. This question remains open. We do not answer this question here, but we address a related question.

Given a CAT(0) space X and a subset $Q \subset X$, we can view Q as a subset of $X \cup \partial X$ under the cone topology. Letting $\overline{Q}$ denote the closure of Q, we call $\overline{Q} \cap \partial X$ the *boundary limit* of Q.

Given G, X, and Y as above, one can use Theorem 1.0.3 to obtain a G-equivariant quasi-isometry from X to Y. One may take the image of a geodesic ray c in X under this quasi-isometry, and examine its boundary limit in Y. We call this set the *boundary image* of c. The question then arises: What are the possible homeomorphism types of the boundary image?

It was observed by K Ruane [22] that such a boundary limit need not be a point. This is in contrast with the case of δ-hyperbolic spaces, where quasi-isometries of the space extend to homeomorphisms of the boundary. It was hoped that all boundary images would be contractible, or have the shape of a point.

It is not hard to see that such boundary images must be compact, connected, and finite dimensional. The main result of this chapter is that these are the only restrictions on the homeomorphism types of such boundary limits:

Theorem 3.1.17. *Let Z be any connected, compact subset of n-dimensional Euclidean space for any n. Then there exists a group G which acts geometrically on two CAT(0) spaces X and Y, a G-equivariant quasi-isometry $f : X \to Y$, and a geodesic c in X, such that the boundary limit of $f(c)$ is homeomorphic to Z.*

In fact in our construction X and Y will be the same space, albeit with different actions of G. Our construction of X, G, and f depends only on n, thus all $Z \subset \mathbb{R}^n$ arise as boundary images obtained from one particular quasi-isometry.

The construction in this chapter is heavily influenced by the constructions of Croke and Kleiner in [9] and its generalizations by C. Mooney in [18]. In fact, one may view the space X as a higher-dimensional analogue of the "blocks" used in their constructions.

For a point p in a metric space and $r > 0$, we will denote by $B_r(p)$ the open ball of radius r centered at p, and for a set Q we will use $N_r(Q)$ to mean the r-neighborhood of the set Q, i.e., $\bigcup_{p \in Q} B_r(p)$. If c is a path in a space, we will occasionally fail to distinguish between c and its image.

3.2 Constructing X, G, and f

We begin by constructing our CAT(0) space X, and the group G which will act on it geometrically. Throughout the chapter we will assume a fixed integer n.

Consider $(n+1)$-dimensional Euclidean space. Consider the group of isometries on this space generated by translations by $e_0, ..., e_n$, the standard basis vectors. The quotient of Euclidean space by this action is T^{n+1}, the standard $(n+1)$-torus. Take n disjoint copies of T^{n+1}, and identify the images of the subspaces spanned by $e_1, ..., e_n$ in each torus. Call the resulting space $\bar{X}$, and the identified image of the origin $\bar{p}$. Denote its universal cover by X. Choose a lift $p \in X$ of $\bar{p}$. The reader may easily verify the following properties:

Proposition 3.2.1. *1. $\bar{X}$ is a space of nonpositive curvature.*

2. $\pi_1(\bar{X}) = F_n \times \mathbb{Z}^n$, where F_n denotes the free group on n generators.

3. X is a CAT(0) space, and is isometric to $T \times \mathbb{R}^n$, where T is a regular tree of degree $2n$.

4. p can be chosen so that $p = (q, \overrightarrow{0}) \in T \times \mathbb{R}^n = X$, where q is a vertex of T.

We then have that the group $G = F_n \times \mathbb{Z}^n$ acts geometrically on the CAT(0) space X. For each of the n tori in $\bar{X}$, the inverse image in X is a disjoint union of Euclidean $(n+1)$-spaces, exactly one of which contains p. Call these spaces $E_1, ..., E_n \subset X$. Denote by $g_1, ..., g_n$ the generators of the F_n factor of G. We may choose these generators so that the subgroup $\langle g_i \rangle \times \mathbb{Z}^n$, acts geometrically on E_i, via translations by $e_0, ..., e_n$, such that g_i translates by e_0 and the standard generators of $\mathbb{Z}^n$ translate by $e_1, ..., e_n$. Indeed, this allows us to assume that the last n coordinates of each E_i correspond exactly to the $\mathbb{R}^n$ factor of X. We will use the symbol $\cdot$ for this group action.

Define a group automorphism $\Phi\colon G \to G$ in the following way: Choosing generators $e_1, ..., e_n$ of $\mathbb{Z}^n$ in the standard way, let $\Phi\big((1, e_i)\big) = (1, e_i)$. Let $\Phi\big((g_i, \overrightarrow{0})\big) = (g_i, e_i)$. The reader may easily check that Φ is an automorphism. This allows us to define a new action of G on X, $\circ$, by $g \circ x = \Phi(g) \cdot x$.

Proposition 3.2.2. *The action $\circ$ is geometric.*

Proof. Since Φ is an automorphism, $\circ$ has the same orbits as $\cdot$, and since both are free actions, it follows that $\circ$ is proper and cocompact. Since each transformation $\Phi(g) \cdot G$ is an isometry, each g acts via $\circ$ by an isometry as well. Thus, $\circ$ is a geometric action. $\square$

For each $i = 1, ..., n$, we define a linear transformation $f_i\colon E_i \to E_i$, which sends e_0 to $e_0 + e_i$ and is the identity on $e_1, ..., e_n$. Note that on each E_i, f_i is G-equivariant in the sense that $f_i(g \cdot x) = g \circ f_i(x)$ for all $g \in \langle g_i \rangle \times \mathbb{Z}^n, x \in E_i$.

Proposition 3.2.3.

1. *X is a union of flats of the form $g \cdot E_i$, with $g \in F_n \times \{\overrightarrow{0}\} \subset G$.*

2. *For $i \neq j$ and all $x \in E_i \cap E_j$, $f_i(x) = f_j(x) = x$.*

3. *There is a unique function $f\colon X \to X$ such that $f\,|_{E_i} = f_i$ for each $i = 1, ..., n$, and such that f is G-equivariant in the sense that $f(g \cdot x) = g \circ f(x)$ for all $g \in G, x \in X$.*

Proof. X is of the form $T \times \mathbb{R}^n$, and so we can write p in the form $(p_T, \overrightarrow{0})$, where p_T is a vertex of T. Since the F_n factor of G acts on T in the standard way, any vertex in T can be written as $g \cdot p_T$ for some $g \in F_n$. The edges coming out of p_T are all contained in one of the E_i, and represent either the e_0 or $-e_0$ direction. Thus, any point q on any edge can be expressed as $g \cdot v$ with $v = (a_0, 0, 0, ..., 0) \in E_i$. Thus, any point $(q, (a_1, ..., a_n)) \in T \times \mathbb{R}^n$ can be written as $g \cdot v'$ with $v' = (a_0, a_1, ..., a_n) \in E_i$, proving (1).

The intersection of E_i and E_j for $i \neq j$ is the connected component of the inverse image of the common n-torus in $\bar{X}$ which contains p. This is isometric to Euclidean n-space, and can be described as the convex hull of the orbit of p under the action of the $\mathbb{Z}^n$ factor of G. This is the space spanned by $e_1, ..., e_n$ in the coordinates of both E_i and E_j, and both f_i and f_j are the identity map on this space, which proves (2).

To see (3), we can simply define $f(g \cdot x) = g \circ f_i(x)$ for $x \in E_i$. (1) and (2) guarantee that this function is well-defined on all of X. (1), together with the equivariance condition, guarantees the uniqueness of f.

□

Defining f as the unique function from Proposition 3.2.3 (3), we record here some of its properties.

Proposition 3.2.4. *For all $(q, v) \in X$, $f\big((q, v)\big) = (q, w)$ for some $w \in \mathbb{R}^n$, that is, f is the identity on the T factor of X.*

Proof. The proposition follows from the definition of f, since the T coordinate of $x \in E_i$ depends only on the e_0 coordinate of x in E_i, which is unchanged by f_i. For points not in any of the E_i, since $f(g \cdot x) = g \circ f_i(x)$, we need only check that $\circ$ and $\cdot$ act identically on the T factor of X. Since $g \circ x = \Phi(g) \cdot x$ and Φ doesn't affect the F_n factor of G, the result follows as the subgroup $\{1\} \times \mathbb{Z}^n \subset G$ acts trivially on the T factor of X. □

We will put coordinates on the flats $g \cdot E_i$, for $g \in F_n \times \{\overrightarrow{0}\}$, however, $g \cdot E_i$ and $g' \cdot E_i$ may be the same flat for distinct g, g'. Since E_i is the convex hull of the orbit of the basepoint under the subgroup $\langle g_i \rangle \times \mathbb{Z}^n$ for some generator g_i, we have $g \cdot E_i = g' \cdot E_i$ if and only if $g' = g g_i^n$ for some $n \in \mathbb{Z}$. Thus, for any such flat, we can always choose a g of minimal length so that our flat is of the form $g \cdot E_i$, and we put coordinates on such a flat by translating the coordinates of E_i via the isometry induced by this minimal g.

Proposition 3.2.5. *For $g \in G$ and $i \in \{1, ..., n\}$, $f(g \cdot E_i) = g \cdot E_i$, and $f\mid_{g \cdot E_i}$ is an affine transformation of $g \cdot E_i$.*

Proof. Choose g so that the origin in $g \cdot E_i$ is $g \cdot p$. Let $A \colon E_i \to E_i$ be defined by $A(x) = g^{-1} \cdot (g \circ x)$. Since $g \circ x = f(g \cdot x)$, it suffices to prove that A is affine.

Since $\circ$ and $\cdot$ have the same effect on the T component of X, we immediately see that A is the identity on the e_0 component of E_i. Let u be the vector given by $g \circ p - g \cdot p$, taken in the coordinates of $g \cdot E_i$. Then for $x \in E_i$, we have $g \circ x = g \cdot f_i(x) + u$ since the coordinates of $g \cdot E_i$ are just translated coordinates of E_i. Letting A_u be translation by $-u$, we then have that $A_u(A(x)) = f_i(x)$.

Thus, after composing A with a translation we have that A is linear, and so A and therefore $f\ |_{g \cdot E_i}$ is an affine transformation.

□

The remainder of the chapter is devoted to proving the following, which in light of Propositions 3.2.2 and 3.2.3(3) implies the main theorem:

Theorem 3.2.6. *Let X and f be defined as above. Given any compact, connected set $Z \subset \mathbb{R}^{n-1}$, there is a geodesic ray $c : [0, \infty) \to X$ such that the boundary limit of $f(c)$ is homeomorphic to Z.*

3.3 Geodesic images under f

This section will use the notation defined in section 3.2.

Consider a geodesic ray $c \colon [0, \infty)$ in the $2n$-regular tree T, with the property $c(0) = p_T$. Since edges in T have length 1, $c(t)$ is a vertex precisely when t is an integer. So, under the standard action of F_n on T, $c(1) = g_{i_1}(p)$ for some generator $g_{i_1} \in F_n$, $c(2) = g_{i_1} g_{i_2}(p)$ for some second generator g_{i_2}, and in general, we have

$$c(k) = w_k(p),$$

where $w_k \in F_n$ is represented by a word of length k in the generators and their inverses. For the remainder of the section, we will assume that c is chosen such that w_k is a *positive* word in the g_i, that is, can be written using only the generators and never their inverses. w_k is a proper initial substring of w_{k+1}, and so there is a unique infinite sequence of generators such that every w_k is an initial subsequence. We define a sequence of integers $\{I_k\}$ such that this infinite sequence is $g_{I_1}, g_{I_2}, g_{I_3}, \ldots$.

For the remainder of the chapter, we will consider T to be a subset of X by identifying it with $T \times \{\overrightarrow{0}\}$. Since $c \subset T$, c lies in the subspace $c \times \mathbb{R}^n$.

Proposition 3.3.1. *$f(c) \subset c \times \mathbb{R}^n$.*

Proof. This is an easy consequence of Proposition 3.2.4. $\square$

The space $c \times \mathbb{R}^n$ is a Euclidean half-space of dimension $n+1$, and is a closed, convex subset of X. We put coordinates on this half space by identifying the point $(c(t), a_1, ..., a_n)$ with $(t, a_1, ..., a_n) \in \mathbb{R}^{n+1}$.

Proposition 3.3.2. *For each integer $k \geq 0$, the space $c \mid_{[k,k+1]} \times \mathbb{R}^n$ lies in exactly one of the flats $g \cdot E_i$. In this space, the coordinates on $c \times \mathbb{R}^n$ and on $g \cdot E_i$ differ by a constant vector v, which is an integer multiple of e_0.*

Proof. The segment $c \mid_{[k,k+1]}$ is an edge of the tree T, and each edge of T is contained in exactly one flat $g \cdot E_i$. But since the flat containing a point depends only on the value of the T factor of X, we have that $c \mid_{[k,k+1]} \times \mathbb{R}^n$ is contained in exactly one $g \cdot E_i$.

The last n coordinates in $c \times \mathbb{R}^n$ are the coordinates from the $\mathbb{R}^n$ factor of X. But $g \cdot E_i$ has coordinates translated from E_i via an element of $F_n \times \{\overrightarrow{0}\}$, which doesn't change the $\mathbb{R}^n$ factor. Since the last n coordinates of E_i are just the coordinates from the $\mathbb{R}^n$ factor of X, we see that the last n coordinates of $g \cdot E_i$ and $c \times \mathbb{R}^n$ agree on their intersection.

If x is a point in $c \mid_{[k,k+1]} \times \mathbb{R}^n$, then the e_0 coordinate in both $c \times \mathbb{R}^n$ and $g \cdot E_i$ differs by a constant from $d\left(\pi_T(x), c(k)\right)$, where π_T is projection onto the T factor. Thus the two coordinates differ by a constant from each other. Furthermore, since the origin (under both coordinate systems) is a vertex of T, they must differ by an integer constant, proving the proposition. $\square$

Definition 3.3.3. Let $V = \{v_1, ..., v_m\}$ be a finite set of vectors in a Euclidean half-space E. A path $c : [0, \infty) \to E$ or $c : [0, a] \to E$ is a *walk in E over V* if $a \in \mathbb{Z}$ and there is a sequence $v_{j_1}, v_{j_2}, v_{j_3}, ...$ of vectors in V such that for each integer k with $[k, k+1]$ in the domain of c, $c(k+1) = c(k) + v_{j_k}$, and $c \mid_{[k,k+1]}$ maps linearly to the line segment connecting $c(k)$ and $c(k) + v_{j_k}$.

Note that a walk is determined by a finite or infinite sequence of vectors in V together with a starting point $c(0)$.

For $i = 1, ..., n$, let $v_i = e_i + e_0$. Recall that the sequence of integers $\{I_k\}$ was chosen so that $c(k) = g_{I_1}...g_{I_k}(p)$.

Lemma 3.3.4. *$f(c)$ is a walk in $c \times \mathbb{R}^n$ over the set $\{v_1, ..., v_n\}$ starting at the origin. The sequence of vectors corresponding to this walk is $v_{I_1}, v_{I_2}, v_{I_3}, ...$.*

Proof. By Proposition 3.3.2, the space $c\mid_{[k,k+1]} \times\mathbb{R}^n$ lies in some $g \cdot E_i$, and the coordinates on the two spaces coincide up to translation by an integer multiple of e_0. In particular, $c([k, k+1])$ is a line segment in the coordinates of $g \cdot E_i$, and since f is affine on this space by Proposition 3.2.5, $f\big(c([k, k+1])\big)$ is a line segment in both coordinate systems.

To show that $f(c)$ is a walk, we thus need only to show that the vector $f(c(k+1)) - f(c(k))$ is one of $v_1, ..., v_n$. This can be done in either coordinate system since they only differ by a translation. Since $c(k+1) - c(k) = e_0$ in both coordinate systems, we then have $f(c(k+1)) - f(c(k)) = e_0 + e_i = v_i$, proving the proposition. $\square$

3.4 V-walks and boundary limits

In this section we will prove Theorem 3.2.6, deferring some technical arguments to the next section. Recall that the *Hausdorff distance* $d_H(S, S')$ between two subsets S and S' of a metric space is the infimum of the set $\{\epsilon > 0 \mid S \subset N_\epsilon(S'), S' \subset N_\epsilon(S)\}$.

Consider a Euclidean half-space $E = \{(x_1, ..., x_n) \in \mathbb{R}^n \mid x_0 \geq 0\}$. E is a CAT(0) space, and its boundary is a closed half-sphere M, which we give the angle metric (this is in fact the Tits metric). Since each point in the boundary is represented by a unique ray emanating from $\overrightarrow{0}$, we will occasionally fail to distinguish between the boundary point and the ray. It will also be convenient to think of points in E as vectors.

Given any $x \neq \overrightarrow{0} \in E$, we can draw the geodesic segment to $\overrightarrow{0}$ and then extend to a ray emanating from $\overrightarrow{0}$. This gives a map $\rho_M \colon E - \{\overrightarrow{0}\} \to M$. We will call the image of a point under this map the *projection of x to M*. Note that, for vectors x and u, we have that $d\big(\rho_M(x), \rho_M(u)\big)$ is exactly the angle between x and u at the origin.

Let V be any set of linearly independent vectors in E. Let $W = \rho_M(V)$, and let $L \subset M$ be the spherical simplex obtained by taking the convex hull of W. The result we will need to prove Theorem 3.2.6 is the following:

Theorem 3.4.1. *Suppose $Z \subset L$ is connected and compact. Then there is a walk in E over V, starting from the origin, whose boundary limit is Z.*

Nonsingular linear transformations induce homeomorphisms of Euclidean space and its boundary under the cone topology. Thus, by applying a change-of-basis transformation and passing to a subspace we may assume that $V = \{e_1, ..., e_n\}$.

Definition 3.4.2. A path $\gamma : [a, a+1] \to L$ is a *W-directed segment* if γ is injective and its image is an initial subsegment of a geodesic segment $[\gamma(a), w]$ for some $w \in W$.

Definition 3.4.3. A path $\gamma : [0, \infty) \to L$ or $\gamma : [0, a] \to L$ is a *W-directed path* if $a \in \mathbb{Z}$, and for each positive integer k in the domain of γ, $\gamma\mid_{[k-1,k]}$ is either constant or a W-directed segment.

The proof of Theorem 3.4.1 will rely on the following two lemmas. The proofs of these lemmas are rather technical and are deferred to the final section.

Lemma 3.4.4. *Let $\gamma : [0, a] \to L$ be any path, and $\epsilon > 0$. Then there is is a W-directed path $\gamma' : [0, a'] \to L$ such that $d_H(\gamma, \gamma') \leq \epsilon$. γ' may be chosen so that $\gamma'(0)$ is any point in $B_\epsilon(\gamma(0))$, and so that $\gamma'(a') \in B_{\frac{\epsilon}{2}}(\gamma(a))$.*

Lemma 3.4.5. *Let $\gamma : [0, a] \to L$ be a W-directed path, and let $\epsilon_1, \epsilon_2 > 0$. Then there is an $R > 0$ such that, for any $x \in E$ with $\|x\| \geq R$ and $d\big(\rho_M(x), \gamma(0)\big) < \frac{\epsilon_1}{2}$, there is a finite walk c in E over V, starting at x, such that $d_H\big(\rho_M(c), \gamma\big) \leq \epsilon_1$. The walk c may be chosen to have arbitrarily long length, and so that the projection of its ending point is within ϵ_2 of $\gamma(a)$.*

We now prove some basic propositions.

Proposition 3.4.6. *Let $c : [0, \infty) \to E$ be a path such that $\lim_{t\to\infty}\|c(t)\| = \infty$. Then the boundary limit of c is*

$$\left\{ \lim_{k\to\infty} \rho_M\big(c(a_k)\big) \middle| a_k \in \mathbb{R}^+, \lim_{k\to\infty} a_k = \infty \right\}.$$

Proof. Let q be a point in the boundary limit of c. Then there is a sequence of points $c(a_1), c(a_2), \ldots$ which converge to q in the cone topology. Recall that a basis of open sets around q consists of sets of the form $N(q, R, \epsilon) = \{x \in E \mid d(x, \overrightarrow{0}) > R, d(\rho_R(x), \rho_R(q)) < \epsilon\}$, where ρ_R is radial projection to the closed ball of radius R centered at $\overrightarrow{0}$.

For any R, all but finitely many of the $c(a_k)$ have $\|c(a_k)\| > R$. It follows that $\lim_{k\to\infty} a_k = \infty$. Since $d\big(\rho_R(c(a_k)), \rho_R(q)\big) < \epsilon$, the angle between $c(a_k)$ and q is less than $\sin^{-1}(\frac{\epsilon}{R})$, thus the angles between $c(a_k)$ and q approach zero. Therefore, $\lim_{k\to\infty} \rho_M\left(c(a_k)\right) = q$.

Conversely, suppose $q = \lim_{k\to\infty} \rho_M\left(c(a_k)\right)$, with $a_k \to \infty$. Then the angle between q and $c(a_k)$ approaches 0 as $k \to \infty$, and thus for any R and ϵ, we can find K such that $d\big(\rho_R(c(a_k)), \rho_R(q)\big) < \epsilon$ for all $k > K$. Since the a_k approach ∞ we can also choose K so that $\|c(a_k)\| > R$ for $k > K$. Thus q is the limit of the points $c(a_k)$ in the cone topology, and the proposition is proved. $\square$

Lemma 3.4.7. *Suppose $\gamma : [0, \infty) \to L$ is any path. Then there is a walk $c : [0, \infty) \to E$ over $V = \{e_1, \ldots, e_n\}$ such that $c(0) = \overrightarrow{0}$ and the boundary limit of c is*

$$\bigcap_{T>0} \overline{\gamma\left([T, \infty)\right)}.$$

Proof. Express γ as a concatenation of paths $\gamma_1 \colon [0, 1] \to L, \gamma_2 \colon [1, 2] \to L$, etc. By Lemma 3.4.4, for each γ_i, we can choose a directed path γ_i' such that $d_H(\gamma_i, \gamma_i') \leq \frac{1}{2^i}$, and we can choose these γ_i' so that the starting point of γ_i' is the ending point of γ_{i-1}'.

Letting $\epsilon_i = \frac{1}{2^i}$, let R_i be the associated value of R needed to approximate γ_i with Lemma 3.4.5. Choose any point $v \in E$ with positive integer coordinates such that $\|v\| > R_1$ and $d\big(\rho_M(v), \gamma_1(0)\big) < \frac{\epsilon_1}{2}$. Then by Lemma 3.4.5, there is a walk c_1 starting at v such that $d_H(\rho_M(c_1), \gamma_1') \leq \epsilon_1$. We can choose c_1 to have arbitrary length. In particular we can choose c_1 so that its ending point has distance greater than R_2 from the origin. We can also choose c_1 so that its ending point, when projected to M, has distance less than $\frac{\epsilon_2}{2}$ from $\gamma_2(0)$.

Lemma 3.4.5 then says there is a walk c_2, starting from the endpoint of c_1, ending at a point of distance at least R_3 from the origin, and such that $d_H(\rho_M(c_2), \gamma_2') \leq \epsilon_2$.

We may choose c_2 so that its ending point has distance less than $\frac{\epsilon_3}{2}$ from $\gamma_3(0)$ when projected to M.

Continuing in this way, we construct a sequence of walks $c_1, c_2, c_3, \ldots$, each ending where the last started, such that $d_H(c_i, \gamma_i') \leq \epsilon_i$. Concatenating all these walks, and concatenating any walk from the origin to v at the beginning gives us a walk $c\colon [0,\infty) \to E$.

By Proposition 3.4.6, the boundary limit of c consists precisely of limits of sequences of the form $\rho_M(c(a_k))$, with $a_k \to \infty$. Given such a sequence, for each k, we can choose a point q_k on some γ'_{i_k} such that $i_k \to \infty$ and $d(\rho_M(c(a_k)), q_k) \to 0$ as $k \to \infty$. We can then choose $b_1, b_2, \cdots \in \mathbb{R}$ such that $b_k \to \infty$, and so that $d(\gamma(b_k), q_k) \to 0$. Thus $d(\rho_M(c(a_k)), \gamma(b_k)) \to 0$ as $k \to \infty$.

The limit of a sequence $\gamma(b_k)$ with $b_k \to \infty$ is precisely a point in $\bigcap_{T>0} \overline{\gamma([T,\infty))}$. Thus we have that the boundary limit is contained in this intersection. But for any sequence $\gamma(b_k)$, we can choose $q_k \in \gamma'_{i_k}$ with $i_k \to \infty$ and $d(\gamma(b_k), q_k) \to 0$. Then we can choose a sequence $a_k \to \infty$ such that $d(\rho_M(c(a_k)), q_k) \to 0$, and so the sequence $\rho_M(c(a_k))$ converges to the same point as $\gamma(b_k)$. This shows the reverse inclusion, proving the proposition. $\square$

We are now in a position to prove Theorem 3.4.1.

Proof of Theorem 3.4.1. Fix any point $q \in Z$. For each positive integer k, let S_k be a finite subset of Z such that $Z \subset N_{\frac{1}{k}}(S_k)$. Since Z is a connected subset of a simplex, any open neighborhood of Z is path-connected. Thus we can choose paths $\gamma_k\colon [0,1] \to N_{\frac{1}{k}}(Z)$ such that $\gamma_k(0) = \gamma_k(1) = q$ and such that $S_k \subset \gamma_k$. Concatenating the paths γ_k gives a path $\gamma\colon [0,\infty) \to L$.

Note that any point of Z lies in the closure of γ, since the S_k get arbitrarily close to every point of Z. Indeed, any point of Z lies in $\overline{\gamma([T,\infty))}$ for any $T \in \mathbb{R}$. For $k \in \mathbb{Z}$, we have $\gamma([k,\infty)) \subset N_{\frac{1}{k}}(Z)$, and thus $Z = \bigcap_{T>0} \overline{\gamma([T,\infty))}$.

By Lemma 3.4.7, there is a walk over V starting at the origin which has $\bigcap_{T>0} \overline{\gamma([T,\infty))} = Z$ as its boundary limit. $\square$

Proof of Theorem 3.2.6. For $i = 1, ..., n$, let $v_i = e_0 + e_i$. Let $V = \{v_1, ..., v_n\}$. Let L be the spherical $(n-1)$-simplex obtained by taking the convex hull of $\rho_M(V)$. Embed Z into L. By Theorem 3.4.1, there is a walk c over V starting at the origin whose boundary limit is Z. Let $v_{I_1}, v_{I_2}, ...$ be the sequence of vectors in V associated to the walk c.

Take the geodesic ray in T which starts at p and passes through $g_{I_1} \cdot p, g_{I_1} g_{I_2} \cdot p$, etc., and call this ray c'. $f(c') \subset c' \times \mathbb{R}^n$ by Proposition 3.2.4. By Lemma 3.3.4, $f(c') = c$ in the coordinates of $c' \times \mathbb{R}^n$. Since $c' \times \mathbb{R}^n$ is a closed, convex subset of X, its boundary is embedded in ∂X. Thus, the boundary limit of $f(c')$ is the boundary limit of c and is homeomorphic to Z. $\square$

As mentioned, Theorem 3.2.6 implies the main theorem.

3.5 Remarks

It is interesting to note that the construction of the group G, the space X, and the map f depended only on the dimension of the space Z. This gives us that, for a fixed n, every compact connected subspace of $\mathbb{R}^{n-1}$ occurs as a boundary image of c for some geodesic c in X. While these boundary images have diverse homeomorphism types, all of their inclusions into ∂X are nullhomotopic. This leaves open whether the boundary image of a geodesic ray under a G-equivariant quasi-isometry is always nullhomotopic in the boundary.

In the proof of the main theorem, only positive geodesics were considered for simplicity's sake. However, the argument of 3.3.4 extends to geodesics in $T \subset X$ which do not represent positive words. Whenever an inverse of a generator g_i^{-1} occurs, the geodesic still passes through some $g \cdot E_i$, but $c(k+1) - c(k)$ is now $-e_0$ in the coordinates of $g \cdot E_i$. This means that a direction in the coordinates of $c \times \mathbb{R}^n$ is the reflection across the $e_0 = 0$ hyperplane of the direction in the coordinates of $g \cdot E_i$. So f sends the vector $c(k+1) - c(k)$ to the reflection of $f_i(-e_0)$ across the hyperplane $e_0 = 0$. Since $f(-e_0) = -e_0 - e_i$, its reflection is $e_0 - e_i$. Calling such a vector v_i', we see that the image of such a geodesic is then just a walk in $\{v_1, ..., v_n\} \cup \{v_1', ..., v_n'\}$.

The geodesics we have constructed are all of the form $c' \times \{0\} \subset T \times \mathbb{R}^n = X$, for some geodesic c' in the tree T. One may also ask what happens to the images of geodesics c which do not stay in the tree T. We answer this with the following proposition:

Proposition 3.5.1. *Let c be any geodesic in $X = T \times \mathbb{R}^n$. Let $c' = \pi_T(c)$, the projection of c to the T factor. If c' is not constant, then up to reparameterization c' is a geodesic, and the boundary image of c is homeomorphic to the boundary limit of c'.*

Proof. c' is a path in a tree which does not retrace itself, thus is a geodesic after reparameterization. Note that c must be contained in the Euclidean half-space $c' \times \mathbb{R}^n$, and as a geodesic ray it corresponds to some vector in this space. This vector must have a strictly positive e_0 coordinate, otherwise c' would be constant. Let v be the vector pointing in this direction, scaled so that the e_0 component of v is 1.

The same argument as in the proof of Proposition 3.3.4 then shows that, in the coordinates of $c' \times \mathbb{R}^n$, $f(c)$ is a walk in the vectors $f_1(v), ..., f_n(v)$. Letting $v = u + e_0$, we then have that $f_i(v) = u + e_i + e_0$. This means the vectors $f_i(v)$ are still linearly independent (one easily checks that adding $u + e_0$ to these vectors produces a basis of $\mathbb{R}^{n+1}$). Thus, $f(c)$ is a walk in $f_1(v), ..., f_n(v)$, and the sequence of these vectors is exactly the sequence of vectors $v_1, ..., v_n$ taken by $f(c')$ with v_i replaced by $f_i(v)$.

Let L be a change-of-basis linear transformation of Euclidean space which takes $\{v_1, ..., v_n\}$ to $\{f_1(v), ..., f_n(n)\}$. This extends to a homeomorphism of Euclidean space together with its boundary sphere. L also takes $f(c')$ to $f(c)$. Thus, the boundary image of $f(c)$ is the image under L of the boundary image of $f(c')$, and so the two are homeomorphic.

□

3.6 Directed paths in a simplex

In this section we will prove Lemmas 3.4.4 and 3.4.5.

Proposition 3.6.1. *It suffices to prove Lemma 3.4.4 for a standard Euclidean simplex K.*

Proof. The projection from the standard Euclidean simplex to the corresponding spherical simplex sends vertices to vertices and geodesics to geodesics. Thus a directed path on K will project to a directed path on the spherical simplex L.

Furthermore, since each point of K is at least distance $\frac{1}{\sqrt{n}}$ from the origin, points on K of distance d apart have angle no more than $2\sin^{-1}(2d\sqrt{n}) \leq 2\pi d\sqrt{n}$. Thus, projection is a Lipschitz map, and if we choose a directed path approximating a path within $\frac{\epsilon}{2\pi\sqrt{n}}$ in K, then the projection will approximate the projected path within ϵ in L. $\square$

In light of the above proposition, we will proceed to prove Lemma 3.4.4 for a Euclidean simplex K. We will still refer to the vertex set of K as W.

Lemma 3.6.2. *Let $q \in K$. Let K' be a face of K, and let q' lie in the interior of K'. Then for any $\epsilon > 0$, there is a W-directed path γ starting at q, ending in $B_\epsilon(q')$, such that $\gamma \subset N_\epsilon\big([q, q']\big)$.*

Proof. The proof will be by induction on the dimension of K'. If the dimension is zero then K' is a vertex, and so the geodesic from q to q' is already W-directed.

Let the dimension d of K' be at least 1, and choose a vertex v. The other vertices span a $(d-1)$–simplex K''. Extend the geodesic segment $[v, q']$ to K'', and denote by q'' the point where this segment intersects K''. Let P denote the 2-dimensional plane containing q, q', and v. Let $\pi_P \colon K \to P$ denote orthogonal projection to P.

Suppose i is an integer with $i > 1$. Since the dimension of K'' is $d-1$, the inductive hypothesis implies that there is a W-directed path β from q, ending within $\frac{\epsilon}{2^i}$ of q'', staying within $\frac{\epsilon}{2^i}$ of the segment $[q, q'']$. If $d(q'', q') \leq \frac{\epsilon}{2}$ then by setting $\gamma = \beta$ we are done, otherwise let x be the point on $[q, q'']$ which is of distance $\frac{\epsilon}{2}$ from $[q, q']$. Let $a \in \mathbb{R}$ be such that $d(\beta(a), x) < \frac{\epsilon}{2^i}$.

Let $\beta' = \beta\mid_{[0,a]}$. Note that β' is still W-directed.

Since q' lies in the interior of K', we have that v and q'' lie on opposite sides of $[q, q']$ in P. By the triangle inequality, this means that $\pi_P(\beta'(a))$ lies on the same side of $[q, q']$ as q''. Thus, the segment $[\beta'(a), v]$, when projected to P, crosses the segment $[q, q']$.

Append to β' the v-directed segment that terminates at the point y which projects to this intersection.

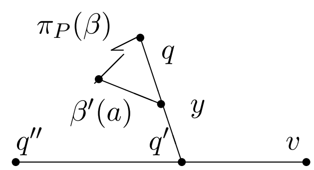

Figure 3.1: The construction projected to the plane P.

The segment $[\beta'(a), y]$ starts at most $\frac{\epsilon}{2^i}$ from the plane P and travels towards a point on P, and thus ends within $\frac{\epsilon}{2^i}$ of P. In particular $d(y, [q, q']) < \frac{\epsilon}{2^i}$. Thus if $d(y, [v, q'']) < \frac{\epsilon}{2}$, we are done. Otherwise, this segment has a length of at least $\frac{\epsilon}{4}$, since its projection starts at least of distance $\frac{\epsilon}{4}$ from $[q, q']$ and ends at distance 0.

Consider the line containing q' and v to be the x-axis in P. Then the line segment $\pi_P\left([\beta(a), y]\right)$ starts at least distance $\frac{\epsilon}{2}$ from the x-axis, and it moves towards a point on the x-axis no more than distance $\operatorname{diam}(K) = \sqrt{2}$ away. If we choose our orientation so that the x values of the segment are increasing, we then have that its slope is less than $-\frac{\epsilon}{2\sqrt{2}}$. Since it is of length at least $\frac{\epsilon}{4}$, this gives a bound on $d\left(\pi_P(\beta'(a)), [v, q'']\right) - d\left(\pi_P(y), [v, q'']\right)$ which is independent of i.

Since β stayed within $\frac{\epsilon}{2^i}$ of the geodesic $[q, q'']$, if i is sufficiently large we have $d\left(\pi_P(\beta(a)), [v, q'']\right) < d(q, [v, q''])$. Thus, assuming we choose i sufficiently large, we have that $d(q, [v, q'']) - d(y, [v, q''])$ is bounded away from zero.

We can now continually repeat this process, replacing q by y and increasing i sufficiently each time. We concatenate the results to obtain a W-directed path γ starting from q. Since at each stage we move no farther than $\frac{\epsilon}{2^i}$ from the geodesic to q' at the previous stage, we stay within ϵ of $[q, q']$ at every point in this process. Further, since each step ends closer to $[v, q'']$ by an amount bounded away from zero, the process eventually terminates within $\frac{\epsilon}{2}$ of $[v, q'']$, and thus within ϵ of q'. □

In particular, setting $K' = K$ in the above Lemma allows us to approximate any line segment by a W-directed path.

Proof of Lemma 3.4.4. Lemma 3.6.2 implies the existence of a W-directed path from any $x \in B_\epsilon(\gamma(0))$ which ends in $B_{\frac{\epsilon}{2}}(\gamma(0))$ and stays within $\epsilon - d\big(\gamma(0), x\big)$ of the segment $[x, \gamma(0)]$. By beginning with such a path, we can assume that our starting point lies in $B_{\frac{\epsilon}{2}}(\gamma(0))$.

Approximate γ by a piecewise-linear path β such that $d_H(\gamma, \beta) < \frac{\epsilon}{2}$. By Lemma 3.6.2 we can find W-directed paths for each linear segment, each starting where the previous ended, ending within $\frac{\epsilon}{2}$ of the next linear segment, and staying within $\frac{\epsilon}{2}$ of the segments connecting their starting and ending points. Thus each stays within ϵ of the corresponding linear segment of β.

Concatenating these W-directed paths gives a single W-directed path γ' with the desired properties. □

We will need the following proposition for the proof of Lemma 3.4.5:

Proposition 3.6.3. *Let $x \in E$ be any nonzero vector with nonnegative coordinates, and let e_i be any elementary basis vector. Let $u = x + e_i$. Then the segment $[x, u]$ projects to a W-directed segment on M, and $d(\rho_M(x), \rho_M(u)) \leq \sin^{-1}\left(\frac{1}{\|x\|}\right)$.*

Proof. The segment $[x, u]$ consists of vectors of the form $x + te_i$, for real values t. All such vectors lie in the plane spanned by x and e_i. The projection of this plane to M is precisely the geodesic containing $\rho_M(x)$ and $\rho_M(e_i)$. But $\rho_M(e_i) \in W$ is a vertex w_i of the simplex L. It is clear that $d(\rho_M(x+te_i), w) \geq d(\rho_M(x+t'e_i), w)$ for $0 \leq t \leq t' \leq 1$. Thus the segment $[x, u]$ projects to a subsegment of $[\rho_M(x), w]$, which is a W-directed segment.

Since v and e_i have fixed lengths, $\rho_M([x, u])$ has maximum length if u is orthogonal to e_i. In this case, its length is the angle between x and u, which is $\sin^{-1}\left(\frac{\|e_i\|}{\|x\|}\right)$. Since $\|e_i\| = 1$, this gives that $d(\rho_M(x), \rho_M(u)) \leq \sin^{-1}\left(\frac{1}{\|x\|}\right)$. □

Proposition 3.6.4. *Given any vector v and basis vector e_i,*

$$\lim_{k \to \infty} d\big(\rho_M(v + ke_i), \rho_M(e_i)\big) = 0.$$

Proof. By the same argument as above, the angle between ke_i and ke_i+v is no more than $\sin^{-1}\left(\frac{\|v\|}{\|ke_i\|}\right)$. Since $\rho_M(ke_i) = \rho_M(e_i)$ and $\|ke_i\| \to \infty$ the proposition follows. □

In particular, Propositions 3.6.3 and 3.6.4 show that we can add k copies of e_i to any vector v, and

$$\lim_{k\to\infty} \rho_M(ke_i + v) = \rho_M(e_i) \in L.$$

Additionally, the linear path from $\rho_M(v)$ to $\rho_M(ke_i + v)$ projects to a $\rho_M(e_i)$-directed segment.

Lemma 3.6.5. *Suppose a, a', and w are colinear in the half-sphere M, and a' is on the geodesic segment from a to w. Further suppose $b \in M$ with $d(b, w) < d(a', w)$, and all distances between these points are $\leq \frac{\pi}{2}$. Then $d(b, a') < d(b, a)$.*

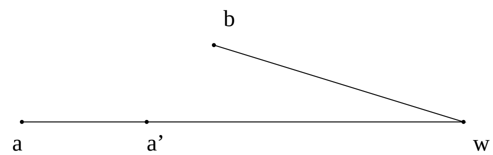

Figure 3.2:

Proof. This is an exercise in spherical geometry. For this proof, we will use the convention $\overline{ab}$ to denote $d(a, b)$. Since all distances here are $\leq \frac{\pi}{2}$, we have that sin and tan are strictly increasing functions, while cos is a strictly decreasing function. Let C be the angle (on the sphere) formed by the geodesic segments $[b, w]$ and $[a, w]$. Since $d(a, w) > d(a', w) > d(b, w)$ and $\cos(C) < 1$, we have

$$\tan\left(\frac{\overline{aw} + \overline{a'w}}{2}\right) > \tan(\overline{bw})\cos(C)$$

$$2\tan\left(\frac{\overline{aw} + \overline{a'w}}{2}\right) > 2\tan(\overline{bw})\cos(C)$$

$$2\cos(\overline{bw})\sin\left(\frac{\overline{aw} + \overline{a'w}}{2}\right) > 2\sin(\overline{bw})\cos\left(\frac{\overline{aw} + \overline{a'w}}{2}\right)\cos(C)$$

We multiply both sides by $\sin\left(\frac{\overline{a'w} - \overline{aw}}{2}\right)$, reversing the inequality since this value is negative, and apply angle difference formulas to obtain

$$\cos(\overline{bw})\big(\cos(\overline{aw}) - \cos(\overline{a'w})\big) < \sin(\overline{bw})\big(\sin(\overline{a'w}) - \sin(\overline{aw})\big)\cos(C)$$

$$\cos(\overline{bw})\cos(\overline{aw}) + \sin(\overline{bw})\sin(\overline{aw})\cos(C) < \cos(\overline{bw})\cos(\overline{a'w}) + \sin(\overline{bw})\sin(\overline{a'w})\cos(C)$$

Applying the spherical law of cosines to the triangles with vertices ABW and $A'BW$, we end up with

$$\cos(\overline{ab}) < \cos(\overline{a'b})$$

$$\overline{ab} > \overline{a'b}$$

□

Let $v = \rho_M(e_n)$. Then v is a vertex of L. Call the span of the other $n-1$ vertices L'. Define $\Psi\colon L - \{v\} \to L'$ such that $\Psi(q)$ is the point of L' obtained by extending the geodesic segment $[v, q]$ to L'.

Proposition 3.6.6. *1. Ψ is induced by orthogonal projection of E to the subspace spanned by $\{e_1, ..., e_{n-1}\}$.*

2. Up to reparameterization, Ψ sends geodesic segments to geodesic segments.

3. Ψ sends W-directed paths to $(W - \{v\})$-directed paths.

4. If $d(v, q) = d(v, q')$, then $d(q, q') \leq d(\Psi(q), \Psi(q'))$.

Proof. Note that an injective path in L is locally a geodesic segment (up to reparameterization) if and only if it lies in $\overline{P} \cap L$ for some 2-plane $P \subset E$ which passes through the origin. Let $q \in L - \{v\}$, let x be any vector such that $\rho_M(x) = q$, and let x' be such that $\rho_M(x') = \Psi(q)$. Since $\rho_M(x') \in L'$, x' lies in the span of $e_1, ..., e_{n-1}$, so x' and e_n are orthogonal.

v, q, and $\Psi(q)$ lie on the same geodesic, thus x must lie in the plane spanned by e_n and x', ie $x = ae_n + bx'$ for scalars a and b. This shows that Ψ is induced by orthogonal projection to the subspace spanned by $e_1, ..., e_{n-1}$, proving (1). This projection is a linear map, thus it sends line segments to line segments, and since geodesic segments on L are projections of line segments in E, we have proved (2).

For $q \in L$ and w a vertex, Ψ sends the geodesic segment $[q, w]$ to the geodesic segment $[\Psi(q), \Psi(w)]$. If $w \neq v$, then $w \in L'$ and so $\Psi(w) = w$. Thus Ψ sends the segment $[q, w]$ to $[\Psi(q), w]$, and sends initial subsegments of the former to initial subsegments of the latter. If $w = v$, a subsegment of $[q, w]$ is sent to a single point, which proves (3).

By (1), we have that for all $q, q' \in L$, $d(\Psi(q), v) = d(\Psi(q'), v) = \frac{\pi}{2}$. Geodesic rays $\alpha(t), \alpha'(t)$ emitting from a common basepoint in a sphere have the property that $d(\alpha(t), \alpha'(t)) < d(\alpha(t'), \alpha'(t'))$ for $0 \le t < t' \le \frac{\pi}{2}$. This proves (4).

$\square$

We require one final lemma:

Lemma 3.6.7. *For any $\epsilon > 0$, there is an $R > 0$ such that, if x is a vector with positive coordinates and $\|x\| > R$, there is a walk $c\colon [0, \infty) \to E$ over V, starting at x, such that $\rho_M(c) \subset B_\epsilon(x)$.*

Proof. We will proceed by induction on the dimension of L. If L is 1-dimensional, then $B_\epsilon(\rho_M(x))$ is an interval. If $\|v\| > \sin^{-1}\left(\frac{2}{\epsilon}\right)$ then Proposition 3.6.3 guarantees that these segments all have length $< \frac{\epsilon}{2}$, thus we can always choose an e_i such that $\rho_M([v, v + e_i]) \subset B_\epsilon(\rho_M(x))$. Thus, setting $R = \sin^{-1}\left(\frac{2}{\epsilon}\right)$ gives the result.

For higher dimensions, set $v = e_n$, and let L' and Ψ be as in Proposition 3.6.6. Let $\rho\colon E \to E$ be orthogonal projection to the subspace spanned by $e_1, ..., e_{n-1}$. By 3.6.6(1), for every vector x, $\Psi(\rho_M(x)) = \rho_M(\rho(x))$. Note that $\|x\| \ge \|\rho(x)\|$.

By the inductive hypothesis, there is an R such that if $\|\rho(x)\| > R$, then there is a walk in $\{e_1, ..., e_{n-1}\}$, starting from $\rho(x)$, whose projection to L' lies within $\frac{\epsilon}{3}$ of $\Psi(\rho_M(x))$. If we use the same sequence of vectors starting from x instead of $\rho(x)$ we obtain a walk c. By Proposition 3.6.6 (1) and (4), we have that $\Psi(\rho_M(c)) \subset B_{\frac{\epsilon}{3}}\big(\Psi(\rho_M(x))\big)$.

Note that since c is a walk in $e_1, ..., e_{n-1}$, we have that for all t, $d\big(\rho_M(c(t)), v\big) \ge d\big(\rho_M(x), v\big)$. Assume R is chosen large enough so that projections of segments $[v, v+e_i]$ have length no more than $\frac{\epsilon}{3}$. We will create another walk c' by inserting copies of e_n into the sequence of vectors for c.

Let a be the smallest integer such that $d\big(\rho_M(c(a)), v\big) - d\big(\rho_M(x), v\big) > \frac{\epsilon}{3}$. We insert copies of e_n at the a^{th} position in the sequence for c. We insert the minimal number

m of such vectors so that $d\big(\rho_M(c(a)+me_n)), v\big) \le d\big(\rho_M(x), v\big)$. Such an m exists by Proposition 3.6.4. After inserting the copies of e_n, our sequence continues with the sequence of vectors from c.

This creates a new sequence of vectors, and we continue in this way, finding the minimal integer a such that $d\big(\rho_M(c(a)), v\big) - d\big(\rho_M(x), v\big) > \frac{\epsilon}{3}$, and inserting e_n's so that $d\big(\rho_M(c(a+m)), v\big) \le d\big(\rho_M(x), v\big)$. Continuing this process to infinity creates a new walk c' with the property that, for all t, $|d\big(\rho_M(c'(t)), v\big) - d\big(\rho_M(x), v\big)| < \frac{2\epsilon}{3}$.

For any t, let q_t be the point on the geodesic $[v, \Psi(\rho_M(x))]$ such that $d(q_t, v) = d\big(\rho_M(c'(t)), v\big)$. Note that since we have only inserted e_n's, we have that $\Psi(\rho_M(c)) = \Psi(\rho_M(c'))$. By Proposition 3.6.6 (4) this implies that $d\big(\rho_M(c'(t)), q_t\big) < \frac{\epsilon}{3}$. Since $d\big(\rho_M(x), q_t\big) < \frac{2\epsilon}{3}$, the triangle inequality implies that $\rho_M(c'(t))$ is always within ϵ of x.

This construction works as long as $\|\rho(x)\| > R$. But the construction using the inductive hypothesis works via orthogonal projection to the span of any $(n-1)$ of our basis vectors. Thus, by replacing R with $2R$, we can guarantee that any x with $\|x\| > R$ has some sufficiently large projection to allow the above construction to work. This completes the proof. $\square$

Proof of Lemma 3.4.5. Let N be the number of directed segments in γ. Let $w_i = \rho_M(e_i)$, so $W = \{w_1, ..., w_n\}$. Let $\delta = \frac{\epsilon_1}{4N}$. Let $R_1 = \csc(\delta)$, let R_2 to be the value obtained from Lemma 3.6.7 which allows the projection of a walk to stay within $\frac{\epsilon_2}{2}$ of its starting point, and let $R = \max(R_1, R_2)$.

To construct our walk c, let $\gamma_1, \gamma_2, ...$ be the W-directed segments of γ. Let w_{i_k} be the vertex that γ_k moves towards, and let T_k be the terminal point of γ_k. We will construct our sequence of vectors c by starting with the empty sequence and appending copies of e_{i_k} to the sequence for each γ_k. Denote by Σ_k the sum of all vectors appended up to the k^{th} step of this process, together with x. To choose how many e_{i_k}'s we append at the k^{th} step, we consider 3 cases:

1. If $T_k = w_{i_k}$, we append enough copies of e_{i_k} so that $d\big(\rho_M(\Sigma_k), w_{i_k}\big) < \frac{\epsilon_1}{2}$.

2. If $d(T_k, w_{i_k}) > d\big(\rho_M(\Sigma_{k-1}), w_{i_k}\big)$, we append no copies of e_{i_k}.

3. Otherwise, we append the minimal number of e_{i_k} required so that $d(T_k, w_{i_k}) > d\big(\rho_M(\Sigma_k), w_{i_k}\big)$.

Cases (1) and (3) are possible by Proposition 3.6.4.

We claim that in each of the above cases,

$$d\big(\rho_M(\Sigma_k), T_k\big) \leq \max\left(\frac{\epsilon_1}{2}, d\big(\rho_M(\Sigma_{k-1}), T_{k-1}\big) + \delta\right).$$

In case 1, the claim is clear since $T_k = w_{i_k}$. In case 2, the claim holds by Lemma 3.6.5.

In case 3, let $A_k \in [T_{k-1}, w_{i_k}]$ and $B_k \in [\rho_M(\Sigma_{k-1}), w_{i_k}]$ be such that $d(A_k, w_{i_k}) = d(B_k, w_{i_k}) = \min\big(d(T_{k-1}, w_{i_k}), d(\rho_M(\Sigma_{k-1}), w_{i_k})\big)$. Let P_k be the point on $[\rho_M(\Sigma_{k-1}), w_{i_k}]$ such that $d(P_k, w_{i_k}) = d(T_k, w_{i_k})$. Then we have

$$d\big(\rho_M(\Sigma_{k-1}), T_{k-1}\big) \geq d(A_k, B_k)$$

$$d(A_k, B_k) \geq d(T_k, P_k)$$

The first inequality holds by Lemma 3.6.5. The second holds because geodesics $\alpha(t)$ and $\alpha'(t)$ emitting from a common basepoint have increasing distance in t for $t < \frac{\pi}{2}$, and $[w_{i_k}, A_k]$ and $[w_{i_k}, B_k]$ are such segments.

Since we added the minimal number of e_{i_k}'s to make $\rho_M(\Sigma_k)$ closer to w_{i_k} than T_k is, we have that $P_k \in \rho_M\left([\Sigma_k - e_{i_k}, \Sigma_k]\right)$. But since $\|\Sigma_k - e_{i_k}\| \geq R = \csc(\delta)$, by Proposition 3.6.3 the length of this segment is no more than $\sin^{-1}\left(\frac{1}{\csc(\delta)}\right) = \delta$. Thus by the triangle inequality we have

$$d(T_k, P_k) \geq d(\rho_M(\Sigma_k), T_k) - \delta$$

and the claim follows.

Thus, for each k, $\rho_M(\Sigma_k)$ gets no more than δ further from T_k. But since $\delta \leq \frac{\epsilon_1}{4N}$, with N the number of segments in γ, we have that $d(\rho_M(\Sigma_N), T_N) < \frac{3\epsilon_1}{4}$, since $d(\rho_M(\Sigma_0), T_0) < \frac{\epsilon_1}{2}$. Thus, it only remains to show that we can append additional vectors causing the walk to end within ϵ_2 of T_N, and to show that we can continue the walk arbitrarily far without leaving $B_{\epsilon_1}(T_N)$.

By Lemma 3.6.2, we can find a W-directed path β in L which starts at $\rho_M(\Sigma_N)$, ends within ϵ_2 of T_N, and stays within ϵ_2 of the geodesic segment $[\Sigma_N, T_N]$. Without loss of generality we may assume $\epsilon_2 < \frac{\epsilon_1}{4}$. Let N' be the number of segments in β, let $\delta' = \frac{\epsilon_2}{4N'}$, and let $R' = \csc(\delta')$.

By Lemma 3.6.7, we can append vectors to our sequence so that the projections stay within $\frac{\epsilon_2}{2}$ of $\rho_M(\Sigma_N)$, until we reach a total sum Σ with $\|\Sigma\| > R'$. We can then repeat the above construction to append a walk whose projection ends within ϵ_2 of T_N and stays within ϵ_2 of β and thus within ϵ_1 of T_N. This gives us our walk which has all the desired properties. $\square$

Chapter 4

The Locally Connected Case

In this chapter we generalize the results of Chapter 3 to include the added assumption that the boundaries in question are locally connected. We will construct examples which prove the following theorem:

Theorem 4.0.8. *For any compact, connected $Z \subset \mathbb{R}^n$, there is a group G which acts geometrically on two CAT(0) spaces X and Y with locally connected boundaries, a G-equivariant map $f\colon X \to Y$, and a geodesic ray c in X such that the boundary limit of $f(c)$ is homeomorphic to Z.*

As before, the construction of G, X, Y, and f depends only on n. X and Y will again be the same space with different actions of G, and their boundaries will be spheres. The examples in this chapter are somewhat analogous to the examples of Buyalo [8] where the G-equivariant quasi-isometries between spaces do not extend to the boundaries.

Throughout this chapter we will assume a fixed integer $n > 1$.

4.1 Constructing X, G, and f

Let Σ_n be a surface of genus n, endowed with a metric of constant negative curvature. Let $\overline{X} = \Sigma_n \times T^n$, where T^n is the standard n-torus. Let X denote the universal cover of $\overline{X}$. The following properties are easy to verify:

Proposition 4.1.1. *1. $\overline{X}$ is a space of nonpositive curvature.*

2. X is a CAT(0) space, and is isometric to $H^2 \times \mathbb{R}^n$, where H^2 denotes the hyperbolic plane.

3. *$\pi_1(\overline{X}) = \Phi \times \mathbb{Z}^n$, where Φ is the Fuchsian group of hyperbolic isometries with the following one-relation presentation:*

$$\Phi = \langle g_1, ..., g_n, h_1, ..., h_n \mid g_1 h_1 g_1^{-1} h_1^{-1} g_2 h_2 g_2^{-1} h_2^{-1} ... g_n h_n g_n^{-1} h_n^{-1} \rangle$$

Setting $G = \Phi \times \mathbb{Z}^n$, we then have a geometric action of G on X, which we denote by $\cdot$. As in Chapter 3, we may choose a basepoint $p \in X$ and choose generators of the $\mathbb{Z}^n$ factor of G such that the generator with a 1 in the i^{th} place and zeros elsewhere translates the $\mathbb{R}^n$ factor of p by the the standard basis vector e_i and fixes the H^2 factor. We choose our generators of Φ so that they fix the $\mathbb{Z}^n$ factor of G. We may assume the metric on H^2 is scaled such that each g_i and h_i translate p by a distance of 1.

Define a function $\psi\colon\ \Phi \to \mathbb{Z}^n$ by defining $\psi(g_i) = e_i, \psi(h_i) = \overrightarrow{0}$. Note that this induces a well-defined homomorphism since the only relation in Φ is always satisfied in an Abelian group. Define an automorphism $\Psi : G \to G$ by considering G as $\Phi \times \mathbb{Z}^n$ and defining $\Psi(g, w) = \big(g, w + \psi(g)\big)$. The reader may verify that Ψ is an automorphism.

Define a new action, $\circ$, by $g \circ x = \Psi(g) \cdot x$ for $g \in G$ and $x \in X$.

Proposition 4.1.2. *The action $\circ$ is geometric.*

Proof. The proof is identical to the proof of Proposition 3.2.2. $\square$

It is well-known that a fundamental domain of the action of Φ on H^2 is a hyperbolic $2n$-gon. Choose such a fundamental domain having p as a vertex, and call it P. We consider H^2 as a subspace of X by identifying it with $H^2 \times \left\{\overrightarrow{0}\right\}$.

We define a function $f\colon\ X \to X$ in the following way: Each vertex v in P is $g \cdot p$ for some $g \in G$, and we define $f(v) = g \circ p$ for these vertices. We then map each edge $[v, w]$ of P linearly to the segment $[f(v), f(w)]$. For points in the interior of P, each lies on a unique geodesic segment connecting p and an opposite edge, and we map such geodesics linearly to the geodesic segments connecting the images of their endpoints.

Any point of H^2 not in P can be written as $g \cdot x$ for some $g \in G$ and $x \in P$. We then define $f(g \cdot x)$ to be $g \circ f(x)$ for all such points. We have now defined f on all of H^2. To extend f to all of X, we simply define $f\big((x, w)\big) = \Big(\pi_{H^2}\big(f(x)\big), \pi_{\mathbb{Z}^n}\big(f(x)\big) + w\Big)$, where π_{H^2} and $\pi_{\mathbb{Z}^n}$ are the projections of X to the H^2 and $\mathbb{Z}^n$ components, respectively.

Proposition 4.1.3. *Using the above definition, f is well-defined and G-equivariant in the sense that $f(g \cdot x) = g \circ f(x)$ for all $g \in G, x \in X$.*

Proof. If $g \in \{1\} \times \mathbb{Z}^n \subset G$, then $g \cdot x = g \circ x$, and the proposition follows for these g. Thus it suffices to prove the proposition for the Φ component of G. Since this component acts on any translate of H^2 by simply preserving the translation, it suffices to prove the proposition for $x \in H^2$.

The proposition holds trivially for points $g \cdot x$ when x lies in the interior of P, since the orbit of such a point has a unique element in P.

If x is a vertex of P, then $x = h \cdot p$ for some $h \in \Phi$. But then any point $g \cdot x$ is just $gh \cdot p$, and $f(g \cdot x) = f(gh \cdot p) = gh \circ p = g \circ f(h \cdot p) = g \circ f(x)$, and we see that G-equivariance holds.

If x lies in the interior of an edge of p, then there is a unique $y \in P$ which lies in the same orbit as x. These points are uniquely defined by their distances along the geodesic segments that make up their respective edges. Since the proposition holds for vertices and both group actions are by isometries, the proposition holds for edges as well. $\square$

It is well-known (see [6] for example) that the boundary of the product of two CAT(0) metric spaces is the join of the two boundaries. Since $\partial H^2 = S^1$ and $\partial \mathbb{R}^n = S^{n-1}$, it follows that $\partial X = \partial\left(H^2 \times \mathbb{R}^n\right) = S^1 * S^{n-1} = S^{n+1}$. In particular, ∂X is locally connected.

The remainder of this section will be devoted to proving the following, which implies Theorem 4.0.8:

Theorem 4.1.4. *Given any compact, connected $Z \subset \mathbb{R}^{n-1}$, there is a geodesic ray $c\colon [0, \infty) \to X$ such that the boundary limit of $f(c)$ is homeomorphic to Z.*

4.2 Useful properties of H^2 and X

In this subsection we set out certain properties which we will need to prove Theorem 4.1.4.

The following theorem is proved in [6]. In fact, the result holds for arbitrary δ-hyperbolic spaces, but we will only need the result for H^2.

Theorem 4.2.1. *In H^2, every quasi-geodesic ray lies in the R-neighborhood of a geodesic ray, for some R.*

This shows that every quasi-geodesic ray in H^2 approaches a unique boundary point in ∂H^2 under the cone topology. We will use this fact by examining the Cayley graph Γ of Φ, using the $g_1, ..., g_n, h_1, ...h_n$ generating set. Recall that by Theorem 1.0.3, there exists a quasi-isometry from Γ to H^2. In particular, geodesic rays in Γ correspond to boundary points in H^2.

We now require a theorem to characterize certain geodesic rays in Γ. Given a sequence of generators $a_1, a_2, ...$ from $\{g_1, ...g_n, h_1...h_n\}$, we can construct a path $c\colon [0, \infty) \to \Gamma$ which starts at the identity and passes through a_1, then a_1a_2, then $a_1a_2a_3$, etc., mapping each interval $[k, k+1]$ isometrically to the edge connecting $a_1...a_{k-1}$ to $a_1...a_k$. Note that such an edge always exists, since at each step we are multiplying on the right by a generator.

Theorem 4.2.2. *Let c be constructed as above. Then c is a geodesic ray.*

To prove this theorem, we first need some terminology. Consider all words of length N, where the alphabet is a chosen set of generators of a finitely-generated group, together with the inverses of such generators. The symmetric group S_N acts on such words by permuting the entries.

Definition 4.2.3. Two words w and w' of length N are *cyclically equivalent* if, for some N-cycle $s \in S_N$, either $w = sw'$ or $w^{-1} = sw'$, where the inverse is taken as an element of the free group on the chosen generators.

Definition 4.2.4. A word v is a *cyclic subword* of a word w if v is a subword of a word w' which is cyclically equivalent to w.

Definition 4.2.5. A word w is *cyclically reduced* if no word cyclically equivalent to w contains a generator followed immediately by its inverse.

We will use the following, which is proved in much greater generality in [19]:

Theorem 4.2.6. *Suppose w is a cyclically-reduced word in $\{g_1, ..., g_n, h_1, ..., h_n\} \cup \{g_1^{-1}, ...g_n^{-1}, h_1^{-1}, ..., h_n^{-1}\}$ which represents the identity in Φ. Then, w contains a cyclic subword of length at least $2n+1$ which is also a cyclic subword of the relator $g_1h_1g_1^{-1}h_1^{-1}...g_nh_ng_n^{-1}h_n^{-1}$.*

Recall that edges in a Cayley graph have a natural labeling by generators of the group. Given a finite-length path γ which starts and ends and vertices in the Cayley graph, we may then form a word w by *reading off the edges* of γ in the following way: Whenever γ traverses an edge from g to gs for a generator s, we add s to the end of w. When γ traverses an edge from gs to g, we add s^{-1} to the end of w.

The following is a well-known fact in geometric group theory:

Proposition 4.2.7. *Suppose γ is a loop in a Cayley graph. Then the word w obtained by reading off the edges of γ represents the identity in the group.*

Proof. Let g be the group element corresponding to the vertex where γ starts and ends. Let w_k denote the subword of w consisting of the first k entries. An easy induction from the definition of reading off edges shows that $\gamma(k) = gw_k$ for each k. Thus, when k is the length of w, we have $g = gw$, showing that w represents the identity. □

We are now in a position to prove Theorem 4.2.2:

Proof of Theorem 4.2.2. It suffices to show that each initial segment of c is a geodesic segment. So suppose that $c\mid_{[0,k]}$ is not a geodesic segment in the graph Γ. Let c' denote a geodesic segment from the identity to $c(k)$. Choose c' so that c and c' share as long an initial segment as possible; call the final vertex of this segment v. Further suppose that c' is chosen so that, given this initial segment, it shares as long a terminal segment as possible with c, and call the starting vertex of this segment v'.

Let w be the word obtained by reading off the edges traversed by c from v to v', and let u be the word obtained by reading the edges traversed by c' from v to v'. Then wu^{-1} is a loop in the Cayley graph, and thus by Proposition 4.2.7 represents the identity in Φ.

By Theorem 4.2.6, there is then a cyclic subword s of wu^{-1} which is also a cyclic subword of the relator $g_1h_1g_1^{-1}h_1^{-1}...g_nh_ng_n^{-1}h_n^{-1}$, and s has length at least $2n+1$. We can then write some cyclic permutation of the relator as st, for another cyclic subword t of length at most $2n-1$. Since st must then represent the identity, we then have that $s = t^{-1}$ in Φ.

w is obtained by reading off the edges of c, and thus it consists only of positive powers of generators and not their inverses. We can see that no cyclic subword of the relator has more than two generators in a row with the same sign, thus at most two letters of s come from w, and at least $2n-1$ come from u^{-1}.

Since u forms a geodesic path, it is of minimal word length. Note that st^{-1} forms a loop, so if two paths have the same starting point and yield s and t^{-1} by reading off their edges, they have the same ending point. If we could replace s with t^{-1} in u we could shorten u's word length without changing the group element it represents, yielding a contradiction, and so at least one letter of s must come from w.

Thus, the edges making up the subword s are contained in both c and c', and thus must pass through either v or v'. Let V be either v or v', whichever they pass through. Consider the path formed by the collection of these edges, and let x and x' be the endpoints of this path which lie in c and c', respectively. Then $1 \leq d(x, V) \leq 2$, and since c' is a geodesic segment, $d(x', V) \geq 2n - 1$.

Since x and x' are connected by a path making the word s, they are also connected by a path making the word t, since st^{-1} forms a loop in the Cayley graph. Thus $d(x, x') \leq 2n - 1$. So we can construct a path from V to x' of length $2n + 1$ by traversing the path from V to x, then from x to x'.

But this means that there is a path from V to x' of minimal length which traverses at least 1 edge identically to c. Replacing the part of c' from V to x' with this path, we find that either $V = v$ and we have extended the initial segment shared by c and c', or $V = v'$ and we have extended the terminal segment shared by c and c' without altering the initial segment. Since the initial and terminal segments were chosen maximally, we have a contradiction. Thus, the only possibility is that $v = v'$, ie, $c = c'$, and thus c is a geodesic segment. $\square$

In X we have a copy of H^2 with an action of Φ and a chosen basepoint p. We can easily embed Γ into H^2 in a Φ-equivariant way: we map the identity to p, we map each vertex g to $g \cdot p$, and we map edges to the geodesic segments connecting the images of their vertices. This is a quasi-isometric embedding. Thus for any geodesic ray c in Γ, c becomes a quasi-geodesic ray in $H^2 \subset X$.

For the following proposition, we will use d_X to denote the distance function in X, and d_Γ to denote the distance function in Γ.

Proposition 4.2.8. *Let c be a geodesic ray in Γ with $c(0) = p$, considered as a quasi-geodesic ray in X. There is a $B > 0$ such that, for all $0 \leq t \leq t'$, $d_X\big(c(t'), p\big) - d_X\big(c(t), p\big) \geq -B$.*

Proof. By Theorem 4.2.1, c lies in the R-neighborhood of some geodesic ray c', for some $R \geq 0$. Since Γ is embedded in H^2 quasi-isometrically, there are values $\lambda \geq 1, K > 0$, such that for any $x, y \in \Gamma$,

$$\frac{1}{\lambda} d_X(x, y) - K \leq d_\Gamma(x, y) \leq \lambda d_X(x, y) + K.$$

We claim that $B = \lambda^2\big(2R + (1 - \frac{1}{\lambda})K\big)$ satisfies the proposition.

To see this, suppose that for some $t' \geq t \geq 0$, we have that $d_X\big(c(t'), p\big) - d_X\big(c(t), p\big) < -B$. Let q be the point on c' closest to $c(t)$. Then there is some $s > t'$ such that $d_X\big(c(s), p\big) > d_X(q, p) + R$, since quasi-geodesic rays get arbitrarily far from any fixed point.

Define $\rho_{c'} \colon X \to c'$ to take points in X to their closest point on c'. Since c' is closed and convex, this is a continuous function. Further, $d_X\big(\rho_{c'}(c(s)), p\big) > d_X(q, p)$, since $d_X\big(c(s), p\big) > d_X(q, p) + R$. Since $\rho_{c'}$ is continuous, this means there is some $t'' > t'$ such that $\rho_{c'}(c(t'')) = q$, and thus $d_X\big(c(t''), q\big) \leq R$.

Since $c(t)$ and $c(t'')$ are both within R of q, the triangle inequality gives us

$$d_X\big(c(t''), c(t)\big) \leq 2R.$$

Also by the triangle inequality and the assumption we have,

$$d_X\big(c(t'), c(t)\big) > B,$$

thus by the quasi-isometry constants we obtain

$$d_\Gamma\big(c(t'), c(t)\big) > \frac{1}{\lambda}(B-K)$$

This gives us that

$$d_\Gamma\big(c(t''), c(t)\big) = t'' - t > t' - t = d_\Gamma\big(c(t'), c(t)\big) > \frac{1}{\lambda}(B-K)$$

Applying the quasi-isometry constants an additional time yields

$$d_X\big(c(t''), c(t)\big) > \frac{1}{\lambda^2}B - (1+\frac{1}{\lambda})K$$

.

Combining with the above gives

$$2R \geq d_X\big(c(t''), c(t)\big) > \frac{1}{\lambda^2}B - (1+\frac{1}{\lambda})K = 2R,$$

a contradiction. □

4.3 Constructing the boundary limit Z

As before, let $\pi_{\mathbb{R}^n}$ and π_{H^2} be the projections from X onto its $\mathbb{R}^n$ and H^2 factors, respectively. We define a function $\mu\colon X \to \mathbb{R}^{n+1}$:

$$\mu(x) = \Big(d\big(\pi_{H^2}(x), p\big), \pi_{\mathbb{R}^n}(x)\Big).$$

Note that $d_X(x,y) \geq d_{\mathbb{R}^{n+1}}\big(\mu(x), \mu(y)\big)$, for all $x, y \in X$. Taking a sequence of generators $a_1, a_2, ...$, we can construct the quasi-geodesic ray c as in the previous section. Denote by C the path in $\mathbb{R}^{n+1}$ given by $C(t) = \mu(f(c(t)))$. Let $v_1, ..., v_n$ be vectors in $\mathbb{R}^{n+1}$ defined by $v_i = e_i + e_0$.

Lemma 4.3.1. *For any walk $b\colon [0,\infty) \to E$ over $\{v_1, ..., v_n\}$ starting from the origin, there is a sequence $a_1, a_2, a_3, ...$ of generators from $\{g_1, ..., g_n, h_1\}$ such that when c and C are constructed as above, there is some $R > 0$ such that $C \subset N_R(b)$ and $b \subset N_R(C)$.*

Proof. Let $v_{i_1}, v_{i_2}, ...$ be the sequence of vectors corresponding to the walk b. As in the previous chapter, we will construct the sequence $a_1, ..., a_n$ by adding on certain elements for each element v_{i_k}.

Since f is G-equivariant, we can immediately determine the $\mathbb{R}^n$ coordinate of $c(j)$ for any integer j. Since $c(j)$ is $a_1...a_jp$, the i^{th} entry in the $\mathbb{R}^n$ component of $f(c(j))$ is precisely the number of times g_i occurs in $a_1, ..., a_j$. Thus, adding on g_i to our sequence causes the e_i component to increase by 1, and all the other components except the e_0 component are unchanged.

$d_X\big(c(j), c(j+1)\big) = 1$ because each edge in Γ maps to a geodesic segment. Thus, adding on a g_i to our sequence causes the e_0 component to change by no more than 1. Further, adding h_1 to our sequence will not change the $e_1, ...e_n$ components, and adding h_1 many times will eventually increase the e_0 component, since the quasi-geodesic ray gets arbitrarily far from p. By Proposition 4.2.8, there is a bound B on how much the e_0 component can decrease from a previous value.

We now construct our sequence as follows: At the k^{th} element v_{i_k} in the sequence corresponding to b, suppose we have already constructed the finite sequence $a_1, ..., a_{j_k}$ corresponding to $v_1, ..., v_{i_k-1}$. We then add g_{i_k} to our sequence. We then add enough copies of h_1 so that the e_0 component of C becomes at least as large as the e_0 component of $b(k)$.

Constructed this way, we can see immediately that for each integer k and each integer $j \in \{j_k, ..., j_{k+1} - 1\}$, the $e_1, ...e_n$ components of $b(k)$ and $C(j)$ are always the same. If we denote π_0 to be projection to the e_0 component, then we have $\pi_0(b(k)) - B - 1 \leq \pi_0(C(j)) \leq \pi_0(b(k)) + 1$. Since b and C travel no more than distance 1 along each coordinate between integer values, this shows that b and C stay within a bounded distance of each other. □

Proposition 4.3.2. *Suppose two paths c and c' in a CAT(0) space X have the property that, for some $R > 0$, $c \subset N_R(c')$ and $c' \subset N_R(c)$. Then c and c' have the same boundary limit.*

Proof. Suppose α is a point on ∂X and $a_1, a_2, ...$ is a sequence of real numbers such that $c(a_1), c(a_2), \ldots$ converges to α in the cone topology. Then choose $b_1, b_2, \ldots$ such that $d\big(c(a_i), c'(b_i)\big) < R$ for each i. Pick a basepoint p.

Suppose α is a point in the boundary limit of c. Take any neighborhood $N(\alpha, T, \epsilon)$ of α. Choose $T' > 0$ large enough so that if q and q' are points on the boundary of a Euclidean ball of radius T' with $d(q, q') < R$, then the distance between their projections to the ball of radius T is no more than $\frac{\epsilon}{2}$.

Choose a_k such that $d\big(c(a_k), p)\big) > T'$, and so that the projection of $c(a_k)$ to $B_T(p)$ is within $\frac{\epsilon}{2}$ of the projection of α. Then the projection of $c'(b_k)$ is within $\frac{\epsilon}{2}$ of the projection of $c(a_k)$ by the CAT(0) inequality, thus $c'(b_k) \in N(\alpha, T, \epsilon)$. Thus α lies in the boundary limit of c' as well.

This shows that the boundary limit of c is a subset of the boundary limit of c', and a completely symmetric argument shows the reverse inclusion, proving the proposition.

□

We now are ready to prove the main theorem of the section:

Proof of Theorem 4.1.4. By Theorem 3.4.1, there is a walk b in $\mathbb{R}^{n+1}$ over $\{v_1, \ldots, v_n\}$, starting at the origin, whose boundary limit is homeomorphic to Z. By Lemma 4.3.1, we can choose a geodesic c in Γ, such that when we consider c as a quasi-geodesic, $\mu(f(c))$ stays within a bounded distance of b. Thus, $\mu(f(c))$ also has boundary limit Z in $\mathbb{R}^{n+1}$.

By Theorem 4.2.1, c stays within a bounded distance of a unique geodesic ray c' emanating from p. c' lies in the Euclidean half-space $c' \times \mathbb{R}^n \subset X$, which is closed and convex. We put coordinates on this half-space by identifying $(c'(t), \overrightarrow{x})$ with the point $(t, \overrightarrow{x}) \in \mathbb{R}^{n+1}$. Note that these coordinates are unchanged by the map μ.

Since c and c' stay within a bounded distance of each other, so do $f(c)$ and $f(c')$. Since μ does not increase distances, $\mu(f(c))$ and $\mu(f(c'))$ stay within a bounded distance of each other, thus the boundary limit of $\mu(f(c'))$ is also Z.

Since $f(c') \subset c' \times \mathbb{R}$, and the coordinates on $c' \times \mathbb{R}$ are unchanged under μ, this means the boundary limit of $f(c')$ is Z in $c' \times \mathbb{R}$. This is a closed, convex subset of X, thus its boundary is embedded in ∂X, and so we have that the boundary image of c' is homeomorphic to Z.

$\square$

References

[1] J Belk. *Thompson's Group F*. PhD thesis, Cornell University, 2004.

[2] I. Benjamini, R. Lyons, Y. Peres, and O. Schramm. Group-invariant percolation on graphs. *Geom. Funct. Anal.*, 9(1):29–66, 1999.

[3] M Bestvina. Local homology properties of boudaries of groups. *Michigan Math J.*, 43(1):123–139, 1996.

[4] J Block and S Weinberger. Aperiodic tilings, positive scalar curvature, and amenability of spaces. *J. Amer. Math. Soc.*, 5:907–918.

[5] P L Bowers and K Ruane. Boundaries of nonpositively curved groups of the form $G \times \mathbb{Z}$.

[6] M R Bridson and A Haefliger. *Metric Spaces of Nonpositive Curvature*. Springer-Verlag, Berlin, 1999.

[7] M Brin and C Squier. Groups of piecewise linear homeomorphisms of the real line. *Invent. Math.*, 79(4):485–498, 1985.

[8] S Buyalo. Geodesics in hadamard spaces. *St. Petersburg Math J.*, 10(2):293–313, 1998.

[9] C. Croke and B. Kleiner. Spaces with nonpositive curvature and their ideal boundaries. *Topology*, 3:549–556, 2000.

[10] C Croke and B Kleiner. The geodesic flow of a nonpositively curved graph manifold. *GAFA*, 12:479–545, 2002.

[11] S Ferry. Geometric topology notes. Available at http://math.rutgers.edu/ sferry/ ps/geotop.pdf.

[12] S Ferry. Homotopy, simple homotopy, and compacta. *Topology*, 19(2):101–110, 1980.

[13] S Ferry. Shape equivalence does not imply CE equivalence. *Proc. Amer. Math Soc.*, 80(1):154–156, 1980.

[14] E Følner. On groups with full Banach mean values. *Math. Scand.*, 3:243–254, 1955.

[15] W J Floyd J W Cannon and W R Parry. Introductory notes on Richard Thompson's groups. *Ensign. Math.*, 2(42):215–256, 1996.

[16] S Mardešlić and J Segal. *Shape Theory*. North-Holland, Amsterdam, 1982.

[17] C Mooney. All CAT(0) boundaries of a group of the form $H \times K$ are CE-equivalent. Available at http://arxiv.org/abs/0707.4316, 2007.

[18] C Mooney. Generalizing the Croke-Kleiner construction. Available at http://arxiv.org/abs/0807.4745, 2008.

[19] T Riley N Brady and H Short. *The Geometry of the Word Problem for Finitely Generated Groups*. Birkhäuser Verlag, Basel-Boston-Berlin, 2007.

[20] I. Namioka. Følner's conditions for amenable semi-groups. *Math. Scand.*, 15:18–28, 1964.

[21] A J Ol'shanskii and M.V. Sapir. Non-amenable finitely presented torsion-by-cyclic groups. *Publ. Math. Inst. Hautes Études Sci.*, 96:43–169, 2002.

[22] K Ruane. PhD thesis, Florida State University, 1996.

[23] D Savchuk. Some graphs related to Thompson's group F. To appear in Trends in Mathematics (Birkhäuser), available at http://arxiv.org/abs/0803.0043, 2008.

[24] R B Sher. Realizing cell-like maps in Euclidean space. *General Topology and Appl.*, 2:75–89, 1972.

[25] D Staley. Thompson's group F and uniformly finite homology. *Algebraic and Geometric Top.*, 9:2349–2360, 2009.

Vita

Daniel Staley

2010	Ph. D. in Mathematics, Rutgers University
2001-05	B. A. in Mathematics from University of California, Berkeley
2001	Graduated from Marin Academy high school.

2005-2007	Henry C. Torrey Fellow, Department of Mathematics, Rutgers University
2007-2009	Teaching assistant, Department of Mathematics, Rutgers University