TREES, DENDRITES, AND THE CANNON-THURSTON MAP

BY

ELIZABETH FIELD

DISSERTATION

Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy in Mathematics
in the Graduate College of the
University of Illinois at Urbana-Champaign, 2020

Urbana, Illinois

Doctoral Committee:

Professor Christopher J. Leininger, Chair
Professor Ilya Kapovich, City University of New York, Director of Research
Professor Nathan Dunfield
Professor Paul Schupp

Abstract

When $1 \to H \to G \to Q \to 1$ is a short exact sequence of three word-hyperbolic groups, Mahan Mitra (Mj) has shown that the inclusion map from H to G extends continuously to a map between the Gromov boundaries of H and G. This boundary map is known as the Cannon-Thurston map. In this context, Mitra associates to every point z in the Gromov boundary of Q an "ending lamination" on H which consists of pairs of distinct points in the boundary of H. We prove that for each such z, the quotient of the Gromov boundary of H by the equivalence relation generated by this ending lamination is a dendrite, that is, a tree-like topological space. This result generalizes the work of Kapovich-Lustig and Dowdall-Kapovich-Taylor, who prove that in the case where H is a free group and Q is a convex cocompact purely atoroidal subgroup of $\mathrm{Out}(F_N)$, one can identify the resultant quotient space with a certain $\mathbb{R}$-tree in the boundary of Culler-Vogtmann's Outer space.

To all of those whose support made this possible.

Acknowledgments

First and foremost, I owe an immeasurable amount of thanks and gratitude to my advisor, Ilya Kapovich. Thank you for always making me feel like I belonged, for believing in me from the start, for your endless patience, and for teaching me what it means to be a mathematician. Thank you for reading countless drafts of my paper and this thesis, and for the detailed feedback you provided. Words cannot express my gratitude to you for all that you have taught me and continue to teach me. I also owe an incredible deal of thanks to Chris Leininger. Thank you for being like a second advisor to me and for introducing me to new and exciting areas of math. Thank you for always having time for me, for your constant support, for listening to and encouraging my ideas, and for taking me seriously as a mathematician. I would also like to thank my prelim and thesis committee members: Nathan Dunfield, Paul Schupp, and Jeremy Tyson. Thank you for being a part of my mathematical community here at UIUC.

I would also like to thank the Southern Connecticut State University Department of Mathematics for providing me with the support and encouragement to pursue a graduate degree in mathematics. I am especially grateful to Therese Bennett, Len Brin, Braxton Carrigan, Joe Fields, Ross Gingrich, Klay Kruczek, and Ray Mugno for giving so generously of their time and energy, encouraging my potential, and helping me to achieve my dreams.

I am deeply grateful to all of my friends, near and far, for their love, friendship, support, and encouragement. Liz, thank you for sharing in so many joys and frustrations. Matthew, thank you for being there through everything. Simone and Hannah, thank you for being such caring and supportive friends and colleagues.

Thank you to my academic siblings, Heejoung, Neha, and Caglar. Heejoung, thank you for always talking math with me. Neha, thank you for being someone I could always look up to. Caglar, thank you for your friendship, your encouragement, and for helping to show me the way. Thank you also to my larger geometric group theory family at UIUC and beyond. Chris, Marissa, and Emily, thanks for always being willing to talk math and participate in seminars with me. Thanks to Yvon for being a wonderful friend and collaborator, and to Jacob for his friendship and many mathematical conversations, especially the semester I visited CUNY.

Thank you also to the many others who have mentored me and supported my mathematical development over the past several years. Thank you to Mark Bell, Spencer Dowdall, Joseph Maher, Jenya Sapir, Sam Taylor, and so many others.

Finally, a heartfelt thanks to my family. Mom and Dad, thank you for your endless support and encouragement. Evan, thank you for never letting me take myself too seriously and for always making me laugh.

Table of Contents

List of Figures

List of Symbols

$\mathbb{H}^n$	Hyperbolic space of dimension n
$\pi_1 S$	Fundamental group of S
F_N	Free group of rank N
Λ_z	Algebraic ending lamination associated to z
Γ_H	Cayley graph of the group H
$[\cdot,\cdot]_X$	Geodesic between two points in the space X
$(\cdot,\cdot)_z$	Gromov product relative to z
Σ_H	Alphabet of the group H
Σ_H^*	Set of all finite words over Σ_H
$\overline{w}$	Group element representing the word $w \in \Sigma_H^*$
$[h]_H$	Conjugacy class of the element $h \in H$
$\|h\|_H$	Length of the element $h \in H$
$[\gamma]$	Equivalence class of the geodesic ray γ
$\partial_z^g X$	Relative geodesic boundary of X with respect to the basepoint x
$\partial^g X$	Geodesic boundary of X
$[(x_n)]$	Equivalence class of the sequence (x_n) converging to infinity
∂X	Gromov boundary of X
$\mathrm{Mod}(S)$	Mapping class group of the surface S
$\mathcal{C}(S)$	Curve complex of the surface S
$\partial^2 H$	Double boundary of H
$\|\|w\|\|_T$	Translation length of the element $w \in F_N$ in the tree T
ϕ_g	Automorphism of H induced by conjugation by g

Chapter 1

Introduction

1.1 Motivation

In [13], Cannon and Thurston showed that when $M = (S \times [0,1])/((x,0) \sim (\phi(x),1))$ is the mapping torus of a closed, hyperbolic surface S by a pseudo-Anosov homeomorphism ϕ of S (so that M is a closed, hyperbolic 3-manifold), the inclusion $i : \pi_1 S \to \pi_1 M$ extends to a continuous, surjective, $\pi_1 S$-equivariant map $\partial i : \partial \pi_1 S \to \partial \pi_1 M$. As $\pi_1 S$ is quasi-isometric to the hyperbolic plane and $\pi_1 M$ is quasi-isometric to $\mathbb{H}^3$ by Milnor-Sv́arc, this map gives a surjective mapping of the circle at infinity bounding $\mathbb{H}^2$ onto the 2-sphere at infinity which bounds $\mathbb{H}^3$. As $\partial \pi_1 S$ is homeomorphic to $\mathbb{S}^1$ and $\partial \pi_1 M$ is homeomorphic to $\mathbb{S}^2$, ∂i is remarkably a space-filling Peano curve.

Associated to the homeomorphism ϕ are two geodesic ending laminations on S, Λ_ϕ^+ and Λ_ϕ^-, the stable and unstable laminations. Lift Λ_ϕ^+ and Λ_ϕ^- to two disjoint copies of the universal cover of S, and glue these two hyperbolic planes together along their boundary circle at infinity. Then, collapse the leaves and complementary components of the lifts of these laminations. Cannon and Thurston showed in [13] that the image of the equator of $\mathbb{S}^2$ along which the two hyperbolic planes were glued under this collapse is precisely the image of $\partial \pi_1 S$ inside $\partial \pi_1 M$ under the Cannon-Thurston map. It turns out that if you consider taking one copy of the hyperbolic plane with its boundary and collapsing the lift of one of the laminations, the resulting quotient space is a *dendrite*, or a tree-like topological space. This dendrite turns out to be the compactification of an $\mathbb{R}$-tree which is dual to the lifted lamination. So, the main result of [13] implies that the 2-sphere boundary of $\pi_1 M$ arises via a gluing of two dendrites.

1.2 History

The work of Cannon and Thurston in [13] has sparked much consideration since its circulation as a preprint in 1984. In modern terminology, if H and G are hyperbolic groups with $H \leq G$ and if the inclusion $i : H \to G$ extends to a continuous boundary map $\partial i : \partial H \to \partial G$, the map ∂i is called the *Cannon-Thurston map*. Such a map automatically exists and is injective when H is a quasiconvex (i.e. undistorted) subgroup of G, since in that case the inclusion map is a quasi-isometric embedding. The 1984 result of Cannon and Thurston gave the first non-trivial example of the existence of such a boundary map.

Mj (formerly Mitra) has since studied the existence of Cannon-Thurston maps in settings which involve distorted subgroups of hyperbolic groups [48, 49, 50]. In particular, Mitra showed in [48] that when

$$1 \to H \to G \to Q \to 1 \tag{*}$$

is a short exact sequence of infinite, hyperbolic groups, the Cannon-Thurston map $\partial i : \partial H \to \partial G$ exists. Since an infinite normal subgroup of infinite index in a word-hyperbolic group G is not quasiconvex [31], this result gives another non-trivial example of the existence of the Cannon-Thurston map. It is known by work of Kapovich and Short [40] that when H is an infinite normal subgroup of a hyperbolic group G, the limit set of H in ∂G is all of ∂G. As this limit set is precisely the image of ∂H under the map ∂i, it follows that the Cannon-Thurston map is surjective in this setting.

In [47], Mitra developed a theory of "algebraic ending laminations" for hyperbolic group extensions to describe when points in ∂H are identified under the Cannon-Thurston map. This work provides an analog of the theory of ending laminations in the context of pseudo-Anosov homeomorphisms of surfaces developed by Thurston [28]. To each point $z \in \partial Q$, Mitra associates an "algebraic ending lamination" on H, $\Lambda_z \subseteq \partial^2 H$, where $\partial^2 H = \{(p, q) \in \partial H \times \partial H \mid p \neq q\}$. The lamination Λ_z is determined by sequences of elements of Q which converge to $z \in \partial Q$ and records limiting information about conjugacy classes which come from applying sequences of automorphisms of H to non-torsion elements $h \in H$. The main result of [48] states that two distinct points $p, q \in \partial H$ are identified under the Cannon-Thurston map if and only if there exists some $z \in \partial Q$ for which (p, q) is a leaf of the ending lamination Λ_z.

If H is a torsion-free, infinite-index, word-hyperbolic, normal subgroup of a word-hyperbolic group G, it follows from combined work of Mosher [57], Paulin [61], Rips-Sela [62], and Bestvina-Feighn [5] that H must be a free product of free groups and surface groups. For a brief explanation of this, see [47]. So, it makes sense to first consider the settings where H is a free group or a surface group. Suppose H is the fundamental group of a closed hyperbolic surface S and Γ is a convex cocompact subgroup of $\mathrm{Mod}(S)$ (and hence Γ is word-hyperbolic [27]). Then, Γ naturally gives rise to a short exact sequence $1 \to H \to E_\Gamma \to \Gamma \to 1$ coming from Birman's short exact sequence for S. Hamenstädt has shown that in this setting, the extension group E_Γ is hyperbolic and the orbit map of Γ into the curve complex of S is a quasi-isometric embedding [33]. Since the boundary of the curve complex consists of ending laminations on S [42], it follows that to each point $z \in \partial\Gamma$, there is an associated ending lamination L_z on the surface S. Mj and Rafi [55] showed that the algebraic ending lamination Λ_z is the same as the diagonal closure of the surface lamination L_z. To each such ending lamination L_z, there is an associated dual $\mathbb{R}$-tree T_z which can be constructed by lifting L_z to $\widetilde{S}$ and collapsing each leaf and complementary component to a point. For more details, see for example [4, 19].

In the free group setting, Mitra's algebraic ending laminations for hyperbolic extensions of free groups are closely related to the theory of algebraic laminations on free groups developed by Coulbois, Hilion, and Lustig in [19]. For any subgroup $\Gamma \leq \mathrm{Out}(F_N)$, the full preimage of Γ under the quotient map $\mathrm{Aut}(F_N) \to \mathrm{Out}(F_N)$, also denoted by E_Γ, fits into the short exact sequence $1 \to F_N \to E_\Gamma \to \Gamma \to 1$. The main result of [23] states that whenever $\Gamma \leq \mathrm{Out}(F_N)$ is a convex cocompact and purely atoroidal subgroup, the extension group, E_Γ, is word-hyperbolic. In [22], Dowdall, Kapovich, and Taylor study the fibers of the Cannon-Thurston map $\partial i : \partial F_N \to \partial E_\Gamma$ in the case where $\Gamma \leq \mathrm{Out}(F_N)$ is convex cocompact and purely atoroidal. Since Γ is convex cocompact, the orbit map to the free factor complex, $\mathcal{F}$, is a quasi-isometric embedding [35] and hence, extends to a continuous embedding $\partial\Gamma \to \partial\mathcal{F}$. By work of Bestvina-Reynolds [8] and Hamenstädt [34], $\partial\mathcal{F}$ consists of equivalence classes of arational F_N-trees. Therefore, there is a class of arational F_N-trees, T_z, associated to each point $z \in \partial\Gamma$. Moreover, each such tree T_z comes equipped with the "dual lamination" $L(T_z)$, defined by Coulbois, Hilion, and Lustig in [19]. A key result of [22] states that for each $z \in \partial\Gamma$, $\Lambda_z = L(T_z)$. This theorem extends the result of Kapovich and Lustig [39] who prove this equality for the specific case where $\Gamma = \langle\varphi\rangle$ is the cyclic group generated by a fully irreducible, atoroidal automorphism of F_N.

Given an $\mathbb{R}$-tree T, Coulbois, Hilion, and Lustig define a suitable topology on $\widehat{T} = \overline{T} \cup \partial T$, where $\overline{T}$ denotes the metric completion of T and ∂T is the Gromov boundary. This topology, known as the "observers' topology", is coarser than the Gromov topology and ensures that $\widehat{T}$ is compact. Recall that a *dendrite* is a compact, connected, locally connected metrizable space which contains no simple closed curves. Coulbois, Hilion, and Lustig show that for any $\mathbb{R}$-tree T, $\widehat{T}$ equipped with the "observers' topology" is a dendrite, as well as a proper, Hausdorff metric space [17]. Dendrites naturally arise from this compactification of simplicial trees, but in general can be much more complicated spaces such as certain Julia sets. Combining the result of [22] with a general result from [17] implies that for each $z \in \partial\Gamma$, for Γ a convex-cocompact and purely atoroidal subgroup of $\mathrm{Out}(F_N)$, $\partial F_N/\Lambda_z$ equipped with the quotient topology is homeomorphic to $\widehat{T}_z$ equipped with the "observers' topology". In particular, $\partial F_N/\Lambda_z$ is homeomorphic to a dendrite. Here, $\partial F_N/\Lambda_z$ means the quotient space of ∂F_N by the equivalence relation on ∂F_N generated by $\Lambda_z \subseteq \partial F_N \times \partial F_N$. The main result of this thesis extends this result to the general setting of hyperbolic group extensions.

1.3 Statement of Main Results

Theorem A. *Let $1 \to H \to G \to Q \to 1$ be a short exact sequence of infinite, finitely generated, word-hyperbolic groups. For each $z \in \partial Q$, let Λ_z denote the algebraic ending lamination on H associated to z. Then for each $z \in \partial Q$, the space $\partial H/\Lambda_z$ is homeomorphic to a dendrite.*

Let $P : \Gamma_G \to \Gamma_Q$ denote the map which is induced by the quotient map $P : G \to Q$, where Γ_G and Γ_Q denote the Cayley graphs of G and Q, respectively. Let $z \in \partial Q$ be arbitrary and take any $z' \in \partial Q$ with $z' \neq z$. Consider a bi-infinite geodesic $\gamma = (z', z) \subseteq \Gamma_Q$ and define the space $X(\gamma)$ to be the subgraph of Γ_G given by $X(\gamma) = P^{-1}(\gamma)$. We show that $X(\gamma)$ satisfies the properties of being a metric graph bundle, as defined by Mj-Sardar [56], and that $X(\gamma)$ is hyperbolic (Proposition 4.5). We go on to show that $X(\gamma)$ also satisfies the properties of being a bi-infinite hyperbolic stack, as defined by Bowditch [10], with fibers being copies of the Cayley graph of H, Γ_H, (Proposition 4.6). We then look at the semi-infinite stack $X(\gamma)^+$ which lies over the geodesic ray $\gamma^+ = [z_0, z)$, where $z_0 \in (z', z)$. We denote the natural "0-th slice" map from $\Gamma_H \to X(\gamma)^+$ by i_γ^+, and also refer to the continuous extension of this map to $\partial i_\gamma^+ : \partial H \to \partial X(\gamma)^+$ as the Cannon-Thurston map. We then show the following.

Theorem B. *Let $1 \to H \to G \to Q \to 1$ be a short exact sequence of infinite, finitely generated, word-hyperbolic groups. Let $z, z' \in \partial Q$ be distinct and let $\gamma \subseteq \Gamma_Q$ be a bi-infinite geodesic in Γ_Q between z and z'. Let $i_\gamma^+ : \Gamma_H \to X(\gamma)^+$ be the inclusion of Γ_H into the semi-infinite stack $X(\gamma)^+$ over $\gamma^+ = [z_0, z)$ for some $z_0 \in \gamma$, and let $i_\gamma : \Gamma_H \to X(\gamma)$ be the inclusion of Γ_H into the bi-infinite stack $X(\gamma)$ over γ. Then,*

1. *the Cannon-Thurston map* $\partial i_\gamma^+ : \partial H \to \partial X(\gamma)^+$ *is surjective; and*

2. *the Cannon-Thurston map* $\partial i_\gamma : \partial H \to \partial X(\gamma)$ *is surjective.*

Using the work of Mitra from [47], we then show that the following holds.

Theorem C. *Let* $1 \to H \to G \to Q \to 1$ *be a short exact sequence of infinite, finitely generated, word-hyperbolic groups. Let* $z, z' \in \partial Q$ *be distinct and let* $\gamma \subseteq \Gamma_Q$ *be a bi-infinite geodesic between* z *and* z'*. Let* $i_\gamma^+ : \Gamma_H \to X(\gamma)^+$ *be the inclusion of* Γ_H *into the semi-infinite stack* $X(\gamma)^+$ *over* $\gamma^+ = [z_0, z)$ *for some* $z_0 \in \gamma$*, and let* $\partial i_\gamma^+ : \partial H \to \partial X(\gamma)^+$ *be the Cannon-Thurston map.*

Then for any distinct $u, v \in \partial H$*, we have* $\partial i_\gamma^+(u) = \partial i_\gamma^+(v)$ *if and only if* (u, v) *is a leaf of the ending lamination* Λ_z*.*

To finish the proof of Theorem A, note that by a general result of Bowditch [10], $\partial X(\gamma)^+$ is a dendrite (Proposition 6.1). Theorem C implies that the Cannon-Thurston map $\partial i_\gamma^+ : \partial H \to \partial X(\gamma)^+$ quotients through to an injective map $\tau_z : \partial H/\Lambda_z \to \partial X(\gamma)^+$. Since ∂i_γ^+ is continuous, the map τ_z is also continuous. By Theorem B, τ_z is also surjective. Thus, $\tau_z : \partial H/\Lambda_z \to \partial X(\gamma)^+$ is a continuous bijection between two compact topological spaces, where $\partial X(\gamma)^+$ is Hausdorff. Therefore, τ_z is a homeomorphism.

1.4 Outline

In Chapter 2, we begin with some preliminary background on hyperbolic metric spaces, groups, and the boundaries of these spaces. We then provide more details and history about the Cannon-Thurston map. This history is followed by preliminaries on laminations, especially algebraic laminations, which provides motivation for the results contained in the remainder of this thesis. This chapter concludes with background on metric graph bundles and stacks of spaces: the two main objects used in the proof of the main results mentioned earlier.

In Chapter 3, we provide proofs of facts about hyperbolic metric spaces and groups which will be used later on. In particular, this chapter contains results about the concatenation of geodesics in hyperbolic metric spaces, as well as results regarding a generalization of the notion of a word in the free group being (almost) cyclically reduced.

The main purpose of Chapter 4 is to give the proof of Theorem B. To do this, we first apply the work of Bowditch [10] and Mj-Sardar [56] to show that the space $X(\gamma) \subseteq \Gamma_G$ is a bi-infinite, hyperbolic stack. We then utilize the hyperbolic stack structure of $X(\gamma)$ to prove Theorem B.

Chapter 5 begins by introducing Mitra's definition of the algebraic ending lamination Λ_z from [47], and then builds up to the proof of Theorem C. Many of the results in this chapter mirror results from [47], but are stated (and proved) in our more specific setting of dealing with the subspace $X(\gamma)$, rather than all of Γ_G.

In Chapter 6, we combine the results from Chapters 4 and 5 to prove our main structural result about the Cannon-Thurston map, Theorem A.

Finally, in Chapter 7, we introduce various open problems about the Cannon-Thurston map which are related to the main result of this thesis.

Chapter 2

Background

2.1 Hyperbolic metric spaces.

In this section, we will discuss some basic definitions and facts about hyperbolic metric spaces. For general references on hyperbolic spaces, see [1, 11, 15, 30, 31, 32]. Let (X, d) be a geodesic metric space. For any $x, y \in X$, we will denote a geodesic between x and y by $[x, y]_X$, or by $[x, y]$ if the space is clear. Given any three points $x, y, z \in X$, the *Gromov product* of x and y relative to z is defined to be

$$(x, y; X)_z := \frac{1}{2}\big(d(x, z) + d(y, z) - d(x, y)\big).$$

If the space X is clear, we will simply write $(x, y)_z$ for $(x, y; X)_z$.

Definition 2.1. Let $\delta \geq 0$. A geodesic metric space (X, d) is called δ-*hyperbolic* if for any $x, y, z \in X$ and any geodesics $[z, x]$ and $[z, y]$ in X, the following holds. Let $x' \in [z, x]$ and $y' \in [z, y]$ be any points such that $d(z, x') = d(z, y') \leq (x, y)_z$. Then, $d(x', y') \leq \delta$. A geodesic metric space (X, d) is said to be *hyperbolic* if it is δ-hyperbolic for some $\delta \geq 0$.

Note that this property implies that for any geodesic triangle $\Delta = [x, y] \cup [y, z] \cup [z, x]$ in X, each side of Δ is contained in the δ-neighborhood of the union of the other two sides. See [1] and [11] for more details and other equivalent definitions of hyperbolicity. Note that in a hyperbolic metric space, the Gromov product $(x, y)_z$ measures how closely the geodesics $[z, x]$ and $[z, y]$ travel.

Definition 2.2. Let (X, d_X) and (Y, d_Y) be metric spaces, and let $\kappa \geq 1$ and $\epsilon \geq 0$. A map $f : X \to Y$ is said to be a (κ, ϵ)-*quasi-isometric embedding* if for all $x_1, x_2 \in X$,

$$\frac{1}{\kappa} d_X(x_1, x_2) - \epsilon \leq d_Y(f(x_1), f(x_2)) \leq \kappa d_X(x_1, x_2) + \epsilon.$$

The map f is said to be a $(\kappa, epsilon)$-*quasi-isometry* if there additionally exists a constant

$C \geq 0$ such that for all $y \in Y$, there exists some $x \in X$ such that $d_Y(y, f(x)) \leq C$. In this case the metric spaces X and Y are said to be *quasi-isometric.* A (κ, ϵ)-*quasigeodesic* in a metric space (X, d) is the image of a (κ, ϵ)-quasi-isometric embedding $f : I \to X$, where $I \subseteq \mathbb{R}$ is a sub-interval. The map f itself is also referred to as a (κ, ϵ)-quasigeodesic.

If X and Y are geodesic metric spaces and $f : X \to Y$ is a quasi-isometry, then it is known that X is hyperbolic if and only if Y is hyperbolic [15, 30].

Definition 2.3. Let (X, d) be a geodesic metric space. A map $e : \mathbb{N} \to \mathbb{R}$ is said to be a *divergence function* for X if for all $T, t \in \mathbb{N}$, all basepoints $x \in X$, and all geodesics $\gamma_1 : [0, a_1] \to X$ and $\gamma_2 : [0, a_2] \to X$ in X with $\gamma_1(0) = \gamma_2(0) = x$, the following condition holds.

If $T + t \leq \min\{a_1, a_2\}$ and $d(\gamma_1(T), \gamma_2(T)) \geq e(0) > 0$, then any path joining $\gamma_1(T + t)$ to $\gamma_2(T + t)$ and lying outside the $(T + t)$-ball around x has length greater than $e(t)$.

It is also known that geodesics "diverge exponentially" in a hyperbolic metric space.

Proposition 2.4 ([11] III.H Proposition 1.25)**.** *Let X be a δ-hyperbolic, geodesic metric space. Then X has an exponential divergence function.*

It turns out that having an exponential divergence function is actually a characterization of hyperbolicity, as it is known that if a geodesic metric space X has an exponential divergence function, then X is hyperbolic (see Proposition 2.20 in [1]). In this thesis, we will use the following more general statement which shows that quasigeodesics which do not necessarily start at the same basepoint also diverge exponentially in a hyperbolic metric space.

Proposition 2.5 (Mitra [47] Proposition 2.4)**.** *Let (X, d) be a δ-hyperbolic, geodesic metric space. Given $K \geq 1$, $\epsilon \geq 0$, and $\alpha \geq 0$, there exist $b > 1$, $A > 0$, and $C > 0$ such that the following holds:*

If r_1 and r_2 are two (K, ϵ)-quasigeodesics in X with $d(r_1(0), r_2(0)) \leq \alpha$ and there exists $T \geq 0$ with $d(r_1(T), r_2(T)) \geq C$, then any path joining $r_1(T+t)$ to $r_2(T+t)$ and lying outside the union of the $\frac{T+t-1}{K+\epsilon}$-balls around $r_1(0)$ and $r_2(0)$ has length greater than Ab^t for all $t \geq 0$.

Let $\mathcal{N}_r(U)$ denote the closed r-neighborhood around a subset U of X. It is also known that in a hyperbolic metric space, any quasigeodesic stays near the geodesic between its endpoints:

Proposition 2.6 ([32] 7.2 A; [15] 3.1.3; [30] 5.6, 5.11)**.** *Let* (X, d) *be a* δ*-hyperbolic metric space and let* $x, y \in X$*. For any* $\kappa \geq 1$ *and* $\epsilon \geq 0$*, there exists* $L = L(\delta, \kappa, \epsilon) \geq 0$ *such that if* α *is a* (κ, ϵ)*-quasigeodesic between* x *and* y*, then for any geodesic* $\beta = [x, y]$*, we have that* $\alpha \subset \mathcal{N}_L(\beta)$ *and* $\beta \subset \mathcal{N}_L(\alpha)$*.*

2.2 Hyperbolic groups.

In the last section, we discussed the notion of a hyperbolic metric space. The first step in our discussion of hyperbolic groups is to define a metric space which is canonically associated to a group and which captures the intrinsic geometry of the group.

Definition 2.7. Let $H = \langle S \mid R \rangle$ be a finitely generated group with finite generating set S. The *Cayley graph* of H with respect to S, denoted $\mathrm{Cay}(H, S)$, is the following graph. There is a vertex for each group element $h \in H$. For all $h_1, h_2 \in H$, there is a directed edge of length 1 from h_1 to h_2 labeled by $s \in S$ if and only if $h_2 = h_1 s$. There is a natural metric on the Cayley graph, denoted by d_H, which results from treating each directed edge of length 1 as an undirected edge of length 1.

A *path* in $\mathrm{Cay}(H, S)$ is a sequence of vertices, $h_1, h_2, \ldots, h_n$, where $h_i \in H$ and $h_{i-1}^{-1} h_i \in S \cup S^{-1}$. Note that edges in a path may be traversed in the opposite orientation. We make the convention that if an edge in a path is traversed in the direction opposite its orientation, then we read the inverse of the label of the edge.

Note that if S is a finite generating set, $\mathrm{Cay}(H, S)$ is a proper, geodesic metric space. For each $h \in H$, left-translation by h induces an action by isometries of $\mathrm{Cay}(H, S)$ which preserves the label and orientation of each edge.

The metric on the Cayley graph induces a natural metric on the group H known as the *word-metric*, where the distance in H between two elements $h, h' \in H$ is exactly the number of elements in $S \cup S^{-1}$ it takes to get from h to h'. This distance is the same as the length of a shortest edge-path in $\mathrm{Cay}(H, S)$ between h and h'. If S' is another finite generating set for H, then it is known that $\mathrm{Cay}(H, S)$ is quasi-isometric to $\mathrm{Cay}(H, S')$. Here, the quasi-isometry constants depend on the number of elements in the generating set S' it takes to express any element in the generating set S. As hyperbolicity is a quasi-isometry invariant for geodesic metric spaces [15, 30], the following notion is well-defined.

Definition 2.8. A finitely generated group H is said to be *word-hyperbolic* if for some, equivalently any, finite generating set of H, there exists $\delta \geq 0$ such that the Cayley graph of H with respect to the word-metric is δ-hyperbolic.

Let H be a word-hyperbolic group and fix a finite, symmetric generating set S_H for H. Denote $\mathrm{Cay}(H, S_H)$ by Γ_H, let d_H, or simply d, denote the word-metric, and let $\Sigma_H := S_H \cup S_H^{-1}$ denote the alphabet of H. A *word* w over the alphabet Σ_H is an expression $s_1 \cdots s_n$, where $s_i \in \Sigma_H$ and $n \geq 0$ (the case $n = 0$ represents the empty word). We will denote the set of all finite words over Σ_H by Σ_H^*, and will think of a word as the label of some (not necessarily geodesic) path in Γ_H.

If $w \in \Sigma_H^*$ is the label of some path in Γ_H from a vertex a to b, then we will denote the group element $a^{-1}b \in H$ representing the word w by $\overline{w}$. Given any element $h \in H$, we will denote the conjugacy class of h in H by $[h]_H$ (or simply by $[h]$ if the ambient group is clear). For a word $w \in \Sigma_H^*$, $|w|_H$ denotes the length of any path labeled by w in Γ_H. The length of an element $h \in H$, also denoted by $|h|_H$, is defined to be the length of any geodesic from the identity 1_H to the vertex h in Γ_H. We will drop the subscript if the group we are working in is clear.

2.3 Boundaries of hyperbolic spaces and groups.

In this section, we will discuss basic definitions and facts about the Gromov boundary of hyperbolic metric spaces and groups. For more details on much of what is discussed here, see [38]. Let (X, d) be a hyperbolic metric space. We begin this section by providing two different descriptions of the boundary of X. More details can be found in [38]. The first definition we give involves equivalence classes of geodesic rays which emanate from a fixed basepoint $x \in X$. If $\gamma_1 : [0, \infty) \to X$ and $\gamma_2 : [0, \infty) \to X$ are two geodesic rays in X, we say that γ_1 and γ_2 are *equivalent*, and denote this by $\gamma_1 \sim \gamma_2$, if there is some $K > 0$ such that for all $t \geq 0$, $d(\gamma_1(t), \gamma_2(t)) \leq K$. Note that for this definition of equivalence, we do not require the rays γ_1 and γ_2 to share a common basepoint.

Definition 2.9. Let (X, d) be a hyperbolic metric space, and fix a basepoint $x \in X$. The *relative geodesic boundary* of X with respect to the basepoint x is defined to be the set

$$\partial_x^g X := \{[\gamma] \mid \gamma : [0, \infty) \to X \text{ is a geodesic ray in } X \text{ with } \gamma(0) = x\}.$$

The *geodesic boundary* of X is defined to be the set

$$\partial^g X := \{[\gamma] \mid \gamma : [0, \infty) \to X \text{ is a geodesic ray in } X\}.$$

Let X be a proper, geodesic, hyperbolic metric space. Given a geodesic ray γ in X with basepoint x, note that $[\gamma]$ is in both $\partial_x^g X$ and $\partial^g X$. As X is proper, for all $[\gamma] \in \partial_x^g X$ and for all $y \in X$, there exists γ' with basepoint y such that $\gamma' \in [\gamma]$. So, there is a bijection $i_{xy} : \partial_x^g X \to \partial_y^g X$ which sends the equivalence class of a ray γ based at x in $\partial_x^g X$ to the equivalence class of a ray γ' based at y in $\partial_y^g X$, where $\gamma \sim \gamma'$. Now, as in [38], define a map $i_x : \partial_x^g X \to \partial^g X$ by $i_x([\gamma]) := [\gamma]$ which sends the equivalence class containing γ in $\partial_x^g X$ to the equivalence class of γ in $\partial^g X$. It is known that when X is proper, this map i_x gives a bijection between the relative geodesic boundary and the geodesic boundary of X (Proposition 2.10 [38]). We will later define a topology on $\partial_x^g X$ and show that when X is a proper, geodesic, hyperbolic metric space, this topology induces a canonical topology on $\partial^g X$. First, we provide another description of a boundary of X which will turn out to be homeomorphic to $\partial^g X$ when X is proper.

Let X be a hyperbolic metric space. A sequence of points $(x_n)_{n\in\mathbb{N}} \in X$ is said to *converge to infinity* if for some basepoint $x \in X$,

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

It is known that this definition is independent of basepoint. Two sequences (x_n) and (y_n) in X which converge to infinity are said to be *equivalent* if

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

If (x_n) and (y_n) are equivalent, we write $(x_n) \sim (y_n)$, and we denote the equivalence class of a sequence (x_n) converging to infinity by $[(x_n)]$. We note again that this equivalence is independent of chosen basepoint.

Definition 2.10. Let (X, d) be a hyperbolic metric space. The *sequential boundary* of X is defined to be

$$\partial X := \{[(x_n)] \mid (x_n) \text{ is a sequence converging to infinity in } X\}.$$

Note that if $\gamma : [0, \infty) \to X$ is a geodesic ray in X, then $(\gamma(n))_{n\geq 0}$ determines a sequence which converges to infinity in X. So, as in [38], define a map $i : \partial^g X \to \partial X$ by $i([\gamma]) := [(\gamma(n))_{n\geq 0}]$. As with the map $i_x : \partial_x^g X \to \partial X$, when X is a proper metric space, the map i gives a bijection between $\partial^g X$ and ∂X (Proposition 2.10 [38]).

We will now give a definition of the topologies on the geodesic and sequential boundaries of X.

Definition 2.11. Let (X, d) be a hyperbolic metric space, and fix a basepoint $x \in X$. For any $p \in \partial_x^g X$ and any $r \geq 0$, define the set

$$\begin{aligned} V_x(p, r) := \{ q \in \partial_x^g X \mid & \text{ for some geodesic rays } \gamma_1, \gamma_2 \text{ based at } x \text{ with} \\ & [\gamma_1] = p \text{ and } [\gamma_2] = q, \text{ we have } \liminf_{t \to \infty} (\gamma_1(t), \gamma_2(t))_x \geq r \}. \end{aligned}$$

The topology on the relative geodesic boundary $\partial_x^g X$ is then generated by the neighborhood basis $\{V_x(p, r) \mid r \geq 0\}$ for any point $p \in \partial_x^g X$. We note that this topology does not depend on the choice of basepoint x.

Note that as the Gromov product tracks how long rays emanating from the same point stay 2δ-close, the set $V(p, r)$ consists of equivalence classes of geodesic rays based at x which travel 2δ-close to a ray representing the point p for (approximately) the distance r. So, two points $a, b \in \partial_x^g X$ are "close" in this boundary at infinity if rays based x to the points a and b stay 2δ-close for a long time. The topology on the sequential boundary is defined in a similar manner.

Definition 2.12. Let (X, d) be a hyperbolic metric space, and fix a basepoint $x \in X$. For any $p \in \partial X$ and $r \geq 0$, define the set

$$\begin{aligned} U_x(p, r) := \{ q \in \partial X \mid & \text{ there exist sequences } (x_n) \text{ and } (y_n) \text{ with } [(x_n)] = p \\ & \text{and } [(y_n)] = q \text{ such that } \liminf_{i,j \to \infty} (x_i, y_j)_x \geq r \}. \end{aligned}$$

The topology on the sequential boundary ∂X is then generated by the neighborhood basis $\{U_x(p, r) \mid r \geq 0\}$. We note again that this topology does not depend on the choice of basepoint x. To get a topology on $\widehat{X} = X \cup \partial X$, define for each $p \in \partial X$ and $r \geq 0$ the additional sets

$$\begin{aligned} U_x'(p, r) := \{ y \in X \mid & \text{ for some sequence } (x_n) \text{ with} \\ & [(x_n)] = p, \text{ we have } \liminf_{i \to \infty} (x_i, y)_x \geq r \}. \end{aligned}$$

For each $p \in \partial X$ put the basis of neighborhoods for $p \in \widehat{X}$ to be $\{U_x(p, r) \cup U_x'(p, r) \mid r \geq 0\}$. For each $y \in X$, we use the same neighborhood basis as in X.

It is known that for a proper, geodesic hyperbolic metric space X, the topologies on $\partial_x^g X$ and ∂X are the same. In particular, for all $x \in X$, the map $i \circ i_x : \partial_x^g X \to \partial X$ is a homeomorphism, provided that X is proper. Further, it is known that if X is proper, then the spaces ∂X and $\widehat{X} = X \cup \partial X$ are compact (Proposition 2.14 [38]).

Definition 2.13. If (X, d) is a proper, geodesic hyperbolic metric space, the boundary ∂X, as defined and topologized above, is called that *Gromov boundary* of X, and $\widehat{X} = X \cup \partial X$ is known as the *Gromov compactification* of X.

2.3.1 More on hyperbolic groups

We note that if G is a hyperbolic group, then the Cayley graph of G is a proper, geodesic metric space. So, the above conclusions hold, in particular, for hyperbolic groups. We noted in the previous section that if S and S' are finite, symmetric generating sets for G, then $\mathrm{Cay}(G, S)$ is quasi-isometric to $\mathrm{Cay}(G, S')$. It is known that a quasi-isometry of hyperbolic metric spaces extends to a homeomorphism between the Gromov boundaries of these spaces. Therefore, if G is a hyperbolic group, we define the *Gromov boundary* of G, denoted by ∂G, to be the Gromov boundary of some (any) Cayley graph of G with respect to a finite generating set.

Let G be a hyperbolic group and let Γ_G be the Cayley graph of G. Each element $g \in G$ acts on Γ_G by left-multiplication: $x \mapsto gx$. This left-translation extends to an action on ∂G which is a homeomorphism. In a hyperbolic group G, any infinite-order element $g \in G$ acts as a *loxodromic* (or *hyperbolic*) isometry on Γ_G. In this case, there are exactly two *poles* in ∂G which are fixed by g: $g^\infty := \lim_{n\to\infty} g^n$ and $g^{-\infty} := \lim_{n\to\infty} g^{-n}$. Additionally, for any $x \in \Gamma_G$, the orbit map $\mathbb{Z} \to \Gamma_G$ given by $n \mapsto g^n x$ is a quasi-isometric embedding with $\lim_{n\to\infty} g^n x = g^\infty$ and $\lim_{n\to\infty} g^{-n} x = g^{-\infty}$, and is called an *axis* of g.

Proposition-Definition 2.14 ([32])**.** *Let G be a hyperbolic group. Then, exactly one of the following holds:*

1. *G is finite, and ∂G is empty;*
2. *G contains $\mathbb{Z}$ as a finite-index subgroup, and ∂G consists of two points; or*
3. *G contains F_2 as a subgroup, and ∂G is an uncountably infinite, compact, metrizable space, with no isolated points.*

If (1) or (2) hold, we say that G is elementary, and if (3) holds, G is said to be non-elementary.

If X is a hyperbolic metric space, a subset $Y \subseteq X$ is said to be *C-quasiconvex* in X if for any $y_1, y_2 \in Y$, any geodesic connecting y_1 to y_2 in X is contained in the C-neighborhood of Y. If G is a hyperbolic group, a subgroup $H \leq G$ is *quasiconvex* in G if some (any) Cayley graph Γ_G of G, the subset of Γ_G corresponding to the subgroup H is a quasiconvex subset of Γ_G. It is known that a finitely generated subgroup H of a hyperbolic group G is quasiconvex in G if any only if for any finite generating sets, the inclusion map $i : \Gamma_H \to \Gamma_G$ is a quasi-isometric embedding [38]. It is also known that if G is a hyperbolic group and $H \leq G$ is a quasiconvex subgroup, then H is finitely generated and hyperbolic.

2.4 The Cannon-Thurston map

Definition 2.15. Let H and G be word-hyperbolic groups with $H \leq G$. If the inclusion map $i : H \to G$ extends to a (necessarily unique and H-equivariant) continuous map between the Gromov boundaries of H and G, $\partial i : \partial H \to \partial G$ such that the extended map $i \cup \partial i : H \cup \partial H \to G \cup \partial G$ is continuous, the map ∂i is called the *Cannon-Thurston map*.

We remark that H and G are hyperbolic groups with H a non-elementary subgroup of G, then if the Cannon-Thurston map $\partial i : \partial H \to \partial G$ exists, then for each $h \in H$ of infinite order, $\partial i(h^\infty) = \lim_{n\to\infty}(i(h))^n$. Additionally, in this setting where H is non-elementary, if $f : \partial H \to \partial G$ is a continuous and H-equivariant map, then f must be the Cannon-Thurston map (Proposition 2.12 [37]).

The term "Cannon-Thurston map" is also widely used in a related sense in the theory of Kleinian groups. Here, we let (S, ρ) be a complete hyperbolic surface of finite volume (possibly with cusps) and let $h : \pi_1 S \to \mathrm{PSL}_2(\mathbb{C}) = \mathrm{Isom}_+(\mathbb{H}^3)$ be a discrete, faithful representation. If the lift $\widetilde{h} : (\widetilde{S}, \widetilde{\rho}) = \mathbb{H}^2 \to \mathbb{H}^3$ extends to a continuous, $\pi_1 S$-equivariant map $j : \partial\mathbb{H}^2 \to \partial\mathbb{H}^3$, the map j is also called a Cannon-Thurston map. In [52], Mj shows that the map j always exists in this setting, provided that h has no "accidental parabolics". In [53], Mj shows that the Cannon-Thurston map exists for arbitrary finitely generated Kleinian groups without parabolics. For other related results and further generalizations about the Cannon-Thurston map in the context of Kleinian groups, see [46], [9], and [43], for example.

The definition of the Cannon-Thurston map in the setting of hyperbolic groups can be naturally extended to the setting where X and Y are hyperbolic spaces for which there is an inclusion map $i : X \to Y$. Note also that if H is a quasiconvex subgroup of a hyperbolic group G, then the Cannon-Thurston map trivially exists (and is injective on ∂H), since geodesic rays with bounded Hausdorff distance in Γ_H will map to quasi-geodesic rays which have bounded Hausdorff distance in Γ_G. In fact, it is known that the Cannon-Thurston map $\partial i : \partial H \to \partial G$ is injective if and only if H is a quasiconvex subgroup of G (see Proposition 2.13 [37] and Lemma 2.1 [51]).

In [47], Mitra gives the following characterization for when the Cannon-Thurston map $\partial i : \partial H \to \partial G$ exists. We include the proof for completeness.

Lemma 2.16 (Mitra [47], Lemma 2.1)**.** *Let H and G be word-hyperbolic groups such that $H \leq G$. The map $i : \Gamma_H \to \Gamma_G$ extends continuously to the Cannon-Thurston map $\partial i : \partial H \to \partial G$ if and only if the following condition is satisfied:*

For all $M > 0$, there exists $N > 0$ such that if $\gamma \subset \Gamma_H$ is a geodesic which lies outside the N-ball around the identity in Γ_H, then any geodesic in Γ_G joining the endpoints of $i(\gamma)$ lies outside the M-ball around the identity in Γ_G.

Proof. Let $p \in \partial H$ and let $\{a_n\}, \{b_n\} \in \Gamma_H$ be two sequences which converge to p and such that for all $n \in \mathbb{N}$, $a_n \neq b_n$. Since $\{a_n\} \to p \in \partial H$ and $b_n \to p \in \partial H$, we may pass to a subsequence $\{a_{n_j}\}$ and $\{b_{n_j}\}$ such that the geodesic segment $\gamma_{n_j} := [a_{n_j}, b_{n_j}]_H$ is outside the j-ball around the identity in Γ_H. Suppose that the above condition is not satisfied. Then there exists some $M > 0$ such that for all $j > 0$, some geodesic in Γ_G between $i(a_{n_j})$ and $i(b_{n_j})$ is contained in the M-ball about the identity in Γ_G. But, this implies that that $\lim_{j\to\infty} i(a_{n_j}) \neq \lim_{j\to\infty} i(b_{n_j})$. Hence if the above condition is not satisfied, then the Cannon-Thurston map does not exist.

Now, suppose that the Cannon-Thurston map does not exist. Then, there exists a point $p \in \partial H$ and sequences $\{a_m\}$, $\{b_m\} \in \Gamma_H$ such that in $\widehat{\Gamma_H}$, $\lim_{m\to\infty} a_m = \lim_{m\to\infty} b_m = p$, but such that in $\widehat{\Gamma_G}$, $\lim_{m\to\infty} i(a_m) = u \in \partial G$ and $\lim_{m\to\infty} i(b_m) = v \in \partial G$, with $u \neq v$. Let $\gamma_m := [a_m, b_m]_H$ denote a geodesic in Γ_H between a_m and b_m. Since $a_m \to p \in \partial H$ and $b_m \to p \in \partial H$, for each $m \in \mathbb{N}$ there exists some constant $N_m > 0$ such that γ_m lies outside the N_m-ball around the identity in Γ_H. Now, as $u, v \in \partial G$ with $u \neq v$, there exists some constant $M > 0$ such that any geodesic between $u, v \in \partial G$ must pass through the M-ball about the identity in Γ_G. As $i(a_m) \to u$ and $i(b_m) \to v$, there exists some index $m_0 > 0$ and some constant $M' > M$ such that for all $m \geq m_0$, any geodesic in Γ_G between $i(a_m)$ and $i(b_m)$ must pass through the (M')-ball about the identity in Γ_G. Since M' does not depend on the index $m \geq m_0$, we have shown that the condition cannot be satisfied if the Cannon-Thurston map does not exist. $\square$

Mitra uses this characterization in [47] to prove that the Cannon-Thurston map exists for hyperbolic group extensions:

Theorem 2.17 (Mitra [47], Theorem 4.3)**.** *Let H, G, and Q be finitely generated, word-hyperbolic groups with*

$$1 \to H \to G \to Q \to 1.$$

Then, the inclusion map $i : \Gamma_H \to \Gamma_G$ extends continuously to the Cannon-Thurston map $\partial i : \partial H \to \partial G$.

If H is a group acting by isometries on a δ-hyperbolic metric space X, then the *limit set* of H in X, $\Lambda H \subseteq \partial X$, is the set of accumulation points of H-orbits of a basepoint $x \in X$. It is known by work of Kapovich and Short [40] that if H is an infinite normal subgroup of a hyperbolic group G, then the limit set of H in G is all of ∂G. Hence, in the above setting of a hyperbolic group extension, the Cannon-Thurston map $\partial i : \partial H \to \partial G$ is surjective when H is infinite. It turns out that when H is non-elementary, the assumption that Q be hyperbolic is redundant for the following reason. Suppose that $1 \to H \to G \to Q \to 1$ is a short exact sequence of finitely generated groups such that G is hyperbolic, H is non-elementary and hyperbolic, and let $P : G \to Q$ be a surjective group homomorphism. Mosher showed in [57] that in this situation, there exists a lift $\sigma : Q \to G$ such that σ is a quasi-isometric embedding, and such that for all $q \in Q$, $P \cdot \sigma(q) = q$. Such a lift σ is called a *quasi-isometric section*. So, when H is non-elementary, the quotient group Q is hyperbolic because it is a quasiconvex subgroup of the hyperbolic group G.

It is known that if G is hyperbolic and H is an infinite normal subgroup of G of infinite index, then H is not quasiconvex in G. Therefore, there is no immediate reason to expect the Cannon-Thurston map to exist in the above setting of a hyperbolic group extension. There are several other settings in which the Cannon-Thurston map is known to exist. In [49], Mitra uses a similar characterization of the existence of the Cannon-Thurston map to show that for a tree of hyperbolic metric spaces X with hyperbolic vertex spaces which quasi-isometrically embed into X, if X itself is hyperbolic, then there exists the Cannon-Thurston map from the boundary of each vertex space to the boundary of X. In particular, this result applies to the setting of a graph of groups where each of the vertex groups are hyperbolic and where the overall fundamental group of the graph of groups is hyperbolic. In [2], Baker and Riley show that the Cannon-Thurston map exists for a family of hyperbolic groups known as *hyperbolic hydra* that contain heavily distorted free subgroups. This result shows that while the Cannon-Thurston map trivially exists for undistorted subgroups, heavy distortion is not an obstruction to the existence of such a map.

It was not until more recently that the first example of the non-existence of the Cannon-Thurston map was found. In [3], Baker and Riley show that there is a hyperbolic group

G with an F_3 subgroup for which the Cannon-Thurston map does not exist. Matsuda and Oguni extend this result in [45] to show that every non-elementary hyperbolic group H occurs as a subgroup of some hyperbolic group G for which the inclusion of H into G does not have a continuous extension to the Cannon-Thurston map.

2.5 Laminations

2.5.1 Surface Laminations

Let S be a closed, hyperbolic surface. A *geodesic* in S is defined to be the image of a complete geodesic in $\mathbb{H}^2 = \widetilde{S}$ under the covering map. A geodesic is said to be *simple* if it has no self-intersections. A non-empty closed subset L of S is said to be a *geodesic lamination* if L consists of a disjoint union of simple geodesics. Each geodesic in L is called a *leaf* of the lamination. The laminations which we will be most interested in are those which are formed from taking the limit of a sequence of longer and longer geodesics. For more background on surface laminations, see [14].

Given a surface S, the *mapping class group* of S, $\mathrm{Mod}(S)$, is defined to be the group of isotopy classes of orientation-preserving homeomorphisms of S:

$$\mathrm{Mod}(S) := \mathrm{Homeo}^+(S)/\mathrm{Homeo}_0(S),$$

where $\mathrm{Homeo}_0(S)$ denotes the connected component of the identity in the group of orientation-preserving homeomorphisms of S, $\mathrm{Homeo}^+(S)$. For general background on mapping class groups, see [26]. If $x \in S$ is a point and $\mathrm{Mod}(S, x)$ denotes the mapping class group of the punctured surface, then the following sequence is known to be exact and is referred to as the Birman exact sequence [26]:

$$1 \to \pi_1 S \to \mathrm{Mod}(S, x) \to \mathrm{Mod}(S) \to 1.$$

A subgroup $\Gamma \leq \mathrm{Mod}(S, g)$ is said to be *convex cocompact* (in the sense of Farb-Mosher [27]) if some orbit of Γ into the Teichmüller space of S, $\mathcal{T}(S)$ is quasiconvex. By work of Kent-Leininger [41] and Hamenstädt [33], this is equivalent to the orbit map of Γ into the curve complex of S, $\mathcal{C}(S)$, being a quasi-isometric embedding.

In the setting where $\Gamma \leq \mathrm{Mod}(S)$ is convex cocompact, Γ naturally gives rise to a short exact sequence $1 \to \pi_1 S \to E_\Gamma \to \Gamma \to 1$, where the extension group E_Γ is hyperbolic

[33]. By work of Masur-Minsky [44], the curve complex of S, $\mathcal{C}(S)$, which is the complex whose simplices are collections of curves which can be realized disjointly on S, is a hyperbolic metric space. Klarreich [42] showed that the boundary of the curve complex, $\partial\mathcal{C}(S)$, consists of ending laminations on S. So, it follows that to each point $z \in \partial\Gamma$, there is an associated ending lamination, L_z, on the surface S. Note that as the orbit map of Γ into $\mathcal{C}(S)$ is a quasi-isometric embedding, we can find a sequence $\{\Phi_i\} \subseteq \Gamma$ and some simple closed curve $c \in S$ such that $\lim_{i\to\infty} \Phi_i \cdot c = z$ in $\mathcal{C}(S) \cup \partial\mathcal{C}(S)$. Let ϕ_i be a representative homeomorphism in the isotopy class of Φ_i, and let $[\alpha]$ denote a geodesic representative of a curve $\alpha \in S$. Let $c_0 := [c]$ denote a geodesic representative of the curve c, set $c_1 := [\phi_1(c_0)]$, and let $c_2 := [\phi_2(c_1)]$. Continuing in this fashion, let $c_i := [\phi_i(c_{i-1})]$ denote a geodesic representative of $\phi_i(c_{i-1})$. Then, the lamination L_z can be realized by taking the limit of the curves c_i on the surface S.

2.5.2 Algebraic Laminations

Given a hyperbolic group H, a non-empty subset $L \subseteq \partial^2 H = \{(x,y) \in \partial H \times \partial H \mid x \neq y\}$ is said to be an *algebraic lamination* on H if L is closed, symmetric (flip-invariant), and H-invariant. Algebraic laminations were defined and studied in the context of free groups by Coulbois, Hilion, and Lustig in [18, 19, 20]. If S is a hyperbolic surface with non-empty boundary, then we can fix an identification of $\pi_1 S$ with F_n. Now, $\pi_1 S$ acts on the universal cover $\widetilde{S} = \mathbb{H}^2$ by covering transformations, which in this case are isometries of the hyperbolic plane. Given a geodesic (surface) lamination L on S, let $\widetilde{L}$ denote the lift of L to $\widetilde{S} = \mathbb{H}^2$. By Milnor-Svarc, $\pi_1 S = F_N$ is quasi-isometric to $\widetilde{S}$, and so the identification of F_N with $\pi_1 S$ induces an F_N-equivariant homeomorphism between the boundary of the fundamental group, ∂F_N, and the boundary of the universal cover, $\partial\widetilde{S}$. Thus, any leaf $l \in \widetilde{L}$ defines a pair of endpoints $(a,b) \in \partial^2\mathbb{H}^2$, which are naturally identified with the points (a,b) and $(b,a) \in \partial^2 F_N$. As the set of all pairs of endpoints of leaves of $\widetilde{L}$ is closed and F_N-invariant, any geodesic lamination on such a surface S naturally defines an algebraic lamination on a free group. Note, however, that not all algebraic laminations on a free group arise from a geodesic lamination on a surface.

Let $\mathcal{A}$ denote a basis of F_N A word $w \in \mathcal{A}^{\pm 1}$ is said to be *reduced* if w does not contain a subword of the form xx^{-1} or $x^{-1}x$, where $x \in \mathcal{A}$. Let $F(\mathcal{A})$ denote the set of all reduced words in $\mathcal{A}^{\pm 1}$. If S is any set of finie, semi-infinite, or bi-infinite reduced words in $\mathcal{A}^{\pm 1}$, we define the *language generated by* S, $\mathcal{L}(S) \subseteq F(\mathcal{A})$, to be the set of all finite subwords of elements of S. A non-empty set $\mathcal{L} \subseteq F(\mathcal{A})$ of finite reduced words in $\mathcal{A}^{\pm 1}$ is called a *laminary language* if

1. $\mathcal{L}$ is closed with respect to inversion;

2. $\mathcal{L}$ is closed with respect to taking subwords; and

3. for any word $u \in \mathcal{L}$, there is some $v \in \mathcal{L}$ for which $v = wuw'$ with $w, w' \in F(\mathcal{A}) \setminus \{1\}$.

If $(p, q) \in \partial^2 F_N$, then we define the *word* corresponding to the leaf (p, q) to be the bi-infinite reduced word $p_{\mathcal{A}}^{-1} q_{\mathcal{A}}$, where $p_{\mathcal{A}}$ and $q_{\mathcal{A}}$ are the reduced infinite words in $\mathcal{A}^{\pm 1}$ that represent p and q, respectively (see Figure 2.1). In [18], Coulbois, Hilion, and Lustig show that a

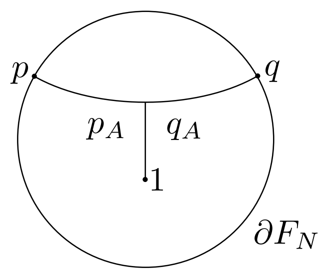

Figure 2.1: Leaf of $\partial^2 F_N$

non-empty set $L \subseteq \partial^2 F_N$ is an algebraic lamination if and only if the language generated by the set of words corresponding to the leafs of L is a laminary language.

2.6 $\mathbb{R}$-trees

An $\mathbb{R}$-*tree* is a geodesic metric space which is 0-hyperbolic (and so every triangle is a degenerate tripod). Note that $\mathbb{R}$-trees are more general than simplicial trees, since they can be infinitely "hairy" (see example 2.9 in [64]). One setting in which $\mathbb{R}$-trees arise naturally is in the context of geodesic laminations on a surface. In particular, let $\mathcal{L}$ be a geodesic lamination of a hyperbolic surface, S, with transverse measure μ. Lift $\mathcal{L}$ and μ to the universal cover $\widetilde{S}$ to get the induced lamination $\widetilde{\mathcal{L}}$ with transverse measure $\widetilde{\mu}$. Define $\widetilde{d} : \widetilde{S} \times \widetilde{S} \to [0, \infty)$ by $\widetilde{d}(x, y) = \inf_\alpha(\widetilde{\mu}(\alpha))$, where α is any arc in $\widetilde{S}$ joining x to y which is transverse to $\widetilde{\mathcal{L}}$. We note that $\widetilde{d}$ is symmetric, since any arc joining x to y is also an arc joining y to x, and that $\widetilde{d}$ satisfies the triangle inequality, since $\widetilde{\mu}$ satisfies the triangle inequality. However, $\widetilde{d}$ is not positive-definite, since if x and y are two distinct points on the same leaf $l \in \widetilde{\mathcal{L}}$ or are in the same complementary component of $\widetilde{\mathcal{L}}$, then $\widetilde{d}(x, y) = 0$. Therefore, $\widetilde{d}$ defines a pseudometric on $\widetilde{S}$.

Now, let $(T_{\mathcal{L}}, d)$ denote the metric space associated to $\widetilde{\mathcal{L}}$ which is obtained from $\widetilde{S}$ by identifying points $(x, y) \in \widetilde{S} \times \widetilde{S}$ with $\widetilde{d}(x, y) = 0$. We claim that $T_{\mathcal{L}}$ is an $\mathbb{R}$-tree. To see this, first suppose arcs in $T_{\mathcal{L}}$ are not unique. To this end, suppose $u, v \in T_{\mathcal{L}}$ are such that $u \neq v$, there are two distinct arcs γ_1, γ_2 joining u and v in $T_{\mathcal{L}}$, and $d(u, v)$ is minimal. Then, consider the preimages γ_1' and γ_2' of these arcs in $\widetilde{S}$. Note that these arcs γ_1' and γ_2' cannot intersect in $\widetilde{S}$, or else this would contradict the minimality of $d(u, v)$. Note also that there must be some leaf $l \in \widetilde{\mathcal{L}}$ which intersects both γ_1' and γ_2'. But, this implies that there are points $a_1 \in \gamma_1'$ and $a_2 \in \gamma_2'$ for which $\widetilde{d}(a_1, a_2) = 0$. This, again, contradicts the minimality of $d(u, v)$, and so arcs in $T_{\mathcal{L}}$ are unique. To finish showing that $T_{\mathcal{L}}$ is an $\mathbb{R}$-tree, we must show that if $a, b, c \in T_{\mathcal{L}}$ are such that $[a, b] \cap [b, c] = \{b\}$, then $[a, b] \cup [b, c] = [a, c]$. Suppose that $\{a, b, c\} \in T_{\mathcal{L}}$ are points which form a minimal triangle in $T_{\mathcal{L}}$. Take a lift of $[a, b]$, $[b, c]$, and $[a, c]$ to $\widetilde{S}$. Then, since these arcs have non-zero measure under $\widetilde{\mu}$, there is some leaf l of $\widetilde{\mathcal{L}}$ which intersects two of these lifts, contradicting the minimality of the triangle in T. Thus, $(T_{\mathcal{L}}, d)$ must be an $\mathbb{R}$-tree. The action of $\pi_1 S$ on $\widetilde{S}$ extends canonically to an isometric $\pi_1 S$-action on this dual tree $T_{\mathcal{L}}$.

2.7 The Dual Algebraic Lamination

Let G be a group acting isometrically on an $\mathbb{R}$-tree, T, and let $\mathrm{Fix}(g)$ denote the subtree of T which is fixed by the element $g \in G$. The action of G on T is said to be *minimal* if T has no proper G-invariant subtree. The action is called *small* if any two elements which pointwise fix a non-trivial arc in T commute. Further, a small G-action is said to be *very small* if for all non-trivial $g \in G$, (a) $\mathrm{Fix}(g)$ does not contain a tripod; and (b) $\mathrm{Fix}(g) = \mathrm{Fix}(g^n)$, for all $g^n \neq 1$. Note that if G acts freely on T, then this action is necessarily a very small action as only the identity element fixes any point in T. In [19], Coulbois, Hilion, and Lustig define a lamination which is canonically dual to a given $\mathbb{R}$-tree T, provided that T has a very small, minimal action of the free group F_N.

Suppose that T is an $\mathbb{R}$-tree with an isometric (left)-action of F_N. Define the *translation length* of an element $w \in F_N$ by

$$||w||_T := \inf\{d_T(P, wP) \mid P \in T\},$$

where d_T denotes the distance in T.

Definition 2.18 ([19])**.** Let T be a very small $\mathbb{R}$-tree with minimal and isometric F_N-action.

The *dual algebraic lamination* associated to T, denoted by $L(T)$, consists of the set of leafs $(p,q) \in \partial^2 F_N$ which satisfy the following:

For any $\epsilon > 0$ and any finite subword of the word corresponding to (p,q), there is some cyclically reduced $w \in F_N$ with translation length $||w||_T \leq \epsilon$ such that v is a subword of w.

For an $\mathbb{R}$-tree T, denote the metric completion by $\overline{T}$, the Gromov boundary by ∂T, and let $\widehat{T} = \overline{T} \cup \partial T$. Note that a point $p \in \partial T$ is determined by a ray $\gamma : [0, \infty) \to T$, and two rays γ and γ' determine the same point in ∂T if any only if their images in T only differ on a compact subset of T. The metric d_T on T extends canonically to a metric on $\widehat{T}$, known as the *metric topology*, though $\widehat{T}$ with this topology is not generally compact. In [17], Coulbois, Hilion, and Lustig define a coarser (weaker) topology on $\widehat{T}$ known as the *observers' topology*. If $p, q \in \widehat{T}$ are distinct points, define the *direction of q at p*, denoted by $\mathrm{dir}_p(q)$, to be the connected component of $\widehat{T} \setminus \{p\}$ which contains q.

Definition 2.19 ([17])**.** The *observers' topology* on $\widehat{T}$ is the topology generated by the set of directions in $\widehat{T}$. Denote $\widehat{T}$ equipped with the observers' topology by $\widehat{T}^{\mathrm{obs}}$.

Coulbois, Hilion, and Lustig [17] show that the metric topology and the observers' topology agree on any finite subtree of $\overline{T}$ and that the restriction of the two topologies to ∂T also agree. They also show that $\widehat{T}^{\mathrm{obs}}$ is compact and Hausdorff. Therefore, $\widehat{T}^{\mathrm{obs}}$ is a dendrite. A key result of [17] shows that the tree $\widehat{T}^{\mathrm{obs}}$ is actually completely determined by the dual algebraic lamination of T, $L(T)$.

Theorem 2.20 (Corollary 2.6 [17])**.** *Let T be an $\mathbb{R}$-tree with a very small, minimal, isometric F_N-action with dense orbits, and let $L(T)$ denote the dual algebraic lamination. Then, there is an F_N-equivariant homeomorphism $\psi : \partial F_N / L(T) \to \widehat{T}^{\mathrm{obs}}$.*

2.8 Free group automorphisms and train track maps

Let F_N denote the free group of rank $N \geq 2$. The *outer automorphism group* of F_N is defined by $\mathrm{Out}(F_N) := \mathrm{Aut}(F_N)/\mathrm{Inn}(F_N)$, where $\mathrm{Inn}(F_N)$ is the subgroup of automorphisms of F_N which arise via conjugation by an element of F_N. An automorphism $\varphi \in \mathrm{Aut}(F_N)$ is called *fully irreducible*, or *irreducible with irreducible powers (iwip)*, if no positive power of φ maps a non-trivial proper free factor of F_N to a conjugate of itself. The map φ is said to be *atoroidal* if it has no non-trivial periodic conjugacy class. For any $\varphi \in \mathrm{Aut}(F_N)$, the group $G_\varphi = F_N \rtimes_\varphi \mathbb{Z} = \langle a_1, \ldots, a_N, t \mid t^{-1}wt = \varphi(w), \forall w \in F_N \rangle$ is known as the *mapping torus* *group* defined by φ. Work of Bestvina-Feighn [5] and Brinkmann [12] gives that the group G_φ is hyperbolic exactly when φ is an atoroidal automorphism. For this reason, atoroidal automorphisms are also known as *hyperbolic* automorphisms.

Let R_N denote the wedge of N circles (called the *rose* of rank N). There is a natural identification of the fundamental group $\pi_1(R_N)$ with the free group F_N, and so each $\phi \in \mathrm{Out}(F_N)$ can be represented by a homotopy equivalence of R_N to itself. Let G be a graph with no vertices of valence 1 or 2, and let $g : R_N \to G$ be a homotopy equivalence (called a *marking*). Then, we say the pair (G, g) is a *marked metric graph* if G has a path metric such that lengths of the edges sum to 1. Given any $\phi \in \mathrm{Out}(F_N)$, there is some marked metric graph G_ϕ such that ϕ can be represented as a cellular map $f_\phi : G_\phi \to G_\phi$. The cellular map f_ϕ is a *train track map* if for every $k > 0$, the map $f_\phi^k : G_\phi \to G_\phi$ is locally injective on the interior of each edge.

In [7], Bestvina and Handel show that every fully irreducible outer automorphism of a free group can be represented by a train track map. Associated to the train track representative of a fully irreducible automorphism φ are two trees, $T_+(\varphi) = T_+$ and $T_-(\varphi) = T_-$, with several nice properties. In particular, both T_+ and T_- are minimal $\mathbb{R}$-trees with an isometric F_N-action which is free and has dense orbits. As such, T_+ and T_- have an associated dual algebraic lamination, $L(T_+)$ and $L(T_-)$. It is these laminations which Kapovich and Lustig show in [39] are the same as the algebraic ending laminations developed by Mitra in [47] which are associated to the points in the boundary of the quotient group $F_N \rtimes_\varphi \mathbb{Z}/F_N = \mathbb{Z} = \langle\varphi\rangle$.

2.9 Metric graph bundles

In [5], Bestvina and Feighn explored the question of when a space which results from the combination of Gromov-hyperbolic spaces will itself be hyperbolic. They introduced the notion of a graph of spaces and provided a "flaring" condition which gives a sufficient condition for the hyperbolicity of a graph of hyperbolic spaces. Mj and Sardar generalized this work in [56] where they introduced the notion of a metric graph bundle and defined the following flaring condition.

Let X and B be connected graphs, each equipped with the path metric where each edge has length 1, and let $p : X \to B$ be a simplicial surjection. For the purpose of this paper, we will consider $\mathbb{N} = \mathbb{Z}_{\geq 0}$.

Definition 2.21. X is said to be a *metric graph bundle* over B if there exists a function $f : \mathbb{N} \to \mathbb{N}$ such that:

(B1) For each vertex $b \in V(B)$, the fiber $F_b := p^{-1}(b)$ is a connected subgraph of X; and for all vertices $u, v \in V(F_b)$, the induced path metric d_b on F_b satisfies $d_b(u, v) \leq f(d_X(u, v))$.

(B2) If $b_1, b_2 \in V(B)$ are any two adjacent vertices and if $x_1 \in V(F_{b_1})$ is any vertex, then there is some vertex $x_2 \in V(F_{b_2})$ adjacent to x_1 in X.

Remark 2.22. Note that if $p : X \to B$ is a metric graph bundle and $W \subseteq B$ is any connected subgraph, then $p : p^{-1}(W) \to W$ is again a metric graph bundle.

Given any metric graph bundle $p : X \to B$ and a connected, closed interval $I \subseteq \mathbb{R}$, a (k, k)-*quasi-isometric lift* of a geodesic $\gamma : I \to B$ is any (k, k)-quasigeodesic $\widetilde{\gamma} : I \to X$ for which $p(\widetilde{\gamma}(n)) = \gamma(n)$ for all $n \in I \cap \mathbb{Z}$.

Definition 2.23. The metric graph bundle $p : X \to B$ is said to satisfy the *flaring condition* if for all $k \geq 1$, there exists $\lambda_k > 1$ and $n_k, M_k \in \mathbb{N}$ such that the following holds: If $\gamma : [-n_k, n_k] \to B$ is any geodesic and $\widetilde{\gamma_1}$ and $\widetilde{\gamma_2}$ are any two (k, k)-quasi-isometric lifts of γ in X which satisfy $d_{\gamma(0)}(\widetilde{\gamma_1}(0), \widetilde{\gamma_2}(0)) \geq M_k$, then we have

$$\lambda_k \cdot d_{\gamma(0)}(\widetilde{\gamma_1}(0), \widetilde{\gamma_2}(0)) \leq \max\{d_{\gamma(-n_k)}(\widetilde{\gamma_1}(-n_k), \widetilde{\gamma_2}(-n_k)), d_{\gamma(n_k)}(\widetilde{\gamma_1}(n_k), \widetilde{\gamma_2}(n_k))\}.$$

The following are two theorems of Mj and Sardar which we will use later. The first is their combination theorem for metric graph bundles, which generalizes the combination theorem of Bestvina-Feighn [5]. The second shows that flaring is a necessary condition for the hyperbolicity of a metric graph bundle.

Theorem 2.24 (Mj-Sardar [56])**.** *Suppose that $p : X \to B$ is a metric graph bundle which satisfies:*

1. *B is a δ-hyperbolic metric space;*
2. *for each $b \in V(B)$, the fiber F_b is δ-hyperbolic with respect to d_b, the path metric induced by X;*
3. *for each $b \in V(B)$, the set of barycenters of ideal triangles in F_b is D-dense; and*
4. *the flaring condition is satisfied.*

Then, X is a hyperbolic metric space.

Theorem 2.25 (Mj-Sardar [56]). *Suppose that $p : X \to B$ is a metric graph bundle which satisfies:*

1. *X is δ-hyperbolic; and*

2. *for each $b \in V(B)$, the fiber F_b is δ-hyperbolic with respect to d_b, the path metric induced by X.*

Then, the metric bundle satisfies the flaring condition.

2.10 Stacks of spaces

In [10], Bowditch defines the notion of a stack of spaces.

Definition 2.26. Let (X, d_X) and (Y, d_Y) be path-metric spaces. A map $f : X \to Y$ is said to be *straight* if there exist functions $F_1, F_2 : [0, \infty) \to [0, \infty)$ such that for all $x, x' \in X$, $F_1(d_X(x, x')) \leq d_Y(f(x), f(x')) \leq F_2(d_X(x, x'))$, where $F_1(t) \to \infty$ as $t \to \infty$. If $X \subseteq Y$, we say that X is a *straight* subspace if the inclusion map $i : X \to Y$ is a straight map with respect to the induced path metric on X.

Definition 2.27. Let $(\mathcal{X}, \rho)$ be a geodesic space, and let $((X_i, \rho_i))_{i \in \mathbb{Z}}$ be a sequence of geodesic subspaces, $X_i \subseteq \mathcal{X}$, called the *sheets* of $\mathcal{X}$ with uniform quasi-isometries $f_i : X_i \to X_{i+1}$. The space $(\mathcal{X}, \rho)$ is said to be a *bi-infinite hyperbolic stack* if it satisfies the conditions (S1)-(S6) stated below.

(S1) Each of the spaces (X_i, ρ_i) are uniformly straight in $\mathcal{X}$, and $\rho(X_i, X_j)$ is bounded away from 0 for $i \neq j$.

(S2) For all $i, j \in \mathbb{Z}$, $\rho(X_i, X_j)$ is bounded below by an increasing linear function of $|i - j|$.

(S3) For all $i \in \mathbb{Z}$, $\mathrm{haus}(X_i, X_{i+1})$ is bounded above.

(S4) The spaces (X_i, ρ_i) are uniformly hyperbolic geodesic spaces.

(S5) The space $(\mathcal{X}, \rho)$ is hyperbolic.

(S6) The union $\bigcup_{i \in \mathbb{Z}} X_i$ is quasidense in $\mathcal{X}$.

Given a bi-infinite stack $\mathcal{X}$, denote by $\mathcal{X}^+$ and $\mathcal{X}^-$ the subsets of $\mathcal{X}$ which consist of the sheets $(X_i)_{i\in\mathbb{N}}$ and $(X_i)_{i\in-\mathbb{N}}$, respectively. Here, $\mathbb{N} = \mathbb{Z}_{\geq 0}$ and $-\mathbb{N} = \mathbb{Z}_{\leq 0}$. We will refer to $\mathcal{X}^+$ and $\mathcal{X}^-$ as *semi-infinite* stacks. Bowditch proves the following about stacks of hyperbolic spaces indexed by any subset $I \subseteq \mathbb{Z}$ of consecutive integers.

Proposition 2.28 (Bowditch [10] Proposition 2.1.7)**.** *Suppose $\mathcal{X}$ is a bi-infinite stack with uniformly hyperbolic sheets $(X_i)_{i\in\mathbb{Z}}$. If $\mathcal{X}$ is hyperbolic, then so is $\mathcal{X}(I)$, where $I \subseteq \mathbb{Z}$ is any set of consecutive integers. In particular, the semi-infinite stacks $\mathcal{X}^+$ and $\mathcal{X}^-$ are hyperbolic whenever $\mathcal{X}$ is hyperbolic.*

Given a (bi-infinite) stack $\mathcal{X}$, Bowditch defines an *r-chain*, $(x_i)_{i\in\mathcal{I}}$, to be a sequence of points, $x_i \in X_i$, such that $\rho(x_i, x_{i+1}) \leq r$ for all $i \in \mathcal{I}$. A *bi-infinite*, *positive*, and *negative* r-chain is defined to be an r-chain indexed by $\mathbb{Z}$, $\mathbb{N}$, and $-\mathbb{N}$, respectively. Bowditch notes that each r-chain interpolates a quasigeodesic in $\mathcal{X}$. If $\mathcal{X}$ is a hyperbolic stack, it comes equipped with its Gromov boundary, $\partial\mathcal{X}$. Thus when $\mathcal{X}$ is a proper, hyperbolic stack, each positive and negative chain determines a point of $\partial\mathcal{X}$. In this setting, there is a fixed r_0 depending on the hyperbolicity constant of $\mathcal{X}$ for which each point in $\mathcal{X}$ is contained in some r_0-chain. Given some r_0 for which each point of $\mathcal{X}$ is contained in an r_0-chain, Bowditch defines $\partial^+\mathcal{X}$ (respectively $\partial^-\mathcal{X}$) to be those subsets of $\partial\mathcal{X}$ which are determined by positive (respectively negative) r_0-chains. Note that the positive chains in $\mathcal{X}^+$ are exactly the positive chains in $\mathcal{X}$, and the negative chains in $\mathcal{X}^-$ are exactly the negative chains in $\mathcal{X}$. Furthermore, two chains determine the same point in $\partial\mathcal{X}^+$ or $\partial\mathcal{X}^-$ if and only if those two chains determine the same point in $\partial\mathcal{X}$. Hence on the level of sets, we can identify $\partial^+\mathcal{X}^+$ with $\partial^+\mathcal{X}$ and $\partial^-\mathcal{X}^-$ with $\partial^-\mathcal{X}$.

Each of the sheets X_i are quasi-isometric to one another, and so we get a homeomorphism from ∂X_i to ∂X_j, for all $i, j \in \mathbb{Z}$. We will let ∂X_0 denote this space which is homeomorphic to ∂X_i for all $i \in \mathbb{Z}$. The notion of the Cannon-Thurston map, as defined earlier between the boundaries of hyperbolic groups, can be extended in the natural way to be defined between the boundaries of hyperbolic spaces. Bowditch proves the following statements about the Cannon-Thurston maps in this setting of stacks of spaces.

Proposition 2.29 (Bowditch [10] see 2.3.2 and 2.3.3)**.** *Let $\mathcal{X}$ be a bi-infinite hyperbolic stack, let $\mathcal{X}^+$ and $\mathcal{X}^-$ be semi-infinite proper hyperbolic stacks, and let ω, ω^+, and ω^- denote the inclusions of X_0 into $\mathcal{X}$, $\mathcal{X}^+$, and $\mathcal{X}^-$, respectively. Then,*

1. *The following continuous Cannon-Thurston maps exist: $\partial\omega : \partial X_0 \to \partial\mathcal{X}$, $\partial\omega^+ : \partial X_0 \to \partial\mathcal{X}^+$, and $\partial\omega^- : \partial X_0 \to \partial\mathcal{X}^-$;*

2. $\partial\mathcal{X} = \partial^+\mathcal{X} \cup \partial^-\mathcal{X} \cup \partial\omega(\partial X_0)$; *and*

3. $\partial\mathcal{X}^+ = \partial^+\mathcal{X} \cup \partial\omega^+(\partial X_0)$ *and* $\partial\mathcal{X}^- = \partial^-\mathcal{X} \cup \partial\omega^-(\partial X_0)$.

Given the Cannon-Thurston maps $\partial\omega$ and $\partial\omega^\pm$, denote by $\widehat{\omega}$ and $\widehat{\omega}^\pm$ the continuous extensions of the inclusion maps. Bowditch defines the maps $\partial\tau^\pm : \partial\mathcal{X}^\pm \to \partial\mathcal{X}$ which extend to continuous maps $\widehat{\tau}^\pm : \widehat{\mathcal{X}^\pm} \to \widehat{\mathcal{X}}$ such that $\widehat{\omega} = \widehat{\tau}^\pm \circ \widehat{\omega}^\pm$. For $y \in \partial^+\mathcal{X}^+ = \partial^+\mathcal{X}$, the map $\partial\tau^+$ is given by $\partial\tau^+(y) = y$; and for $a \in \partial X_0$, we have that $\partial\tau^+ \circ \partial\omega^+(a) = \partial\omega(a)$. The map $\partial\tau^-$ is defined similarly. Bowditch proves that $\partial\tau^\pm$ are continuous maps. Using this structure, Bowditch shows the following.

Lemma 2.30 (Bowditch [10] see 2.3.5, 2.3.6, 2.3.7, and 2.3.9)**.** *Let $\mathcal{X}$ be a bi-infinite, proper, hyperbolic stack.*

1. *Suppose $a \in \partial X_0$ and $y \in \partial^+\mathcal{X}$. Then, $\partial\omega(a) = y$ if and only if there is a sequence $(\underline{x}^n)_{n\in\mathbb{N}}$ of positive chains, $\underline{x}^n = (x_i^n)_{i\in\mathbb{N}}$, each converging to y, and with x_0^n converging to $a \in \partial X_0$.*

2. *Given $a \in \partial X_0$ and $y \in \partial^\pm\mathcal{X}$, we have $\partial\omega^\pm(a) = y$ if and only if $\partial\omega(a) = y$.*

3. *Suppose $a, b \in \partial X_0$ are distinct. If $\partial\omega^+(a) = \partial\omega^+(b) = y$, then $y \in \partial^+\mathcal{X}$; and if $\partial\omega^-(a) = \partial\omega^-(b) = y$, then $y \in \partial^-\mathcal{X}$.*

4. *If $a, b \in \partial X_0$ and $\partial\omega(a) = \partial\omega(b)$, then either $\partial\omega^+(a) = \partial\omega^+(b)$ or $\partial\omega^-(a) = \partial\omega^-(b)$.*

Chapter 3

Results about hyperbolic metric spaces and groups

In this chapter, we prove basic results about hyperbolic metric spaces and groups which will be needed later in this thesis.

3.1 Paths in hyperbolic spaces.

Proposition 3.1. *Let (X, d) be a δ-hyperbolic metric space and let $A \geq 0$. If $x, y, z \in X$ are such that $(x, z)_y \leq A$, then $[x, y] \cup [y, z]$ is a $(1, 2A)$-quasigeodesic.*

Proof. Suppose that x, y, and z are such that $(x, z)_y \leq A$. We need to show that for all $p \in [x, y]$ and $q \in [y, z]$, $d(p, y) + d(y, q) \leq d(p, q) + 2A$. By the triangle inequality,

$$\begin{aligned}(p, q)_y &= \frac{1}{2}(d(p, y) + d(q, y) - d(p, q)) \\ &\leq \frac{1}{2}(d(p, y) + d(q, y) - (d(p, z) - d(q, z))) \\ &= \frac{1}{2}(d(p, y) + d(y, z) - d(p, z)) = (p, z)_y.\end{aligned}$$

Similarly, $(p, z)_y \leq (x, z)_y$. Therefore, $(p, q)_y \leq A$ by hypothesis, and so $d(p, q) + d(y, q) = d(p, q) + 2(p, q)_y \leq d(p, q) + 2A$. Hence, $[x, y] \cup [y, z]$ is a $(1, 2A)$-quasigeodesic. □

The next proposition says that geodesic quadrilaterals in hyperbolic metric spaces must either be "tall and thin" or "short and long".

Proposition 3.2. *Let (X, d) be a δ-hyperbolic metric space and let $x, y, z, w \in X$. Then, either there are points $a \in [x, y]$ and $a' \in [z, w]$ with $d(a, a') \leq 2\delta$, or there are points $b \in [x, w]$ and $b' \in [y, z]$ with $d(b, b') \leq 2\delta$.*

Proof. Consider the geodesic quadrilateral with sides $[x, y]$, $[y, z]$, $[z, w]$, and $[x, w]$. Draw in the diagonal $[y, w]$ and consider the two triangles $xyw = [x, y] \cup [y, w] \cup [w, x]$ and $ywz = [y, w] \cup [w, z] \cup [z, y]$. Mark internal points $p \in [x, y]$, $q \in [x, w]$, and $r \in [y, w]$ such that

$d(x,p) = d(x,q)$, $d(w,q) = d(w,r)$, and $d(y,p) = d(y,r)$. Similarly, mark internal points $q' \in [y,z]$, $p' \in [z,w]$, and $r' \in [y,w]$ such that $d(z,q') = d(z,p')$, $d(w,p') = d(w,r')$, and $d(y,q') = d(y,r')$. Note that since X is δ-hyperbolic, we have that $\max\{d(p,q), d(q,r), d(p,r)\} \leq \delta$ and $\max\{d(p',q'), d(q',r'), d(p',r')\} \leq \delta$. There are two cases to consider.

First, suppose that $d(y,r) \leq d(y,r')$. In this case, there exists some point $s \in [y,w]$ between r and r' such that $d(s,[x,w]) \leq \delta$ and $d(s,[y,z]) \leq \delta$. Hence, there exist $b \in [x,w]$ and $b' \in [y,z]$ such that $d(b,b') \leq d(b,s) + d(s,b') \leq 2\delta$.

Now, suppose that $d(y,r) > d(y,r')$. In this case, there is some point $s' \in [y,w]$ between r' and r such that $d(s',[x,y]) \leq \delta$ and $d(s',[z,w]) \leq \delta$. So, there is some $a \in [x,y]$ and $a' \in [z,w]$ with $d(a,a') \leq 2\delta$.

□

Proposition 3.3. *Let (X,d) be a δ-hyperbolic metric space and let $A \geq 0$. If $x,y,z,w \in X$ are such that $(x,z)_y \leq A$, $(y,w)_z \leq A$, and $d(y,z) > 10\delta + 2A$, then $[x,y] \cup [y,z] \cup [z,w]$ is a $(1, 4\delta + 4A)$-quasigeodesic.*

Proof. Fix $x,y,z,w \in X$ such that $(x,z)_y \leq A$, $(y,w)_z \leq A$, and $d(y,z) > 10\delta + 2A$. We need to show that for all $p,q \in [x,y] \cup [y,z] \cup [z,w]$, the distance between p and q along $[x,y] \cup [y,z] \cup [z,w]$ is at most $d(p,q) + 4\delta + 4A$. This statement is certainly true if p and q are on the same geodesic segment, and the proof of Proposition 3.1 shows that it also holds if p and q are on adjacent segments. So, it remains to show that if $p \in [x,y]$ and $q \in [z,w]$, then $d(p,y) + d(y,z) + d(z,q) \leq d(p,q) + 4\delta + 4A$.

So, fix $p \in [x,y]$ and $q \in [z,w]$ and let $[p,q]$ denote the geodesic segment between p and q. Since $d(y,z) > 10\delta + 2A$, there exists a point $r \in [y,z]$ such that $d(r,y) > 5\delta + A$ and $d(r,z) > 5\delta + A$. As geodesic quadrilaterals are 2δ-thin, there exists some $r' \in [y,p] \cup [p,q] \cup [q,z]$ at distance at most 2δ from r. We claim that this point $r' \in [p,q]$. Suppose instead that $r' \in [p,y]$. Then since $(p,r)_y \leq (x,z)_y \leq A$, we have that $d(y,[p,r]) \leq A + \delta$. So, $d(z,y) \leq d(p,x) + A + \delta$. But then,

$$\begin{aligned}
d(p,x) + A + \delta - d(x,y) &\leq d(p,r') + d(r',x) + A + \delta - d(x,r') - d(r',y) \\
&= d(p,r') + A + \delta - d(r',y) \\
&\leq d(p,r') + A + \delta - [d(y,p) - d(p,r')] \\
&= 2d(p,r') + A + \delta - d(y,p) \\
&< 4\delta + A + \delta - (5\delta + A) = 0,
\end{aligned}$$

which is a contradiction. Similarly, we cannot have that $r' \in [z, q]$ and hence our claim that $r' \in [p, q]$ must be true.

As $(p, r)_y \leq A$ and $(q, r)_z \leq A$, we have that $d(p, y) + d(y, r) \leq d(p, r) + 2A$ and $d(r, z) + d(z, q) \leq d(r, q) + 2A$. By the triangle inequality, $d(p, r) \leq d(p, r') + 2\delta$ and $d(q, r) \leq d(q, r') + 2\delta$. Therefore, $d(p, y) + d(y, z) + d(z, q) \leq d(p, q) + 4\delta + 4A$. □

Lemma 3.4. *Let (X, d) be a δ-hyperbolic metric space and let $x, y, z, w \in X$. If there exist points $a \in [x, w]$ and $b \in [y, z]$ such that $d(a, b) \leq 2\delta$, then $[x, y] \cup [y, z] \cup [z, w]$ is a $(1, 4\delta + 4d(y, z))$-quasigeodesic.*

Proof. Let $x, y, z, w \in X$ be as above and consider the geodesic quadrilateral with edges $[x, y]$, $[y, z]$, $[z, w]$, and $[x, w]$. Note that both $(x, z)_y$ and $(y, w)_z$ are bounded by $d(y, z)$. So, the proof of Proposition 3.3 shows that if $p \in [x, y]$ and $q \in [y, z]$, then $d(p, y) + d(y, q) \leq d(p, q) + 2d(y, z)$. If $p \in [x, y]$ and $q \in [z, w]$, then there exist points $u \in [p, q]$ and $v \in [y, z]$ with $d(u, v) \leq 2\delta$. Thus, $d(p, v) \leq d(p, u) + 2\delta$ and $d(q, v) \leq d(q, u) + 2\delta$. Additionally, $d(p, y) + d(y, v) \leq d(p, v) + 2d(y, z)$ and $d(q, z) + d(z, v) \leq d(q, v) + 2d(y, z)$. Therefore, we have that

$$\begin{aligned} d(p, y) + d(y, z) + d(z, q) &\leq d(p, v) + d(q, v) + 4d(y, z) \\ &\leq d(p, q) + 4\delta + 4d(y, z). \end{aligned}$$

□

3.2 Words in hyperbolic groups.

For the remainder of this chapter, we assume that H is a word-hyperbolic group with a fixed finite generating set S_H. We will also usually abbreviate $|h|_H$ by $|h|$ for $h \in H$. The following definitions generalize the notion of words in a free group being cyclically and almost cyclically reduced to the context of a general word-hyperbolic group.

Definition 3.5. Let $\kappa \geq 0$. An element $h \in H$ is said to be *κ-almost conjugacy minimal* in H if $|h|_H \leq |h'|_H + \kappa$ for all $h' \in [h]_H$. If $\kappa = 0$, then h is said to be *conjugacy minimal.* A geodesic $[a, aw] \subseteq \Gamma_H$ is said to be a *κ-almost conjugacy minimal representative* if $\overline{w} \in H$ is κ-almost conjugacy minimal. If $[a, aw] \subseteq \Gamma_H$ is a κ-almost conjugacy minimal representative, then we will also refer to the word w labeling this geodesic as a *κ-almost conjugacy minimal representative.* If $\kappa = 0$, then $[a, aw]$ and w are said to be *conjugacy minimal representatives.*

Lemma 3.6. *Fix an element $h \in H$ and any constant $\kappa \geq 0$. If h is κ-almost conjugacy minimal, then $(1, hh)_h \leq \frac{\kappa+\delta}{2}$.*

Proof. Let $h \in H$ be κ-almost conjugacy minimal and suppose the geodesic $[1, h] \subseteq \Gamma_H$ is labeled by $\alpha h' \beta$, where $|\alpha| = |\beta| = (1, hh)_h$ and $\beta\alpha = s$, with $|s| \leq \delta$. Then $h =_H \alpha h' s \alpha^{-1}$, and so $h's \in [h]_H$. As h is κ-almost conjugacy minimal, we have that $|h| \leq |h's| + \kappa \leq |h'| + \delta + \kappa$. Finally, $|h| = |\alpha| + |h'| + |\beta| = 2(1, hh)_h + |h'|$, and so we have that $(1, hh)_h \leq \frac{\kappa+\delta}{2}$. □

Lemma 3.7. *There exists a constant $C \geq 0$ such that for any element $h \in H$, the following holds. Suppose that $u, c \in H$ are such that $h = c^{-1}uc$, where $u \in [h]$ is conjugacy minimal and $|c|$ is the smallest element conjugating h to any conjugacy minimal representative. Then, the path $[c, 1] \cup [1, u] \cup [u, uc] \subseteq \Gamma_H$ is a $(1, C)$-quasigeodesic.*

Proof. Let $h, u, c \in H$ be as in the hypothesis above and consider the quadrilateral in Γ_H with vertices 1, c, u, uc, and edges $[1, u]$ labeled by u, $[c, uc]$ labeled by h, and $[1, c]$ and $[u, uc]$ both labeled by c. We want to show the path $\gamma = [c, 1] \cup [1, u] \cup [u, uc]$ is a $(1, C)$-quasigeodesic, for some constant $C \geq 0$.

Let $p \in [1, c]$ and $q \in [1, u]$ be such that $d(1, p) = d(1, q) = (c, u)_1$. As $d(p, q) \leq \delta$, we must have that $d(1, p)$ is also at most δ. Otherwise, $q^{-1}c$ would be a shorter word conjugating h to a cyclic conjugate of u, contradicting the minimality of $|c|$. Similarly, we have that $(1, uc)_u \leq \delta$. If $|u| > 12\delta$, then by Proposition 3.3, γ is a $(1, 8\delta)$-quasigeodesic.

If $|u| \leq 12\delta$, then since H is finitely generated, there are only finitely many possibilities for such u. Hence, there are only finitely many cases to consider and the result holds by taking C to be, for instance, the length of the longest path γ that we get in this setting. □

Corollary 3.8. *For any $\kappa \geq 0$, there exists a constant $M > 0$ such that if $h \in H$ is κ-almost conjugacy minimal, then there is an element $c \in H$ with $|c| \leq M$ and a conjugacy minimal element $u \in [h]$ such that $h = c^{-1}uc$.*

Proof. Let $c \in H$ be a shortest length element conjugating h to any conjugacy minimal element in $[h]$. By Lemma 3.7, there exists some constant $C > 0$ such that $[c, 1] \cup [1, u] \cup [u, uc]$ is a $(1, C)$-quasigeodesic. So, $2|c| + |u| \leq |h| + C$. Since h is κ-almost conjugacy minimal, we have that $|h| \leq |u| + \kappa$. Thus, $|c| \leq \frac{C+\kappa}{2}$. □

Lemma 3.9. *For any $\kappa \geq 0$ there exists a constant $A \geq 0$ such that if $h \in H$ satisfies $(1, hh)_h \leq A$, then h is κ-almost conjugacy minimal.*

Proof. Let $h \in H$ be such that $(1, hh)_h \leq A$. Then by Proposition 3.1, the path $[1, h] \cup [h, hh]$ is a $(1, 2A)$-quasigeodesic. Additionally, all subpaths of $[1, h] \cup [h, hh]$ are $(1, 2A)$-quasigeodesics. In particular, any (non-reduced) edge-path representing a cyclic conjugate of h is a $(1, 2A)$-quasigeodesic. Choose a cyclic conjugate, h' of h such that $ch'c^{-1} = u$, where $u \in [h]_H$ is conjugacy minimal and $|c|$ is smallest.

Consider the points 1, c, u, and uc; geodesics $[1, c]$, $[1, u]$, and $[u, uc]$; and the $(1, 2A)$-quasigeodesic path between c and uc, call it γ', labeled by the (non-reduced) word h'. By Lemma 3.7, there is some constant C for which $\gamma = [c, 1] \cup [1, u] \cup [u, uc]$ is a $(1, C)$-quasigeodesic. As γ' and γ are quasigeodesics sharing the same endpoints, Proposition 2.6 implies that γ' and γ live in a D-neighborhood of each other for some constant $D \geq 0$ depending only on the quasi-isometry constants and δ.

We will now show that $|c|$ is bounded. If $|u| \leq 12\delta$, then there are only finitely many cases to check and we can take maximum length we get in these cases. So, suppose that $|u| > 12\delta$. Note that the distance between any point on $[1, c]$ must be at least $|u|$ from a point on $[u, uc]$ as otherwise we would get a contradiction with u being conjugacy minimal. So by Proposition 3.2, there must exist points $x \in [1, u]$ and $x' \in [c, uc]$ such that $d(x, x') \leq 2\delta$. Let x_0 denote the point along γ' where the two paths labeled by h meet. As the triangle with vertices c, x_0, and uc is δ-thin, there must exist a point $x'' \in [c, x_0] \cup [x_0, uc] = \omega'$ such that $d(x', x'') \leq \delta$. Therefore $d(x, x'') \leq 3\delta$. Now, consider the word c' which labels the path from x to x'' and note that c' conjugates a cyclic conjugate of u to a cyclic conjugate of h'. Therefore, by the minimality of c, we must have in this case that $|c| \leq |c'| \leq 3\delta$.

We now want to show that the distance between x_0 and $[1, u]$ is bounded. Consider the point $y_0 \in \gamma$ which is closest to x_0. Without loss of generality, we may assume that either $y_0 \in [1, u]$ or $y_0 \in [c, 1]$. If $y_0 \in [1, u]$, then $d(x_0, [1, u]) \leq M$. If $y_0 \in [c, 1]$, then $d(x_0, [1, u]) \leq M + |c|$. As $|c|$ is bounded by some constant, we have that the distance between x_0 and $[1, u]$ is also bounded by some constant. Therefore, h is κ-almost conjugacy minimal for some $\kappa \geq 0$ independent of h. □

For the purpose of this paper, if X is a graph, then we will assume that any quasi-isometry or quasi-isometric embedding takes vertices to vertices and edges to edge-paths. The following lemma follows from Proposition 2.6.

Lemma 3.10. *Let $K \geq 1$ and $C \geq 0$. Then, for any $\kappa \geq 0$, there exists $\kappa' \geq 0$ such that if $w \in \Sigma_H^*$ is a κ-almost conjugacy minimal representative and $\psi : \Gamma_H \to \Gamma_H$ is any (K, C)-quasi-isometry, then $\psi(w)$ is a κ'-almost conjugacy minimal representative.*

Chapter 4

Application of bundles and stacks

For the remainder of this thesis, we make the following convention:

Convention 4.1. Let $1 \to H \overset{i}{\to} G \overset{P}{\to} Q \to 1$ be a short exact sequence of three infinite, word-hyperbolic groups. Fix finite, symmetric generating sets S_H, S_G, and S_Q for H, G, and Q, respectively, so that $i(S_H) \subseteq S_G$ and $S_Q := P(S_G)$. Let Γ_H, Γ_G, and Γ_Q denote the Cayley graphs with respect to these generating sets. Let $P : \Gamma_G \to \Gamma_Q$ also denote the map on the Cayley graphs induced by $P : G \to Q$ which is given as follows. If $v \in \Gamma_G$ is a vertex labeled by the element $g \in G$, then v will get sent to the vertex in Γ_Q labeled by the element $P(g) \in Q$. Suppose $e = [g_1, g_2] \in \Gamma_G$ is an edge between adjacent vertices $g_1, g_2 \in \Gamma_G$. If g_1 and g_2 are in the same coset of H in G, then e will get collapsed to the vertex $P(g_1) = P(g_2)$ in Γ_Q. Otherwise, e will get mapped to the edge between $P(g_1)$ and $P(g_2)$ in Γ_Q labeled by $P(g_1^{-1}g_2) \in S_Q$.

4.1 Hyperbolicity of subgraphs of the Cayley graph

The goal of this section is to show that any subgraph of Γ_G which consists of copies of Γ_H which live over a geodesic ray in the quotient group Q is a hyperbolic metric space. We make the following convention:

Convention 4.2. Suppose $\gamma = (z', z)$ is a bi-infinite geodesic in Γ_Q between $z', z \in \partial Q$ with $z' \neq z$; and let $z_0 \in V(\gamma)$ be a vertex of γ which minimizes $d_Q(1, \gamma)$. Label the sequence of vertices in order along the portion of γ from z_0 to z by $z_0, z_1, z_2, \ldots$; and similarly, label the sequence of vertices in order along the portion of γ from z_0 to z' by $z_0, z_{-1}, z_{-2}, \ldots$. Let $\gamma^+ = [z_0, z)$ and $\gamma^- = (z', z_0]$.

Definition 4.3. The *subgraph of* Γ_G *corresponding to* γ is

$$X(\gamma) := P^{-1}(\gamma).$$

Note that we can think of $X(\gamma)$ as the subgraph of Γ_G with vertical fibers that are copies of Γ_H corresponding to the cosets g_iH, where $g_i \in P^{-1}(z_i)$ for each $z_i \in V(\gamma)$. Since $S_Q = P(S_G)$, there are edges between adjacent cosets g_iH and $g_{i+1}H$ between any vertex g_ih and the vertex $g_ihP^{-1}([z_i, z_{i+1}])$, where $[z_i, z_{i+1}]$ is the edge in γ between z_i and z_{i+1}. Let $P_\gamma : X(\gamma) \to \gamma$ denote the restriction of P to $X(\gamma)$.

Mj and Sardar showed in [56] that $P : \Gamma_G \to \Gamma_Q$ is a metric graph bundle. The same reasoning shows that the restricted map P_γ is a metric graph bundle as well. We include the argument below for completeness.

Proposition 4.4. *Given $P : \Gamma_G \to \Gamma_Q$ as in Convention 4.1, the map $P : \Gamma_G \to \Gamma_Q$ and the restricted map $P_\gamma : X(\gamma) \to \gamma$ are metric graph bundles.*

Proof. For each vertex $q \in V(\Gamma_Q)$, $P^{-1}(q) = F_q$ is a copy of Γ_H, and so the induced path metric d_q is equal to d_H for all q. Hence, condition (B1) is satisfied by the function $f(n) := \max\{d_H(1, g) \mid d_G(1, g) \leq n\}$. Now, suppose $q_1, q_2 \in \Gamma_Q$ are adjacent vertices where $P(g_1H) = q_1$ and $P(g_2H) = q_2$. Since P maps edges between distinct cosets of Γ_H in Γ_G isometrically onto edges in Γ_Q, there exist some $h_1, h_2 \in H$ such that g_1h_1 and g_2h_2 are adjacent in Γ_G. Therefore $s = (g_1h_1)^{-1}g_2h_2 \in S_G$. Hence, for all $x_1 = g_1h \in V(F_{q_1})$, x_1 is adjacent to $x_1s = g_1hs = (g_1hh_1^{-1}g_1^{-1})g_2h_2$ in Γ_G. This element is contained in the coset $g_2H = F_{q_2}$ since H is normal in G, and so condition (B2) is satisfied. By Remark 2.22, $P_\gamma : X(\gamma) \to \gamma$ is also a metric graph bundle. □

Condition (B2) says that if we choose any lift g_0 of z_0, there exists $g_1 \in P^{-1}(z_1)$ such that $d_{X(\gamma)}(g_0, g_1) = 1$. Continuing in this fashion, we get a lift $\sigma : \gamma \to X(\gamma)$, where $\sigma(z_i) = g_i$, such that $d_{X(\gamma)}(g_i, g_{i+1}) = 1$ for all i. By the triangle inequality and the fact that γ is a geodesic in Γ_Q, we have that $d_{X(\gamma)}(g_i, g_j) \leq d_Q(Pg_i, Pg_j) = d_Q(z_i, z_j)$. But, as every path in Γ_G projects to a path in Γ_Q of no greater length and as $X(\gamma) \subseteq \Gamma_G$, we also have that $d_Q(Pa, Pb) \leq d_G(a, b) \leq d_{X(\gamma)}(a, b)$. Hence, for all $g_i, g_j \in \sigma(\gamma)$, $d_{X(\gamma)}(g_i, g_j) = d_G(g_i, g_j) = d_Q(z_i, z_j)$.

Proposition 4.5. *The space $X(\gamma)$ is hyperbolic.*

Proof. By Theorem 2.25, Γ_G satisfies the flaring condition since Γ_G and Γ_H are both hyperbolic and for each $q \in \Gamma_Q$, $F_q := p^{-1}(q)$ is a copy of Γ_H. Suppose that σ is a (K, C)-quasi-isometric lift of γ to $X(\gamma)$. Note that for all $a, b \in \gamma$, $d_Q(a, b) = d_Q(P \cdot \sigma(a), P \cdot \sigma(b)) \leq d_G(\sigma(a), \sigma(b))$. Also, since $X(\gamma) \subseteq \Gamma_G$, $d_G(\sigma(a), \sigma(b)) \leq d_{X(\gamma)}(\sigma(a), \sigma(b))$. So, any quasi-isometric lift of a portion of γ to $X(\gamma)$ is also a quasi-isometric lift when considered as a path in Γ_G. Thus, we have that $X(\gamma)$ satisfies the flaring condition. Additionally, the barycenters of ideal triangles in Γ_H are dense since the H-orbit of the barycenter of any ideal triangle in Γ_H is dense in Γ_H. Therefore by Theorem 2.24, we have that $X(\gamma)$ is hyperbolic. □

We now apply the work of Bowditch in [10] to our setting of hyperbolic group extensions. Let γ be as in Convention 4.2, and recall that $P : \Gamma_G \to \Gamma_Q$ is the projection map and $X(\gamma) := P^{-1}(\gamma)$.

Proposition 4.6. *The space $X(\gamma)$ with the induced path metric $d_{X(\gamma)}$ from Γ_G is a hyperbolic stack.*

Proof. We need to show that $X(\gamma)$ satisfies conditions (S1)-(S6). For each vertex $z_i \in \gamma$, choose some $g_i \in G$ such that $P(g_i) = z_i$. For each $i \in \mathbb{Z}$, the sheet X_i of $X(\gamma)$ is the copy of Γ_H which corresponds to the coset $g_i\Gamma_H$ of H in G. Since X_i and X_j represent different cosets of Γ_H in Γ_G for $i \neq j$, we have that $d_G(X_i, X_j) \leq d_{X(\gamma)}(X_i, X_j)$ is bounded away from 0 for $i \neq j$. Now, for all $i \in \mathbb{Z}$, let $\beta_i(n) := \max\{d_{X_i}(a, b) \mid d_{X(\gamma)}(a, b) \leq n\}$. Then, $\beta_i^{-1}(d_{X_i}(a, b)) \leq d_{X(\gamma)}(a, b) \leq d_{X_i}(a, b)$, and so condition (S1) is satisfied.

We see that condition (S2) is satisfied since $d_{X(\gamma)}(X_i, X_j) \geq d_Q(z_i, z_j) = |i-j|$. Similarly, we have that the Hausdorff distance between X_i and X_{i+1} in $X(\gamma)$ is at most 2, and so condition (S3) is satisfied. As each X_i is a copy of Γ_H which is δ-hyperbolic, we have that (S4) holds. Additionally, $\bigcup_{i\in\mathbb{Z}} X_i$ is in the 1-neighborhood of $X(\gamma)$, and so (S6) is satisfied. Finally, we have by Proposition 4.5 that condition (S5) is satisfied. Therefore, we have that $X(\gamma)$ is a bi-infinite hyperbolic stack. □

Let $X(\gamma)^+ := P^{-1}(\gamma^+)$ and $X(\gamma)^- := P^{-1}(\gamma^-)$. As the bi-infinite stack $X(\gamma)$ is hyperbolic by Proposition 4.5, applying Proposition 2.28 yields the following corollary.

Corollary 4.7. *The one-sided stacks $X(\gamma)^+$ and $X(\gamma)^-$ are hyperbolic.*

4.2 Surjectivity of the Cannon-Thurston map

Recall, as in Convention 4.2, that z_0 denotes a point on γ closest to the identity in Γ_Q, the vertices along γ between z_0 and z are labeled by $z_1, z_2, \ldots$, and the vertices along γ between z_0 and z' are labeled by $z_{-1}, z_{-2}, \ldots$. Then, for all $x_i \in X_i$, $Px_i = z_i$. Since $X(\gamma)$ satisfies property (B2) of being a metric graph bundle, every vertex in $X(\gamma)$ is contained in some 1-chain. This can also be seen for the following reason. Each consecutive fiber of $X(\gamma)$ is quasi-isometric Let $y \in \partial X(\gamma)$, and let $y_n \in X(\gamma)$ be a sequence of vertices in $X(\gamma)$ which converge to y. As every vertex in $X(\gamma)$ is contained in some 1-chain, for each $n \in \mathbb{N}$ we can construct a 1-chain $\underline{x}^n = (x_i^n)_{i=0}^{m_n}$ in $X(\gamma)$ with terminal point $x_{m_n}^n := y_n$ as follows. Without loss of generality, assume that $y_n \in X(\gamma)^+$. Then, there exists some $m_n \in \mathbb{N}$ and some $h \in H$ such that $y_n = g_{m_n} h \in X_{m_n} = g_{m_n}\Gamma_H$, where $g_i = \sigma(z_i)$. Set $x_{m_n}^n := y_n$ and define $x_{m_n-1}^n := g_{m_n} h g_{m_n}^{-1} g_{m_n-1}$. Given the point $x_{m_n-j}^n$, where $j \in \{1, 2, \ldots, m_n - 1\}$, set $x_{m_n-j-1}^n := x_{m_n-j}^n g_{m_n-j}^{-1} g_{m_n-j-1}$. Note that for each i, $x_i^n \in X_i = g_i\Gamma_H$, and so $\underline{x}^n$ defined in this fashion is a 1-chain in $X(\gamma)$ with terminal point y_n.

We now have a sequence of 1-chains $\underline{x}^n$ with terminal points y_n converging to $y \in \partial X(\gamma)$. Passing to a subsequence, we may assume that x_0^n converges to $x_0 \in X_0 \cup \partial X_0$. Suppose first that $x_0 \in X_0$. Then, since the points x_1^n remain in a compact subset of X_1, they subconverge on a point $x_1 \in X_1$ with $d_{X(\gamma)}(x_0, x_1) = 1$. Continuing on in this fashion, we can pass to a subsequence of our partial chains to get an infinite 1-chain $\underline{x} = \{x_0, x_1, \ldots, x_n, \ldots\}$ in $X(\gamma)$, where x_i^n converges to $x_i \in X_i$ for all i. Note that for large enough n, x_i^n remains uniformly close to x_i for arbitrarily many i. Hence, we must have that the terminal points of the chains $\underline{x}^n$ converge to the terminal point of $\underline{x}$ in $X(\gamma) \cup \partial X(\gamma)$. Since the chains $\underline{x}^n$ each have terminal point y_n, we therefore have that $y \in \partial X(\gamma)$ is the terminal point of a 1-chain in $X(\gamma)$, and so $y \in \partial^+ X(\gamma)$. Suppose now that $x_0 \in \partial X_0$. Then, by Lemma 2.30 (1), we have that $\partial\omega_\gamma(x_0) = y$, where $\partial\omega_\gamma : \partial X_0 \to \partial X(\gamma)$. Therefore, for all $y \in \partial X(\gamma)$, either y is the endpoint of a 1-chain in $X(\gamma)$, or $y \in \omega_\gamma(\partial X_0)$.

So, suppose that $(x_i)_{i \in \mathbb{Z}}$ is a 1-chain. We have that for all $i, j \in \mathbb{Z}$ with $i < j$,

$$\begin{aligned} d_Q(Px_i, Px_j) &\leq d_{X(\gamma)}(x_i, x_j) \\ &\leq d_{X(\gamma)}(x_i, x_{i+1}) + d_{X(\gamma)}(x_{i+1}, x_{i+1}) + \cdots + d_{X(\gamma)}(x_{j-1}, x_j) \\ &= d_Q(Px_i, Px_{i+1}) + \cdots + d_Q(Px_{j-1}, Px_j) \\ &= d_Q(z_i, z_j). \end{aligned}$$

Hence, every 1-chain in $X(\gamma)$ interpolates a geodesic in $X(\gamma)$ which is an isometric lift of γ. Furthermore, as for all $i, j \in \mathbb{Z}$, $d_Q(Px_i, Px_j) \leq d_G(x_i, x_j) \leq d_{X(\gamma)}(x_i, x_j)$, we have that every 1-chain interpolates a geodesic in Γ_G as well. Therefore, if $y \in \partial^+ X(\gamma)$ is the terminal point in $X(\gamma)$ of the positive 1-chain $(y_i)_{i \in \mathbb{N}}$, then the terminal point of this chain in Γ_G will determine a point of ∂G as well. As the only r-chains we will be considering in $X(\gamma)$ are 1-chains, all 1-chains in this space will now simply be referred to as chains.

Convention 4.8. Given $P : \Gamma_G \to \Gamma_Q$ as in Convention 4.1 and γ as in Convention 4.2, let $\sigma : \gamma \to \Gamma_G$ denote an isometric lift of γ such that for all $z_i \in \gamma$, $P(\sigma(z_i)) = z_i$ and set $g_i := \sigma(z_i)$. Let $X(\gamma) := P^{-1}(\gamma)$ and $X(\gamma)^+ := P^{-1}(\gamma^+)$ be the one-sided stacks which consist of the sheets $X_i = g_i\Gamma_H$ for all $i \in \mathbb{Z}$ and $i \in \mathbb{N}$, respectively. Denote by $\omega_\gamma : X_0 \to X(\gamma)$ and $\omega_\gamma^+ : X_0 \to X(\gamma)^+$ the inclusions of the sheet $X_0 = g_0\Gamma_H$ into $X(\gamma)$ and $X(\gamma)^+$ respectively. Define $i_{X_0} : \Gamma_H \to X_0$ as follows. Set $i_{X_0}(h) := g_0 \cdot g_0^{-1}hg_0 = hg_0$ for all vertices $h \in \Gamma_H$. Extend i_{X_0} to a map on all of Γ_H by sending an edge $[a, b]$ to a shortest path between ag_0 and bg_0. Now let $i_\gamma : \Gamma_H \to X(\gamma)$ be given by $i_\gamma := \omega_\gamma \circ i_{X_0}$ and $i_\gamma^+ : \Gamma_H \to X(\gamma)^+$ be given by $i_\gamma^+ := \omega_\gamma^+ \circ i_{X_0}$. Note that if $1 \in \gamma$, then i_{X_0} and i_γ^+ are simply the identity inclusion map $i : \Gamma_H \to \Gamma_G$.

Lemma 4.9. *The maps $i_\gamma : \Gamma_H \to X(\gamma)$ and $i_\gamma^+ : \Gamma_H \to X(\gamma)^+$ as given in Convention 4.8 extend continuously to the maps $\widehat{i}_\gamma : \widehat{\Gamma_H} \to \widehat{X(\gamma)}$ and $\widehat{i}_\gamma^+ : \widehat{\Gamma_H} \to \widehat{X(\gamma)^+}$, respectively.*

Proof. Given $\omega_\gamma : X_0 \to X(\gamma)$ and $\omega_\gamma^+ : X_0 \to X(\gamma)^+$ as in Convention 4.8, note that Proposition 2.29 gives that the Cannon-Thurston maps $\partial\omega_\gamma : \partial X_0 \to \partial X(\gamma)$ and $\partial\omega_\gamma^+ : \partial X_0 \to \partial X(\gamma)$ both exist. Let $\widehat{\omega_\gamma} : \widehat{X_0} \to \widehat{X(\gamma)}$ and $\widehat{\omega_\gamma^+} : \widehat{X_0} \to \widehat{X(\gamma)^+}$ denote the continuous extensions of ω_γ and ω_γ^+. For all $g \in G$, conjugation by g gives an automorphism of H which takes $h \in H$ to $g^{-1}hg$. This automorphism is a quasi-isometry from Γ_H to itself. So, $i_{X_0} : \Gamma_H \to X_0$ is a quasi-isometry from Γ_H to $X_0 = g_0\Gamma_H$, and so extends to a homeomorphism $\partial i_{X_0} : \partial H \to \partial X_0$. Hence, $\partial i_\gamma := \partial\omega_\gamma \circ \partial i_{X_0}$ and $\partial i_\gamma^+ : \partial\omega_\gamma^+ \circ \partial i_{X_0}$ exist, are continuous, and extend i_γ and i_γ^+ continuously to the maps $\widehat{i_\gamma}$ and $\widehat{i_\gamma^+}$, respectively. ☐

Lemma 4.9 allows us to now refer to the maps ∂i_γ and ∂i_γ^+ as Cannon-Thurston maps. The goal of the remainder of this section is to show that the maps ∂i_γ and ∂i_γ^+ are surjective. We will first show surjectivity for the case where the geodesic γ lives over the identity in Γ_Q.

Convention 4.10. Let $\gamma = (z', z)$ be as in Convention 4.2 and let $\gamma' := z_0^{-1} \cdot \gamma = (z_0^{-1}z', z_0^{-1}z)$. Note that $1 \in V(\gamma')$. For each $z_i \in \gamma$, let $z_i' := z_0^{-1} \cdot z_i$. Given $\sigma : \gamma \to \Gamma_G$ as in Convention 4.8, let $\sigma' : \gamma' \to \Gamma_G$ be such that $\sigma' := g_0^{-1} \cdot \sigma$. Set $g_i' := \sigma'(z_i')$, and denote the sheet $g_i'\Gamma_H$ by X_i'. Note that the sheet X_0' is the identity coset $1 \cdot \Gamma_H$, and so the map $i_{X_0'} : \Gamma_H \to X_0'$ is the identity map.

In a similar manner as Bowditch [10], we define a map $\widehat{\tau}_{\gamma'} : \widehat{X(\gamma')} \to \widehat{\Gamma_G}$ with $\widehat{i} = \widehat{\tau}_{\gamma'} \circ \widehat{i}_{\gamma'}$ and will later show that $\partial\tau_{\gamma'} : \partial X(\gamma') \to \partial G$ is continuous. Let $\tau_{\gamma'} := \widehat{\tau}_{\gamma'}|_{X(\gamma')}$ be the identity inclusion of $X(\gamma')$ into Γ_G given by $\tau_{\gamma'}(g) = g$. Note that for all $h \in H$, $\tau_{\gamma'} \circ i_{\gamma'}(h) = h = i(h)$. As the map $i_{X_0'}$ is the identity map, $\partial\omega_{\gamma'} = \partial i_{\gamma'}$. So by Proposition 2.29, we have that $\partial X(\gamma') = \partial i_{\gamma'}(\partial H) \cup \partial^{\pm} X(\gamma')$. If $(y_i)_{i\in\mathbb{N}}$ is a positive 1-chain in $X(\gamma')$ with endpoint $y \in \partial^+ X(\gamma')$, then $(\tau_{\gamma'}(y_i))_{i\in\mathbb{N}}$ interpolates a geodesic ray in Γ_G with the same label as the geodesic ray interpolated by (y_i) in $X(\gamma')$. Denote the endpoint of this geodesic ray in Γ_G by $\overline{y} \in \partial G$, and for all $y \in \partial^{\pm} X(\gamma')$ define $\partial\tau_{\gamma'}(y) := \overline{y}$. Finally, for all $a \in \partial H$, define $\partial\tau_{\gamma'}(\partial i_{\gamma'}(a)) := \partial i(a)$. Note that if $(x_i)_{i\in\mathbb{N}}$ and $(y_i)_{i\in\mathbb{N}}$ are distinct but equivalent 1-chains in $X(\gamma')$, then the geodesic rays interpolated by these chains are Hausdorff close in both $X(\gamma')$ and Γ_G. Hence, $\tau_{\gamma'}$ is well-defined on equivalence classes of chains. To finish showing that $\widehat{\tau}_{\gamma'}$ is well-defined, we need the following lemma.

Lemma 4.11. *Let γ' be as in Convention 4.10. Suppose $(\underline{x}^n)_{n\in\mathbb{N}}$ is a sequence of positive chains in $X(\gamma')$, where $\underline{x}^n = (x^n_i)_{i\in\mathbb{N}}$ is a positive chain with terminal point $y_n \in \partial^+ X(\gamma')$. Suppose also that in $\widehat{X(\gamma')}$, $y_n \to y \in \partial X(\gamma')$ and in $\widehat{X'_0}$, $x^n_0 \to \partial i_{X'_0}(a) \in \partial X'_0$. Then in $\widehat{\Gamma_G}$, $\widehat{\tau}_{\gamma'}(y_n) \to \widehat{i}(a)$.*

Proof. Let $f(n) = \max\{d_{X'_0}(a,b) \mid d_G(\tau_{\gamma'}(a), \tau_{\gamma'}(b)) \le n\}$. Note that since Γ_G is finitely generated, such a maximum exists and that $f(n) \to \infty$ as $n \to \infty$. For each $n \in \mathbb{N}$, there exists $a_n \in \Gamma_H$ such that $x^n_0 = i_{X'_0}(a_n) = a_n$. As $x^n_0 \to \partial i_{X'_0}(a)$, this implies that $a_n \to a \in \partial H$ in $\widehat{\Gamma_H}$. Let $\lambda_n = [\widehat{\tau}_{\gamma'}(x^n_0), \widehat{\tau}_{\gamma'}(y_n))_G = [i(a_n), \widehat{\tau}_{\gamma'}(y_n))_G$ be the geodesic ray in Γ_G interpolated by $(\tau_{\gamma'}(x^n_i))_{i\in\mathbb{N}}$ for each $n \in \mathbb{N}$.

Suppose that in $\widehat{\Gamma_G}$, $\lim_{n\to\infty} \widehat{\tau}_{\gamma'}(y_n) \neq \lim_{n\to\infty} i(a_n)$. Then, there exist constants $R, N > 0$ such that for all $n \ge N$, $d_G(1, \lambda_n) \le R$. So, for each $n \ge N$ there exists some point $x^n_{i_n}$ in the chain $\underline{x}^n$ such that $d_G(1, \tau_{\gamma'}(x^n_{i_n})) \le R$. Then, we have that

$$\begin{aligned} d_G(1, \tau_{\gamma'}(x^n_{i_n})) &\ge d_Q(P \cdot 1, P \cdot \tau_{\gamma'}(x^n_{i_n})) \\ &= d_Q(1, P \cdot x^n_{i_n}) \\ &= |i_n|. \end{aligned}$$

As $d_G(1, \tau_{\gamma'}(x^n_{i_n})) \le R$, this means that $|i_n| \le R$. Note that since $\underline{x}^n$ is a 1-chain, we have that $d_G(\tau_{\gamma'}(x^n_0), \tau_{\gamma'}(x^n_{i_n})) = d_{X(\gamma')}(x^n_0, x^n_{i_n}) = |i_n| \le R$ So,

$$\begin{aligned} d_{X'_0}(1, x^n_0) &\le f(d_G(\tau_{\gamma'}(1), \tau_{\gamma'}(x^n_0))) \\ &= f(d_G(1, \tau_{\gamma'}(x^n_0))) \\ &\le f(d_G(1, \tau_{\gamma'}(x^n_{i_n})) + d_G(\tau_{\gamma'}(x^n_{i_n}), \tau_{\gamma'}(x^n_0))) \\ &\le f(2R). \end{aligned}$$

But, $d_{X'_0}(1, x^n_0) \to \infty$ as $n \to \infty$ since $x^n_0 \to \partial i_{X'_0}(a) \in \partial X'_0$, and so we have a contradiction. Therefore, $d_G(1, \lambda_n) \to \infty$ as $n \to \infty$. Hence, in $\widehat{\Gamma_G}$, $\lim_{n\to\infty} \tau_{\gamma'}(x^n_0) = \lim_{n\to\infty} \widehat{\tau}_{\gamma'}(y_n)$. As $\tau_{\gamma'}(x^n_0) = i(a_n)$ and $i(a_n) \to \widehat{i}(a)$ as $n \to \infty$, we have that $\widehat{\tau}_{\gamma'}(y_n) \to \widehat{i}(a)$ as desired. □

Lemma 4.12. *The map* $\widehat{\tau}_{\gamma'} : \widehat{X(\gamma')} \to \widehat{\Gamma_G}$ *is well-defined and satisfies* $\widehat{\tau}_{\gamma'} \circ \widehat{i}_{\gamma'} = \widehat{i} : \widehat{\Gamma_G} \to \widehat{\Gamma_H}$

Proof. If $x \in \Gamma_H$, then $\tau_{\gamma'} \circ i_{\gamma'}(x) = x = i(x)$. Similarly, if $a \in \partial H$, then $\partial\tau_{\gamma'} \circ \partial i_{\gamma'}(a) = \partial i(a)$. So, $\widehat{i} = \widehat{\tau}_{\gamma'} \circ \widehat{i}_{\gamma'}$. Now, it suffices to show that $\partial\tau_{\gamma'} : \partial X(\gamma') \to \partial G$ is well-defined. First, we need to show that if $y \in \partial^+ X(\gamma')$ and $a \in \partial H$ are such that $\partial i_{\gamma'}(a) = y$, then $\partial\tau_{\gamma'}(y) = \partial\tau_{\gamma'}(\partial i_{\gamma'}(a))$. So, suppose that $y \in \partial^+ X(\gamma')$ and $a \in \partial H$ are such that $\partial i_{\gamma'}(a) = y$. Since $\partial i_{\gamma'}(a) = y$ and $\partial i_{X_0'}$ is the identity, this implies that $\partial i_{\gamma'}(a) = \partial\omega_{\gamma'} \circ \partial i_{X_0'}(a) = \partial\omega_{\gamma'}(a) = y$. By Lemma 2.30 (1), there exists a sequence $(\underline{x}^n)_n$ of positive chains, each converging to y, with x_0^n converging to $a = \partial i_{X_0'}(a) \in \partial X_0'$. By Lemma 4.11, the existence of such a sequence of chains implies that $\partial\tau_{\gamma'}(y) = \partial i(a)$. Hence, $\partial\tau_{\gamma'}(y) = \partial\tau_{\gamma'}(\partial i_{\gamma'}(a))$.

Now, suppose that $a, b \in \partial H$ with $a \neq b$ are such that $\partial i_{\gamma'}(a) = \partial i_{\gamma'}(b)$. Since $\partial i_{X_0'}$ is the identity, this implies that $\partial\omega_{\gamma'}(a) = \partial\omega_{\gamma'}(b)$. By Lemma 2.30 (4), we may assume without loss of generality that $\partial\omega^+_{\gamma'}(a) = \partial\omega^+_{\gamma'}(b)$. Since a and b are distinct, we have by Lemma 2.30 (3) that $\partial\omega^+_{\gamma'}(a) = \partial\omega^+_{\gamma'}(b) = y \in \partial^+ X(\gamma')$. By Lemma 2.30 (2), we now have that $\partial\omega_{\gamma'}(a) = \partial\omega_{\gamma'}(b) = y$. So by the same reasoning as above, Lemma 2.30 (1) and Lemma 4.11 give that $\partial\tau_{\gamma'}(\partial i_{\gamma'}(a)) = \partial\tau_{\gamma'}(\partial i_{\gamma'}(b)) = \partial\tau_{\gamma'}(y) = \overline{y}$. □

Corollary 4.13. *If* $a, b \in \partial H$ *are such that* $\partial i^+_{\gamma'}(a) = \partial i^+_{\gamma'}(b)$ *then* $\partial i(a) = \partial i(b)$.

Proof. Suppose $a, b \in \partial H$ are such that $\partial i^+_{\gamma'}(a) = \partial i^+_{\gamma'}(b)$. If $a = b$, then $\partial i(a) = \partial i(b)$. So, suppose $a \neq b$. As $\partial i_{X_0'}$ is the identity, we have that $\partial\omega^+_{\gamma'}(a) = \partial\omega^+_{\gamma'}(b)$. By Lemma 2.30 (3), there exists $y \in \partial^+ X(\gamma')$ such that $\partial\omega^+_{\gamma'}(a) = \partial\omega^+_{\gamma'}(b) = y$. By Lemma 2.30 (2), this implies that $\partial\omega_{\gamma'}(a) = \partial\omega_{\gamma'}(b)$. So, $\partial i_{\gamma'}(a) = \partial\omega_{\gamma'} \circ \partial i_{X_0'}(a) = \partial\omega_{\gamma'} \circ \partial i_{X_0'}(b) = \partial i_{\gamma'}(b)$. As $\partial\tau_{\gamma'}$ is well-defined by Lemma 4.12, we have that $\partial\tau_{\gamma'}(\partial i_{\gamma'}(a)) = \partial\tau_{\gamma'}(\partial i_{\gamma'}(b))$, and so $\partial i(a) = \partial i(b)$. □

The goal of the remainder of this section is to use this work of Bowditch to prove that the Cannon-Thurston map $\partial i^+_\gamma : \partial H \to \partial X(\gamma)^+$ is surjective.

Lemma 4.14. *Fix* $\gamma = (z', z) \subseteq \Gamma_Q$ *as in Convention 4.2 and let* $X(\gamma)^+$ *be as described above. Let* $(y_n)_{n\in\mathbb{N}}$ *be a* 1*-chain in* $X(\gamma)^+$ *and denote the word which labels the geodesic from* y_0 *to* y_n *in* $X(\gamma)^+$ *by* α_n*. Fix some* $h \in H$ *of infinite order and let* ρ_n *denote any path in* $X(\gamma)^+$ *which is the concatenation of a path labeled by* α_n *followed by a path labeled by* h *and finally a path labeled by* α_n^{-1}*. Then, there exists some constant* $C \geq 0$ *independent of* n *(but dependent on* h*) such that for all* n*,* ρ_n *is a* $(1, C)$*-quasigeodesic in* $X(\gamma)^+$*.*

Proof. Let $(y_n)_{n\in\mathbb{N}}$ be a 1-chain in $X(\gamma)^+$, and for each $n \geq 0$ let α_n denote the word which labels the geodesic from y_0 to y_n. Given $h \in H$, let β denote any quasigeodesic in $X(\gamma)^+$

labeled by h. Since $h \in H$ is fixed, there exists some constant $C' \geq 0$ such that β is a $(1, C')$-quasigeodesic. For each $n \geq 0$, let $[x_n = y_0, y_n]$ be the geodesic in $X(\gamma)^+$ from $x_n = y_0$ to y_n labeled by α_n, let $z_n \in X(\gamma)^+$ be a point such that β is a quasigeodesic in $X(\gamma)^+$ from y_n to z_n, and let $[z_n, w_n]$ be the geodesic in $X(\gamma)^+$ labeled by α_n^{-1}. Denote by δ_n the label of the geodesic in $X(\gamma)^+$ between x_n and w_n.

For each $n \geq 0$, consider the quadrilateral in $X(\gamma)^+$ with vertices $y_0 = x_n$, y_n, z_n, w_n, and with sides $[x_n, y_n]$ labeled by α_n, β labeled by h, $[z_n, w_n]$ labeled by α_n^{-1}, and $[x_n, w_n]$ labeled by δ_n. Unless otherwise specified, we will denote $d_{X(\gamma)^+}$ simply by d, and all geodesic and quasigeodesic segments considered are geodesics or quasigeodesics in $X(\gamma)^+$.

As before, we need to show that if p and q are arbitrary points on $\rho_n = [x_n, y_n] \cup \beta \cup [z_n, w_n]$, then the distance between p and q along ρ_n is at most $d(p, q) + C$. There are two cases to consider. By Proposition 3.2, either there is a point on $[x_n, w_n]$ at most distance 2δ in $X(\gamma)^+$ from a point on $[y_n, z_n]$, or there is a point on the side $[x_n, y_n]$ at most distance 2δ in $X(\gamma)^+$ from a point on the side $[z_n, w_n]$. If there is some point on the side $[x_n, w_n]$ within 2δ of a point on the side $[y_n, z_n]$, then Lemma 3.4 gives that $[x_n, y_n] \cup [y_n, z_n] \cup [z_n, w_n]$ is a $(1, 4\delta + 4d(y_n, z_n))$-quasigeodesic. Since β is a $(1, C')$-quasigeodesic between y_n and z_n, this gives that ρ_n is a $(1, C)$-quasigeodesic for some $C \geq 0$.

So, suppose now that the two sides labeled by α_n and α_n^{-1} come within 2δ of each other in $X(\gamma)^+$. We make the following claim:

Claim: If $a \in [x_n, y_n]$ and $a' \in [z_n, w_n]$ are the furthest points in $X(\gamma)^+$ from y_n and z_n, respectively, such that $d(a, a') \leq 2\delta$, then there is some constant $K > 0$ dependent on h but independent of n such that $\max\{d(a, y_n), d(a', z_n)\} \leq K$.

Assuming this claim, we will now show that ρ_n is a $(1, C)$-quasigeodesic in $X(\gamma)^+$. First fix $p \in [x_n, y_n]$ and $q \in \beta$. Since in $X(\gamma)^+$, $(p, q; X(\gamma)^+)_{y_n}$ is bounded by $|\beta|_{X(\gamma)^+} \leq |h|_H$, we have that

$$\begin{aligned} d(p, y_n) + d(y_n, q) &= d(p, q) + 2(p, q; X(\gamma)^+)_{y_n} \\ &\leq d(p, q) + 2|h|_H. \end{aligned}$$

So, suppose $p \in [x_n, y_n]$ and $q \in [z_n, w_n]$. If $p \in [a, y_n]$ and $q \in [a', z_n]$, then $d(p, y_n) + |\beta|_{X(\gamma)^+} + d(z_n, q) \leq d(p, q) + |h|_H + 2K$. Now suppose $p \in [x_n, a]$ and $q \in [a', z_n]$. Since

$d(a, a') \leq 4\delta$, we have by the triangle inequality that

$$\begin{aligned} d(q, a) &\leq d(q, a') + d(a', a) \\ &\leq K + 2\delta, \text{ and} \end{aligned}$$

$$d(p, a) \leq d(p, q) + d(q, a).$$

Therefore,

$$\begin{aligned} d(p, a) + d(a, y_n) + |\beta|_{X(\gamma)^+} + d(z_n, q) &\leq d(p, q) + d(q, a) + K + |h|_H + K \\ &\leq d(p, q) + 3K + 2\delta + |h|_H. \end{aligned}$$

The final case to consider is when $p \in [x_n, a]$ and $q \in [w_n, a']$. In this case, there must be a point $u \in [p, q]$ and $v \in [a, a']$ such that $d(u, v) \leq 2\delta$. This is because by choice of a and a', there are no points at which $[q, a']$ is within a distance of 2δ of $[p, a]$ in $X(\gamma)^+$. So, we have that $d(q, v) \leq d(q, u) + d(u, v)$ and $d(p, v) \leq d(p, u) + d(u, v)$. Additionally, $d(p, a) \leq d(p, v) + d(v, a)$ and $d(q, a') \leq d(q, v) + d(v, a')$. Hence, we have that

$$\begin{aligned} & d(p, a) + d(a, y_n) + |\beta|_{X(\gamma)^+} + d(z_n, a') + d(a', q) \\ &\quad \leq d(p, a) + K + |h|_H + K + d(a', q) \\ &\quad \leq d(p, v) + d(v, a) + 2K + |h|_H + d(q, v) + d(v, a') \\ &\quad \leq d(p, v) + d(q, v) + 2K + |h|_H + 2\delta \\ &\quad \leq d(p, u) + d(q, u) + 2d(u, v) + 2K + |h|_H + 2\delta \\ &\quad \leq d(p, q) + 2K + |h|_H + 6\delta. \end{aligned}$$

Proof of Claim: Suppose to the contrary that there is no such bound on the how long the sides labeled by α_n and α_n^{-1} stay uniformly close in $X(\gamma)^+$. Let S_Q be the generating set for Q and let $L = \{w \in \Sigma_Q^* \mid w \text{ a geodesic in } Q\}$. Since Q is a hyperbolic group, the language L of geodesic words is a regular language for Q (see [25]) which is accepted by some finite state automaton, $\mathcal{A}$, with start state s_0. Then, $\gamma^+ = [z_0, z) \subseteq \Gamma_Q$ gives an infinite path from s_0 in $\mathcal{A}$ such that all states are accept states. Let γ_n denote the initial portion of the path γ^+ of length n, i.e., $\gamma_n := P([y_0 = x_n, y_n])$.

For each n, assume without loss of generality that the side of ρ_n labeled by α_n begins at the vertex y_0 and ends at the vertex y_n. Let y_{i_n} denote the vertex along the side α_n where the side labeled by α_n and the side labeled by α_n^{-1} begin to be 2δ close. Note that after the point y_{i_n}, the sides labeled by α_n and α_n^{-1} will continue to travel within a distance of $|h|_{X(\gamma)^+}$ of each other in $X(\gamma)^+$. Project the $X(\gamma)^+$-geodesic $[y_{i_n}, y_n]$ to Q and feed this geodesic, $P([y_{i_n}, y_n])$, into $\mathcal{A}$. Note that by assumption, the length of these geodesics go to infinity as $n \to \infty$. So, there will be some $n > 0$ for which some state in $\mathcal{A}$ repeats more times than the number of words in G of length at most $|h|_{X(\gamma)^+}$. Note that the label of any loop in $\mathcal{A}$ is a periodic Q-geodesic word. Since there is a state that repeats more times than the number of words in G of length at most $|h|_{X(\gamma)^+}$, it follows that there is some subpath of $[y_{i_n}, y_n]$ labeled by a word $v \in \Sigma_Q^*$ which has infinite order in Q and some word $m \in \Sigma_G^*$ of length at most $|h|_{X(\gamma)^+}$ such that in G, $P^{-1}(v)m(P^{-1}(v))^{-1} = m$ and such that h is conjugate to m in G. As h has infinite order in G and h is conjugate to m, it follows that m is infinite order in G as well. As $P^{-1}(v)$ and m commute in G, this implies that $(P^{-1}(v))^p = m^q$, for some $p, q \neq 0$. But then $v^p = 1$ in Q, because h projects to the identity in Q which means that m projects to the identity in Q as well. The fact that $v^p = 1$ contradicts v being a periodic geodesic in Q. This completes the proof of the claim and the lemma. □

Theorem B. *Let $1 \to H \to G \to Q \to 1$ be a short exact sequence of infinite, finitely generated, word-hyperbolic groups. Let $z, z' \in \partial Q$ be distinct and let $\gamma \subseteq \Gamma_Q$ be a bi-infinite geodesic in Γ_Q between z and z'. Let $i_\gamma^+ : \Gamma_H \to X(\gamma)^+$ be the inclusion of Γ_H into the semi-infinite stack $X(\gamma)^+$ over $\gamma^+ = [z_0, z)$, and let $i_\gamma : \Gamma_H \to X(\gamma)$ be the inclusion of Γ_H into the bi-infinite stack $X(\gamma)$, as in Convention 4.8. Then,*

1. *the Cannon-Thurston map $\partial i_\gamma^+ : \partial H \to \partial X(\gamma)^+$ is surjective; and*
2. *the Cannon-Thurston map $\partial i_\gamma : \partial H \to \partial X(\gamma)$ is surjective.*

Proof. Let $\gamma = (z', z) \subseteq \Gamma_Q$ be as in Convention 4.2 and let $\gamma' := z_0^{-1} \cdot \gamma$ be as in Convention 4.10. We will first show that the Cannon-Thurston maps $\partial i_{\gamma'}^+ : \partial H \to \partial X(\gamma')^+$ and $\partial i_{\gamma'} : \partial H \to \partial X(\gamma')$ are surjective.

Consider first the map $\partial i_{\gamma'}^+ : \partial H \to \partial X(\gamma)^+$. Since $\partial i_{\gamma'}^+ = \partial \omega_{\gamma'}^+ \circ \partial i_{X_0'}$ and $\partial i_{X_0'}$ is the identity, it suffices to show that $\partial \omega_{\gamma'}^+$ is surjective. By Proposition 2.29 (3), we need only show that if $y \in \partial^+ X(\gamma')^+$, then there exists $a \in \partial X_0'$ such that $\partial i_{\gamma'}^+(a) = y$. So, suppose that $y \in \partial^+ X(\gamma')^+$ is the endpoint of the chain (y_n) and fix some $h \in H$ of infinite order. Let α_n be the word which labels the path from y_0 to y_n in $X(\gamma')^+$, and consider the path ρ_n in $X(\gamma')^+$ which is labeled by the word $\alpha_n h \alpha_n^{-1}$.

By Lemma 4.14, ρ_n is a $(1, C)$-quasigeodesic in $X(\gamma')^+$ for some C independent of n. Let h_n be the word which labels the geodesic in X_0' between the endpoints of ρ_n. Since

$|h_n|_H \to \infty$, there exists a subsequence h_{n_i} such that $y_0 h_{n_i} \to a \in \partial X_0'$. Since $\partial\omega_{\gamma'}^+$ is a continuous extension of $\omega_{\gamma'}^+$, we have that

$$\lim_{n_i\to\infty} \omega_{\gamma'}^+(y_0 h_{n_i}) = \partial\omega_{\gamma'}^+ \lim_{n_i\to\infty} y_0 h_{n_i} = \partial\omega_{\gamma'}^+(a).$$

Since $y_{n_i} \to y$ and since ρ_{n_i} is a quasigeodesic and $y_{n_i} \in \rho_{n_i}$, it follows that $\lim_{n_i\to\infty} y_{n_i} = \lim_{n_i\to\infty} \omega_\gamma^+(y_0 h_{n_i}) = y$ in $\widehat{X(\gamma')^+}$.

To see that $\partial i_{\gamma'} : \partial H \to \partial X(\gamma')$ is surjective, note that by Proposition 2.29 (2), $\partial X(\gamma') = \partial^+ X(\gamma') \cup \partial^- X(\gamma') \cup \partial i_{\gamma'}(\partial H)$. Note that the map $\widehat{i}_{\gamma'} : \widehat{H} \to \widehat{X(\gamma')}$ is defined in the same way as $\widehat{i}_{\gamma'}^+$. So, to show the surjectivity of $\partial i_{\gamma'}$, it suffices to note that in the above argument, we can replace $y \in \partial^+ X(\gamma')^+$ with $y' \in \partial^- X(\gamma')^-$. As the same reasoning holds, we have that $\partial i_{\gamma'} : \partial H \to \partial X(\gamma')$ is surjective as well.

Now, let $t_{g_0}^H : \Gamma_H \to g_0\Gamma_H$, $t_{g_0} : X(\gamma') \to X(\gamma)$, and $t_{g_0}^+ : X(\gamma')^+ \to X(\gamma)^+$ denote the maps induced by left-translation of the vertices of Γ_H, $X(\gamma')$, and $X(\gamma')^+$, respectively, by the element $g_0 = \sigma(z_0)$. Note that for all $h \in H$, $\omega_\gamma \circ t_{g_0}^H(h) = t_{g_0} \circ i_{\gamma'}(h)$ and $\omega_\gamma^+ \circ t_{g_0}^H(h) = t_{g_0}^+ \circ i_{\gamma'}^+(h)$. Since $t_{g_0}^H$, t_{g_0}, and $t_{g_0}^+$ are isometries, these maps extend continuously to the boundary maps $\partial t_{g_0}^H : \partial H \to \partial g_0 H$, $\partial t_{g_0} : \partial X(\gamma') \to \partial X(\gamma)$, and $\partial t_{g_0}^+ : \partial X(\gamma')^+ \to \partial X(\gamma)$, respectively, which are homeomorphisms. Hence, we have that for all $a \in \partial H$, $\partial\omega_\gamma \circ \partial t_{g_0}^H(a) = \partial t_{g_0} \circ \partial i_{\gamma'}(a)$ and $\partial\omega_\gamma^+ \circ \partial t_{g_0}^H(a) = \partial t_{g_0}^+ \circ \partial i_{\gamma'}^+(a)$. As $\partial i_{\gamma'}$ and $\partial i_{\gamma'}^+$ are surjective by the above argument and as $\partial t_{g_0}^H$, ∂t_{g_0}, and $\partial t_{g_0}^+$ are homeomorphisms, this implies that $\partial\omega_\gamma$ and $\partial\omega_\gamma^+$ are surjective.

As noted previously, each $g \in G$ gives rise to an automorphism ϕ_g of H with $\phi_g(h) = g^{-1}hg$. This automorphism of H induces a quasi-isometry of Γ_H taking an edge $[u, v]$ to a shortest edge path between $\phi_g(u)$ and $\phi_g(v)$. As $\phi_g : \Gamma_H \to \Gamma_H$ is a quasi-isometry, it extends to a homeomorphism $\partial\phi_g : \partial H \to \partial H$. Recall that $i_\gamma = \omega_\gamma \circ t_{g_0}^H \circ \phi_{g_0}$ and $i_\gamma^+ = \omega_\gamma^+ \circ t_{g_0}^H \circ \phi_{g_0}$. So, $\partial i_\gamma = \partial\omega_\gamma \circ \partial t_{g_0}^H \circ \partial\phi_{g_0}$ and $\partial i_\gamma^+ = \partial\omega_\gamma^+ \circ \partial t_{g_0}^H \circ \partial\phi_{g_0}$. As $\partial\omega_\gamma$ and $\partial\omega_\gamma^+$ are surjective, and as $\partial t_{g_0}^H$ and $\partial\phi_{g_0}$ are homeomorphisms, we have that ∂i_γ and ∂i_γ^+ are surjective. □

Recall that given the maps $\partial i_\gamma^+ : \partial H \to \partial X(\gamma)^+$ and $\partial i_\gamma : \partial H \to \partial X(\gamma)$, Bowditch defines a map $\partial\tau^+ : \partial X(\gamma)^+ \to \partial X(\gamma)$ with $\partial i_\gamma = \partial\tau^+ \circ \partial i_\gamma^+$. This map is given by $\partial\tau^+(y) = y$ for all $y \in \partial^+ X(\gamma)^+$, and $\partial\tau^+ \circ \partial i_\gamma^+(a) = \partial i_\gamma(a)$ for all $a \in \partial H$. We can now show the following about the map $\partial\tau^+$.

Corollary 4.15. *The map $\partial\tau^+ : \partial X(\gamma)^+ \to \partial X(\gamma)$ as defined above is surjective.*

Proof. By Proposition 2.29 (2) and Theorem B (2), we have that $\partial X(\gamma) = \partial i_\gamma(\partial H)$. Suppose $y \in \partial X(\gamma)$. By Theorem B (2), there exists $a \in \partial H$ such that $\partial i_\gamma(a) = y$. Then by definition of τ^+, we have that $\tau^+(\partial i_\gamma^+(a)) = \partial i_\gamma(a) = y$. □

Chapter 5

Ending laminations

Recall that by Convention 4.1 we have fixed a short exact sequence $1 \to H \to G \to Q \to 1$ of three infinite word-hyperbolic groups with Cayley graphs Γ_H, Γ_G, and Γ_Q, respectively. For each $g \in G$, conjugation by g gives an automorphism ϕ_g of H defined by $\phi_g(h) = g^{-1}hg$. Note that ϕ_g provides a bijection of the vertices of Γ_H which is a quasi-isometry of Γ_H with parameters depending on $|g|$. As such, ϕ_g extends to a homeomorphism of ∂H that coincides with the action of left-multiplication by g^{-1}. We will also denote this homeomorphism by ϕ_g. When $\lambda = [a, b]$ is a geodesic segment in Γ_H, we will denote a geodesic in Γ_H between $\phi_g(a)$ and $\phi_g(b)$ by λ_g. Similarly, if $\lambda = (u, v)$ is a bi-infinite geodesic in Γ_H with endpoints in ∂H, then $\lambda_g = (\phi_g(u), \phi_g(v)) = (g^{-1}u, g^{-1}v)$ also denotes the bi-infinite geodesic in Γ_H between the images of the endpoints of λ under the homeomorphism ϕ_g.

Given $\kappa \geq 1$ and $\epsilon \geq 0$, define a (κ, ϵ)-*quasi-isometric section* to be a (κ, ϵ)-quasi-isometric embedding $\sigma : \Gamma_Q \to \Gamma_G$ such that $P \cdot \sigma$ is the identity map on Γ_Q. The existence of such a quasi-isometric section in the setting of Convention 4.1 is guaranteed by Mosher [57]. If $\gamma \subseteq \Gamma_Q$ is a bi-infinite geodesic or a geodesic ray, we will also refer to a (κ, ϵ)-quasi-isometric embedding $\sigma : \gamma \to \Gamma_G$ as a quasi-isometric section. All sections we consider in this paper are assumed to take vertices to vertices and edges to edge-paths.

Definition 5.1. An *algebraic lamination* on H is defined to be a non-empty subset L of the double boundary $\partial^2 H$ which is closed, symmetric (flip-invariant), and H-invariant. If $L \subseteq \partial^2 H$ is an algebraic lamination, an element $(p, q) \in L$ will be referred to as a *leaf* of the lamination. As each point $(p, q) \in \partial^2 H$ can be represented by a bi-infinite geodesic λ in Γ_H from p to q, we will sometimes refer to the geodesic λ as a leaf of the lamination as well.

In [47], Mitra describes a set of algebraic ending laminations on Γ_H associated to the hyperbolic group extension (*) which are parametrized by points in the Gromov boundary of Γ_Q. These algebraic ending laminations are defined below.

Convention 5.2. Fix $\kappa \geq 1$ and $\epsilon \geq 0$, and let $\sigma : \Gamma_Q \to \Gamma_G$ be a quasi-isometric section of Γ_Q into Γ_G. For a fixed $z \in \partial Q$, let $[1, z) \subseteq \Gamma_Q$ be a geodesic ray from the identity to z. Denote the n^{th} vertex along $[1, z)$ by z_n, and set $g_n := \sigma(z_n)$.

Definition 5.3 (Mitra, [47])**.** Let $z \in \partial Q$.

1. Let $h \in H$ be an element of infinite order. Choose a geodesic $[1, z) \subseteq \Gamma_Q$ as in Convention 5.2. Define $R_{z,h}$ to be the set of all pairs $(a, aw) \in H \times H$ such that there is some $n \geq 0$ for which $w \in [g_n h g_n^{-1}]_H$ and w is a conjugacy minimal representative of $g_n h g_n^{-1}$ in H. Let $\overline{R_{z,h}}$ denote the closure of $R_{z,h}$ in $\widehat{H} \times \widehat{H}$, and set

$$\Lambda_{z,h} := \overline{R_{z,h}} \cap \partial^2 H.$$

   So, $\Lambda_{z,h}$ consists of all points $(p, q) \in \partial^2 H$ for which there exists a sequence $(a_{n_i}, a_{n_i} w_{n_i})$ in $H \times H$ such that $(a_{n_i}, a_{n_i} w_{n_i})$ converges to (p, q) in $\widehat{H} \times \widehat{H}$ as $n_i \to \infty$, where w_{n_i} is some conjugacy minimal representative of $g_{n_i} h g_{n_i}^{-1}$ in H.

2. The *algebraic ending lamination* corresponding to z is

$$\Lambda_z := \bigcup_{\substack{h \in H, \\ h \text{ infinite order}}} \Lambda_{z,h}.$$

3. The *algebraic ending lamination* for the short exact sequence (*) is

$$\Lambda := \bigcup_{z \in \partial Q} \Lambda_z.$$

Remark 5.4. We note the following about Definition 5.3 and the laminations Λ_z and Λ.

1. The lamination Λ_z is H-invariant and non-empty.

2. While $\Lambda_{z,h}$ is not necessarily symmetric as defined, $\Lambda_{z,h} \cup \Lambda_{z,h^{-1}}$ is symmetric.

3. As $\Lambda_{z,h}$ is a closed subset of $\partial^2 H$ for each $h \in H$ by construction, this shows that Λ_z is closed. Therefore, Λ_z is an algebraic lamination of H.

4. Theorem C also implies that Λ_z is a closed subset of $\partial^2 H$.

5. Mitra explained in [47] that in Definition 5.3 (2), it suffices to choose a finite collection of elements $h \in H$.

6. In Definition 5.3, the quasi-isometric section σ only needs to be defined on the ray $[1, z)$ rather than on all of Γ_Q.

7. The lamination Λ_z is independent of choice of quasi-isometric section, since if $\sigma : [1, z)$ and $\sigma' : [1, z)$ are two quasi-isometric sections, $[\sigma(z_n)h\sigma(z_n)^{-1}]_H = [\sigma'(z_n)h\sigma'(z_n)^{-1}]_H$.

8. The lamination Λ_z is independent of geodesic ray $[1, z)$ by Mitra's Lemma 3.3 of [47].

9. The definitions of Λ_z and Λ are independent of the choice of generating set for Q. This follows from the proof of Lemma 3.3 [47] which can be adapted to show that Λ_z is actually independent of quasigeodesic ray from 1 to z.

10. Fix $z_0 \in \Gamma_Q$, $z \in \partial Q$, and let $\gamma = [z_0, z)$ be a geodesic ray in Γ_Q with vertices $z'_n \in \gamma$ such that $d_Q(z_0, z_n) = n$. Let $\sigma' : [z_0, z) \to \Gamma_G$ be a quasi-isometric section with $\sigma'(z'_n) = g'_n$ and let Λ'_z be the algebraic ending lamination obtained by considering conjugacy minimal representatives of $g'_n h(g'_n)^{-1}$. The proof of Lemma 3.3 [47] also shows that $\Lambda_z = \Lambda'_z$. So, when defining Λ_z, we can consider a geodesic ray from any basepoint $z_0 \in \Gamma_Q$ converging to $z \in \partial Q$.

The next proposition shows how leaves of the lamination Λ_z behave under the action of conjugation by elements of G.

Proposition 5.5. *Let $1 \to H \to G \to Q \to 1$ be as in Convention 4.1 and let $P : \Gamma_G \to \Gamma_Q$ be the induced map. Then for all $g \in G$, $z \in \partial Q$, and $(u, v) \in \partial^2 H$, we have that (u, v) is a leaf of Λ_z if and only if $(g^{-1}u, g^{-1}v)$ is a leaf of $\Lambda_{P(g)^{-1}z}$.*

Proof. Fix $z \in \partial Q$, $g \in G$, and set $q_0 := P(g)$. Let $\lambda = (u, v)$ be a leaf of Λ_z. If $[1, z)$ is a geodesic ray in Γ_Q with vertices $1, z_1, z_2, \ldots$, then $q_0^{-1} \cdot [1, z) = [q_0^{-1}, q_0^{-1}z)$ is a geodesic ray in Γ_Q with vertices $q_0^{-1}, q_0^{-1}z_1, q_0^{-1}z_2, \ldots$. Since Λ_z is independent of quasi-isometric section, we may assume that σ is a quasi-isometric section with $\sigma(q_0) = g$. As in Convention 5.2, we will denote $\sigma(z_i)$ by g_i.

Since $(u, v) \in \Lambda_z$, there is some sequence $(a_i, a_i w_i) \in H \times H$ such that $w_i \in [g_{n_i} h g_{n_i}^{-1}]_H$ is a conjugacy minimal representative of $g_{n_i} h g_{n_i}^{-1}$ in H for some $n_i \geq 0$ and such that $a_i \to u$ and $a_i w_i \to v$ in $\widehat{\Gamma_H}$ as $i \to \infty$. Note that the sequence $(\phi_g(a_i), \phi_g(a_i w_i)) = (\phi_g(a_i), \phi_g(a_i)\phi_g(w_i))$ converges to $(\phi_g(u), \phi_g(v)) = (g^{-1}u, g^{-1}v)$ in $\widehat{H} \times \widehat{H}$.

Since $w_i \in [g_{n_i} h g_{n_i}^{-1}]_H$, we have that $\phi_g(w_i) \in [g^{-1} g_{n_i} h g_{n_i}^{-1} g]_H$. As mentioned earlier, there exist constants $K \geq 1$ and $C \geq 0$ such that ϕ_g is a (K, C)-quasi-isometry. Since for each $i \geq 0$ we have that w_i is a conjugacy minimal representative, Lemma 3.10 implies that there exists some $\kappa \geq 0$ such that for all $i \geq 0$, $\phi_g(w_i)$ is a κ-almost conjugacy minimal representative of $[g^{-1}g_{n_i}hg_{n_i}^{-1}g]_H$ in H. So, for each $i \geq 0$, there exists some $c_i \in H$ with $|c_i|_H \leq \kappa$ such that $c_i^{-1}\phi_g(w_i)c_i$ is a conjugacy minimal representative of $[g^{-1}g_{n_i}hg_{n_i}^{-1}g]_H$. As $(\phi_g(a_i), \phi_g(a_i)\phi_g(w_i)) \to (g^{-1}u, g^{-1}v)$ and $|c_i| \leq \kappa$ for all $i \geq 0$, we must also have that $(\phi_g(a_i)c_i, \phi_g(a_i)\phi_g(w_i)c_i) = (\phi_g(a_i)c_i, \phi_g(a_i)c_ic_i^{-1}\phi_g(w_i)c_i) \to (g^{-1}u, g^{-1}v)$.

For each $n_i \geq 0$, the element $g^{-1}g_{n_i}$ is in the same coset of H in G as $\sigma(q_0^{-1}z_{n_i})$. So, $c_i^{-1}\phi_g(w_i)c_i$ is a conjugacy minimal representative of $[\sigma(q_0^{-1}z_{n_i})h\sigma(q_0^{-1}z_{n_i})^{-1})]_H$. Therefore, by definition of $\Lambda_{q_0^{-1}z}$ and Remark 5.4 (5), we have that $\lambda_g = (g^{-1}u, g^{-1}v)$ is a leaf of $\Lambda_{q_0^{-1}z}$.

Now, suppose that $\lambda_g = (g^{-1}u, g^{-1}n)$ is a leaf of $\Lambda_{P(g)^{-1}z}$. Let $g^{-1}u = u'$, $g^{-1}v = v'$, and let $\lambda' = (u', v')$. Then the forward direction of this proposition shows that $\lambda'_{g^{-1}} \in \Lambda_{P(g^{-1})^{-1}P(g)^{-1}z} = \Lambda_z$. As $\lambda'_{g^{-1}} = (u, v) = \lambda$, the reverse direction of this proposition follows. □

The main result of Mitra in [47] is the following.

Theorem 5.6 (Mitra [47], Theorem 4.11)**.** *Suppose that $1 \to H \to G \to Q \to 1$ is as in Convention 4.1 and $\partial i : \partial H \to \partial G$ is the Cannon-Thurston map. Then for distinct points $u, v \in \partial H$, $\partial i(u) = \partial i(v)$ if and only if $(u, v) \in \Lambda$.*

The goal of the remainder of this section is to prove Theorem C. We first show that if $\lambda = (u, v)$ is a leaf of Λ_z, then ∂i_γ^+ identifies the endpoints u and v.

Proposition 5.7. *Let $1 \to H \to G \to Q \to 1$ be as in Convention 4.1, γ be as in Convention 4.2, i_γ^+ be as in Convention 4.8, and let $\partial i_\gamma^+ : \partial H \to \partial X(\gamma)^+$ denote the Cannon-Thurston map. If $\lambda = (u, v)$ is a leaf of Λ_z, then $\partial i_\gamma^+(u) = \partial i_\gamma^+(v)$.*

Proof. Let $\lambda = (u, v) \in \Lambda_z$ and suppose that $h \in H$ is such that λ is a leaf of $\Lambda_{z,h}$. By Remark 5.4, we can consider $\Lambda_{z,h}$ defined by the geodesic ray $[z_0, z)$. If $\sigma' : \Gamma_Q \to \Gamma_G$ is any quasi-isometric section, then $[\sigma'(z_i)h\sigma'(z_i)^{-1}]_H = [g_ihg_i^{-1}]_H$ for all $z_i \in [z_0, z)$. Hence, there exist elements $a_i \in H$ and conjugacy minimal representatives $w_i \in [g_{n_i}hg_{n_i}^{-1}]_H$ for some $n_i \geq 0$ such that $a_i \to u$ and $a_iw_i \to v$ as $i \to \infty$. Note that since w_i is conjugacy minimal, we have that $[a_iw_i^{-1}, a_i] \cup [a_i, a_iw_i] \cup [a_iw_i, a_iw_i^2]$ is a $(1, C_1)$-quasigeodesic for $C_1 = C_1(\delta)$ by Lemma 3.6 and Proposition 3.3. So, we have that $a_iw_i^{-1} \to u$ and $a_iw_i^2 \to v$ as well.

Suppose for each $i \geq 0$ that $g_{n_i}hg_{n_i}^{-1} =_H c_i^{-1}w_i'c_i$, where $c_i \in H$ is a minimal length element conjugating $g_{n_i}hg_{n_i}^{-1}$ to a cyclic conjugate w_i' of w_i. Mark vertices p_i on $[a_iw_i^{-1}, a_i]$ and q_i on $[a_iw_i, a_iw_i^2]$ where the path labeled by $(w_i')^2$ begins and ends. Let $x_i = p_ic_i$ and $y_i = q_ic_i$ denote the vertices at the end of the paths labeled by c_i which start at p_i and

q_i, respectively. Now, as in the proof of Lemma 3.7 the minimality of $|c_i|$ requires that $(x_i, q_i; \Gamma_H)_{p_i} \leq \delta$ and $(p_i, y_i; \Gamma_H)_{q_i} \leq \delta$. Hence, by Proposition 3.3, $[x_i, p_i] \cup [p_i, q_i] \cup [q_i, y_i]$ is a $(1, 8\delta)$-quasigeodesic in Γ_H. So, we must have that $x_i \to u$ and $y_i \to v$ in $\widehat{\Gamma_H}$. Note that the geodesic in Γ_H between x_i and y_i is labeled by the word $c_i^{-1}(w_i')^2 c_i =_H g_{n_i} h^2 g_{n_i}^{-1}$.

Recall that $i_\gamma^+(x_i) = x_i g_0$ and $i_\gamma^+(y_i) = y_i g_0$. So, the geodesic between $i_\gamma^+(x_i)$ and $i_\gamma^+(y_i)$ in $X(\gamma)$ is labeled by a word representing the element $g_0^{-1} g_{n_i} h^2 g_{n_i}^{-1} g_0$. To show that $\partial i_\gamma^+(u) = \partial i_\gamma^+(v)$, we must show that in $X(\gamma)^+$, the distance between some fixed point and the geodesic between $i_\gamma^+(x_i)$ and $i_\gamma^+(y_i)$, goes to infinity as $i \to \infty$. For each $i > 0$, consider the path $\rho_i \in X(\gamma)^+$ from $i_\gamma^+(x_i)$ to $i_\gamma^+(y_i)$ which consists of the geodesic $[x_i g_0, x_i g_{n_i}]$ labeled by $g_0^{-1} g_{n_i}$, followed by the quasigeodesic from $x_i g_{n_i}$ to $x_i g_{n_i} h^2$ labeled by h^2, followed by the geodesic $[x_i g_{n_i} h^2, y_i g_0]$ labeled by $g_{n_i}^{-1} g_0$. This path is a $(1, C)$-quasigeodesic in $X(\gamma)^+$ by Lemma 4.14 for some constant $C \geq 0$ independent of i.

So, take an arbitrary point $p \in \rho_n$. We will show that p is far from g_0 in $X(\gamma)^+$, and so the distance in $X(\gamma)^+$ between a quasigeodesic between $i_\gamma^+(x_n)$ and $i_\gamma^+(y_n)$ and g_0 goes to infinity as n goes to infinity. Note that since $d_H(1, x_n) \to \infty$, we must have that $d_{X(\gamma)^+}(g_0, i_\gamma^+(x_n)) \to \infty$.

Suppose first that the point p belongs to the initial part of ρ_n which is labeled by $g_0^{-1} g_n$. In this case, $p = i_\gamma^+(x_n) g_0^{-1} g_j$, $0 \leq j \leq n$. There are two cases for us to consider:

1. If $j \leq \frac{1}{2} d_{X(\gamma)^+}(g_0, i_\gamma^+(x_n))$, then

$$\begin{aligned} d_{X(\gamma)^+}(p, g_0) &= d_{X(\gamma)^+}(i_\gamma^+(x_n) g_0^{-1} g_j, g_0) \\ &\geq d_{X(\gamma)^+}(i_\gamma^+(x_n), g_0) - j \\ &\geq \frac{1}{2} d_{X(\gamma)^+}(g_0, i_\gamma^+(x_n)). \end{aligned}$$

2. If $j > \frac{1}{2} d_{X(\gamma)^+}(g_0, i_\gamma^+(x_n))$, then

$$\begin{aligned} d_{X(\gamma)^+}(p, g_0) &= d_{X(\gamma)^+}(i_\gamma^+(x_n) g_j, g_0) \\ &\geq j \\ &> \frac{1}{2} d_{X(\gamma)^+}(g_0, i_\gamma^+(x_n)). \end{aligned}$$

In both cases, $d_{X(\gamma)^+}(p, g_0) \to \infty$ as $n \to \infty$. The case where p belongs to the terminal part of ρ_n which is labeled by $g_n^{-1} g_0$ is handled similarly.

Finally, if p is a vertex in the portion of ρ_n which is labeled by h^2, then since $h^2 \in H$ is fixed, in $X(\gamma)^+$, p must lie a bounded distance away from the element $i^+_\gamma(x_n)g_0^{-1}g_n$. In this case, we have that $d_{X(\gamma)^+}(p, g_0) \geq d_{X(\gamma)^+}(i^+_\gamma(x_n)g_0^{-1}g_n, g_0) - |h^2|_H$, and

$$\lim_{n\to\infty} d_{X(\gamma)^+}(i^+_\gamma(x_n)g_0^{-1}g_n, g_0) - |h^2|_H = \infty.$$

Therefore, the distance between $[i_\gamma(x_n)^+, i^+_\gamma(y_n)]_{X(\gamma)^+}$ and g_0 in $X(\gamma)^+$ goes to infinity as $n \to \infty$. Hence, $\partial i^+_\gamma(u) = \partial i^+_\gamma(v)$. □

The following several lemmas from Mitra [47] will allow us to show that certain geodesics are conjugacy minimal representatives. We have stated and proved these results in the setting where $\gamma = (z', z) \subseteq \Gamma_Q$ does not necessarily go through the identity. However, we will apply these results in a simpler setting where γ does go through the identity. We have included the more general statements here to illuminate what happens in the general setting. The following three results are used to prove Corollary 5.12, which is key to the proof of Theorem C.

Lemma 5.8 (Cf. Mitra [47], Lemma 4.2 [47])**.** *There exists $\kappa \geq 0$ such that for any $(u, v) \in \partial^2 H$ with $\partial i(u) = \partial i(v)$, any geodesic subsegment $[p, q]$ of $\lambda = (u, v)$ has an extension $[r, q]$ in λ with $d_H(p, r)$ equal to 0 or 1 such that $[r, q]$ is a κ-almost conjugacy minimal representative.*

The next lemma is proved in a similar manner to Mitra's Lemma 4.3 in [47].

Lemma 5.9. *Given $\kappa \geq 0$, there exists $C \geq 1$ such that for any distinct $z, z' \in \partial Q$ and for any geodesic $\gamma = (z', z) \subset \Gamma_Q$ with $z_0 \in \gamma$ the following holds:*

If $\lambda = [1, h] \subseteq \Gamma_H$ and λ_{g_0} is a κ-almost conjugacy minimal representative for some $g_0 \in P^{-1}(z_0)$, then there exists a $(C, 0)$-quasi-isometric section σ_0 of (z', z) into $X(\gamma)$ containing g_0 such that for all $g \neq g_0$ in $\sigma_0((z', z))$, λ_g is a conjugacy minimal representative.

Proof. Let $\gamma = (z', z)$ be as in Convention 4.2. Let $\sigma : (z', z) \to X(\gamma)$ be an isometric lift of (z', z) into $X(\gamma)$ with $\sigma(z_0) = g_0$ and such that λ_{g_0} is a κ-almost conjugacy minimal representative for some $\kappa \geq 0$. We will construct the quasi-isometric section σ_0 satisfying the conclusions of the lemma inductively.

Set $\sigma_0(z_0) = g_0$. For each $n \geq 0$ set $s_n := \sigma(z_n)^{-1}\sigma(z_{n+1})$, and for each $n \leq 0$ set $s_{n-1} = \sigma(z_n)^{-1}\sigma(z_{n-1})$. Note that since σ is an isometric embedding, $|s_n| = 1$ for all n. So, there exists some $K_1 \geq 1$ and $\epsilon_1 \geq 0$ such that $\phi_{s_n} : \Gamma_H \to \Gamma_H$ is a (K_1, ϵ_1)-quasi-isometry for all $n \geq 0$. As λ_{g_0} is a κ-almost conjugacy minimal representative, there exists $\kappa' \geq 0$ such that $\phi_{s_0}(\lambda_{g_0}) = \lambda_{g_0 s_0}$ is a κ'-conjugacy minimal representative by Lemma 3.10.

By Corollary 3.8, there exists $c_0 \in H$ and $M' \geq 0$ with $|c_0|_H \leq M'$ such that $\lambda_{g_0 s_0 c_0}$ is a conjugacy minimal representative. Set $\sigma_0(z_1) := g_0 s_0 c_0$. We can similarly define $\sigma_0(z_{-1})$.

Suppose that $\sigma_0(z_j)$ has been constructed satisfying the conclusions of the lemma for all $-m \leq j \leq n$. By assumption, $\lambda_{\sigma_0(z_n)}$ is a conjugacy minimal representative, and so by Lemma 3.10 there exists $\kappa'' \geq 0$ such that $\lambda_{\sigma_0(z_n)s_n}$ is a κ''-almost conjugacy minimal representative. Then by Corollary 3.8, there exists $c_n \in H$ and $M'' \geq 0$ with $|c_n|_H \leq M''$ such that $\lambda_{\sigma_0(z_n)s_n c_n}$ is a conjugacy minimal representative. Set $\sigma_0(z_{n+1}) := \sigma_0(z_n)s_n c_n$. We can similarly define $\sigma_0(z_{-m-1})$. Note that $d_{X(\gamma)}(\sigma_0(z_i), \sigma_0(z_{i+1})) \leq \max\{M', M''\}$, and so σ_0 is a $(C, 0)$-quasi-isometric section, where $C := \max\{M', M''\}$ and λ_g is a conjugacy minimal representative for all $g \neq g_0$ in $\sigma_0((z', z))$. □

The following corollary is obtained from the previous lemma by translating the quasi-isometric section by an element of G. Here, we choose the quasi-isometric section σ_0 to go through the point $g_0 \in \Gamma_G$ rather than the identity.

Corollary 5.10 (Cf. Mitra [47] Corollary 4.4)**.** *Given $\kappa \geq 0$, there exists $C \geq 1$ such that for any geodesic ray $[z_0, z)$ in Γ_Q and any $g \in P^{-1}([z_0, z))$ the following holds:*

If $\lambda = [1, h] \subseteq \Gamma_H$ and λ_{g_0} is a κ-almost conjugacy minimal representative for some $g_0 \in P^{-1}(z_0)$, then there exists a $(C, 0)$-quasi-isometric section σ_0 of $[z_0, z)$ into Γ_G containing $g \in \Gamma_G$ such that for all $g' \neq g$ in $\sigma_0([z_0, z))$, $\lambda_{g_0 g^{-1} g'}$ is a conjugacy minimal representative.

Proof. By Lemma 5.9, there exists a $(C, 0)$-quasi-isometric section $\sigma' : (z', z) \to X(\gamma)$ with $\sigma'(z_0) = g_0$ such that for all $g' \neq g_0$ in $\sigma'((z', z))$, $\lambda_{g'}$ is a conjugacy minimal representative. Suppose that $g \in P^{-1}(z_n)$ and set $\sigma_0(z_n) := g$. For each integer i with $i \geq -n$, set $\sigma_0(z_{n+i}) := t_{gg_0^{-1}} \cdot \sigma'(z_i)$. Now, $\sigma_0 : [z_0, z) \to X(\gamma)^+$ is a $(C, 0)$-quasi-isometric section since it is a left-translate of σ' by $gg_0^{-1} \in G$. Also, note that for all $g' \neq g$ in $\sigma([z_0, z))$, we have that $g' = t_{gg_0^{-1}} \cdot \sigma'(z_i)$ for some $i \geq -n$ with $i \neq 0$. Then, $\lambda_{g_0 g^{-1} g'} = \lambda_{g_0 g^{-1} g g_0^{-1} \sigma'(z_i)} = \lambda_{\sigma'(z_i)}$ is a conjugacy minimal representative by Lemma 5.9. □

The following lemma will allow us to reduce to the simpler setting where $\gamma = (z', z) \subseteq \Gamma_Q$ passes through the identity in Γ_Q.

Lemma 5.11. *Suppose $\gamma = (z', z)$ is as in Convention 4.2 and let $\gamma' := z_0^{-1} \cdot \gamma = (z_0^{-1} z', z_0^{-1} z)$ be as in Convention 4.10. Let $X(\gamma)$ and $X(\gamma')$ be the stacks as in Convention 4.8 where the section $\sigma : \gamma \to X(\gamma)$ is such that $\sigma(z_0) = g_0$ and $\sigma' : \gamma' \to X(\gamma')$ is chosen so that $\sigma' = g_0^{-1} \cdot \sigma$. Let i_γ^+ and $i_{\gamma'}^+$ be as in Convention 4.8, and let $\partial i_\gamma^+ : \partial H \to \partial X(\gamma)^+$ and $\partial i_{\gamma'}^+ : \partial H \to \partial X(\gamma')^+$ be the Cannon-Thurston maps.*

Then for any two distinct points $u, v \in \partial H$, $\partial i^+_\gamma(u) = \partial i^+_\gamma(v)$ if and only if $\partial i^+_{\gamma'}(\phi_{g_0}(u)) = \partial i^+_{\gamma'}(\phi_{g_0}(v))$, where $g_0 \in P^{-1}(z_0)$.

Proof. Let $\gamma = (z', z)$ be as in Convention 4.2, let $\gamma' := z_0^{-1} \cdot \gamma = (z_0^{-1}z', z_0^{-1}z)$, and fix some $g_0 \in P^{-1}(z_0)$. Recall that i^+_γ is given by $i^+_\gamma(h) = t_{g_0} \cdot \phi_{g_0}(h) = hg_0$ and $i^+_{\gamma'}$ is given by $i_{\gamma'}(h) = h$. Suppose first that $u, v \in \partial H$ are distinct points such that $\partial i^+_\gamma(u) = \partial i^+_\gamma(v)$. Then, for any sequences $(u_n), (v_n) \in \Gamma_H$ with $u_n \to u$ and $v_n \to v$ in $\widehat{\Gamma_H}$, we have that in $\widehat{X(\gamma)^+}$, $\lim_{n\to\infty} i^+_\gamma(u_n) = \lim_{i\to\infty} i^+_\gamma(v_n)$. So, in $\widehat{X(\gamma)^+}$ we have that $\lim_{i\to\infty} u_n g_0 = \lim_{i\to\infty} v_n g_0$. Note that $X(\gamma)^+ = g_0 X(\gamma')^+$ and so left-translation by g_0^{-1} gives an isometry from $X(\gamma)^+$ to $X(\gamma')^+$. Therefore, in $\widehat{X(\gamma')^+}$ we have that $\lim_{i\to\infty} g_0^{-1} u_n g_0 = \lim_{i\to\infty} g_0^{-1} v_n g_0$. So by definition of $i^+_{\gamma'}$, we have that $\lim_{i\to\infty} i^+_{\gamma'}(\phi_{g_0}(u_n)) = \lim_{i\to\infty} i^+_{\gamma'}(\phi_{g_0}(v_n))$ in $\widehat{X(\gamma')^+}$. Since in $\widehat{\Gamma_H}$ $\phi_{g_0}(u_n) \to \phi_{g_0}(u)$ and $\phi_{g_0}(v_n) \to \phi_{g_0}(v)$ as $n \to \infty$, we have that $\partial i^+_{\gamma'}(\phi_{g_0}(u)) = \partial i^+_{\gamma'}(\phi_{g_0}(v))$ by the continuity of $\widehat{i^+_{\gamma'}}$ (Lemma 4.9). The reverse implication follows in the same manner by noting that left-translation by g_0 gives an isometry from $X(\gamma')^+$ to $X(\gamma)^+$. ☐

The following result follows directly from Lemma 5.8 and Corollary 5.10. This corollary will be used in the proof of Theorem C to construct a sequence of conjugacy minimal representatives which converge to some bi-infinite geodesic $\lambda \subseteq \partial^2 H$ whose endpoints are identified by ∂i^+_γ.

Corollary 5.12 (Cf. Mitra [47] Lemma 4.5)**.** *There exists C' such that for any $\lambda = (u, v)$, $u, v \in \partial H$ with $\partial i^+_\gamma(u) = \partial i^+_\gamma(v)$, any geodesic ray $[z_0, z)$ in Γ_Q, and any geodesic subsegment $[p, q]$ of λ_g for some $g \in P^{-1}([z_0, z))$ the following holds:*

There exists an extension $[r, q] = \mu$ of $[p, q]$ in λ_g with $d_H(p, r)$ equal to 0 or 1 and a $(C', 0)$-quasi-isometric section $\sigma : [z_0, z) \to X(\gamma)$ such that $gr \in \sigma([z_0, z))$ and $\mu_{g_0 r^{-1} g^{-1} g'}$ is a conjugacy minimal representative for all $g' \neq gr$ in $\sigma([z_0, z))$.

Proof. Let $\lambda = (u, v)$ be such that $\partial i^+_\gamma(u) = \partial i^+_\gamma(v)$, let $[z_0, z) \in \Gamma_Q$ be a geodesic ray, let $g \in P^{-1}([z_0, z))$, and let $[p, q]$ be any geodesic subsegment of $\lambda_g = (\phi_g(u), \phi_g(v))$. By Lemma 5.11, $\partial i^+_{\gamma'}(\phi_{g_0}(u)) = \partial i^+_{\gamma'}(\phi_{g_0}(v))$. So by Corollary 4.13, we have that $\partial i(\phi_g(u)) = \partial i(\phi_g(v))$. So by Lemma 5.8, there exists an extension $[r, q] = \mu$ of $[p, q]$ in λ_g with $d_H(p, r)$ equal to 0 or 1 and such that $[r, q]$ is a κ-almost conjugacy minimal representative for some $\kappa \geq 0$. Let $\mu' = [1, r^{-1}q]$ and note that μ' is also a κ-almost conjugacy minimal representative since it has the same label as μ. By Lemma 3.10, μ_{g_0} and μ'_{g_0} are κ'-almost conjugacy minimal representatives for some $\kappa' \geq 0$ depending on g_0. So by Corollary 5.10, there exists $C' \geq 1$ and a $(C', 0)$-quasi-isometric section $\sigma : [z_0, z) \to \Gamma_G$ containing $gr \in \Gamma_G$ such that for all

$g' \neq gr \in \sigma([z_0, z))$, $\mu'_{g_0 r^{-1} g^{-1} g'}$ is a conjugacy minimal representative. Therefore, $\mu_{g_0 r^{-1} g^{-1} g'}$ is also a conjugacy minimal representative. □

For the next portion of this chapter, we will assume that the bi-infinite geodesic $\gamma = (z', z) \subseteq \Gamma_Q$ goes through the identity in Q, and so $\gamma^+ = [1, z)$. Note that several of the previous lemmas simplify in this case. We now make the following convention.

Convention 5.13. Let $\gamma = (z', z)$ be a bi-infinite geodesic in Γ_Q between $z', z \in \partial Q$ with $z' \neq z$ and assume that $1 \in \gamma$. Label the sequence of vertices in order along the portion of γ from 1 to z by $1 = z_0, z_1, z_2, \ldots$. Similarly, label the sequence of vertices in order along the portion of γ from 1 to z' by $1 = z_0, z_{-1}, z_{-2}, \ldots$. Let $\sigma_0 : \gamma \to \Gamma_G$ denote an isometric lift of γ through the identity in Γ_G, i.e. such that $\sigma_0(1) = 1$, and set $g_i := \sigma_0(z_i)$. Let $X(\gamma)$ and $X(\gamma)^+$ denote the stacks over $\gamma = (z', z)$ and $\gamma^+ = [1, z)$, respectively. Finally, let $i_\gamma : \Gamma_H \to X(\gamma)$ and $i_\gamma^+ : \Gamma_H \to X(\gamma)^+$ be the respective inclusion maps given by $i_\gamma(h) = h$ and $i_\gamma^+(h) = h$ for all $h \in H$.

Before proving Theorem C, we will first introduce some necessary terminology as well as some lemmas which were first stated by Mitra in [47].

Given a (finite or infinite) geodesic $\lambda \subset \widehat{\Gamma_H}$ with endpoints $a, b \in \widehat{\Gamma_H}$ and an element $g \in G$, recall that $\lambda_g \subset \widehat{\Gamma_H}$ denotes the geodesic joining $\phi_g(a) = g^{-1} a g$ and $\phi_g(b) = g^{-1} b g$. For any quasi-isometric section $\sigma : \Gamma_Q \to \Gamma_G$ and geodesic λ, Mitra defines the set

$$B(\lambda, \sigma) := \bigcup_{g \in \sigma(Q)} t_g \cdot i(\lambda_g),$$

where t_g denotes left-translation by the element $g \in G$. For our purposes, we will consider the subset of $B(\lambda, \sigma)$ which lives in $X(\gamma)^+$:

$$B_{\gamma^+}(\lambda, \sigma) = \bigcup_{g \in \sigma([1,z))} t_g \cdot i_\gamma^+(\lambda_g).$$

Note that $B_{\gamma^+}(\lambda, \sigma) = B(\lambda, \sigma) \cap P^{-1}([1, z))$ and that if λ is a bi-infinite geodesic, then $B_{\gamma^+}(\lambda, \sigma)$ is independent of quasi-isometric section σ for the same reason Mitra uses to show $B(\lambda, \sigma)$ is independent of quasi-isometric section [47].

On the vertices of Γ_H, define the map $\pi_{g,\lambda} : \Gamma_H \to \lambda_g$ by sending $h \in H$ to a closest vertex on λ_g. We will now define a projection map to the set $B_{\gamma^+}(\lambda, \sigma)$. As σ is a quasi-isometric section, for each $g' \in X(\gamma)^+$, there is a unique $g \in \sigma([1, z))$ and $h \in H$ such that

$g' = t_g \cdot i_\gamma^+(h)$. So, define

$$\Pi_\lambda^\sigma(g') = \Pi_\lambda^\sigma \cdot t_g \cdot i_\gamma^+(h) := t_g \cdot i_\gamma^+ \cdot \pi_{g,\lambda}(h).$$

The following statements are versions of the analogous statements from Mitra [47] which apply to the setting in which we are working. In most cases, the proofs that Mitra provided go through with no changes to the reasoning. We provide details of the necessary modifications where they are needed.

The same proof of Mitra's Theorem 3.7 of [48] verifies the following statement. In particular, this lemma will be used to show that if $\sigma : [1, z) \to X(\gamma)^+$ is a (K, ϵ)-quasi-isometric section, then the projection of σ to $B_{\gamma^+}(\lambda, \sigma)$ is also a quasi-isometric section.

Lemma 5.14 (Cf. Mitra [47], Theorem 4.6)**.** *For all $K \geq 1$ and $\epsilon \geq 0$, there exists a constant $C \geq 1$ such that if $\sigma : [1, z) \to X(\gamma)^+$ is any (K, ϵ)-quasi-isometric section and $\lambda \subseteq \Gamma_H$ is any bi-infinite geodesic, then for all $x, y \in X(\gamma)^+$, $d_{X(\gamma)^+}(\Pi_\lambda^\sigma(x), \Pi_\lambda^\sigma(y)) \leq C d_{X(\gamma)^+}(x, y)$.*

Lemma 5.15 (Cf. Mitra [47], Lemma 4.7)**.** *For all $K \geq 1$ and $\epsilon \geq 0$ there exists $A \geq 1$ such that if $\sigma : [1, z) \to X(\gamma)^+$ is a (K, ϵ)-quasi-isometric section, then for all $p, q \in \sigma([1, z))$ and $x \in t_p \cdot i_\gamma^+(\lambda_p)$ there exists $y \in t_q \cdot i_\gamma^+(\lambda_q)$ such that $d_{X(\gamma)^+}(x, y) \leq A d_Q(Px, Py) = A d_Q(Pp, Pq)$.*

Proof. Let $\sigma : [1, z) \to X(\gamma)^+$ be a (K, ϵ)-quasi-isometric section, $p, q \in \sigma([1, z))$, $x \in t_p \cdot i_\gamma^+(\lambda_p)$, and set $y = \Pi_\lambda^\sigma(xp^{-1}q)$. Note that $y \in t_q \cdot i_\gamma^+(\lambda_q)$. Then by Lemma 5.14, there exists a constant $C \geq 1$ such that $d_{X(\gamma)^+}(\Pi_\lambda^\sigma(x), \Pi_\lambda^\sigma(xp^{-1}q)) = d_{X(\gamma)^+}(x, y) \leq C d_{X(\gamma)^+}(x, xp^{-1}q)$. Since $p, q \in \sigma([1, z))$ and σ is a (K, ϵ)-quasi-isometric section, we have that $|p^{-1}q| \leq K d_Q(Pp, Pq) + \epsilon$. Therefore, $d_{X(\gamma)^+}(x, xp^{-1}q) = |p^{-1}q| \leq K d_Q(Pp, Pq) + \epsilon$. So, let $A = C(K + \epsilon)$. As $Px = Pp$ and $Py = Pq$, we have finally that $d_{X(\gamma)^+}(x, y) \leq A d_Q(Px, Py) = A d_Q(Pp, Pq)$ as required. □

The following is the version of Lemma 4.8 [47] that we need for our purposes. It is proved by an argument similar to the one given by Mitra using the previous lemma.

Lemma 5.16 (Cf. Mitra [47], Lemma 4.8)**.** *For all $K \geq 1$ and $\epsilon \geq 0$ there exists $M \geq 0$ such that the following holds. Suppose λ is a bi-infinite geodesic in Γ_H and a is a vertex on λ splitting λ into semi-infinite geodesics λ^- and λ^+. Suppose further that $\sigma : [1, z) \to X(\gamma)^+$ is a (K, ϵ)-quasi-isometric section such that $\sigma([1, z)) \subseteq B_{\gamma^+}(\lambda, \sigma)$ and $i_\gamma^+(a) \in \sigma([1, z))$. Then, any geodesic in $X(\gamma)^+$ joining a point in $B_{\gamma^+}(\lambda^-, \sigma)$ to a point in $B_{\gamma^+}(\lambda^+, \sigma)$ passes through an M-neighborhood of $\sigma([1, z))$.*

Lemma 5.17 (Cf. Mitra [47], Corollary 4.10)**.** *Given $K \geq 1$, $\epsilon \geq 0$, there exists α such that if $\lambda = (u, v)$ is such that $\partial i_\gamma^+(u) = \partial i_\gamma^+(v)$ then the following is satisfied:*

If σ and σ' are (K, ϵ)-quasi-isometric sections such that $B_{\gamma^+}(\lambda, \sigma) = B_{\gamma^+}(\lambda, \sigma')$ and σ, σ' are contained in $B_{\gamma^+}(\lambda, \sigma)$, then there exists $N \geq 0$ such that for all $n \geq N$,

$$d_{X(\gamma)^+}(\sigma(z_n), \sigma'(z_n)) \leq \alpha.$$

Proof. Let $\lambda = (u, v)$ be such that $\partial i_\gamma^+(u) = \partial i_\gamma^+(v)$ and let σ and σ' be (K, ϵ)-quasi-isometric sections satisfying the hypotheses of the lemma. Let (p_n) and (q_n) be a sequence of vertices on λ such that $p_n \to u$ and $q_n \to v$ as $n \to \infty$. For each $n \geq 0$, Lemma 5.16 guarantees there exist points $z_{n'}, z_{n''} \in [1, z)$ such that any geodesic in $X(\gamma)^+$ joining $i_\gamma^+(p_n)$ to $i_\gamma^+(q_n)$ passes through an M-neighborhood of both $\sigma(z_{n'})$ and $\sigma'(z_{n''})$. Since $\partial i_\gamma^+(u) = \partial i_\gamma^+(v)$ and $\widehat{i_\gamma^+}$ is continuous, we must have that the sequences $\{i_\gamma^+(p_n)\}$, $\{i_\gamma^+(q_n)\}$, $\{\sigma(z_{n'})\}$, and $\{\sigma'(z_{n''})\}$ all converge to the same point in $\partial X(\gamma)^+$. Since σ and σ' are quasi-isometric sections of $[1, z)$ into $X(\gamma)^+$ and as $d_{X(\gamma)^+}(1, [i_\gamma^+(p_n), i_\gamma^+(q_n)]) \to \infty$, we must have that $z_{n'} \to z$ and $z_{n''} \to z$. Therefore, $\sigma([1, z))$ and $\sigma'([1, z))$ are asymptotic quasigeodesic rays in $X(\gamma)^+$ and we have that for all $n \geq N$,

$$\max\{d_{X(\gamma)^+}(\sigma(z_n), \sigma'([1, z))), d_{X(\gamma)^+}(\sigma([1, z)), \sigma'(z_n))\} \leq \alpha'.$$

But since σ and σ' are (K, ϵ)-quasi-isometric sections, if $z_{n'}$ is such that

$$d_{X(\gamma)^+}(\sigma(z_n), \sigma'(z_{n'})) \leq \alpha',$$

then we have that

$$\begin{aligned} d_{X(\gamma)^+}(\sigma(z_n), \sigma'(z_n)) &\leq d_{X(\gamma)^+}(\sigma(z_n), \sigma'(z_{n'})) + d_{X(\gamma)^+}(\sigma'(z_{n'}), \sigma'(z_n)) \\ &\leq \alpha' + K|n - n'| + \epsilon \\ &\leq \alpha' + K\alpha' + \epsilon = \alpha \end{aligned}$$

Thus for all $n \geq N$, $d_{X(\gamma)^+}(\sigma(z_n), \sigma'(z_n)) \leq \alpha$. □

We are now ready to prove the main theorem of this chapter which is reminiscent of Mitra's Theorem 4.11 [47].

Theorem C. *Let $1 \to H \to G \to Q \to 1$ be a short exact sequence of infinite, finitely generated, word-hyperbolic groups. Let $z, z' \in \partial Q$ be distinct and let $\gamma \subseteq \Gamma_Q$ be a bi-infinite* *geodesic in Γ_Q between z and z'. Let $i_\gamma^+ : \Gamma_H \to X(\gamma)^+$ be the inclusion of Γ_H into the semi-infinite stack $X(\gamma)^+$ over $\gamma^+ = [z_0, z)$ as in Convention 4.8, and let $\partial i_\gamma^+ : \partial H \to \partial X(\gamma)^+$ be the Cannon-Thurston map.*

Then for any distinct $u, v \in \partial H$, we have $\partial i_\gamma^+(u) = \partial i_\gamma^+(v)$ if and only if (u, v) is a leaf of the ending lamination Λ_z.

Proof. Suppose first that $\gamma = (z', z)$ is as in Convention 5.13 with $1 \in \gamma$. By Proposition 5.7, it suffices to show that if $\partial i_\gamma^+(u) = \partial i_\gamma^+(v)$, then $\lambda = (u, v) \in \Lambda_z$. So, let $u, v \in \partial H$ be distinct points such that $\partial i_\gamma^+(u) = \partial i_\gamma^+(v)$. As the set of leaves of $\partial^2 H$ whose endpoints are identified under ∂i_γ^+ is H-invariant, we may assume that $\lambda = (u, v)$ passes through $1 \in \Gamma_H$.

Let $\sigma_0 : [1, z) \to X(\gamma)^+$ be the isometric lift of γ^+ into Γ_G through the identity as in Convention 5.13. Let $\sigma_e := \Pi_\lambda^{\sigma_0} \cdot \sigma_0$ be the projection of σ_0 onto $B_{\gamma^+}(\lambda, \sigma_0)$ and set $g_n' := \sigma_e(z_n)$. By Lemma 5.14, σ_e is a $(C, 0)$-quasi-isometric section of $[1, z)$ into $B_{\gamma^+}(\lambda, \sigma_0)$ for some $C \geq 1$.

By Corollary 5.12, there exists $C' \geq 1$ such that for any $g \in \sigma_0([1, z))$ and any $[p, q] \subseteq \lambda_g$, there exists an extension $[r, q] =: \mu$ of $[p, q]$ in λ_g with $d_H(p, r) = 0$ or 1 and a $(C', 0)$-quasi-isometric section σ such that $gr \in \sigma([1, z))$ and $\mu_{r^{-1}g^{-1}g'}$ is a conjugacy minimal representative for all $g' \neq gr$ in $\sigma([1, z))$. Projecting σ to $B_{\gamma^+}(\lambda, \sigma_0)$ yields, by Lemma 5.14, a $(C_2, 0)$-quasi-isometric section for some $C_2 \geq 1$.

If σ' is any $(C_2, 0)$-quasi-isometric section, Lemma 5.17 gives that there is some $\alpha > 0$ such that if $\sigma' \subseteq B_{\gamma^+}(\lambda, \sigma_0)$, then there exists some $N \geq 0$ such that for all $n \geq N$, $d_{X(\gamma)^+}(g_n', \sigma'(z_n)) \leq \alpha$. Given this α, Proposition 2.5 guarantees there are some $b > 1$, $A > 0$, and $\eta > 0$ depending on α and C_2 such that if $\sigma'([1, z))$ is a $(C_2, 0)$-quasi-isometric section of $[1, z)$ into $X(\gamma)^+$ with $d_{X(\gamma)^+}(\sigma'(z_n), g_n') \geq \eta$, then any path in $i_\gamma^+(\Gamma_H)$ joining $\sigma'(1)$ and $\sigma_e(1)$ has length greater than or equal to Ab^n.

Now, let λ^+ and λ^- denote the two closures of the components of $\lambda \setminus \{1\}$. Note that for each $n > 0$, $g_n' \in t_{g_n} \cdot i_\gamma^+(\lambda_{g_n})$. Hence, for all $n > 0$ there exists $p_n \in \lambda_{g_n}^-$ and $q_n \in \lambda_{g_n}^+$ such that $d_{X(\gamma)^+}(t_{g_n} \cdot i_\gamma^+(p_n), g_n') = d_{X(\gamma)^+}(g_n p_n, g_n') = \eta + 1$ and $d_{X(\gamma)^+}(t_{g_n} \cdot i_\gamma^+(q_n), g_n') = d_{X(\gamma)^+}(g_n q_n, g_n') = \eta$. By Corollary 5.12, for each $n > 0$ there exists $r_n \in \lambda_{g_n}^-$ with $d_H(r_n, p_n) = 0$ or 1 and a $(C', 0)$-quasi-isometric section σ_n of $[1, z)$ into $X(\gamma)^+$ satisfying the following two conditions:

1. $g_n r_n = \sigma_n(z_n)$
2. If $\mu^{(n)}$ is the subsegment of λ_{g_n} in Γ_H joining r_n and q_n, then $\mu^{(n)}_{r_n^{-1} g_n^{-1} \sigma_n(z_m)}$ is a conjugacy minimal representative for all $z_m \neq z_n$.

For each $n > 0$, define a new quasi-isometric section $\tau_n(z_i) := t_{g_n q_n r_n^{-1} g_n^{-1}} \cdot \sigma_n(z_i)$ which is obtained by left-translating σ_n to go through the point $g_n q_n \in t_{g_n} \cdot i_\gamma^+(\lambda_{g_n})$. We will now project σ_n and τ_n to the set $B_{\gamma^+}(\lambda, \sigma_0)$ to get new quasi-isometric sections which satisfy the hypotheses of Lemma 5.17. Denote these new $(C_2, 0)$-quasi-isometric sections by $\sigma_n' := \Pi_\lambda^{\sigma_0} \cdot \sigma_n$ and $\tau_n' := \Pi_\lambda^{\sigma_0} \cdot \tau_n$.

By Lemma 5.17, there is some α such that for every index $n > 0$, $d_{X(\gamma)^+}(g_k', \sigma_n'(z_k)) \leq \alpha$ as long as $k \geq N$ for some constant $N = N(n)$. So, the $(C_2, 0)$-quasigeodesic rays interpolated by σ_n' and σ_e satisfy the hypotheses of Proposition 2.5 since there is some point along these rays where $d_{X(\gamma)^+}(g_k', \sigma_n'(z_k)) \leq \alpha$ and since the rays were defined so that $d_{X(\gamma)^+}(g_n', \sigma_n'(z_n)) = d_{X(\gamma)^+}(g_n', g_n r_n) \geq \eta$. As any path in $i_\gamma^+(\Gamma_H)$ is distance at least n/C_2 from any path in $t_{g_n} \cdot i_\gamma^+(\Gamma_H)$, we have that there exists $b > 1$ and $A > 0$ such that the portion of $i_\gamma^+(\lambda)$ between $\sigma_n'(1)$ and $\sigma_e(1) = g_0'$ has length at least Ab^n. As the same holds true for the quasigeodesic rays interpolated by τ_n' and σ_e, the portion of $i_\gamma^+(\lambda)$ between $\tau_n'(1)$ and $\sigma_e(1) = g_0'$ also has length greater than or equal to Ab^n.

Note that for all $n \geq 0$, $\sigma_n(1)$, $\sigma_n'(1)$, $\tau_n(1)$, and $\tau_n'(1)$ all lie in $i_\gamma^+(\Gamma_H)$. Let $[\sigma_n'(1)^*, \tau_n'(1)^*]$ denote the subsegment of λ joining $(i_\gamma^+)^{-1} \cdot \sigma_n'(1)$ and $(i_\gamma^+)^{-1} \cdot \tau_n'(1)$. Then, the sequence $\{[\sigma_n'(1)^*, \tau_n'(1)^*]\}$ converges to λ in $\widehat{\Gamma_H}$.

Since $d_{X(\gamma)^+}(g_n r_n, g_n q_n) \leq 2\eta + 2$, there exists $\rho > 0$ such that $r_n^{-1} q_n$ is an element of H with $|r_n^{-1} q_n|_H \leq \rho$. Since there are only finitely many of these, we may pass to a subsequence n_j such that $r_{n_j}^{-1} q_{n_j} = h$ where h is some fixed element of H. Note that the subsequence $\{[\sigma_{n_j}'(z_0)^*, \tau_{n_j}'(z_0)^*]\}$ also converges to λ in $\widehat{\Gamma_H}$.

Let $[\sigma_n(1)^*, \sigma_n'(1)^*]$ denote a geodesic segment in Γ_H joinging $(i_\gamma^+)^{-1} \cdot \sigma_n(1)$ and $(i_\gamma^+)^{-1} \cdot \sigma_n'(1)$ and define $[\tau_n'(1)^*, \tau_n(1)^*]$ similarly. Since $\sigma_n'(1)^* = (i_\gamma^+)^{-1} \cdot \Pi_\lambda^{\sigma_0} \cdot i_\gamma^+(\sigma_n(1))$, we must have that in Γ_H, $(\sigma_n(1)^*, \tau_n'(1)^*)_{\sigma_n'(1)^*} \leq 2\delta$. Otherwise, there would be a point on $i_\gamma^+(\lambda)$ closer to $\sigma_n(1)$ than $\sigma_n'(1)$, contradicting the definition of $\sigma_n'(1)$ as the projection of $\sigma_n(1)$ to $i_\gamma^+(\lambda)$. For a similar reason, $(\sigma_n'(1)^*, \tau_n(1)^*)_{\tau_n'(1)^*} \leq 2\delta$. So by Proposition 3.3, we have that for all n sufficiently large (so that $d_H(\sigma_n'(1)^*, \tau_n'(1)^*) > 14\delta$), $[\sigma_n(1)^*, \sigma_n'(1)^*] \cup [\sigma_n'(1)^*, \tau_n'(1)^*] \cup [\tau_n'(1)^*, \tau_n(1)^*]$ is a $(1, 12\delta)$-quasigeodesic. Thus, for all n sufficiently large, there is some constant $B > 0$ depending only on δ such that $[\sigma_n(1)^*, \sigma_n'(1)^*] \cup [\sigma_n'(1)^*, \tau_n'(1)^*] \cup [\tau_n'(1)^*, \tau_n(1)^*]$ lies in a B-neighborhood of the geodesic $[\sigma_n(1)^*, \tau_n(1)^*]$ in Γ_H.

As the sequence $\{[\sigma_{n_j}'(1)^*, \tau_{n_j}'(1)^*]\}$ converges to λ, we must also have that the sequence

$$\{[\sigma_{n_j}(1)^*, \sigma_{n_j}'(1)^*] \cup [\sigma_{n_j}'(1)^*, \tau_{n_j}'(1)^*] \cup [\tau_{n_j}'(1)^*, \tau_{n_j}(1)^*]\}$$

converges to λ. In particular, $\{\sigma_{n_j}(1)^*\}$ and $\{\tau_{n_j}(1)^*\}$ must converge to the endpoints of λ

in Γ_H. Recall that σ_n was chosen so that, in particular, $\mu^{(n)}_{r_n^{-1}g_n^{-1}\sigma_n(1)}$ is a conjugacy minimal representative. Since $\mu^{(n)}_{r_n^{-1}g_n^{-1}\sigma_n(1)}$ is the label of the geodesic in Γ_H between $\sigma_n(1)^*$ and $(g_n q_n r_n^{-1} g_n^{-1}\sigma_n(1))^* = \tau_n(1)^*$, we have that $\{[\sigma_{n_j}(1)^*, \tau_{n_j}(1)^*]\}$ is a sequence of conjugacy minimal representatives of $\phi_{r_{n_j}^{-1}g_{n_j}^{-1}\sigma_{n_j}(1)}(h)$.

Let $\sigma'' : [1, z) \to \Gamma_G$ be any quasi-isometric section. Note that for all $n \geq 0$, $\sigma''(z_n)$ and $\sigma_n(1)^{-1}g_n r_n$ are in the same coset of H in G. Therefore, $\phi_{r_{n_j}^{-1}g_{n_j}^{-1}\sigma_{n_j}(1)}(h)$ and $\phi_{(\sigma''(z_n))^{-1}}(h)$ have the same conjugacy minimal representatives. Hence, $\lambda = (u, v) \in \Lambda_{z,h} \subseteq \Lambda_z$.

Finally, suppose that $\gamma = (z', z)$ goes through $z_0 \in \Gamma_Q$ rather than the identity. Then, $\gamma' := z_0^{-1}\gamma = (z_0^{-1}z', z_0^{-1}z)$ does go through the identity. If $\partial i^+_\gamma(u) = \partial i^+_\gamma(v)$, Lemma 5.11 implies that $\partial i^+_{\gamma'}(\phi_{g_0}(u)) = \partial i^+_{\gamma'}(\phi_{g_0}(v))$. By the above, this implies that

$$\lambda_{g_0} = (\phi_{g_0}(u), \phi_{g_0}(v)) \in \Lambda_{z_0^{-1}z}.$$

Finally by Proposition 5.5, this implies that $\lambda \in \Lambda_z$ as desired. $\square$

Chapter 6

Structure of the Cannon-Thurston map

We can now prove the main result of this thesis, Theorem A from Chapter 1. Recall that a *dendrite* is a compact, connected, locally connected metrizable space which contains no simple closed curves.

Proposition 6.1 (Bowditch [10] c.f. 2.5.2)**.** *Let $\mathcal{X}$ be a bi-infinite hyperbolic stack and let $\mathcal{X}^+$ be the corresponding semi-infinite stack. Then, the Gromov boundary $\partial\mathcal{X}^+$ is a dendrite.*

If $L \subseteq \partial^2 H$ is an algebraic lamination on H, then $\partial H/L$ denotes the quotient space of ∂H by the equivalence relation generated by $L \subseteq \partial^2 H$.

Theorem A. *Let $1 \to H \to G \to Q \to 1$ be a short exact sequence of infinite, finitely generated, word-hyperbolic groups and choose $z \in \partial Q$. Then, the space $\partial H/\Lambda_z$ is homeomorphic to a dendrite.*

Proof. Let $\gamma = (z', z)$ be as in Convention 4.2 and let $X(\gamma)^+$ and $i_\gamma^+ : \Gamma_H \to X(\gamma)^+$ be as in Convention 4.8. Denote by $\pi_z : \partial H \to \partial H/\Lambda_z$ the quotient map. If $a, b \in \partial H$ are such that $\pi_z(a) = \pi_z(b)$, then $\partial i_\gamma^+(a) = \partial i_\gamma^+(b)$ by Proposition 5.7. So, the Cannon-Thurston map $\partial i_\gamma^+ : \partial H \to \partial X(\gamma)^+$ quotients through to a map $\tau_z : \partial H/\Lambda_z \to \partial X(\gamma)^+$ with $\partial i_\gamma^+ = \tau_z \circ \pi_z$. We will show that τ_z is a continuous bijection from a compact topological space to a Hausdorff topological space, and thus is a homeomorphism.

Note that the Gromov boundary of a proper hyperbolic space is compact and metrizable (see for instance [38]), and so $\partial X(\gamma)^+$ is compact Hausdorff and $\partial H/\Lambda_z$ is compact. As ∂i_γ^+ is continuous by virtue of being a Cannon-Thurston map (Lemma 4.9) and the quotient map π_z is also continuous, the map τ_z must be continuous. By Theorem B, ∂i_γ^+ is surjective and so τ_z must also be surjective. If $a', b' \in \partial H/\Lambda_z$ are such that $\tau_z(a') = \tau_z(b') = u \in \partial X(\gamma)^+$, then since ∂i_γ^+ is surjective, there must exist $a, b \in \partial H$ such that $\partial i_\gamma^+(a) = \partial i_\gamma^+(b) = u$. But by Theorem C, this implies that $(a, b) \in \Lambda_z$, and so τ_z is injective. It now follows that $\tau_z : \partial H/\Lambda_z \to \partial X(\gamma)^+$ is a homeomorphism. Therefore by Proposition 6.1, $\partial H/\Lambda_z$ is a dendrite. $\square$

Chapter 7

Open Problems

I would like to end by posing some open problems related to this result.

7.1 Multiplicity of the Cannon-Thurston map

In [22], Dowdall, Kapovich, and Taylor show that if $\Gamma \leq \mathrm{Out}(F_N)$ is purely atoroidal and convex cocompact, then the size of the fibers of the Cannon-Thurston map $\partial i : \partial F_N \to \partial E_\Gamma$ is bounded by twice the rank of the free group. This result generalizes work of Kapovich-Lustig [39], who prove the $2N$ bound for the case where $\Gamma = \langle \varphi \rangle$ is the cyclic group generated by a fully irreducible, atoroidal automorphism. In [29], Ghosh generalizes this result of Kapovich-Lustig to the case where φ is any atoroidal automorphism. Here, $E_\Gamma = G_\varphi = F_N \rtimes_\varphi \mathbb{Z}$. Ghosh shows that given the short exact sequence $1 \to F_N \to G_\varphi \to \langle \varphi \rangle \to 1$, the CT map $\partial i : \partial F_N \to \partial G_\varphi$ is finite-to-one, with the cardinality of the preimage of each point in ∂G_φ bounded by some function of N. However, Ghosh does not give an explicit bound in terms of N. One question to investigate is the following:

Question 7.1. If $\varphi \in \mathrm{Out}(F_N)$ is an atoroidal automorphism and $\partial i : \partial F_N \to \partial G_\varphi$ is the Cannon-Thurston map, then for each $p \in \partial G_\varphi$, is $|(\partial i)^{-1}(p)| \leq 2N$?

The proof in [39] relies on the fully irreducible and atoroidal assumption on φ, as then the attracting and repelling trees are suited to apply the "index theory" from [16]. In the setting of a general atoroidal φ, such trees are no longer available. One approach to this question may be to relate the structure of Mitra's lamination Λ_φ in this setting to the attracting laminations which appear when working with relative train track representatives of φ.

Question 7.2. Let $1 \to H \to G \to Q \to 1$ be a short exact sequence of torsion-free, finitely generated, infinite, hyperbolic groups. Is the Cannon-Thurston map $\partial i : \partial H \to \partial G$ finite-to-one? Is there some bound depending on the structure of H?

As H is a free product of surface groups and a free group in this setting, the answer to the above question ultimately reduces to understanding the case where $H = F_N$ and then interpolating the free group case with known results from surface theory. A place to start might be to examine the case where $H = F_N$ and Q is a purely atoroidal but not convex cocompact subgroup of $\mathrm{Out}(F_N)$ such that G is hyperbolic and determine whether the size of the fibers of the Cannon-Thurston map is also bounded by $2N$.

7.2 Fibers of the Cannon-Thurston map

While there are several settings in which the Cannon-Thurston map is known to exist, there is only an explicit algebraic description of this boundary map for the setting of a hyperbolic group extension. One area to explore would be to develop an algebraic description of the Cannon-Thurston map in other settings where such a map is known to exist. For instance, Mitra shows in [49] that the Cannon-Thurston map exists in the setting of a hyperbolic graph of groups where the vertex and edge graphs are infinite, hyperbolic, and where the defining edge monomorphisms are quasi-isometric embeddings. We say that an endomorphism of the free group F_N is *expanding* if there is some $k \geq 1$ such that for all $w \in F_N$, $|\varphi^k(w)| \geq 2|w|$, where $|w|$ denotes the length of $w \in F_N$.

Question 7.3. Let φ be an injective, but not surjective, expanding endomorphism of F_N. Let $G_\varphi = \langle F_N, t \mid t^{-1}wt = \varphi(w), w \in F_N \rangle$ be the ascending HNN-extension of F_N along φ such that G_φ is hyperbolic. What is an algebraic description of the points in ∂F_N which are identified by the CT map $\partial i : \partial F_N \to \partial G_\varphi$?

In [58], Mutanguha gives conditions on the endomorphism φ which guarantee the mapping torus G_φ will be hyperbolic. In this setting, there is an associated Bass-Serre tree $\mathcal{X}$ for the HNN-extension G_φ. The boundary of $\mathcal{X}$ consists of two parts: an uncountable "backward" boundary, and a single-point "forward" boundary. As φ is not surjective, there should be no lamination associated to the single forward boundary point since φ^{-1} cannot be iterated arbitrarily many times. To each point z in the backward boundary, there is a naturally associated algebraic lamination on F_N, Λ_φ, which depends only on φ (and not on z). An algebraic description of the Cannon-Thurston map in this setting may have applications toward determining when subgroups of F_N are undistorted in G_φ, extending work of Scott-Swarup [63], Mj-Rafi [55], Dowdall-Kent-Leininger [21], and Dowdall-Taylor [23].

In[54], Mj-Pal extend the work of Mitra [49] to the setting where the vertex and edge groups are relatively hyperbolic. Additionally, Pal gives conditions in [60] for the existence of the Cannon-Thurston map in the setting of relatively hyperbolic group extensions. Another question to explore is the following:

Question 7.4. How can the fibers of the Cannon-Thurston map be described in terms of some algebraic ending lamination on H in the setting of a relatively hyperbolic group extension when such a map exists?

7.3 Geometric structure of the Cannon-Thurston map

In [24], Dowdall and Taylor show that when $\Gamma \leq \mathrm{Out}(F_N)$ is purely atoroidal and convex cocompact, the extension group is hyperbolic. However, the converse does not hold. For instance, any automorphism $\phi \in \mathrm{Out}(F_N)$ which is atoroidal but not fully irreducible will give rise to a hyperbolic extension by Brinkman [12], but is not convex cocompact by Bestvina-Feighn [6]. A subgroup $\Gamma \leq \mathrm{Out}(F_N)$ is *nonelementary* if Γ contains two independent fully-irreducible elements. Uyanik [65] gives examples of nonelementary subgroups of $\mathrm{Out}(F_N)$ which are not convex cocompact but which give rise to hyperbolic extensions. One open area of exploration would be to analyze the dendrites that arise from these examples. In particular, let μ be a probability measure on $\mathrm{Out}(F_N)$ with finite first moment whose support generates a subgroup $\Gamma \leq \mathrm{Out}(F_N)$ which is nonelementary. Let $\hat{\mu}$ be the associated exit measure on $\partial\Gamma$. Then, it is known by work of [36] and [59] that for $\hat{\mu}$-almost every $z \in \partial\Gamma$, there exists a class of $\mathbb{R}$-tree $[T_z] \in \partial CV_N$ which is free, arational, and uniquely ergodic. The following question is based on this result and the results of [22] and [39].

Question 7.5. Let $\Gamma \leq \mathrm{Out}(F_N)$ be nonelementary and such that the corresponding extension E_Γ of F_N is hyperbolic, and let μ be a probability measure on $\mathrm{Out}(F_N)$ with finite first moment whose support generates Γ. Then for $\hat{\mu}$-almost every $z \in \partial\Gamma$, is $\partial F_N/\Lambda_z$ homeomorphic to $\widehat{T}_z$?

A more open-ended question to explore would be to determine the geometric structure of the dendrite $\partial H/\Lambda_z$ for a general hyperbolic group extension.

Question 7.6. Let $1 \to H \to G \to Q \to 1$ be a short exact sequence of torsion-free, finitely generated, infinite, hyperbolic groups, and let $z \in \partial Q$. What is a geometric description of the dendrite $\partial H/\Lambda_z$?

References

[1] J. M. Alonso, T. Brady, D. Cooper, V. Ferlini, M. Lustig, M. Mihalik, M. Shapiro, and H. Short. Notes on word hyperbolic groups. In *Group theory from a geometrical viewpoint*, pages 3 – 63. World Scientific, 1991.

[2] O. Baker and T. Riley. Cannon-Thurston maps, subgroup distortion, and hyperbolic hydra. *arXiv e-prints*, page arXiv:1209.0815, Sep 2012.

[3] O. Baker and T. R. Riley. Cannon-Thurston maps do not always exist. *Forum of Mathematics, Sigma*, 1, 2013.

[4] M. Bestvina. $\mathbb{R}$-trees in topology, geometry, and group theory. In *Handbook of Geometric Topology*, pages 55 – 91. Elsevier Science & Technology Books, 2002.

[5] M. Bestvina and M. Feighn. A combination theorem for negatively curved groups. *Journal of Differential Geometry*, 35(1):85–101, 1992.

[6] M. Bestvina and M. Feighn. Hyperbolicity of the complex of free factors. *Advances in Mathematics*, 256:104 – 155, 2014.

[7] M. Bestvina and M. Handel. Train tracks and automorphisms of free groups. *The Annals of Mathematics*, 135(1):1–51, 1992.

[8] M. Bestvina and P. Reynolds. The boundary of the complex of free factors. *Duke Mathematical Journal*, 164(11):2213–2251, 2015.

[9] B. Bowditch. The Cannon-Thurston map for punctured-surface groups. *Mathematische Zeitschrift*, 255(1):35–76, 2007.

[10] B. Bowditch. Stacks of hyperbolic spaces and ends of 3-manifolds. In *Geometry and Topology Down Under*, pages 65–138. American Mathematical Society, Providence, Rhode Island, 2013.

[11] M. R. Bridson and Haefliger A. *Metric spaces of non-positive curvature*. Springer, 2010.

[12] P. Brinkmann. Hyperbolic automorphisms of free groups. *Geometric and Functional Analysis*, 10(5):1071–1089, 2000.

[13] J. W. Cannon and W. P. Thurston. Group invariant Peano curves. *Geometry & Topology*, 11(3):1315–1355, 2007.

[14] A. J. Casson and S. A. Bleiler. *Automorphisms of Surfaces after Nielsen and Thurston.* London Mathematical Society Student Texts. Cambridge University Press, 1988.

[15] M. Coornaert, T. Delzant, and A. Papadopoulos. *Géométrie et théorie des groupes: les groupes hyperboliques de Gromov.* Springer-Verlag, 1990.

[16] T. Coulbois and A. Hilion. Rips induction: index of the dual lamination of an $\mathbb{R}$-tree. *Groups, Geometry, and Dynamics*, 8(1):97 – 134, 2014.

[17] T. Coulbois, A. Hilion, and M. Lustig. Non-unique ergodicity, observers' topology and the dual algebraic lamination for $\mathbb{R}$-trees. *Illinois J. Math.*, 51(3):897–911, 07 2007.

[18] T. Coulbois, A. Hilion, and M. Lustig. $\mathbb{R}$-trees and laminations for free groups I: Algebraic laminations. *Journal of the London Mathematical Society*, 78(3):723–736, 2008.

[19] T. Coulbois, A. Hilion, and M. Lustig. $\mathbb{R}$-trees and laminations for free groups II: The dual lamination of an $\mathbb{R}$-tree. *Journal of the London Mathematical Society*, 78(3):737–754, 2008.

[20] T. Coulbois, A. Hilion, and M. Lustig. $\mathbb{R}$-trees and laminations for free groups III: currents and dual $\mathbb{R}$-tree metrics. *Journal of the London Mathematical Society*, 78(3):755–766, 08 2008.

[21] S. Dowdall, R. Kent IV, and C. Leininger. Pseudo-Anosov subgroups of fibered 3-manifold groups. *Groups, Geometry, and Dynamics*, 8(4):1247 – 1282, 2014.

[22] S. Dowdall, I. Kapovich, and S. J. Taylor. Cannon-Thurston maps for hyperbolic free group extensions. *Israel Journal of Mathematics*, 216(2):753–797, 2016.

[23] S. Dowdall and S. Taylor. Hyperbolic extensions of free groups. *Geometry & Topology*, 22(1):517–570, 2018.

[24] S. Dowdall and S. Taylor. Hyperbolic extensions of free groups. *Geometry & Topology*, 22(1):517–570, 2018.

[25] D. B. A. Epstein, M. S. Paterson, J. W. Cannon, D. F. Holt, S. V. Levy, and W. P. Thurston. *Word Processing in Groups.* A. K. Peters, Ltd., Natick, MA, USA, 1992.

[26] B. Farb and D. Margalit. *A Primer on Mapping Class Groups (PMS-49).* Princeton University Press, 2012.

[27] B. Farb and L. Mosher. Convex cocompact subgroups of mapping class groups. *Geometry & Topology*, 6(1):91–152, 2002.

[28] A. Fathi, F. Laudenbach, and V. Poénaru. *Thurston's Work on Surfaces (MN-48)*. Mathematical Notes - Princeton University Press. Princeton University Press, 2012.

[29] P. Ghosh. Limits of conjugacy classes under iterates of Hyperbolic elements of $\mathsf{Out}(\mathbb{F})$. *ArXiv e-prints*, September 2017.

[30] É. Ghys and P. de La Harpe, editors. *Sur les groupes hyperboliques d'après Mikhael Gromov*. Progress in mathematics. Birkhäuser, 1990.

[31] E. Ghys, A. Haefliger, and A. Verjovsky. *Group Theory from a Geometrical Viewpoint*. World Scientific, 1991.

[32] M. Gromov. Hyperbolic groups. In *Essays in Group Theory*, pages 75 – 263. Springer New York, New York, NY, 1987.

[33] U. Hamenstädt. Word hyperbolic extensions of surface groups. *arXiv preprint math/0505244*, 2005.

[34] U. Hamenstädt. The boundary of the free splitting graph and the free factor graph. *arXiv preprint arXiv:1211.1630*, 2012.

[35] U. Hamenstädt and S. Hensel. Stability in outer space. *Groups, Geometry, and Dynamics*, 12(1):359–398, 2018.

[36] C. Horbez. The Poisson boundary of $\mathrm{Out}(F_N)$. *Duke Math. J.*, 165(2):341–369, 02 2016.

[37] W. Jeon, I. Kapovich, C. Leininger, and K. Ohshika. Conical limit points and the cannon-thurston map. *Conformal Geometry and Dynamics of the American Mathematical Society*, 20(4):58–80, 2016.

[38] I. Kapovich and N. Benakli. Boundaries of hyperbolic groups. *Contemporary Mathematics Combinatorial and Geometric Group Theory*, page 39–93, 2002.

[39] I. Kapovich and M. Lustig. Cannon-Thurston fibers for iwip automorphisms of F_N. *Journal of the London Mathematical Society*, 91(1):203–224, 2015.

[40] I. Kapovich and H. Short. Greenbergs theorem for quasiconvex subgroups of word hyperbolic groups. *Canadian Journal of Mathematics*, 48(6):1224–1244, 1996.

[41] R. P. Kent and C. J. Leininger. Shadows of mapping class groups: Capturing convex cocompactness. *Geometric and Functional Analysis*, 18(4):1270–1325, Jun 2008.

[42] E. Klarreich. The boundary at infinity of the curve complex and the relative teichmüller space, 2018.

[43] C. Leininger, M. Mj, and S. Schleimer. The universal Cannon-Thurston map and the boundary of the curve complex. *Commentarii Mathematici Helvetici*, pages 769–816, 2011.

[44] H. A. Masur and Y. N. Minsky. Geometry of the complex of curves I: Hyperbolicity. *Inventiones Mathematicae*, 138, 03 1999.

[45] Y. Matsuda and S. Oguni. On Cannon-Thurston maps for relatively hyperbolic groups. *Journal of Group Theory*, 17(1):41–47, 2014.

[46] Y. N. Minsky. On rigidity, limit sets, and end invariants of hyperbolic 3-manifolds. *Journal of the American Mathematical Society*, 7(3):539–588, 1994.

[47] M. Mitra. Ending laminations for hyperbolic group extensions. *Geometric and Functional Analysis*, 7(2):379–402, 1997.

[48] M. Mitra. Cannon-Thurston maps for hyperbolic group extensions. *Topology*, 37(3):527–538, 1998.

[49] M. Mitra. Cannon-Thurston maps for trees of hyperbolic metric spaces. *Journal of Differential Geometry*, 48(1):135–164, 1998.

[50] M. Mitra. Coarse extrinsic geometry: a survey. *The Epstein Birthday Schrift*, 1998.

[51] M. Mitra. On a theorem of Scott and Swarup. *Proceedings of the American Mathematical Society*, 127(6):1625–1631, 1999.

[52] M. Mj. Cannon-Thurston maps for surface groups. *Annals of Mathematics*, 179(1):1–80, Jan 2014.

[53] M. Mj. Cannon-Thurston maps for Kleinian groups. *Forum of Mathematics, Pi*, 5, 2017.

[54] M. Mj and A. Pal. Relative hyperbolicity, trees of spaces and Cannon-Thurston maps. *Geometriae Dedicata*, 151(1):59 – 78, 2011.

[55] M. Mj and K. Rafi. Algebraic ending laminations and quasiconvexity. *Algebraic & Geometric Topology*, 18(4):1883–1916, 2018.

[56] M. Mj and P. Sardar. A combination theorem for metric bundles. *Geometric and Functional Analysis*, 22(6):1636–1707, 2012.

[57] L. Mosher. Hyperbolic extensions of groups. *Journal of Pure and Applied Algebra*, 110(3):305 – 314, 1996.

[58] J. P. Mutanguha. Irreducible Nonsurjective Endomorphisms of F_n are Hyperbolic. *ArXiv e-prints*, page arXiv:1908.08214, Aug 2019.

[59] H. Namazi, A. Pettet, and P. Reynolds. Ergodic decompositions for folding and unfolding paths in Outer space. *ArXiv e-prints*, 2014.

[60] A. Pal. Relatively hyperbolic extensions of groups and Cannon-Thurston maps. *Proceedings - Mathematical Sciences*, 120(1):57 – 68, Feb 2010.

[61] F. Paulin. Outer automorphisms of hyperbolic groups and small actions on $\mathbb{R}$-trees. *Mathematical Sciences Research Institute Publications Arboreal Group Theory*, page 331–343, 1991.

[62] E. Rips and Z. Sela. Structure and rigidity in hyperbolic groups I. *Geometric & Functional Analysis GAFA*, 4(3):337–371, May 1994.

[63] G. P. Scott and G. A. Swarup. Geometric finiteness of certain Kleinian groups. *Proceedings of the AMS*, 109(3):765–768, 1990.

[64] P. B. Shalen. Dendrology of groups: An introduction. *Essays in Group Theory Mathematical Sciences Research Institute Publications*, page 265–319, 1987.

[65] C. Uyanik. Hyperbolic extensions of free groups from atoroidal ping-pong. *Algebr. Geom. Topol.*, 19(3):1385–1411, 2019.