Semiconjugacies between Kleinian Group Actions on the Riemann Sphere

Erica Klarreich

Abstract

We consider the problem of characterizing topologically the action of a Kleinian group on the Riemann sphere $\hat{\mathbf{C}}$. We prove that certain geometrically infinite Kleinian group actions on $\hat{\mathbf{C}}$ can be obtained from geometrically finite ones by a semiconjugacy that is determined by the end invariants of the geometrically infinite group. This turns out to be related to the problem of continuous extensions of maps of hyperbolic 3-space $\mathbf{H^3}$ to maps of its boundary at infinity, $\hat{\mathbf{C}}$. Along the way we consider the general problem of extending maps to the boundary at infinity in Gromov-hyperbolic metric spaces. We give criteria for extending to the boundary a map $h : X \to Y$ between hyperbolic spaces if it extends to the boundary of certain subsets of X.

1 Introduction

1.1 Kleinian Groups

If Γ is a Kleinian group (a discrete group of orientation-preserving isometries of $\mathbf{H^3}$), then the action of Γ extends naturally to an action by conformal homeomorphisms (Möbius transformations) on the boundary at infinity $\hat{\mathbf{C}}$ of $\mathbf{H^3}$. A fundamental problem is to understand the topology of this action and its interplay with the geometry of the hyperbolic manifold (or orbifold) $\mathbf{H^3}/\Gamma$.

The Kleinian groups for which this interplay is best understood are the *geometrically finite* groups without parabolics (also known as *convex cocompact* groups); for these groups, each end of the quotient manifold corresponds to a collection of components of the domain of discontinuity in $\hat{\mathbf{C}}$, which are shuffled by the group action. Moreover, the limit set of such a group is naturally homeomorphic to the Gromov boundary of the group (see [8, 18]), a purely group-theoretical structure. If a Kleinian group Γ is not geometrically finite, we may try to model its action on $\hat{\mathbf{C}}$ by the action of an isomorphic geometrically finite group Γ_0, via a semiconjugacy (that is, a continuous map from $\hat{\mathbf{C}}$ to $\hat{\mathbf{C}}$ that is equivariant with respect to the actions of Γ_0 and Γ and a given isomorphism $i : \Gamma_0 \to \Gamma$). Thurston has conjectured (see [1]):

Conjecture *If Γ is a finitely generated Kleinian group then there exists an isomorphic geometrically finite Kleinian group Γ_0 and a semiconjugacy from the action of Γ_0 on $\hat{\mathbf{C}}$ to the action of Γ on $\hat{\mathbf{C}}$.*

We will prove this conjecture in a special case, and show that the identifications that the semiconjugacy makes on $\hat{\mathbf{C}}$ are completely determined by the *end invariants* of the quotient manifold $\mathbf{H^3}/\Gamma$, which are all either Riemann surfaces or geodesic laminations called *ending laminations* (see [4, 30]). This proposed model for the topology of the action has been shown by Cannon and Thurston to be correct when the quotient manifold is the infinite cyclic cover of a compact surface bundle that fibers over the circle (see [12]), and more generally by Minsky ([24]) for any surface group whose quotient manifold admits a uniform lower bound on injectivity radius. We will apply Minsky's theorem to certain subgroups of Γ to obtain maps from portions of $\hat{\mathbf{C}}$ to $\hat{\mathbf{C}}$ that piece together to provide the desired semiconjugacy.

In this paper we will study finitely generated torsion-free Kleinian groups that satisfy the following two conditions:
C1. Γ is freely indecomposable; that is, Γ cannot be written as a non-trivial free product.
C2. There is a uniform lower bound on the injectivity radius at every point of $\mathbf{H^3}/\Gamma$.

If Γ satisfies $C1$ then Bonahon has shown that Γ is topologically tame, that is, $\mathbf{H^3}/\Gamma$ is the interior of a compact manifold (see [4]). By Thurston's geometrization theorem (see [26, 22]) there exists a geometrically finite group Γ_0 such that the quotient manifolds $\mathbf{H^3}/\Gamma_0$ and $\mathbf{H^3}/\Gamma$ are homeomorphic. It is this group Γ_0 that we use to model the action of Γ on $\hat{\mathbf{C}}$. We will show the following:

Theorem A *Let Γ_0 and Γ be finitely generated, torsion-free Kleinian groups with homeomorphic quotients, that satisfy C1 and C2 and such that Γ_0 is geometrically finite. Then there exists a homeomorphism $g : \mathbf{H^3}/\Gamma_0 \to \mathbf{H^3}/\Gamma$ in the same homotopy class as the given homeomorphism, whose lift $\tilde{g} : \mathbf{H^3} \to \mathbf{H^3}$ to the universal covers extends to a continuous, surjective map $\tilde{g} : \hat{\mathbf{C}} \to \hat{\mathbf{C}}$ that semiconjugates the action of Γ_0 to the action of Γ.*

When we assume that Γ satisfies $C2$, the geometrically finite group Γ_0 given by Thurston's geometrization theorem is in fact convex cocompact, so that Γ may be regarded as a hyperbolic group in the sense of Gromov. Let $\partial\Gamma$ denote the Gromov boundary of Γ, and let $\Lambda_\Gamma \subset \hat{\mathbf{C}}$ denote the limit set of Γ. An immediate consequence of Theorem A is

Corollary *Let Γ be a finitely generated torsion-free Kleinian group that satisfies $C1$ and $C2$. Then there is a continuous, surjective map $\partial\Gamma \to \Lambda_\Gamma$ that is equivariant with respect to the action of Γ.*

If Γ_0 is as in Theorem A then the components of the domain of discontinuity Ω of Γ_0 are topological disks. If Ω_α is a component of Ω that corresponds to a degenerate end of $\mathbf{H^3}/\Gamma$, then Ω_α can be equipped with a lamination $\tilde{\lambda}_\alpha$ that is a lift of the corresponding ending lamination of $\mathbf{H^3}/\Gamma$. The semiconjugacy constructed in Theorem A identifies points in the following way:

Theorem B *Let $\tilde{g} : \hat{\mathbf{C}} \to \hat{\mathbf{C}}$ be the map constructed in Theorem A. If x and y are in $\hat{\mathbf{C}}$ then $\tilde{g}(x) = \tilde{g}(y)$ iff for some Ω_α, x and y lie on the closure of some leaf or complementary component of $\tilde{\lambda}_\alpha$.*

A similar construction appears in holomorphic dynamics; certain filled Julia sets may be obtained from the unit disk by collapsing along the leaves of a lamination (see e.g. [13]).

If $i : \Gamma_0 \to \Gamma$ is the isomorphism determined by the homeomorphism between $\mathbf{H^3}/\Gamma_0$ and $\mathbf{H^3}/\Gamma$

then any semiconjugacy with respect to i must agree with the semiconjugacy of Theorem A on the limit set of Γ_0. Thus, when Γ satisfies $C1$ and $C2$, Theorems A and B give an alternate description of the ending laminations of $\mathbf{H^3}/\Gamma$: they are the laminations determined by the identifications any semiconjugacy from Γ_0 to Γ makes on the limit set of Γ_0. This point of view is different from that of the original definition of ending laminations by Thurston, which relies heavily on the internal geometry of the quotient 3-manifold $\mathbf{H^3}/\Gamma$.

The main strategy in constructing the semiconjugacy $\tilde{g}$ is as follows: in our setting, if Ω_α is a component of the domain of discontinuity of Γ_0 then the closure of Ω_α in $\hat{\mathbf{C}}$ is the boundary at infinity of a topological half-space H_α in $\mathbf{H^3}$; the half-spaces H_α are the components of the complement of the convex hull (in $\mathbf{H^3}$) of the limit set of Γ_0. Using Minsky's theorem, for each set H_α we can construct a map $\tilde{g}_\alpha : H_\alpha \to \mathbf{H^3}$ that extends continuously to the boundary and identifies points according to the ending lamination of Ω_α (if it has one). We piece these maps together by "filling in" on the convex hull, to obtain a map $\tilde{g} : \mathbf{H^3} \to \mathbf{H^3}$ defined on all of $\mathbf{H^3}$. The map $\tilde{g}$ extends continuously to the boundary of each half-space H_α. However, the typical point in the limit set of Γ_0 will in general not lie on the boundary of a set H_α, but rather will be an accumulation point of some sequence of the sets H_α; in fact, the Hausdorff dimension of the limit set of Γ_0 is strictly greater than the Hausdorff dimension of the portion of the limit set that lies in the boundary of the sets H_α (see [7]). The problem then is to control the map at those points where the half-spaces accumulate.

1.2 The general extension problem in hyperbolic spaces

Given a map $h : X \to Y$ between Gromov hyperbolic spaces, we can ask what is a sufficient condition for h to extend continuously to a map from the boundary at infinity $\partial_\infty X$ of X to the boundary at infinity $\partial_\infty Y$ of Y (see Section 3 for definitions). One such sufficient condition is that h is a quasi-isometry (see [18]); but it is certainly not a necessary condition. In this paper we give a more general condition, motivated by the picture in the Kleinian groups setting. A different sufficient condition is given by Mitra in [25].

The proof that the map $\tilde{g} : \mathbf{H^3} \to \mathbf{H^3}$ constructed in the Kleinian groups problem extends to the boundary uses only coarse properties of hyperbolic space, and can be generalized to the setting of Gromov hyperbolic spaces. Consider a pair $(\Delta, \{S_\alpha\})$ where Δ is a proper, geodesic Gromov-hyperbolic space and $\{S_\alpha\}$ is a collection of closed, disjoint, path-connected subsets of Δ. If we have a map h from Δ to another hyperbolic space whose restriction to each set S_α extends continuously to the boundary, we may ask whether the map extends continuously to all of $\partial_\infty\Delta$. To ensure that such a map will extend to the boundary we must mimic some of the regularity that the Kleinian group actions on $\mathbf{H^3}$ give to the geometry and arrangement of the collections $\{H_\alpha\}$ and $\{\tilde{g}(H_\alpha)\}$ in $\mathbf{H^3}$. We will consider the following conditions, each of which holds in the Kleinian groups setting:

($*_1$) The complement $\mathcal{C}$ of the sets S_α is open and path connected.
($*_2$) There is some real number $q \geq 0$ for which the sets S_α are all q-quasiconvex.
($*_3$) There exists a real number $c > 0$ such that $d(S_\alpha, S_\beta) > c$ for all α and β.

It is useful to consider the space E_Δ obtained from Δ by identifying each set S_α to a point; after Farb ([15]), we will call this space the *electric space* of Δ. The metric on Δ induces an *electric metric* on E_Δ. We will show that the space E_Δ, although generally not a proper metric space, is hyperbolic in the sense of Gromov; thus, it can be equipped with a natural boundary at infinity. We will show that the boundary of E_Δ can be identified with the subset of $\partial_\infty \Delta$ consisting of those points that do not lie on the boundary of any of the sets S_α.

If $h : X \to Y$ is a map between proper, geodesic Gromov-hyperbolic spaces and there is a collection of subsets H_α of X such that the pairs $(X, \{H_\alpha\})$ and $(Y, \{h(H_\alpha)\})$ satisfy $(*_1) - (*_3)$, then h induces a map $h_E : E_X \to E_Y$ on the electric spaces. If h_E is a quasi-isometry (which we can arrange to be true in the Kleinian groups setting) then it will extend continuously to a map from $\partial_\infty E_X \to \partial_\infty E_Y$ (see e.g. [17]); this will give us information about the extendibility of h at those points of $\partial_\infty X$ that do not lie on the boundary of any set H_α. This will enable us to show:

Theorem C *Let X and Y be proper, geodesic Gromov-hyperbolic spaces, $\{H_\alpha\}$ a collection of closed, disjoint path-connected subsets of X, and $h : X \to Y$ a quasi-Lipschitz map such that for every H_α, $h|_{H_\alpha}$ extends continuously to a continuous map $h : \partial_\infty H_\alpha \to \partial_\infty Y$. If $(X, \{H_\alpha\})$ and $(Y, \{h(H_\alpha)\})$ satisfy $(*_1) - (*_3)$ and the induced map $h_E : E_X \to E_Y$ is a quasi-isometry then h extends continuously to a continuous map $h : \partial_\infty X \to \partial_\infty Y$.*

(Note that in this context, where maps are not usually assumed to be continuous, we say that a map $f : X \to Y$ extends continuously to a map from $\partial_\infty X$ to $\partial_\infty Y$ if for every $\xi \in \partial_\infty X$, if (x_n) is a sequence in X that converges to ξ, then $f(x_n)$ converges to a point $\eta \in \partial_\infty Y$ that is uniquely determined by ξ.)

Since $h_E : E_X \to E_Y$ is a quasi-isometry its extension to the boundary is in fact injective (see [18]). This leads to the following result about where h can be non-injective, which will be used to prove Theorem B:

Theorem D *Let $h : \partial_\infty X \to \partial_\infty Y$ be the extension map constructed in Theorem C. If ξ and ρ are points in $\partial_\infty X$ such that $h(\xi) = h(\rho)$ then for some α and β, $\xi \in \partial_\infty H_\alpha$ and $\rho \in \partial_\infty H_\beta$.*

If we make the stronger assumption that $h_E : E_X \to E_Y$ is a bi-Lipschitz homeomorphism (this is true in the Kleinian groups setting) then we can say the following about the identifications of the map $h : \partial_\infty X \to \partial_\infty Y$:

Theorem E *Let X, Y, $\{H_\alpha\}$ and $h : X \to Y$ be as in Theorem C, and suppose that the induced map $h_E : E_X \to E_Y$ on electric spaces is a bi-Lipschitz homeomorphism. If ξ and ρ are points in*

$\partial_\infty X$ such that $h(\xi) = h(\rho)$ then there is a finite chain of electric sets $\{H_0, ...H_m\}$ such that

(1) $\xi \in \partial_\infty H_0$ and $\rho \in \partial_\infty H_m$.

(2) For $i = 0, ..., m-1$ there is a point $p_i \in \partial_\infty H_i \cap \partial_\infty H_{i+1}$ such that $h(p_i) = h(\xi) = h(\rho)$.

In the Kleinian groups setting, we will see that such a chain cannot consist of more than one set H_α; thus two points on $\hat{\mathbf{C}}$ can only be identified by the semiconjugacy if they lie in the closure of the same component of the domain of discontinuity of Γ_0. This gives us Theorem B.

In Section 2 we review background material on Kleinian groups, and discuss the history of Thurston's conjecture on semiconjugacies. We use a theorem of Minsky and the structure of geometrically finite Kleinian groups satisfying $C1$ to show that Theorem A is a special case of Theorem C. Section 3 consists of a review of the basic ideas of Gromov-hyperbolic spaces. Section 4 consists of the proof of Theorem C. In Section 5 we study the boundary of the space E_Δ, and prove Theorems D and E; we use Theorem E and some additional information about ending laminations to prove Theorem B.

Acknowledgments. I would like to thank my advisor, Dr. Yair Minsky, for many helpful and enjoyable discussions. I would also like to thank the State University of New York at Stony Brook, where most of this work was completed, for its hospitality.

2 Kleinian Groups

A Kleinian group is a discrete subgroup Γ_0 of the group of orientation-preserving isometries of $\mathbf{H^3}$. If Γ_0 has no torsion (elliptic elements) then $N_0 = \mathbf{H^3}/\Gamma_0$ is a complete, orientable hyperbolic 3-manifold, and $\Gamma_0 \cong \Pi_1(N_0)$. The group Γ_0 acts by conformal homeomorphisms (Möbius transformations) on the Riemann sphere $\hat{\mathbf{C}}$, the boundary at infinity of $\mathbf{H^3}$. This action partitions $\hat{\mathbf{C}}$ into two Γ_0-invariant sets: the *limit set* Λ_{Γ_0}, which is the set of accumulation points of any orbit under Γ_0, and its complement Ω_{Γ_0}, the *domain of discontinuity*, an open subset of $\hat{\mathbf{C}}$. Γ_0 acts freely and properly discontinuously on $\mathbf{H^3} \cup \Omega_{\Gamma_0}$; the quotient space $\overline{N_0} = (\mathbf{H^3} \cup \Omega_{\Gamma_0})/\Gamma_0$ is a manifold-with-boundary.

2.1 The ends of a hyperbolic manifold

Let X be a Hausdorff locally compact topological space. If $K \subset X$ is compact, define

$\varepsilon(K) = \{E : E$ a connected component of $X \setminus K$, $\overline{E}$ not compact$\}$.

If there is a compact set $K \subset X$ such that for every compact set $K' \supset K$ the inclusion induces a bijection $\varepsilon(K') \to \varepsilon(K)$ then we call each element of $\varepsilon(K)$ an *end* of X; this definition is essentially independent of K, since if K and K' are two such compact sets then there is a natural bijection between $\varepsilon(K)$ and $\varepsilon(K')$. Heuristically, the ends of X are those parts of X that go off to infinity.

If for some such compact set K the open set E is a connected component of $X \setminus K$, we say that E is a neighborhood of the corresponding end.

Geometrically finite ends. Let $N_0 = \mathbf{H^3}/\Gamma_0$ be a manifold that satisfies $C1$ and that has no cusps (that is, Γ_0 has no parabolic elements), and let $CC(N_0)$ denote the *convex core* of N_0, which is the smallest closed convex submanifold of N_0 whose inclusion is a homotopy equivalence. The convex core may be obtained as the quotient by Γ_0 of the convex hull $CH(\Gamma_0)$ in $\mathbf{H^3}$ of the limit set of Γ_0 (see [14]). An end of N_0 is called *geometrically finite* if it has a neighborhood that is disjoint from $CC(N_0)$. And end that is not geometrically finite is called *degenerate*, or *geometrically infinite*. We say that N_0 is geometrically finite if all of its ends are geometrically finite; equivalently, N_0 is geometrically finite if its convex core is compact. If N_0 has cusps, we say that N_0 is geometrically finite if its convex core has finite volume.

For geometrically finite groups without cusps (that is, convex cocompact groups), Theorem A has long been known: if $N_0 = \mathbf{H^3}/\Gamma_0$ and $N = \mathbf{H^3}/\Gamma$ are homeomorphic geometrically finite manifolds without cusps then there is a homeomorphism $g : N_0 \to N$ whose lift $\tilde{g} : \mathbf{H^3} \to \mathbf{H^3}$ is a quasi-isometry, which (necessarily) extends to a quasiconformal homeomorphism $\tilde{g} : \hat{\mathbf{C}} \to \hat{\mathbf{C}}$ that conjugates the action of Γ_0 to the action of Γ.

2.2 Parabolic Pinching

The question of modelling the action of a Kleinian group on $\hat{\mathbf{C}}$ by the action of a convex cocompact group has largely been answered for geometrically finite groups with parabolics. One class of groups that has been extensively studied are the *regular b-groups*, that is, geometrically finite groups that preserve a simply connected component of the domain of discontinuity (see [3, 20, 2] for discussions of regular b-groups). If Γ is a regular b-group then it is a *surface group*, that is, it is isomorphic to the fundamental group of a surface; if Γ is isomorphic to the fundamental group of a closed surface, we may try to model its action on $\hat{\mathbf{C}}$ by the action of an isomorphic Fuchsian group (a Fuchsian group is a geometrically finite Kleinian surface group whose limit set is a round circle). Concretely, let Δ be the invariant component of the domain of discontinuity of Γ, and let $\phi : D \to \Delta$ be a conformal homeomorphism, where D is the unit disk in $\hat{\mathbf{C}}$. The group $\Gamma_0 = \phi^{-1} \circ \Gamma \circ \phi$ is a convex cocompact Fuchsian group, and we may ask whether the map ϕ extends to an equivariant map defined on all of $\hat{\mathbf{C}}$. In contrast to the situation considered in the previous subsection when Γ had no parabolic elements and was geometrically finite, we cannot expect to construct an actual conjugacy between the actions of Γ_0 and Γ on $\hat{\mathbf{C}}$; certain loxodromic elements of Γ_0 correspond to parabolic elements of Γ (called accidental parabolics), and if x and y are the two fixed points in $\hat{\mathbf{C}}$ of such an element of Γ_0 then any semiconjugacy must map both x and y to the single fixed point of the corresponding parabolic element of Γ.

Abikoff has shown (see [2]) that the limit set of Γ $(= \partial\Delta)$ is locally connected, hence by classical complex analysis the map $\phi : D \to \Delta$ extends continuously to a map $\overline{D} \to \overline{\Delta}$. Let D' denote the component $\hat{\mathbf{C}} \setminus \overline{D}$ of the domain of discontinuity of Γ_0; equip D' with a Poincaré metric. Work of Floyd ([16]) shows that we may define an extension of ϕ to all of $\hat{\mathbf{C}}$ to obtain a semiconjugacy from the action of Γ_0 to the action of Γ, with the following properties:

(1) If $\gamma_0 \in \Gamma_0$ corresponds to a parabolic element $\gamma \in \Gamma$ and l is the axis of γ_0 in D' then ϕ maps all of l (including endpoints) onto the fixed point of γ (see Figure 1).

(2) The collapsings in (1) are the only identifications that ϕ induces.

For more general geometrically finite Kleinian groups with parabolics, it is generally accepted that a similar pinching picture exists; it seems probable that the techniques used in the proofs of Theorems A and C of this paper would also work in the setting of geometrically finite groups with parabolics.

2.3 The Cannon-Thurston theorems

The first progress towards proving Thurston's semiconjugacy conjecture for geometrically infinite groups was made by Cannon and Thurston, who studied the following class of examples of Kleinian surface groups: let S be a closed surface of genus ≥ 2, and let $\phi : S \to S$ be a homeomorphism. Let M be the manifold obtained from $S \times [0,1]$ by the identification $(x,0) \sim (\phi(x),1)$. M is called the *mapping torus* of ϕ, and it fibers over the circle S^1. Thurston has shown that if ϕ is pseudo-Anosov then M admits a hyperbolic structure (see [31]). Let N be the infinite cyclic cover of M corresponding to the fiber group $\Pi_1(S)$. N is a hyperbolic manifold homeomorphic to $S \times \mathbf{R}$, and both of its ends are geometrically infinite. If Γ_0 is an isomorphic Fuchsian group, Cannon and Thurston ([12]) have shown that there is a homeomorphism from $N_0 = \mathbf{H^3}/\Gamma_0$ to N whose lift to $\mathbf{H^3}$ extends continuously to $\hat{\mathbf{C}}$ to give a semiconjugacy from the action of Γ_0 to the action of Γ. Such a map must send the limit set of Γ_0 onto the limit set of Γ; in this case, the limit set of Γ_0 is a circle and the limit set of Γ is the entire Riemann sphere, so we obtain an equivariantly parametrized Peano curve.

Cannon and Thurston show the existence of a semiconjugacy by constructing a model manifold N_{mod} that is homeomorphic and quasi-isometric to N, and whose geometry is completely determined by the isotopy type of the pseudo-Anosov homeomorphism ϕ. Given a pseudo-Anosov homeomorphism $\phi : S \to S$ we can associate to S a pair of measured singular foliations (Φ_x, dx) and (Φ_y, dy), called the *stable* and *unstable foliations of* ϕ that satisfy the following properties (see [10]):

(1) $\phi(\Phi_x) = \Phi_x$ and $\phi(\Phi_y) = \Phi_y$.

(2) There is some multiplier $k > 1$ such that $\phi(dx) = k dx$ and $\phi(dy) = \frac{1}{k} dy$.

Thus, in the fibred manifold M, each time we go around S^1 we have scaled in the x-direction by k and in the y-direction by $\frac{1}{k}$. The metric on the model manifold $N_{mod} \cong S \times \mathbf{R}$ is defined by

$$ds^2 = k^{2t} dx^2 + k^{-2t} dy^2 + (\log k)^2 dt^2$$

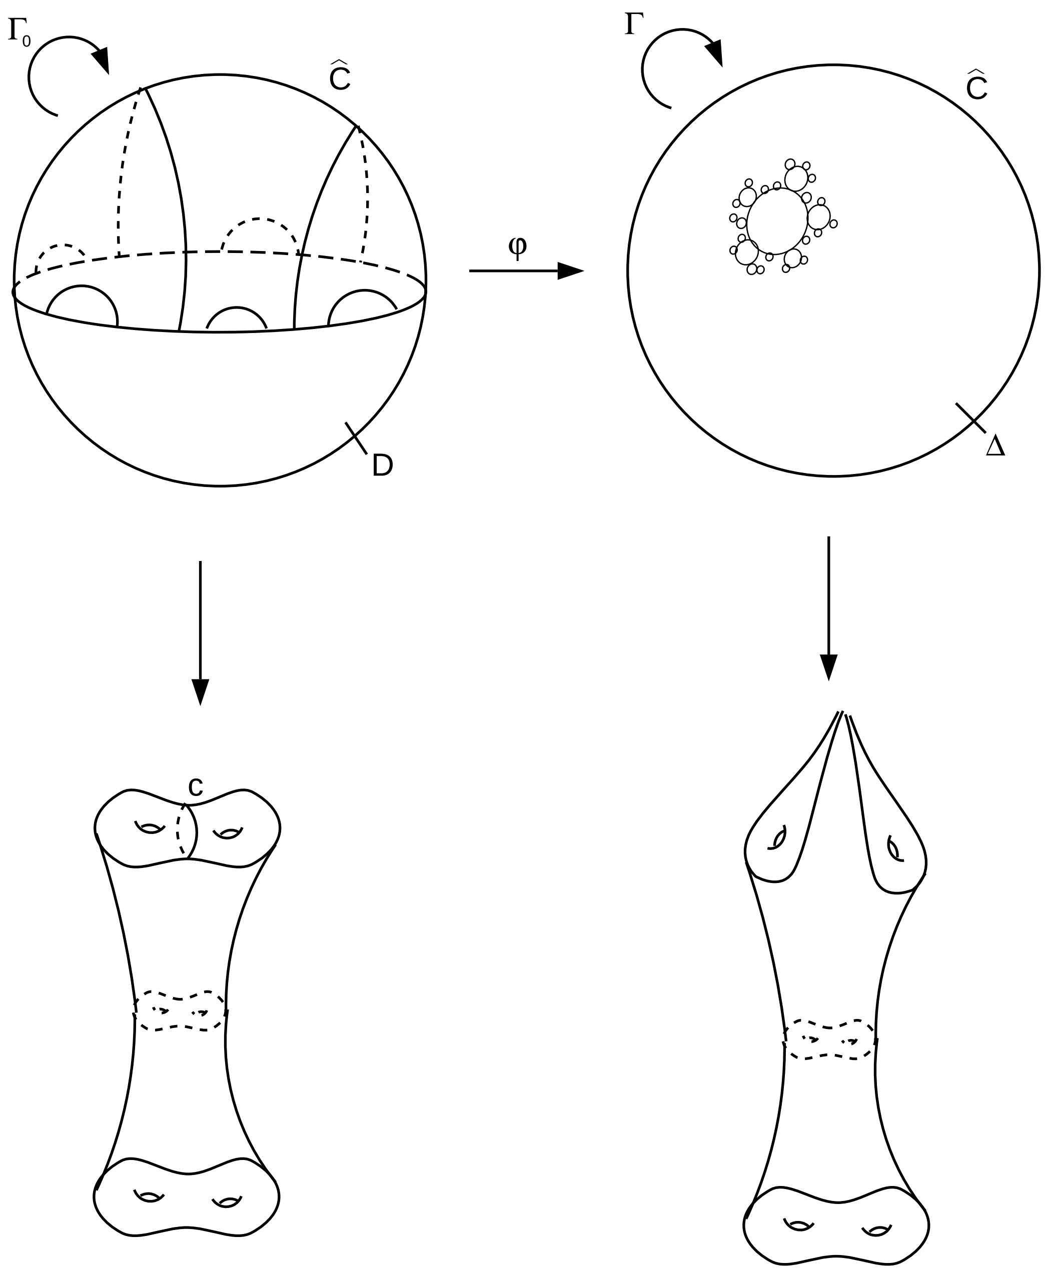

Figure 1: *If the curve c corresponds to a cusp in $\mathbf{H^3}/\Gamma$ then each lift of c to $\hat{\mathbf{C}}$ is pinched by ϕ to a point.*

where dt^2 is the metric on $\mathbf{R}$. Let $\mathbf{H^2} \times \mathbf{R}$ be the universal cover of N_{mod}. The homeomorphism from $N_{mod} = S \times \mathbf{R}$ to $N = S \times \mathbf{R}$ is a quasi-isometry, and its lift $g : \mathbf{H^2} \times \mathbf{R} \to \mathbf{H^3}$ to the universal covers is also a quasi-isometry (see [12]).

Let Γ_0 be a Fuchsian group isomorphic to Γ. Γ_0 preserves a geodesic plane in $\mathbf{H^3}$, giving a homeomorphism of $\mathbf{H^3}$ with $\mathbf{H^2} \times \mathbf{R}$; the action of Γ_0 preserves the product structure of $\mathbf{H^2} \times \mathbf{R}$, giving a product structure to $N_0 = \mathbf{H^3}/\Gamma_0$ (that is, giving a specific homeomorphism of N_0 with $S \times \mathbf{R}$). The induced map from N_0 to N_{mod} lifts to a homeomorphism $f : \mathbf{H^3} \to \mathbf{H^2} \times \mathbf{R}$.

Let $F = g \circ f : \mathbf{H^3} \to \mathbf{H^3}$; Cannon and Thurston have shown that F extends continuously to $\hat{\mathbf{C}}$. Since F is the lift of a homeomorphism from N_0 to N, it is equivariant with respect to the actions of Γ_0 and Γ, and its extension to $\hat{\mathbf{C}}$ gives a surjective semiconjugacy from the action of Γ_0 on $\hat{\mathbf{C}}$ to the action of Γ. The proof that F extends to $\hat{\mathbf{C}}$ is fairly involved; however, it is not hard to show that, in analogy to the situation in the case of accidental parabolics in b-groups, F collapses certain portions of $\hat{\mathbf{C}}$ that correspond to the foliations Φ_x and Φ_y. Lift Φ_x and Φ_y to foliations $\tilde{\Phi}_x$ and $\tilde{\Phi}_y$ on the universal cover $\mathbf{H^2}$ of S. Let L be a leaf of either $\tilde{\Phi}_x$ or $\tilde{\Phi}_y$, say $\tilde{\Phi}_x$; then the topological disk $L \times \mathbf{R}$ is totally geodesic in $\mathbf{H^2} \times \mathbf{R}$ (see e.g. [24]). The metric ds^2 restricted to $L \times \mathbf{R}$ has the form

$$k^{-2t} dy + (\log k)^2 dt^2;$$

this is a hyperbolic metric, in which curves of the form $L \times \{t\}$ correspond to horoballs with a common boundary point. The topological disk $P = f^{-1}(L \times \mathbf{R})$ meets $\hat{\mathbf{C}} = \partial_\infty \mathbf{H^3}$ in a topological circle and its intrinsic geometry is that of the hyperbolic plane. The map f restricted to P takes the geodesic $f^{-1}(L) \times \{0\}$ and its equidistant curves to nested horocycles in $L \times \mathbf{R}$ (see Figure 2). Thus the extension of f to the boundary circle of P collapses the entire upper semicircle to a point,

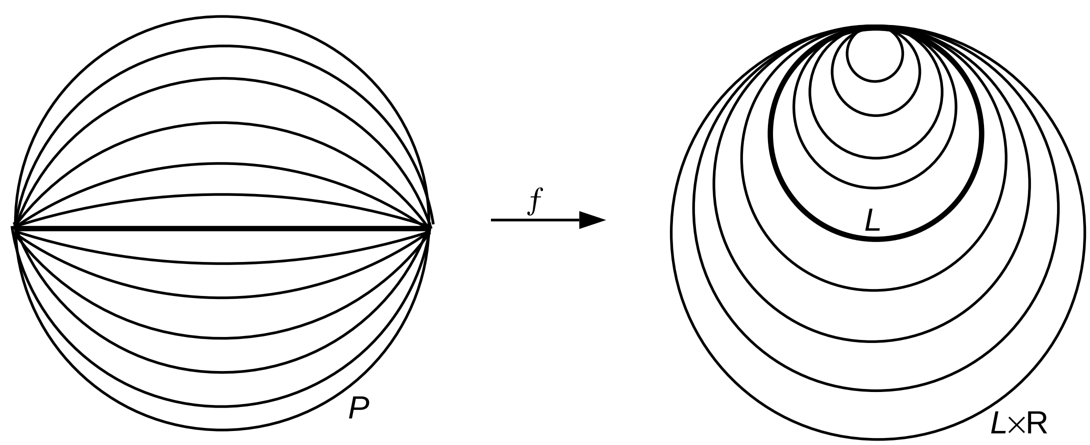

Figure 2: *The restriction of f to P takes the geodesic $t = 0$ and its equidistant curves to nested horocycles in $L \times \mathbf{R}$.*

and stretches the lower semicircle around the boundary circle of $L \times \mathbf{R}$. The map g restricted to $L \times \mathbf{R}$ is a quasi-isometry so it extends continuously to an injective map of the boundary at infinity $\partial_\infty(L \times \mathbf{R})$ (see e.g. [18]); hence the map $F = g \circ f$ identifies points in $\partial_\infty P$ in the same way as the map f. If L is a leaf of $\tilde{\Phi}_y$ then we get a similar picture, but with the lower semicircle being collapsed, rather than the upper semicircle.

Thus, the semiconjugacy $F : \hat{\mathbf{C}} \to \hat{\mathbf{C}}$ identifies points in the following way: Let Ω_+ and Ω_- denote the two components of the domain of discontinuity of Γ_0 (with Ω_+ corresponding to $\mathbf{H^2} \times \{\infty\}$ and Ω_- corresponding to $\mathbf{H^2} \times \{-\infty\}$ with respect to our chosen identification of $\mathbf{H^3}$ with $\mathbf{H^2} \times \mathbf{R}$). We can pull back the foliations $\tilde{\Phi}_x$ and $\tilde{\Phi}_y$, via the map f, to foliations $\tilde{\Phi}_x^+$ and $\tilde{\Phi}_y^-$ on Ω_+ and Ω_-, respectively. The semiconjugacy $F : \hat{\mathbf{C}} \to \hat{\mathbf{C}}$ will collapse each leaf of $\tilde{\Phi}_x^+$ and each leaf of $\tilde{\Phi}_y^-$ to a point; Cannon and Thurston have shown that these are the only identifications that occur. Thus, we have a complete topological description of the semiconjugacy $F : \hat{\mathbf{C}} \to \hat{\mathbf{C}}$.

Cannon and Thurston also studied a second class of Kleinian groups: closed surface groups Γ for which the quotient manifold, again homeomorphic to $S \times \mathbf{R}$, has one geometrically finite end and one end that is quasi-isometric to an end of the manifold we have just discussed, the infinite cyclic cover of a pseudo-Anosov mapping torus. For such a group, the domain of discontinuity consists of a single topological disk, and the limit set is a dendrite. Again, the existence of a semiconjugacy from a Fuchsian group action on $\hat{\mathbf{C}}$ is proved using a model manifold, homeomorphic to $S \times \mathbf{R}$; the metric on $S \times \mathbf{R}$ models the geometry of a geometrically finite end on, say, $S \times (-\infty, 0]$ and is similar to the metric ds^2 on $S \times [0, \infty)$. If Ω_+ and Ω_- are the components of the domain of discontinuity of the Fuchsian group and $\tilde{\Phi}_x^+$ is the lift of the stable foliation to Ω_+ then the semiconjugacy from $\hat{\mathbf{C}}$ to $\hat{\mathbf{C}}$ maps Ω_- homeomorphically onto the domain of discontinuity of Γ and maps $\overline{\Omega}_+$ onto the limit set of Γ, collapsing each leaf of $\tilde{\Phi}_x^+$ to a single point (see Figure 3). In particular, the existence of a semiconjugacy implies that the limit set of such a group Γ is locally connected, since it is the continuous image of a circle.

2.4 Ending laminations and foliations, and semiconjugacies for surface groups with lower bounds on injectivity radius

In constructing semiconjugacies for more general Kleinian groups, the correct generalization both of the simple closed curves that control the collapsing in the accidental parabolics case and of the stable and unstable foliations in the Cannon-Thurston examples is the *ending lamination* or *ending foliation* (these are equivalent constructions). Let S be a closed surface and let σ be a hyperbolic metric on S. A *geodesic lamination* on (S, σ) is a closed disjoint union of simple geodesics in S, called the *leaves* of the lamination. We will also think of laminations as topological objects: a lamination is an equivalence class whose members are closed subsets of S that are ambient isotopic to a geodesic lamination; the equivalence is under isotopy of S. Let $\mathcal{GL}(S)$ denote the space of geodesic laminations on (S, σ). Let $\mathcal{ML}(S)$, the measured lamination space, be the space of

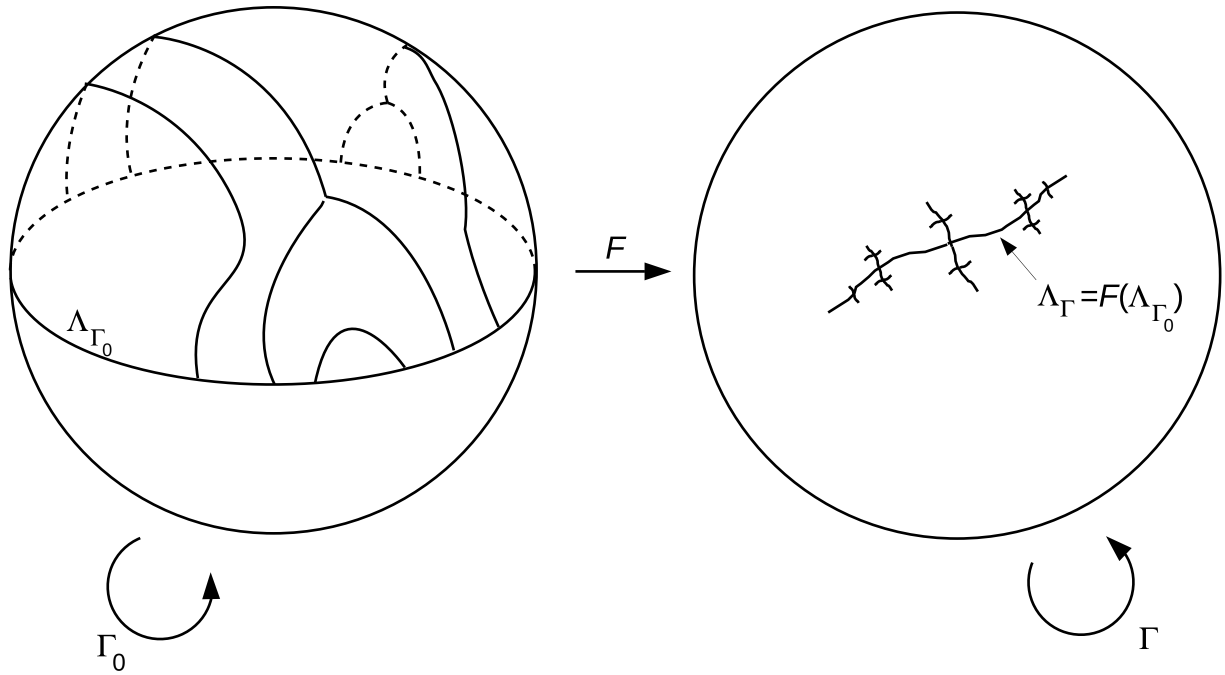

Figure 3: *When there is only one degenerate end, the semiconjugacy* $F : \hat{\mathbf{C}} \to \hat{\mathbf{C}}$ *collapses each leaf of the foliation* $\tilde{\Phi}_x^+$ *to a point in the limit set of* Γ. F *sends the limit set* Λ_{Γ_0} *of* Γ_0 *onto the limit set* Λ_Γ *of* Γ, *and maps the bottom hemisphere homeomorphically onto the domain of discontinuity of* Γ.

geodesic laminations on (S, σ) equipped with transverse measures. $\mathcal{ML}(S)$ has a natural topology (see [9, 10, 30]), such that if σ and τ are two different metrics on S then there is a canonical homeomorphism between the measured lamination spaces of (S, σ) and (S, τ). If μ is a measured lamination then the *support* of μ is the lamination on which the transverse measure of μ is supported.

Denote by $\mathcal{MF}(S)$ the space of equivalence classes of topological foliations on S with three- or more-pronged saddle singularities, that support transverse measures; two foliations are equivalent if they are isotopic up to collapsing segments joining singularities (Whitehead moves). There is a natural homeomorphism between $\mathcal{MF}(S)$ and $\mathcal{ML}(S)$, obtained by straightening leaves of a foliation to their geodesic representatives.

Let $N = \mathbf{H}^3/\Gamma$ be a hyperbolic 3-manifold that satisfies $C1$ and that has no cusps. Thurston and Bonahon have shown that a manifold that satisfies $C1$ is topologically tame, that is, it is the interior of a compact manifold (see [4, 30]); so for each end e of N a boundary surface S_e may be associated to e. We will see that if e is geometrically finite, S_e will have a canonical Riemann surface structure, induced by the conformal structure of the domain of discontinuity of Γ. If e is degenerate then S_e will not have a canonical Riemann surface structure; instead, Thurston has defined an ending invariant for e that is a lamination on S_e.

Definition 2.1 *An end e of a manifold N that satisfies C1 and has no cusps is called* **simply degenerate** *if there is a sequence $\{\gamma_i\}$ of closed curves on S_e whose geodesic representatives γ_i^* in N are eventually contained in any neighborhood of e. We say that the curves γ_i^** **exit the end**.

Bonahon ([4]) has shown that if N satisfies $C1$ then all of its ends are either geometrically finite or simply degenerate. Thurston has shown the following (see [4, 5] for a proof):

Proposition 2.2 *Let e be a simply degenerate end of a hyperbolic 3-manifold without cusps, that satisfies C1. There is a unique $\lambda_e \in \mathcal{GL}(S_e)$, called the* **ending lamination** *of e, such that any sequence (γ_i) of simple closed curves in S_e whose geodesic representatives (γ_i^*) exit the end e will accumulate in $\mathcal{ML}(S)$ onto a set of measured laminations with support λ_e.*

The ending lamination λ_e satisfies the following properties (see e.g. [24]), which will be used in the proof of Theorem B:

(1) λ_e is maximal (i.e. there is no lamination $\lambda \in \mathcal{GL}(S_e)$ such that $\lambda_e \subset \lambda$ and $\lambda_e \neq \lambda$).
(2) Every half-leaf of λ_e is dense in λ_e.
(3) λ_e contains no closed geodesics.

For closed surface groups that satisfy $C2$, Minsky has shown (see [24]) that the ending laminations (or foliations) of the group control the geometry of the quotient manifold, as did the stable and unstable foliations in the Cannon-Thurston examples (in fact, for these examples the stable and unstable foliations *are* the ending foliations). Let G be a Kleinian group that satisfies $C2$ and that is isomorphic to the fundamental group of a closed surface S, and let G_0 be an isomorphic Fuchsian or quasi-Fuchsian group (note that a surface group automatically satisfies $C1$). The quotient manifolds $\mathbf{H^3}/G$ and $\mathbf{H^3}/G_0$ are both homeomorphic to $S \times \mathbf{R}$ (see [4, 30]). A surface group without parabolics must fall into one of the following categories (this classification is due to Bers, Marden, Maskit, Bonahon and Thurston (see [3, 20, 4, 30])):

(1) The group is geometrically finite, its limit set is a Jordan curve (so that its domain of discontinuity is the union of two topological disks), and the convex core of the quotient manifold is homeomorphic to $S \times [-1, 1]$ (unless the group is Fuchsian, in which case the convex core is homeomorphic to $S \times \{0\}$). Such groups where the limit set is not a round circle are called *quasi-Fuchsian*.

(2) The quotient manifold has one degenerate end, the domain of discontinuity consists of a single topological disk, and the convex core of the quotient manifold is homeomorphic to $S \times [0, \infty)$.

(3) The quotient manifold has two degenerate ends, the limit set of the group is all of $\hat{\mathbf{C}}$, and the convex core of the quotient manifold is the entire manifold.

If G_0 and G are as above, let e_+ and e_- denote the two ends of $S \times \mathbf{R}$, and S_+ and S_- denote the surfaces $S \times \{\infty\}$ and $S \times \{-\infty\}$. Let Ω_+ and Ω_- denote the two components of the domain of discontinuity of G_0, which cover S_+ and S_-, respectively. If G has one degenerate end (say e_+), let λ_+ denote the ending lamination of G. If G has two degenerate ends, let λ_+ and λ_- denote the ending laminations. We can lift $\lambda_\pm$ to laminations $\tilde{\lambda}_\pm$ on $\Omega_\pm$. Using techniques similar to those of Cannon and Thurston, Minsky has shown that for surface groups satisfying $C2$, the ending laminations determine the geometry of the degenerate ends and the topology of the action on $\hat{\mathbf{C}}$ (see [24]). The idea is as follows: represent λ_+ and λ_- as measured foliations (Φ_x, dx) and (Φ_y, dy), and again construct a model manifold $N_{mod} \cong S \times \mathbf{R}$, with the following metric:

$$ds^2 = e^{2t}dx^2 + e^{-2t}dy^2 + dt^2.$$

If (γ_i^*) is a sequence of simple closed geodesics in $N = \mathbf{H^3}/G$ that exit the end e_+, say, we may construct a sequence of pleated surfaces σ_i in N, each containing the curve γ_i^*, that also exit the end e_+. In [23], Minsky uses the lower bound on injectivity radius in N and the fact that the curves (γ_i^*) converge to λ_+ to show that the surfaces σ_i have similar geometry to the level surfaces $S \times \{i\}$ in N_{mod}. These techniques can be used to show that N_{mod} and N are quasi-isometric; then techniques similar to those used in the Cannon-Thurston examples give the following result, showing the existence and describing the collapsing of semiconjugacies on $\hat{\mathbf{C}}$ (formulated here in terms of laminations, not foliations); we will use this result in the proofs of Theorems A and B.

Theorem 2.3 *(Minsky, [24]) Let $f : \mathbf{H^3}/G_0 \to \mathbf{H^3}/G$ be a homeomorphism, where G_0 is quasi-Fuchsian and G satisfies $C2$. Then there is a homotopic map $g : \mathbf{H^3}/G_0 \to \mathbf{H^3}/G$ such that the lift $\tilde{g} : \mathbf{H^3} \to \mathbf{H^3}$ to the universal covers is quasi-Lipschitz and extends continuously to a map $\tilde{g} : \hat{\mathbf{C}} \to \hat{\mathbf{C}}$ that collapses leaves and complementary components of $\tilde{\lambda}_\pm$ to points. That is, if G has one degenerate end then for x and y in $\hat{\mathbf{C}}$, $\tilde{g}(x) = \tilde{g}(y)$ if and only if x and y lie on the closure of the same leaf or complementary component of $\tilde{\lambda}_+$. If both ends of G are degenerate then $\tilde{g}(x) = \tilde{g}(y)$ also if x and y lie on the closure of the same leaf or complementary component of $\tilde{\lambda}_-$. If G is quasi-Fuchsian then $\tilde{g} : \hat{\mathbf{C}} \to \hat{\mathbf{C}}$ is a homeomorphism.*

By work of Lott, the map g may be chosen to be a homeomorphism such that $\tilde{g}$ is Lipschitz on $\mathbf{H^3}$ (see [19]).

2.5 Construction of the homeomorphism g of Theorem A

We return to the setting of Theorem A: Let Γ_0 be a geometrically finite Kleinian group that satisfies $C1$ and $C2$, and let $N_0 = \mathbf{H^3}/\Gamma_0$. Let $f : N_0 \to N$ be a homeomorphism, where $N = \mathbf{H^3}/\Gamma$ is a hyperbolic manifold that satisfies $C2$. The homeomorphism f induces an isomorphism $f_* : \Gamma_0 \to \Gamma$. We wish to replace f by a homotopic map g whose lift to the universal covers extends continuously to the boundary and identifies boundary points according to the ending laminations of N; to do this, we will modify f on each end of N_0 to obtain a map $g : N_0 \to N$ whose restriction to each end of N_0 looks like the map constructed by Minsky in Theorem 2.3.

If N_0 has finite volume then Mostow's rigidity theorem says that f is homotopic to an isometry, whose lift necessarily extends to a Möbius transformation on $\hat{\mathbf{C}}$ (see [27, 28]); so in this case there is nothing to prove. Otherwise, N_0 has at least one end; let e be an end of N_0, and let $K_e \subset N_0$ be the corresponding component of the boundary of the convex core. Since Γ_0 is freely indecomposable K_e is incompressible in N_0 (see [4]). So K_e lifts to a union of disks in $\mathbf{H^3}$, each of which is a boundary component of $CH(\Gamma_0)$. If $\tilde{K_\alpha}$ is one of these disks then $\tilde{K_\alpha}$ bounds a topological half-space H_α on its nonconvex side that meets $\hat{\mathbf{C}}$ in the closure of a component Ω_α of Ω (see Figure 4). Let $\Gamma_\alpha \subset \Gamma_0$ be the stabilizer of Ω_α; Γ_α is called the *component subgroup* of Γ_0 with respect to Ω_α. Orthogonal projection from $int(H_\alpha) \cup \Omega_\alpha$ to $\tilde{K_\alpha}$ induces a foliation of $int(H_\alpha) \cup \Omega_\alpha$ which gives a homeomorphism of $int(H_\alpha) \cup \Omega_\alpha$ with $\tilde{K_\alpha} \times (0, \infty]$; this foliation is invariant under the action of Γ_α so it descends to the corresponding component of $\overline{N}_0 \setminus CC(N_0)$, which is then homeomorphic to $K_e \times (0, \infty]$. Write $S_e = K_e \times \{\infty\}$, so that Ω_α covers S_e.

Given a component Ω_α of the domain of discontinuity Ω of Γ_0, the stabilizer $\Gamma_\alpha \subset \Gamma_0$ of Ω_α is a surface group since Ω_α is a topological disk. The map f lifts to the intermediate cover $\mathbf{H^3}/\Gamma_\alpha$ to give a homeomorphism $\hat{f_\alpha} : \mathbf{H^3}/\Gamma_\alpha \to \mathbf{H^3}/f_*(\Gamma_\alpha)$. We may apply Minsky's theorem to this homeomorphism, for the group Γ_α is geometrically finite; this is a consequence of a theorem of Thurston, that any finitely generated subgroup of an infinite-covolume geometrically finite group is geometrically finite (see e.g. [26]). Thus, by Minsky's theorem the map $\hat{f_\alpha}$ is homotopic to a map $\hat{g}_\alpha$ whose lift $\tilde{g}_\alpha$ to the universal cover extends continuously to the boundary.

For an end e, let n_e denote the corresponding component of the complement of $CC(N_0)$, so that the boundary of n_e in N_0 is K_e. Lifting K_e and n_e to the intermediate cover $\mathbf{H^3}/\Gamma_\alpha = S_e \times (-\infty, \infty)$ gives sets $\hat{K_e}$ and $\hat{n}_e$, such that $\hat{K_e} = \tilde{K_\alpha}/\Gamma_\alpha$ and $\hat{n}_e = H_\alpha/\Gamma_\alpha$. Since for every $\gamma \in \Gamma_0 \setminus \Gamma_\alpha$, $\gamma(H_\alpha) \cap H_\alpha = \emptyset$, we have $H_\alpha/\Gamma_\alpha = H_\alpha/\Gamma$; hence $\hat{K_e}$ and $\hat{n}_e$ project homeomorphically to K_e and n_e, respectively. If $\hat{g}_\alpha : S_e \times \mathbf{R} \to S_e \times \mathbf{R}$ is the map given by Minsky's theorem then we may adjust $\hat{g}_\alpha$ by a homotopy on a compact set so that $\hat{g}_\alpha$ is still a homeomorphism, and so that $\hat{g}_\alpha(\hat{n}_e) = \hat{f_\alpha}(\hat{n}_e)$ and $\hat{g}_\alpha$ and $\hat{f_\alpha}$ agree on $\hat{K_e}$; this does not change the behavior at infinity of the lift $\tilde{g}_\alpha : \mathbf{H^3} \to \mathbf{H^3}$ of $\hat{g}_\alpha$.

With this adjustment $\hat{g}_\alpha|_{\hat{n}_e}$ descends to a homeomorphism g_e from n_e to a subset of N, and g_e agrees with f on K_e. Note that if Γ_α and Γ_β are two component subgroups of Γ_0 that correspond to the same end e then the maps they induce on n_e are the same (if we have adjusted $\hat{g}_\alpha$ and $\hat{g}_\beta$ in the same way), since their lifts $\tilde{g}_\alpha$ and $\tilde{g}_\beta$ are related by a Möbius transformation that conjugates Γ_α to Γ_β; thus the map g_e did not depend on our choice of component subgroup. We will piece together the maps g_e by "filling in" on the convex core, to obtain the map $g : N_0 \to N$ of Theorem A.

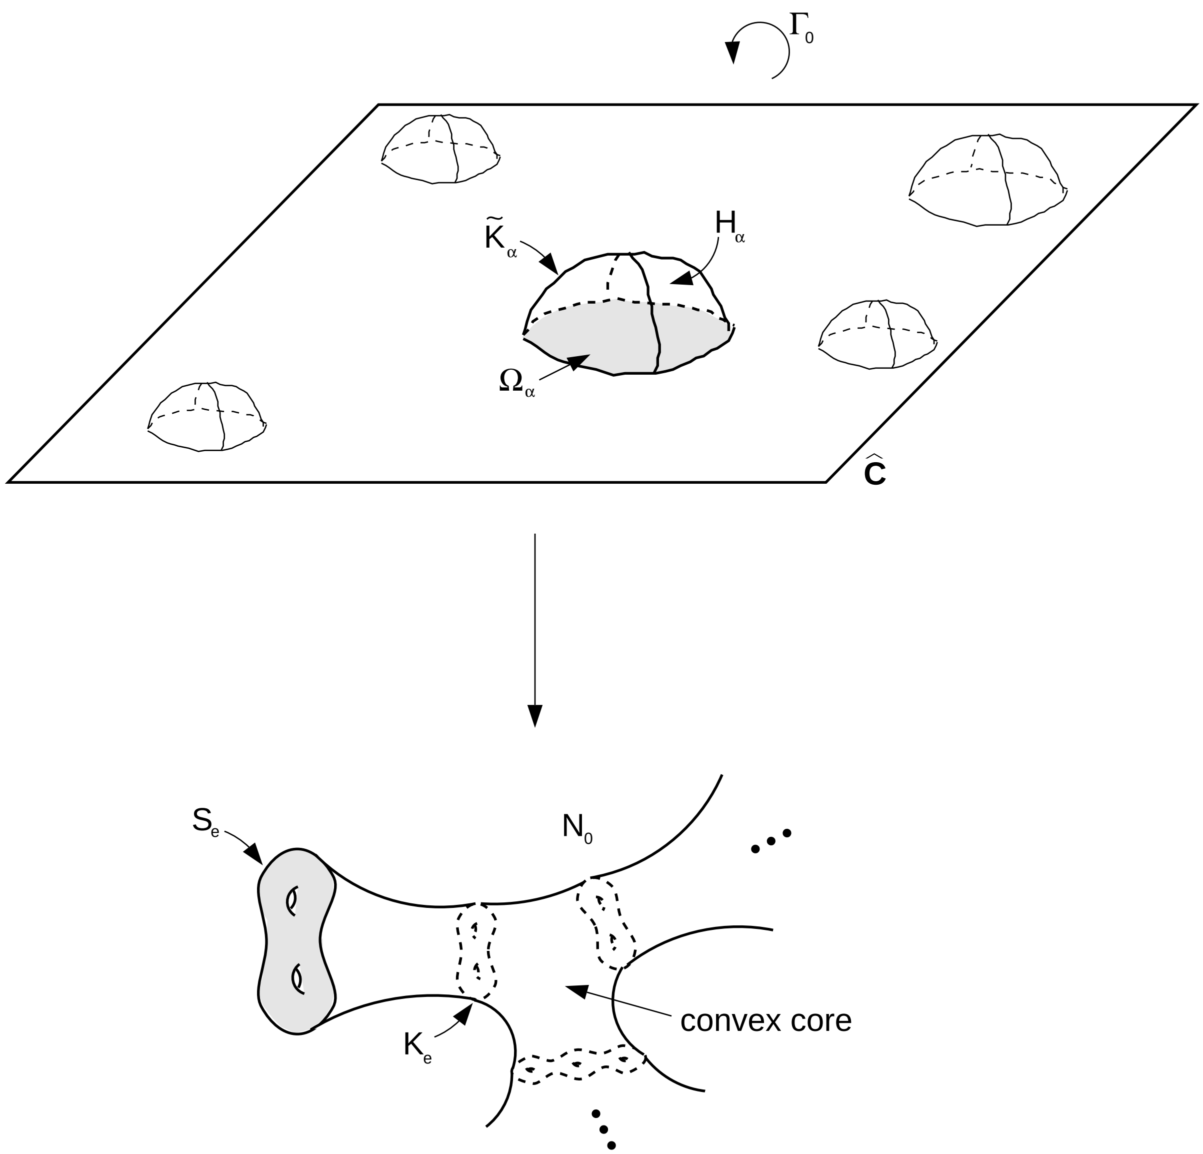

Figure 4: *The convex hull of* Γ_0 *is the closure of the complement in* $\mathbf{H^3}$ *of the topological half-spaces* H_α*. The quotient of* H_α *by* Γ_α *(equivalently, by* Γ_0*) is the complementary component of the convex core of* N_0 *whose boundary surface is* K_e*.*

Define a map $g : N_0 \to N$ in the following way:
(1) $g|_{CC(N_0)} = f|_{CC(N_0)}$; and
(2) If e is an end of N_0, $g|_{n_e} = g_e|_{n_e}$.

The maps given by Minsky's theorem are Lipschitz, so by standard arguments we can adjust g by a homotopy supported on a compact set so that g is a homeomorphism that is Lipschitz on each piece n_e and is a bi-Lipschitz map from $CC(N_0)$ to $g(CC(N_0))$, with respect to their intrinsic path metrics. Note that this is not the same as saying that the restriction of g to $CC(N_0)$ is a bi-Lipschitz map; since $g(CC(N_0))$ is generally not convex, its intrinsic path metric is strictly larger than the hyperbolic metric that $g(CC(N_0))$ inherits from N. If we adjust g in this way then the lift of g to the universal covers is a Lipschitz homeomorphism $\tilde{g} : \mathbf{H^3} \to \mathbf{H^3}$ such that $\tilde{g}|_{CH(\Gamma_0)}$ is a bi-Lipschitz homeomorphism with respect to the intrinsic path metrics of $CH(\Gamma_0)$ and $\tilde{g}(CH(\Gamma_0))$; another way of saying this last statement is that the induced map on the electric spaces of $(\mathbf{H^3}, \{H_\alpha\})$ and $(\mathbf{H^3}, \{\tilde{g}(H_\alpha)\})$ is a bi-Lipschitz homeomorphism, where the sets H_α are the closures in $\mathbf{H^3}$ of the complementary components of the convex hull of Γ_0 (recall that the electric spaces of $(\mathbf{H^3}, \{H_\alpha\})$ and $(\mathbf{H^3}, \{\tilde{g}(H_\alpha)\})$ are the spaces we get from $\mathbf{H^3}$ by collapsing the sets H_α and $\tilde{g}(H_\alpha)$, respectively, to points; the metric on the electric spaces is the inherited path metric from $\mathbf{H^3}$). Since we have only adjusted g by a homotopy on a compact set we have not changed the behavior at infinity of its lift $\tilde{g} : \mathbf{H^3} \to \mathbf{H^3}$. We will show that this lift extends continuously to the boundary.

2.6 Reduction of Theorem A to Theorem C

Let Γ_α be a component subgroup of Γ_0. Write $\hat{N}_0^\alpha = \mathbf{H^3}/\Gamma_\alpha$ and $\hat{N}^\alpha = \mathbf{H^3}/f_*(\Gamma_\alpha)$, so that $\hat{N}^\alpha$ is a covering space of N. We will begin by showing that when the covering $\hat{N}^\alpha \to N$ is finite-to-one, Theorem A is a simple extension of Minsky's result, Theorem 2.3. When the covering $\hat{N}^\alpha \to N$ is infinite-to-one, we will show that Theorem A is a result of Theorem C. We need the following:

Lemma 2.4 *Let Γ_0 be a geometrically finite Kleinian group that satisfies $C1$ and $C2$, and such that the domain of discontinuity does not have 0 or 2 components. Then if Ω_α is a component of the domain of discontinuity, Ω_α has infinitely many translates in $\hat{\mathbf{C}}$.*

Proof:

The domain of discontinuity Ω cannot have exactly one component, since then Γ_0 would be a surface group and by the classification of surface groups that satisfy $C2$ it would have a degenerate end.

Suppose Ω has more than two components. Let Ω_1, Ω_2 and Ω_3 be components of Ω and let Γ_1, Γ_2 and Γ_3 be their component subgroups. We will show that Ω_1 has infinitely many translates. The limit sets of Γ_1, Γ_2 and Γ_3 are all Jordan curves so no two can be the same. Hence, there is a (loxodromic) element of Γ_2 whose attracting fixed point does not lie in the limit set of Γ_1. The sets $\gamma^n(\Omega_1)$ for $n \in \mathbf{Z}$ are all distinct, and each is a component of Ω. $\square$

Now, to reduce the proof of Theorem A to that of Theorem C, we will consider two cases:

Case 1: the covering $\hat{N}^\alpha \to N$ is finite-to-one. As previously stated, for this case we will show that the result of Theorem A is essentially already known; that is, it is an easy extension of Minsky's result. If the covering is finite-to-one then the corresponding cover $\hat{N}_0^\alpha \to N_0$ is also finite-to-one, and so Γ_α has finite index in Γ_0. Then the component Ω_α of the domain of discontinuity of Γ_0 stabilized by Γ_α must have only finitely many translates in $\hat{\mathbf{C}}$; hence by Lemma 2.4, the domain of discontinuity of Γ_0 has either 0 or 2 components. If Ω is empty then N_0 is compact; in this setting we have seen that Theorem A is a consequence of Mostow's rigidity theorem, which tells us that the map $\tilde{g}$ extends to a Möbius transformation of $\hat{\mathbf{C}}$ (and Theorem B just turns into the fact that a Möbius transformation from $\hat{\mathbf{C}}$ to $\hat{\mathbf{C}}$ is injective). If Ω has two components then there are two possibilities for Γ_0 and N_0:

(1) If Γ_0 preserves the two components of Ω then Γ_0 is a quasi-Fuchsian group, and $N_0 \cong S \times \mathbf{R}$; this is the case of Theorems A and B that has already been proved by Minsky.

(2) Otherwise, Γ_0 contains an index two subgroup G_0 that is quasi-Fuchsian (namely the subgroup of Γ_0 that preserves the components of Ω). Maskit calls Γ_0 an *extended quasi-Fuchsian group* (see [21]). In this case, N_0 has one end and is doubly covered by the manifold $\mathbf{H^3}/\ G_0$. The map $g : N_0 \to N$ we have constructed lifts to a homeomorphism $g_0 : \mathbf{H^3}/G_0 \to \mathbf{H^3}/f_*(G_0)$, and it is easily checked that the map g_0 agrees with Minsky's map up to a homotopy on a compact set; thus its lift to $\mathbf{H^3}$ extends to the boundary in the same way as Minsky's map.

Case 2: the covering $\hat{N}^\alpha \to N$ is infinite-to-one. We will show that in this case, Theorem A is a result of Theorem C; that is, that the map $\tilde{g} : \mathbf{H^3} \to \mathbf{H^3}$ constructed in Section 2.5 satisfies the hypotheses of Theorem C. So, let $\tilde{g}_\alpha : \mathbf{H^3} \to \mathbf{H^3}$ be the map given by Minsky's result (Theorem 2.3), for the groups Γ_α and $f_*(\Gamma_\alpha)$. Let e'_+ and e'_- denote the two ends of $\hat{N}^\alpha = \mathbf{H^3}/f_*(\Gamma_\alpha)$; according to Minsky's theorem, the identifications $\tilde{g}_\alpha$ makes on $\hat{\mathbf{C}}$ are determined by the geometry of e'_+ and e'_-. Let e' denote the end of N corresponding to e in N_0. One of the two ends of $\hat{N}^\alpha$, say e'_+, has a neighborhood that covers a neighborhood of e' homeomorphically and isometrically (see Figure 5); we must understand the geometry of the other end, e'_-. We have assumed that a neighborhood of e'_- covers N in an infinite-to-one manner; a theorem of Thurston (see [6] for a proof) tells us that in this case, e'_- is geometrically finite. Hence $\hat{N}^\alpha$ has at most one degenerate end. If the end e'_+ is degenerate then its ending lamination agrees with the ending lamination λ_α of e'. Minsky's theorem then gives us the following about the identifications of $\tilde{g}_\alpha$ on $\hat{\mathbf{C}}$:

Given a component Ω_α of the domain of discontinuity of Γ_0, if the end e' of N corresponding to Ω_α is geometrically finite then the map $\tilde{g}_\alpha : \hat{\mathbf{C}} \to \hat{\mathbf{C}}$ is a homeomorphism. If e' is degenerate then for x and y in $\hat{\mathbf{C}}$, $\tilde{g}_\alpha(x) = \tilde{g}_\alpha(y)$ if and only if x and y lie on the closure of the same leaf or complementary component of $\tilde{\lambda}_\alpha$.

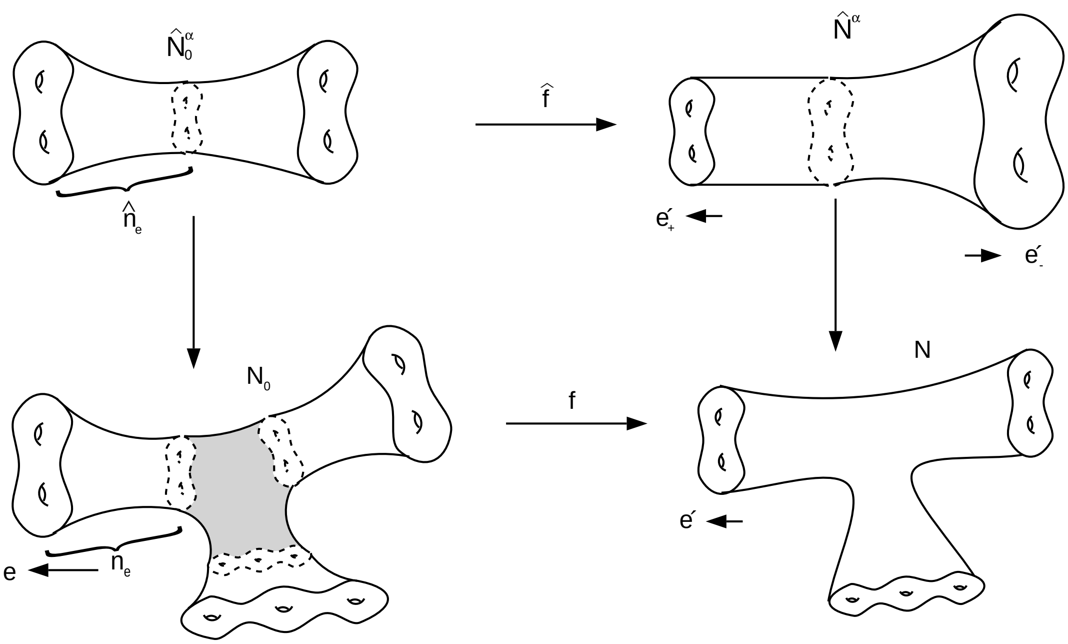

Figure 5: *The intermediate covers $\hat{N}_0^\alpha$ and $\hat{N}^\alpha$.*

We have now constructed a homeomorphism $\tilde{g} : \mathbf{H^3} \to \mathbf{H^3}$ that satisfies the following:

(1) $\tilde{g}$ is Lipschitz.

(2) On the electric level (with respect to $(\mathbf{H^3}, \{H_\alpha\})$ and $(\mathbf{H^3}, \{\tilde{g}(H_\alpha)\})$), $\tilde{g}$ is a quasi-isometry (see Section 3 for definitions); in fact, it is bi-Lipschitz.

(3) For each set H_α, $\tilde{g}|_{H_\alpha}$ extends continuously to the boundary $\overline{\Omega}_\alpha$, and either is a homeomorphism on Ω_α or identifies points on $\overline{\Omega}_\alpha$ according to the ending lamination $\tilde{\lambda}_\alpha$.

Thus, if we can show that the systems $(\mathbf{H^3}, \{H_\alpha\})$ and $(\mathbf{H^3}, \{\tilde{g}(H_\alpha)\})$ satisfy conditions $(*_1)$ – $(*_3)$ then we will have shown that the map $\tilde{g} : \mathbf{H^3} \to \mathbf{H^3}$ satisfies the hypotheses of Theorem C. We have

Theorem 2.5 *Let Γ_0 and Γ be as in Theorem A, and suppose that the domain of discontinuity of Γ_0 consists of more than two components. Then the map $\tilde{g} : \mathbf{H^3} \to \mathbf{H^3}$ defined above satisfies the conditions of Theorem C.*

Proof:

We must show that $(*_1) - (*_3)$ hold for the systems $(\mathbf{H^3}, \{H_\alpha\})$ and $(\mathbf{H^3}, \{\tilde{g}(H_\alpha)\})$.

For $(*_1)$, simply note that $\mathbf{H^3} \setminus \cup_\alpha H_\alpha$ is the interior of the convex hull of Λ_{Γ_0}, the limit set of Γ_0, so it is open and path-connected. Since $\tilde{g}$ is a homeomorphism $\mathbf{H^3} \setminus \cup_\alpha \tilde{g}(H_\alpha)$ is also open and path-connected.

For $(*_2)$, we must show that there exists some q such that the sets H_α and $\tilde{g}(H_\alpha)$ are all q-quasiconvex (see Section 3 for a definition of quasiconvexity). If we can show that the sets are all quasiconvex then the uniformity of the quasiconvexity constant follows immediately: two such sets that correspond to the same end of N_0 or N have the same quasiconvexity constant since one is the image of the other by a Möbius transformation, and there are only finitely many ends. Fix α; we will show that $\tilde{g}(H_\alpha)$ is quasiconvex (the proof we give also works for H_α).

Let $\hat{N}^\alpha$ be the intermediate cover of N corresponding to Γ_α. We have seen that $\hat{N}^\alpha$ is homeomorphic to $S_e \times \mathbf{R}$. As in the previous subsection, let e'_+ and e'_- denote the two ends of $\hat{N}^\alpha$, and e' the end of N corresponding to Γ_α. We have seen that one end, e'_+, has a neighborhood that projects homeomorphically to a neighborhood of e', and that the other end, e'_-, is geometrically finite. Since e'_- is geometrically finite, it has a neighborhood that is disjoint from the convex core C_α of $\hat{N}^\alpha$ (see Figure 6). The lift $\tilde{C_\alpha}$ of C_α to $\mathbf{H^3}$ is a convex set. The projections of $\tilde{C_\alpha}$ and $\tilde{g}(H_\alpha)$ to $\hat{N}^\alpha$ are both neighborhoods of e'_+, and their boundaries are homotopic. If we choose a homotopy $H : S_e \times [0,1] \to \hat{N}^\alpha$ between the boundaries of the projections of $\tilde{C_\alpha}$ and $\tilde{g}(H_\alpha)$, the lengths of the paths $H_x : [0,1] \to \hat{N}^\alpha$ defined by $H_x(t) = H(x,t)$ are uniformly bounded, by compactness of S_e; so lifting to $\mathbf{H^3}$ we have a homotopy $\tilde{H} : \tilde{S_e} \times [0,1] \to \mathbf{H^3}$ between the boundaries of $\tilde{C_\alpha}$ and $\tilde{g}(H_\alpha)$, such that there is a uniform bound of the length of each path $\tilde{H}_{\tilde{x}} : [0,1] \to \mathbf{H^3}$. Hence the boundaries of $\tilde{g}(H_\alpha)$ and $\tilde{C_\alpha}$ are within a bounded Hausdorff distance of each other. So since $\tilde{C_\alpha}$ is convex, $\tilde{g}(H_\alpha)$ must be quasiconvex.

For $(*_3)$ we must show that there is a constant c such that the sets H_α (resp. $\tilde{g}(H_\alpha)$) are all separated from each other by at least c. Let H_α and H_β be distinct half-spaces; we will show that there is a lower bound (independent of α and β) on the length of any path from H_α to H_β. So let p be a path from x to y, where $x \in H_\alpha$ and $y \in H_\beta$; without loss of generality we may assume that $x \in \partial H_\alpha$ and $y \in \partial H_\beta$. Now the path p descends (via the projection $\mathbf{H^3} \to \mathbf{H^3}/\Gamma_0$) to a path $proj(p)$ in $N_0 = \mathbf{H^3}/\Gamma_0$ whose endpoints a and b lie on the boundary of the convex core of N_0.

Suppose first that a and b lie on different boundary components of the convex core. The convex core of N_0 has finitely many boundary components, and each is an embedded surface; so there is a lower bound on the distance in N_0 between any two of these components. Hence there is a lower bound on the length of $proj(p)$, which is the same as the length of p.

Otherwise a and b lie on the same boundary component K of the convex core; then since H_α and H_β are distinct, $proj(p)$ is not homotopic relative endpoints to a path in K. K is an embedded surface in N_0, so if $u \in K$ then for every $\epsilon > 0$ there is a $\delta(u,\epsilon) > 0$ such that if $v \in K$ and

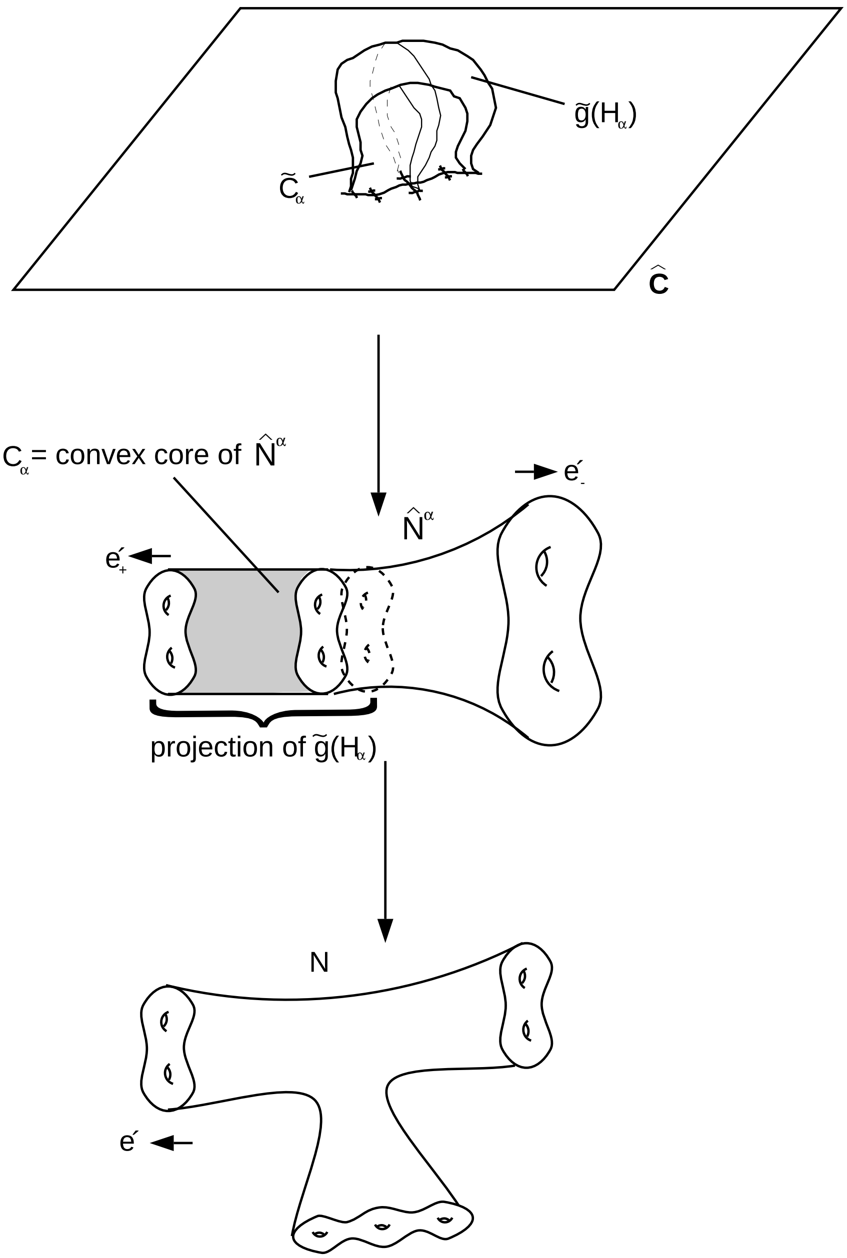

Figure 6: *The boundaries of $\tilde{C}_\alpha$ and $\tilde{g}(H_\alpha)$ are within finite Hausdorff distance of each other, so since $\tilde{C}_\alpha$ is convex, $\tilde{g}(H_\alpha)$ must be quasiconvex.*

$d(u,v) \leq \delta(u,\epsilon)$ then u and v are within ϵ of each other in the intrinsic metric on K; since K_e is compact, one $\delta(\epsilon)$ may be chosen for all $u \in K$. Let ϵ_0 denote the injectivity radius of the convex core of N_0. Suppose that $l(p) \leq \delta(\epsilon_0)$. Then there is a path r in K_e from a to b whose length is at most ϵ_0. But the path $r \cdot (proj(p))$ is a homotopically nontrivial loop, so its length is at least $2\epsilon_0$. So the length of p is at least ϵ_0.

So we have shown that $(\mathbf{H^3}, \{H_\alpha\})$ satisfies $(*_3)$. The only thing we used in the proof was the topological structure of the convex core of N_0, so since $\tilde{g}$ is a homeomorphism the same proof works for $(\mathbf{H^3}, \{\tilde{g}(H_\alpha)\})$. $\square$

To summarize, we have shown the following:

Theorem 2.6 *Let Γ_0 and Γ be as in Theorem A. If the domain of discontinuity of Γ_0 has 0 or two components then Theorem A holds for Γ_0 and Γ. In all other cases, Theorem A is a special case of Theorem C.*

3 Gromov-hyperbolic spaces

In this section we will present an overview of some of the basic theory of Gromov-hyperbolic spaces. References for the material in this section are [17], [18], [11] and [8].

Let (Δ, d) be a metric space. If Δ is equipped with a basepoint 0, define the *Gromov product* $(x|y)$ of the points x and y in Δ to be

$$(x|y) = (x|y)_0 = \frac{1}{2}(d(x,0) + d(y,0) - d(x,y)).$$

Definition 3.1 *Let $\delta \geq 0$ be a real number. The metric space Δ is δ-**hyperbolic** if*

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

for every $x, y, z \in \Delta$ and for every choice of basepoint.

We say that Δ is hyperbolic in the sense of Gromov if Δ is δ-hyperbolic for some δ.

For example:

1. Every bounded metric space is hyperbolic.

2. Every real tree is 0-hyperbolic (a real tree is a space Δ such that for any two points x and y in Δ there is a unique topological segment joining x and y, and the length of that segment is $d(x,y)$.)

3. Hyperbolic space $\mathbf{H^n}$ is δ-hyperbolic with $\delta = \log 3$.

A metric space Δ is *geodesic* if any two points in Δ can be joined by a geodesic segment (not necessarily unique). If x and y are in Δ we write $[x, y]$, ambiguously, to denote some geodesic from x to y.

Heuristically, a δ-hyperbolic space is "tree-like"; more precisely, if we define an ϵ-*narrow* geodesic polygon to be one such that every point on each side of the polygon is at distance $\leq \epsilon$ from a point in the union of the other sides, then we have

Proposition 3.2 *In a geodesic δ-hyperbolic metric space, every n-sided polygon ($n \geq 3$) is $4(n-2)\delta$-narrow.*

If $S \subset \Delta$ is a set and $x \in \Delta$, then the *projection* of x to S is the subset $proj_S x$ of S consisting of all points $y \in S$ such that $d(x, y)$ is minimal over all points in S. If $T \subset \Delta$ then the projection of T to S is the union of all sets $proj_S x$ where $x \in T$. The following is an easy consequence of the above proposition:

Lemma 3.3 *For every positive integer N, there exists a constant $K(N, \delta) > 0$ such that the following holds:*
Let l be a geodesic segment. Let g be a path in Δ that is the union of at most N geodesic segments $[a_0, a_1], [a_1, a_2], ..., [a_{N-1}, a_N]$, and assume that g lies outside of the K-neighborhood of l. Then the diameter of the projection of g to l is bounded above by 32δ.

In a geodesic hyperbolic space, the Gromov product of two points x and y is roughly the distance from 0 to $[x, y]$; we have

Proposition 3.4 *Let Δ be a geodesic, δ-hyperbolic space and let $x, y \in \Delta$. Then*

$$d(0, [x, y]) - 4\delta \leq (x|y) \leq d(0, [x, y])$$

for every geodesic segment $[x, y]$.

3.1 The boundary of a hyperbolic space

If Δ is a hyperbolic space, Δ can be equipped with a boundary in a natural way. We say that a sequence (x_n) of points in Δ *converges at infinity* if we have $\lim_{m,n\to\infty}(x_m|x_n) = \infty$; note that this definition is independent of the choice of basepoint, by Proposition 3.4. Given two sequences (x_m) and (y_n) that converge at infinity, say that (x_m) and (y_n) are *equivalent* if $\lim_{m,n\to\infty}(x_m|y_n) = \infty$. Since Δ is hyperbolic, it is easily checked that this is an equivalence relation. Define the *boundary at infinity* $\partial_\infty\Delta$ of Δ to be the set of equivalence classes of sequences that converge at infinity. If

$\xi \in \partial_\infty \Delta$ then we say that a sequence of points in Δ *converges to* ξ if the sequence belongs to the equivalence class ξ. Write $\overline{\Delta} = \Delta \cup \partial_\infty \Delta$.

If $r : [0, \infty) \to \Delta$ is a geodesic ray then there exists a point $\xi \in \partial_\infty \Delta$ such that $r(t_n) \to \xi$ for every sequence (t_n) of positive real numbers such that $t_n \to \infty$. Write $\xi = r(\infty)$. The following statements are easily proved:

Proposition 3.5 *Let Δ be a geodesic δ-hyperbolic space.*

(a) Let (x_n) and (y_n) be sequences of points in Δ. Suppose that (x_n) converges to a point $\xi \in \partial_\infty \Delta$ and there is a real number B such that $d(x_n, y_n) \leq B$ for every n. Then (y_n) also converges to ξ.

(b) Let $r_1, r_2 : [0, \infty) \to \Delta$ be geodesic rays such that $r_1(\infty) = r_2(\infty)$. Then every point on r_1 is at distance $\leq d(r_1(0), r_2(0)) + 8\delta$ from r_2, and there exists a real number $T \geq 0$ such that for every $t \geq T$, $r_1(t)$ is at distance $\leq 8\delta$ from r_2.

The Gromov product can be naturally extended to $\overline{\Delta} \times \overline{\Delta}$, by

$$(a|b) = \inf[\liminf_{i,j \to \infty} (a_i|b_j)]$$

where the infimum is taken over all sequences (a_i), (b_j) in Δ such that (a_i) converges to a and (b_j) converges to b. We have $(a|b) = \infty$ iff $a = b \in \partial_\infty \Delta$, and

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

for all x, y and z in $\overline{\Delta}$.

3.2 The visual metric on $\overline{\Delta}$

Let Δ be a proper geodesic δ-hyperbolic space equipped with a basepoint 0 (a metric space is *proper* if its closed metric balls are all compact). Gromov constructs a "visual metric" on $\overline{\Delta}$ by scaling the metric on Δ in the following way: Let $w > 1$ be a real number. Let $f_w : \Delta \to \mathbf{R}$ be the function $f(x) = w^{-d(0,x)}$. If σ is a path in Δ, define $l_w(\sigma)$ to be the integral of f along σ. If x and y are points of Δ, define

$$|x - y|_w = \inf(l_w(\sigma))$$

where the infimum is taken over all continuous paths σ from x to y. $|\cdot|_w$ gives a metric on Δ, such that the induced topology is the same as the original topology on Δ.

Gromov has shown that there is a constant $w_0(\delta) > 1$ such that for every real number w strictly between 1 and w_0, the following hold:

(1) The identity map of Δ extends to a bijection from $\overline{\Delta}$ to the completion of Δ with respect to $|\cdot|_w$, so that $|\cdot|_w$ induces a metric on $\overline{\Delta}$, which we will call the *visual metric of $\overline{\Delta}$ with respect to w*. The visual metrics all induce the same topology on $\overline{\Delta}$, and $\overline{\Delta}$ is compact. If $\xi \in \partial_\infty \Delta$ then a sequence (x_n) of points in Δ converges to ξ in the sense of the previous subsection iff it converges to ξ in the metric topology.

(2) Let ξ and η be in $\overline{\Delta}$ and let $[\xi, \eta]$ be a geodesic from ξ to η. Then

$$|\xi - \eta|_w \leq \nu w^{-d(0,[\xi,\eta])},$$

where $\nu \geq 1$ is a constant that depends on Δ and w. Furthermore, if ξ and η are in $\partial_\infty \Delta$ and $\xi \neq \eta$ then

$$\frac{1}{\nu} w^{-d(0,[\xi,\eta])} \leq |\xi - \eta|_w.$$

Note that the diameter of $\overline{\Delta}$ in the metric $|\cdot|_w$ is at most $\frac{2}{\log(w)}$.

Example: if $\Delta = \mathbf{H^n}$ then $\overline{\Delta}$ is homeomorphic to the closed unit ball in R^n, and the topology induced by the visual metrics agree with the topology inherited from R^n.

3.3 Quasiconvex sets and their visual diameter

If S is a subset of a geodesic metric space Δ and $q > 0$ is a real number, write

$$S^q = \{x \in \Delta : d(x, S) \leq q\}.$$

Definition 3.6 *A subset S of a geodesic metric space is q-***quasiconvex** *if whenever x and y are points in S, every geodesic segment $[x, y]$ is contained in S^q.*

For the remainder of this subsection Δ will be a proper, geodesic δ-hyperbolic space. If $S \subset \Delta$, we will write $\partial_\infty S$ for the subset of $\partial_\infty \Delta$ consisting of those points that are accumulation points of elements of S (to be distinguished from the boundary of S in Δ). We have

Proposition 3.7 *Let $S \subset \Delta$ be q-quasiconvex. If $x \in S$ and $\xi \in \partial_\infty S$, then any geodesic ray $g : [0, \infty) \to \Delta$ from x to ξ is contained in $S^{q+8\delta}$.*

Proof:

Let (x_n) be a sequence of points in S that converges to ξ, and let $r_n : [0, \infty) \to \Delta$ be the map that sends the interval $[0, d(x, x_n)]$ isometrically onto a geodesic $[x, x_n]$ and maps $[d(x, x_n), \infty)$ onto the point x_n. Since Δ is proper, by Ascoli's theorem (r_n) has a subsequence (which we will continue to call (r_n)) that converges to a map $r : [0, \infty) \to \Delta$, in the topology of uniform convergence on

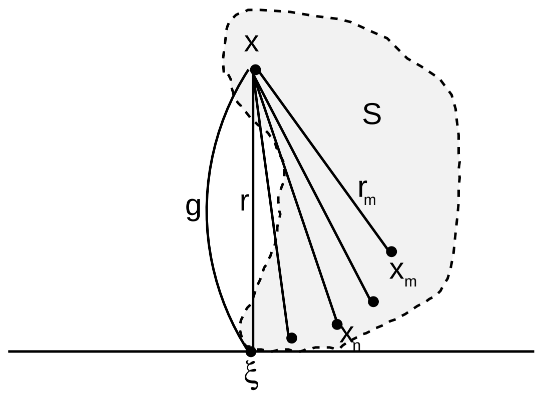

Figure 7

compact sets (see Figure 7). The map r is a geodesic ray from x to ξ, so by Proposition 3.5 *(b)*, every point on g is at distance $\leq 8\delta$ from r. At the same time, every point on r is at distance $\leq q$ from S; hence each point on g is within $q + 8\delta$ of a point of S. □

We will also need the following simple fact:

Proposition 3.8 *Let $S \subset \Delta$ be closed and q-quasiconvex and let T be a subset of Δ that satisfies $S \subset T \subset S^R$, where R is a real number. Then T is Q-quasiconvex, where $Q = q + R + 8\delta$.*

Quasiconvex sets that are far from the basepoint 0 of Δ have small diameters in the visual metrics on $\overline{\Delta}$. We have the following, which will be crucial for controlling the behavior of the sets H_α of Theorem C:

Proposition 3.9 *Let $S \subset \Delta$ be q-quasiconvex. Then the diameter of S in the metric $|\cdot|_w$ is bounded above by $\nu w^{-(d(0,S)-q)}$.*

Proof:
Let x and y be in S. Every point on $[x, y]$ is within q of a point in S, so $d(0, [x, y]) \geq d(0, S) - q$. Hence $|x - y|_w \leq \nu w^{-d(0,[x,y])} \leq \nu w^{-(d(0,S)-q)}$. □

3.4 Quasi-isometries and quasi-geodesics

Let Δ_0 and Δ be two metric spaces. Let $k \geq 1$ and $\mu \geq 0$ be real numbers. We say that a map $f : \Delta_0 \to \Delta$ is (k, μ)-*quasi-Lipschitz* if $d(f(x_1), f(x_2)) \leq k \cdot d(x_1, x_2) + \mu$ for all x_1 and x_2 in Δ_0.

A map $f : \Delta_0 \to \Delta$ is a (k, μ)-*quasi-isometry* if

$$\frac{1}{k} d(x_1, x_2) - \mu \le d(f(x_1), f(x_2)) \le k d(x_1, x_2) + \mu$$

for all x_1 and x_2 in Δ_0.

A quasi-isometry between two δ-hyperbolic spaces extends continuously to the boundary, in the following sense:

Theorem 3.10 *Let Δ_0 and Δ be Gromov-hyperbolic, and let $h : \Delta_0 \to \Delta$ be a quasi-isometry. For every sequence (x_n) of points in Δ_0 that converges to a point ξ in $\partial_\infty \Delta_0$, the sequence $(h(x_n))$ converges to a point in $\partial_\infty \Delta$ that depends only on ξ, so that h defines a continuous map from $\partial_\infty \Delta_0$ to $\partial_\infty \Delta$. The map $h : \partial_\infty \Delta_0 \to \partial_\infty \Delta$ is injective.*

Note that Theorems C and D of this paper can be thought of as a generalization of this theorem, since if there are no electric sets then Theorems C and D together reduce to this statement. The proof of the above statement is, among other things, an important step in the proof of Mostow's rigidity theorem.

If Δ is a metric space, a (k, μ)-*quasigeodesic* is a rectifiable path $p : I \to \Delta$, where I is an interval in $\mathbf{R}$, such that for all s and t in I,

$$\frac{1}{k} l(p|_{[s,t]}) - \mu \le d(p(s), p(t)) \le k \cdot l(p|_{[s,t]}) + \mu.$$

Note that if a path $p : I \to \Delta$ is parametrized by arc length then it is a quasigeodesic if and only if it is a quasi-isometry. The following theorem tells us that in a hyperbolic space (the image of) a quasigeodesic is quasiconvex.

Theorem 3.11 *Let Δ be a geodesic, δ-hyperbolic space.*

(a) Let $I = [a, b]$ be a closed interval in $\mathbf{R}$, and let $p : I \to \Delta$ be a (k, μ)-quasigeodesic. Let l be a geodesic segment in Δ whose endpoints are $p(a)$ and $p(b)$. Then there is a constant H that depends only on δ, k and μ, such that the Hausdorff distance between the images of p and l is at most H.

(b) Suppose that Δ is proper. Let $p : I \to \Delta$ be a (k, μ)-quasigeodesic, where $I = [0, \infty)$ or $I = (-\infty, \infty)$. Then there exists, respectively, a geodesic ray or an infinite geodesic l, such that the Hausdorff distance between the images of p and l is at most H, where again H is a constant that depends only on δ, k and μ.

A consequence of the preceding theorem is the following key property of quasigeodesics, which makes them useful tools for showing that quasi-isometries extend continuously to the boundary:

Proposition 3.12 *Let Δ be a proper geodesic δ-hyperbolic space and let $p : [0, \infty) \to \Delta$ be a quasigeodesic ray. Then p terminates at a unique point ξ in $\partial_\infty \Delta$; that is, as $t \to \infty$, $p(t) \to \xi$.*

Proposition 3.12 is the first step in showing that a quasi-isometry of hyperbolic spaces extends continuously to the boundary: if the spaces in question are proper and geodesic, it shows that the restriction of the quasi-isometry to any geodesic ray extends continuously to the boundary point of the ray, since the image of the ray under the quasi-isometry is a quasigeodesic. We will use a similar strategy to show that a map satisfying the conditions of Theorem C extends to the endpoints of geodesic rays.

4 Continuity of the Extension

In this section we will prove Theorem C. If $h : X \to Y$ is a map that satisfies the hypotheses of Theorem C, we want to show that h extends continuously to a map from $\partial_\infty X$ to $\partial_\infty Y$. We know that on each set H_α, h extends to the boundary. Moreover, h is assumed to be a quasi-isometry away from the sets H_α; so since quasi-isometries extend to the boundary, it is not difficult to show that h extends to those points of $\partial_\infty X$ that have a neighborhood that is disjoint from all of the sets H_α. Thus, the main problem will be to control h where the sets H_α accumulate. Suppose ξ is a point in $\partial_\infty X$ that is the limit of a sequence of sets (H_n). A first step in showing that h extends to ξ is to show that the sets $h(H_n)$ converge to a unique point in $\partial_\infty Y$. To do this, we must be able to show in particular that the visual diameters of the sets $h(H_n)$ approach 0 as n approaches infinity (in all of what follows we will assume that we have fixed a basepoint 0 in X and Y such that $h(0) = 0$, and a visual metric d_{vis} on X and Y, where $d_{vis} = | \cdot |_w$ for some constant $w > 1$, as in Section 3). Since the sets $h(H_n)$ are uniformly quasiconvex, we can get an upper bound on their visual diameters if we know how far they are from 0, by Proposition 3.9. Our main tool for showing that a set $h(H_n)$ is far from 0 is the*electric metric* (using terminology coined by Farb in [15]):

Definition 4.1 *If $(\Delta, \{S_\alpha\})$ is a pair satisfying $(*_1) - (*_3)$, and $\mathcal{C}$ denotes the complement in Δ of the sets S_α, then the* **electric length** *l_{el} of a rectifiable path p in Δ is defined by $l_{el}(p) = l(p \cap \mathcal{C})$. The* **electric distance** *d_{el} on Δ is the path metric induced by l_{el}. We will call the sets S_α* **electric sets***, and the space E_Δ obtained from Δ by collapsing each set S_α to a point, the* **electric space of** *Δ.*

Note that d_{el} is not an actual distance function, since if x and y are distinct points in one of the electric sets S_α then $d_{el}(x, y) = 0$. However, the electric distance descends to a genuine metric on E_Δ; we will use d_{el} to denote this metric as well as the electric metric on Δ.

In the setting of Theorem C, the electric distance enables us to show that certain sets $h(H_n)$ in Y are far from 0 because of the following: one hypothesis of Theorem C is that the map

$h_E : E_X \to E_Y$ is an (L, κ)-quasi-isometry, for some constants L and κ. This hypothesis translates upstairs (that is, in X and Y) to the condition that

$$\frac{1}{L} d_{el}(x, y) - \kappa \leq d_{el}(h(x), h(y)) \leq L \cdot d_{el}(x, y) + \kappa$$

for all x and y in X. The electric distance is bounded above by the standard distance, so one consequence of the above inequality is that for all x and y in X,

$$\frac{1}{L} d_{el}(x, y) - \kappa \leq d(h(x), h(y)).$$

Thus, if two points in X are far apart in the electric metric then their images under h are far apart in the standard metric on Y.

We will begin by studying electric quasigeodesics, which will play a similar role in the proof of Theorem C as quasigeodesics play in the proof that a quasi-isometry of hyperbolic spaces extends continuously to the boundary. We will prove a generalization of Theorem 3.11 that gives us some control over the distance between an electric quasigeodesic and a standard geodesic with the same endpoints. We will show that a (standard) geodesic ray is an electric quasigeodesic ray exactly when its endpoint at infinity does not lie on the boundary of any electric set. This will enable us to extend the map h to the endpoints of such rays, and then to the entire boundary.

4.1 Electric quasigeodesics

If $(\Delta, \{S_\alpha\})$ satisfies $(*_1) - (*_3)$, let $\Pi_\Delta : \Delta \to E_\Delta$ denote the projection map. If $e : I \to \Delta$ is a rectifiable path, we say that e is an *electric* (k, μ)*-quasigeodesic* if

$$\frac{1}{k} l_{el}(e|_{[s,t]}) - \mu \leq d_{el}(e(s), e(t)) \leq k \cdot l_{el}(e|_{[s,t]}) + \mu$$

for all $s, t \in I$ with $s < t$. Note that a path $e : I \to \Delta$ is an electric (k, μ)-quasigeodesic if and only if $\Pi_X \circ e : I \to E_\Delta$ is a quasigeodesic in the electric metric.

Since d_{el} is a path metric, if $e : I \to \Delta$ is any path from a to b then we have $d_{el}(a, b) \leq l_{el}(e)$; so for e to be a (k, μ)-quasigeodesic it is in fact sufficient for e to satisfy

$$\frac{1}{k} l_{el}(e|_{[s,t]}) - \mu \leq d_{el}(e(s), e(t))$$

for all s and t in I.

We have the following relationship between standard quasigeodesics and electric quasigeodesics:

Proposition 4.2 *Let $e : I \to \Delta$ be an electric (k, μ)-quasigeodesic and let J be a subinterval of I such that $e(J) \subset \mathcal{C}$. Then $e|_J$ is a (non-electric) (k, μ)-quasigeodesic.*

Proof:
For $s, t \in J$ we have $l_{el}(e|_{[s,t]}) = l(e|_{[s,t]})$, and $d_{el}(e(s), e(t)) \leq d(e(s), e(t))$. Hence

$$\frac{1}{k} l(e|_{[s,t]}) - \mu \leq d_{el}(e(s), e(t)) \leq d(e(s), e(t)).$$

So $e|_J$ is a (k, μ)-quasigeodesic. $\square$

We have not yet shown that electric quasigeodesics actually exist; but since d_{el} is defined as an infimum of path lengths, given any $\epsilon > 0$ and points x and y in Δ, we can find a path p from x to y such that $l_{el}(p_\epsilon) - \epsilon \leq d_{el}(x, y)$. It is easily checked that such a path is an electric $(1, \epsilon)$-quasigeodesic.

Since the electric sets are separated from each other by at least c, every time a path p in Δ travels from one electric set to another, at least c is contributed to the electric length of p. Say that p *does not backtrack* if it passes through each electric set at most once. Thus, if p does not backtrack and it enters electric sets n times then the electric length of p is at least $c(n-1)$.

If $e : I \to \Delta$ is an electric (k, μ)-quasigeodesic then it can be modified so that it does not backtrack, as follows: choose some α such that e passes through S_α, and let s and t be the first and last points of I, respectively, that are mapped by e into S_α. Replace e on $[s, t]$ by a path from $e(s)$ to $e(t)$ entirely contained in S_α. Since e has finite electric length it jumps from one electric set to another a finite number of times, so after a finite number of modifications we will have produced a path that does not backtrack. In the process, we have if anything reduced the electric length of e, so the new path is again an electric (k, μ)-quasigeodesic. Note that we have not changed e on the endpoints of I.

We will begin by examining the relationship between electric quasigeodesics and standard geodesics in Δ. If $S \subset \Delta$, let $EN_0(S)$ denote the 0-neighborhood of S in the electric metric, that is, the union of S with those electric sets that S intersects. We can make the following generalization of Theorem 3.11 (Farb has proved a similar statement in [15] when the electric sets are horoballs):

Proposition 4.3 *Let e be an electric (k, μ)-quasigeodesic in Δ, with endpoints x and y. Then for every point z in $[x, y]$ there is a point in $EN_0(e)$ that is within (non-electric) distance D of z, where $D = D(\delta, k, \mu)$ depends only on k, μ and the hyperbolicity constant δ of Δ. Moreover, for every point z' in e there is a point in $[x, y]$ that is within electric distance D of z'.*

Proof:
Choose $R > 0$, and suppose z is a point in $[x, y]$ for which no point of $EN_0(e)$ comes within R of z. Let (a, b) be the maximal subinterval of $[x, y]$ containing z and such that no point of $EN_0(e)$

comes within R of (a,b). We will show that if R is chosen sufficiently large then there is an upper bound m on the length of all such intervals (a,b). Then we have that $d(z, EN_0(e)) \leq R+m$, and the first statement is proved. The idea of the argument is as follows: if (a,b) is "too long" then by Lemma 3.3 the portion of e that lies outside the R-neighborhood of $[x,y]$ must travel through many electric sets; but each time e travels from one electric set to another a definite amount is contributed to its electric length, making it inefficient, which contradicts the fact that it is an electric quasigeodesic.

So, let a' and b' be points in $EN_0(e)$ such that $d(a,a') \leq R$ and $d(b,b') \leq R$. We can assume that a' and b' actually lie on e; let e' be a subsegment of e whose endpoints are a' and b'. As per the discussion preceding this proposition, we can assume, after removing some portion of e', that e' does not backtrack.

Let e'_{str} be the path obtained by straightening the portions of e' that pass through the sets S_α and those that pass through $\mathcal{C}$; that is, if a component of $e' \cap S_\alpha$ or $e' \cap \mathcal{C}$ has endpoints s and t, replace that component with $[s,t]$. Then since the sets S_α are all q-quasiconvex and the components of $e' \cap \mathcal{C}$ are all (k,μ)-quasigeodesics, it is easy to see that e'_{str} lies outside of the $(R-q-H)$-neighborhood of (a,b), where $H = H(\delta,k,\mu)$ is the constant from Theorem 3.11.

Choose $N \in \mathbf{Z}_+$ such that $\frac{cN}{2} > 32\delta k$ (where c is the minimum separation between any two of the electric sets). Let R be sufficiently large that $R-q-H > K(N)$, where $K(N)$ is the constant from Lemma 3.3.

Let $j = [\frac{l([a,b])}{32\delta}]$; by Lemma 3.3, e'_{str} must consist of at least $N \cdot j$ geodesic arcs, so e' must travel through at least $\frac{N \cdot j}{2}$ of the electric sets S_α. Every time e' passes from one electric set to another, at least c is added to its electric length; so we have

$$l_{el}(e') \geq (\frac{N \cdot j}{2} - 1)c \geq [\frac{N}{2} \cdot (\frac{l([a,b])}{32\delta} - 1) - 1]c.$$

But since e' is an electric (k,μ)-quasigeodesic we also have that

$$l_{el}(e') \leq k \cdot d_{el}(a',b') + \mu \leq k(l([a,b]) + 2R) + \mu.$$

Combining these two inequalities and simplifying, we have

$$(\frac{Nc}{2(32\delta)} - k)l([a,b]) \leq 2Rk + (\frac{N}{2} + 1)c + \mu.$$

By our choice of N, $(\frac{Nc}{2(32\delta)} - k)$ is positive, so the above inequality gives an upper bound on $l([a,b])$ and proves the first statement of the proposition.

To prove the second statement, assume that z' is a point on e such that $EN_0(z')$ does not come within hyperbolic distance R of $[x,y]$. Let e' now be the maximal subsegment of e containing z' such that $EN_0(e')$ does not come within R of $[x,y]$. Let a' and b' be the endpoints of e'. There is a point in $EN_0(a')$ and a point in $EN_0(b')$ that come within R of points a and b, respectively, of $[x,y]$; without loss of generality we may assume that a' and b' themselves satisfy $d(a',a) = R$ and $d(b',b) = R$. By the same reasoning as above, there is an upper bound m on $l([a,b])$ that depends only on δ, k and μ. Now we have

$$l_{el}(e') \leq k \cdot d_{el}(a',b') + \mu \leq k(l([a,b]) + 2R) + \mu \leq k(m + 2R) + \mu.$$

Hence

$$d_{el}(z',[x,y]) \leq d_{el}(z',a') + d_{el}(a',a) \leq l_{el}(e') + R \leq k(m + 2R) + \mu + R.$$

$\square$

4.2 Extension of h to the endpoints of rays

Returning to the setting of Theorem C, if $r : [0,\infty) \to X$ is a geodesic ray whose endpoint ξ at infinity lies on one of the sets $\partial_\infty H_\alpha$ then it is fairly easy to see that the restriction of h to the ray r extends continuously to ξ. It is more difficult to extend h along a ray whose endpoint does not lie on the boundary of any electric set. We will use the following terminology, given a system $(\Delta, \{S_\alpha\})$:

Definition 4.4 *If $\xi \in \partial_\infty \Delta$ does not lie on the boundary of any set S_α then we call ξ a* **residual point**. *We denote by $\partial_\infty \Delta_{res}$ the set of residual points in $\partial_\infty \Delta$.*

The following result shows in particular that if $r : [0,\infty) \to \Delta$ is a geodesic ray whose endpoint at infinity is a residual point then r is an electric quasigeodesic.

Theorem 4.5 *There exist constants a, b, a' and A such that if r is a geodesic segment or ray in Δ starting at 0 then there exists a sequence (χ_n) of sets that are either electric sets or single points on r, such that the following hold:*

(1) $\cup_n \chi_n^A$ covers r;

(2) $an - b \leq d_e(0,\chi_n) \leq a'n$; and

(3) $a|m-n| - b \leq d_e(\chi_m,\chi_n) \leq a'|m-n|$.

The sequence (χ_n) is infinite if the endpoint at infinity of r is a residual point, and is finite otherwise.

Proof:

Note first that up to some adjustment of constants, (2) is a consequence of (3); thus we need only show that (1) and (3) hold.

We will build our sequence (χ_n) out of electric sets that remain close to r for a long time; wherever there is no such electric set, we will instead include a point on the ray r. If we make sure that these electric sets and points appear at regular intervals (that is, that any two adjacent ones are not too far apart) then it is easy to show the right-hand side of the inequality in (3). To see that the left-hand side of the inequality holds, we show that because of the way we have chosen the members of the sequence, any electric quasigeodesic from χ_n to χ_m must either travel through many electric sets or it must spend a long time in the complement of the electric sets; in either case, we get a lower bound on its electric length.

So, set $x_0 = 0$ and consider the set Σ_0 of those S_α that come within D of x_0, where $D = D(\delta, 1, 1)$ is the constant from Proposition 4.3. If Σ_0 is empty then take $\chi_0 = \{x_0\}$. If Σ_0 is nonempty then it is a finite set since the electric sets are all separated from each other by at least c. So if $k_\alpha = \max\{t : d(r(t), S_\alpha) \leq D\}$ then there is an electric set S_{α_0} in Σ_0 for which k_{α_0} is maximal over all elements of Σ_0 (note that k_{α_0} may be infinity). Take $\chi_0 = S_{\alpha_0}$. Let y_0 denote the last point in $\{p \in r : d(p, S_{\alpha_0}) \leq D\}$ when traveling away from 0, that is, $y_0 = r(k_{\alpha_0})$ (note that there will be such a point y_0 except when the endpoint at infinity of r is in S_{α_0} or perhaps when r is a finite segment, in which cases the sequence (χ_n) terminates here). If Σ_0 is empty then set $y_0 = x_0$.

We wish to construct a sequence (χ_n) of sets that are close enough together that for some A the union of their A-neighborhoods covers r, but far enough apart that each χ_n contributes a definite amount to the electric length of any path connecting two of the sets χ_j and χ_k such that $j < n < k$.

Accordingly, let x_1 be the point $2D + 2c$ units farther along r than y_0 (again traveling away from 0). Let Σ_1 be the (at most finite) set of those electric sets that come within D of x_1. Note that Σ_0 and Σ_1 are disjoint. If Σ_1 is empty then take $\chi_1 = \{x_1\}$. Otherwise, let S_{α_1} be an element of Σ_1 for which k_{α_1} is maximal over all elements of Σ_1. Take $\chi_1 = S_{\alpha_1}$. Let y_1 be the last point in $\{p \in r : d(p, S_{\alpha_1}) \leq D\}$ (again, there will be such a last point except when the endpoint of r is in S_{α_1} or perhaps when r is finite, in which cases the sequence terminates here). If Σ_1 is empty then take $y_1 = x_1$. Let x_2 be the point $2D + 2c$ units farther along r than y_1. Continue constructing the sequences (χ_n), (x_n) and (y_n) in this manner (see Figure 8). Note that the sets Σ_n are all disjoint.

To show that the sequence (χ_n) satisfies (1), observe first that $\chi_n^{D+q+8\delta}$ contains the interval $[x_n, y_n]$. This is trivial if $\chi_n = x_n = y_n$. If χ_n is an electric set then by construction there are points a_n and b_n in χ_n such that $d(x_n, a_n) \leq D$ and $d(y_n, b_n) \leq D$. This implies that the segment $[x_n, y_n]$ lies within $D + 8\delta$ of the geodesic segment $[a_n, b_n]$, since the quadrilateral $[x_n, y_n] \cup [x_n, a_n] \cup [a_n, b_n] \cup [b_n, y_n]$ is 8δ-narrow; also, every point in $[a_n, b_n]$ lies within q of a point on χ_n, since χ_n is q-quasiconvex. Thus the sets $\chi_n^{D+q+8\delta}$ cover all of r except possibly the intervals $[y_n, x_{n+1}]$; these intervals all have length $2D + 2c$, so if we set $A = 3D + q + 8\delta + 2c$ then $\cup_n \chi_n^A$ covers r.

So it remains to show that the left-hand inequality in *(3)* holds, that is, there exist constants a and b such that for all n and m, $a|m - n| - b \leq d_{el}(\chi_m, \chi_n)$. We will show that there is a constant a such that $d_{el}(x_m, x_n) \geq a|m - n| - a - 1$; the triangle inequality then gives us that

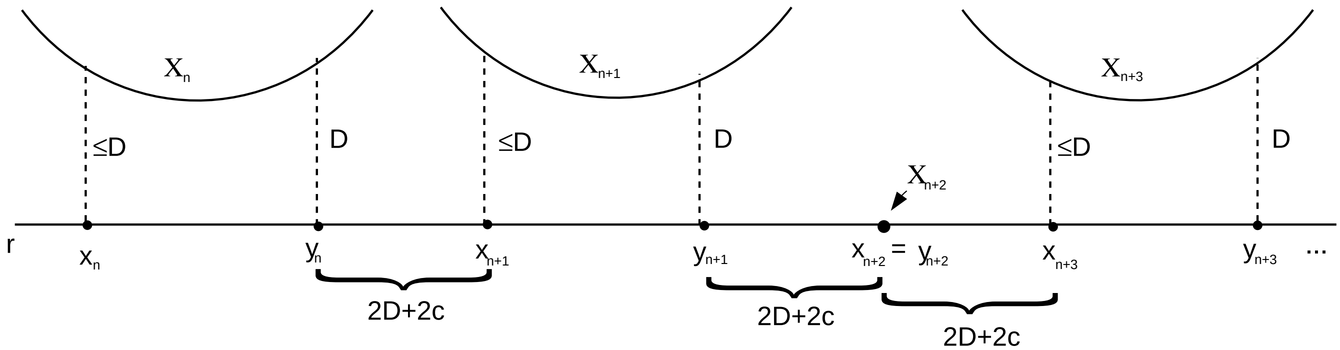

Figure 8: *A possible configuration of some of the sets in the sequence* (χ_n).

$d_{el}(\chi_m, \chi_n) \geq a|m-n| - a - 1 - 2D$.

Let e_{mn} be an electric $(1,1)$-quasigeodesic path from x_m to x_n, so that $d_{el}(x_m, x_n) \geq l_{el}(e_{mn}) - 1$. For each i strictly between m and n we will construct a segment of e_{mn} of a definite electric length; if we can control the extent to which these segments overlap then we can get a lower bound on $l_{el}(e_{mn})$ in terms of $|m-n|$.

By Proposition 4.3, for every i strictly between m and n there is a point x'_i in $EN_0(e_{mn})$ such that $d(x_i, x'_i) \leq D$. If x'_i is in the complement of the sets S_α then it lies on a segment of $e_{mn} \cap \mathcal{C}$ whose length is at least c. Choose a subsegment s_i such that s_i contains x'_i and $l(s_i) = c$. Setting $a = \frac{1}{2}c$, each such segment s_i contributes $2a$ to $l_{el}(e_{mn})$, and if $j \neq k$ then $s_j \cap s_k$ is at most a point since by our construction $d(x'_j, x'_i) \geq 2c$.

If x'_i lies in some electric set S_i then there is a segment of e_{mn} of length at least c between S_i and the next electric set e_{mn} enters; choose t_i to be the subsegment of that segment that abuts S_i, and whose length is c. Each segment t_i contributes $2a$ to $l_{el}(e_{mn})$, and if $j \neq k$ then $t_j \cap t_k = \emptyset$ since by construction $S_j \neq S_k$ (since $S_j \in \Sigma_j$ and $S_k \in \Sigma_k$, and $\Sigma_k \cap \Sigma_j = \emptyset$).

Thus, we have constructed $n-m-1$ segments of e_{mn}, each of which contributes $2a$ to $l_{el}(e_{mn})$. It is possible for s_j and t_k to overlap, but by counting all the lengths of the segments s_j and t_k we have at worst counted some pieces twice, so we have

$$l_{el}(e_{mn}) \geq \frac{1}{2}[\sum_j l(s_j) + \sum_k l(t_k)] = \frac{1}{2}(2a(|n-m|-1)) = a(|n-m|) - a$$

and hence $d_{el}(x_m, x_n) \geq a(|n-m|) - a - 1$. □

Returning to the setting of Theorem C, if two points x and y in X are far from each other in the electric metric on X then we have seen that $h(x)$ and $h(y)$ are far apart in the standard metric on Y. If $r : [0, \infty) \to X$ is a geodesic ray whose endpoint at infinity is a residual point then the previous result gives us a sequence (χ_n) of sets close to r whose electric distance from 0 is getting large linearly with respect to n; thus, the sets $h(\chi_n)$ are also getting far from 0 linearly with respect to n. The sets $h(\chi_n)$ are uniformly quasiconvex, so by Proposition 3.9 their visual diameters are getting small exponentially fast with respect to n; hence the sum of their diameters is finite. This enables us to show the following:

Theorem 4.6 *Let $r : [0, \infty) \to X$ be a geodesic ray from 0 to a point ξ on $\partial_\infty X$. Then $h|_r$ extends continuously to ξ.*

Proof:

Let L and κ denote the quasi-Lipschitz constants of h and the quasi-isometry constants of h_E.

If ξ lies on the boundary of some electric set H_α then since H_α is quasiconvex, by Propositions 3.5*(b)* and 3.7 for large values of t there are points p_t in H_α such that each p_t is within a bounded distance of $r(t)$. As $t \to \infty$, the points $h(p_t)$ converge to a unique point η in $\partial_\infty Y$, since $h|_{H_\alpha}$ extends continuously to the boundary; so since h is quasi-Lipschitz, $h(r(t))$ must also tend towards η as $t \to \infty$.

Otherwise ξ is a residual point. In this case Theorem 4.5 gives an infinite sequence (χ_n) of sets such that $\cup_n \chi_n^A$ covers r and $d_{el}(0, \chi_n^A) \geq a \cdot n - b - A$; so $d(0, h(\chi_n^A)) \geq \frac{1}{L}(a \cdot n - b - A) - \kappa$. The sets $h(\chi_n)$ are all q-quasiconvex (since each χ_n is either a point or one of the sets H_α), and since h is (L, κ)-quasi-Lipschitz we have $h(\chi_n^A) \subset (h(\chi_n))^{LA+\kappa}$; hence by Proposition 3.8 the sets $h(\chi_n^A)$ are all Q-quasiconvex where $Q = q + LA + \kappa + 8\delta$. Thus we have for every n,

$$diam_{vis}(h(\chi_n^A)) \leq \nu w^{-[\frac{1}{L}(a \cdot n - b - A) - \kappa] + Q},$$

by Proposition 3.9 (where $diam_{vis}$ denotes the diameter in the visual metric $|\cdot|_w$ on Y). Thus the series

$$\sum_n diam_{vis}(h(\chi_n^A))$$

is dominated by the series

$$\sum_n \nu w^{-[\frac{1}{L}(a \cdot n - b - A) - \kappa] + Q},$$

which converges since $w > 1$. Therefore, by going out far enough in the sequence (χ_n^A) we can make the tail end of the first series as small as we want. As we travel along r towards ξ we move out farther and farther in the sequence (χ_n^A) (see Figure 9); so if we choose T large enough we can make the visual diameter of the set $\{h(r(t)) : t \geq T\}$ as small as we want. Hence $h(r(t))$ must converge to a unique point as $t \to \infty$. $\square$

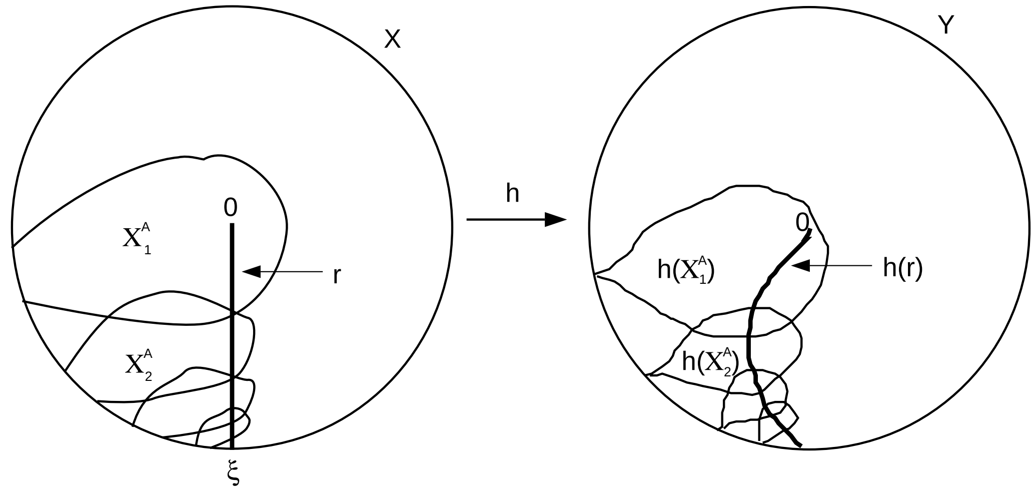

Figure 9: *The path $h(r)$ must terminate at a unique point in $\partial_\infty Y$.*

4.3 Extension of h to the boundary

In order to extend h to the endpoint of a geodesic ray r, we showed that for the sequence of sets (χ_n) associated to r by Theorem 4.5, the series

$$\sum_{n=1}^{\infty} diam_{vis}(h(\chi_n^A))$$

is convergent. The estimates we used to show that the series is convergent involve constants that do not depend of the choice of ray r. This observation enables us to prove the following uniformity property of these series: we can make the tail end of such a series as small as we like by going out past those sets χ_n such that $h(\chi_n^A)$ intersects a fixed ball in Y around 0; the radius of this ball is independent of the choice of ray r and associated sequence (χ_n). This is the main fact that we will need in order to show that h extends to all of $\partial_\infty X$.

Lemma 4.7 *For every $\epsilon > 0$, there is an $N \in \mathbf{Z}_+$ such that the following holds: Let r be a geodesic ray in X starting at 0, and let (χ_n) be the sequence of sets associated to r by Theorem 4.5. Then*

$$\sum_{j:d(0,h(\chi_j^A))\geq N} diam_{vis}(h(\chi_j^A)) < \frac{\epsilon}{2}.$$

Sketch of proof:

Let $\sigma_k = \{\chi_j^A : k < d(0, h(\chi_j^A)) \le k+1\}$. The idea of the proof is as follows: if χ_j^A is in σ_k then by Proposition 3.9 we have that $diam_{vis}(h(\chi_j^A)) \le \nu w^{-k+Q}$, where Q is the quasiconvexity constant of the sets $h(\chi_j^A)$. Note that the series in the statement of the lemma is the same as the series

$$\sum_{k=N}^{\infty}[\sum_{\chi_j^A \in \sigma_k} diam_{vis}(h(\chi_j^A))];$$

hence it is dominated by the series

$$\sum_{k=N}^{\infty} |\sigma_k| \nu w^{-k+Q}.$$

Because the sets χ_m and χ_n are at least $a \cdot |n-m| - b$ apart in the electric metric, their images are also far apart, and we get an upper bound on the number of sets χ_n^A in each σ_k, roughly proportional to k; concretely, it is easy to show that there exist constants c_1 and c_2, independent of the choice of ray r and sequence (χ_n), such that $|\sigma_k| \le c_1 k + c_2$. Thus each term of the series

$$\sum_{k=0}^{\infty} \sum_{\sigma_k} diam_{vis}(h(\chi_j^A))$$

is bounded above by the corresponding term of the convergent series

$$\sum_{k=0}^{\infty} (c_1 k + c_2) \nu w^{-k+Q}.$$

So by going out far enough in the first series (where "far enough" is independent of r and (χ_n)) we can make its tail end as small as we like. $\square$

We are now able to prove

Theorem C *Let X and Y be proper, geodesic δ-hyperbolic spaces, $\{H_\alpha\}$ a collection of closed, disjoint path-connected subsets of X, and $h : X \to Y$ a quasi-Lipschitz map such that for every H_α, $h|_{H_\alpha}$ extends continuously to a continuous map $h : \partial_\infty H_\alpha \to \partial_\infty Y$. If $(X, \{H_\alpha\})$ and $(Y, \{h(H_\alpha)\})$ satisfy $(*_1) - (*_3)$ and the induced map $h_E : E_X \to E_Y$ is a quasi-isometry then h extends continuously to a continuous map $h : \partial_\infty X \to \partial_\infty Y$.*

Proof:

Let $\xi \in \partial_\infty X$. By Theorem 4.6, if $r : [0, \infty) \to X$ is a geodesic ray from 0 to ξ then the image under h of r approaches a unique point η in $\partial_\infty Y$. The Hausdorff distance between any two rays from 0 to ξ is finite by Proposition 3.5*(b)*, so since h is quasi-Lipschitz the images of these rays under h must both approach the same point as $t \to \infty$, by Proposition 3.5*(a)*; thus, η does not depend on the choice of r.

Given $\epsilon > 0$, let V_ϵ denote the open ball of radius ϵ around η in the visual metric on $\overline{Y}$. We will show that for every $\epsilon > 0$ there is a neighborhood of ξ in $\overline{X}$ that is mapped by h into V_ϵ.

To construct this neighborhood, we will start with a certain open set W in X that is mapped by h into $V_{\frac{\epsilon}{2}}$, and consider the set S consisting of all points x in $\overline{X}$ such that some geodesic segment $[0,x]$ intersects W. We will see that W can be chosen so that S will contain a neighborhood of ξ in $\overline{X}$ (see Figure 10). If x is a point in S then let q_x be a point of $W \cap [0,x]$. The point $h(q_x)$ is in

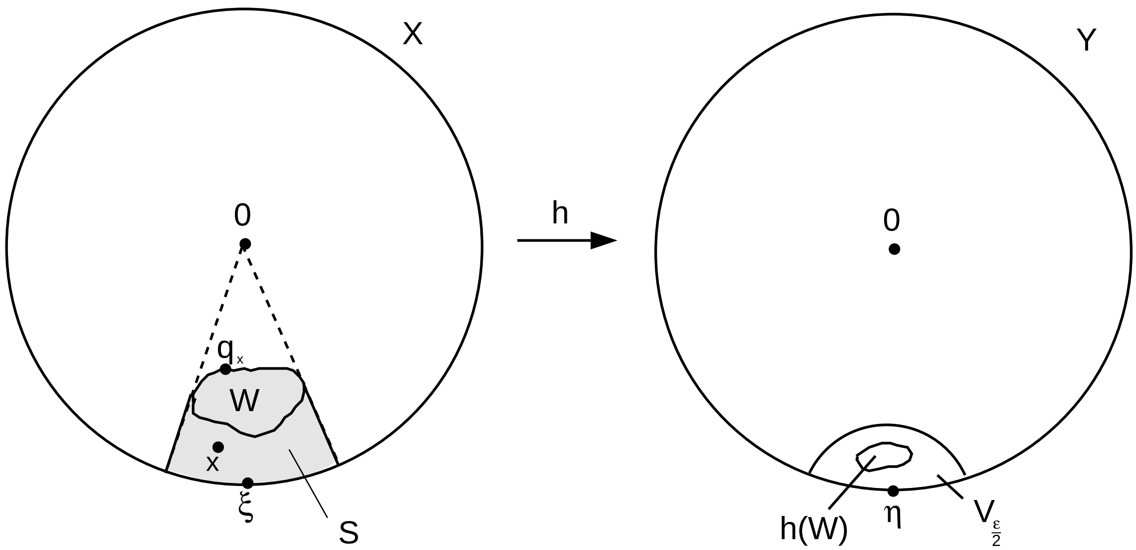

Figure 10

$V_{\frac{\epsilon}{2}}$, so if we could show that $h(x)$ and $h(q_x)$ are within $\frac{\epsilon}{2}$ of each other in the visual metric then we would have that $h(x) \in V_\epsilon$, which is what we would like. In order to make this true we will have to modify S, in such a way that the resulting set still contains a neighborhood of ξ.

If (χ_n) is the sequence of sets associated to $[0,x]$ by Theorem 4.5, consider the subsequence (χ_{n_k}) consisting of those sets that come within A of $[q_x,x]$; note that the sets $(\chi^A_{n_k})$ cover $[q_x,x]$. Let N be the constant from Lemma 4.7. If for all n_k we were to have $h(\chi^A_{n_k}) \cap B(0_Y,N) = \emptyset$ then the sum of the visual diameters of the sets $h(\chi^A_{n_k})$ would be at most $\frac{\epsilon}{2}$; so since $h([q_x,x])$ is contained in the union of these sets, we would have $d_{vis}(h(x),h(q_x)) \leq \frac{\epsilon}{2}$. But we cannot guarantee that this will be the case; there is no reason to expect that the sets $h(\chi^A_{n_k})$ will not intersect $B(0_Y,N)$. To prevent this from occurring, we will carve out of S as much as possible of those sets H^A_α for which $h(H^A_\alpha)$ intersects $B(0_Y,N)$, and, since some of the sets χ_n may consist of single points in $\mathcal{C}$ (the complement of the sets H_α), we will also remove some parts of $\mathcal{C}$ whose image under h intersects $B(0_Y,N)$.

So, let Σ denote the collection of those sets $\overline{H_\alpha^A}$ such that $h(H_\alpha^A)$ intersects $B(0,N)$ in Y. If $h(H_\alpha^A)$ intersects $B(0,N)$ then $h(H_\alpha)$ must intersect $B(0, N+LA+\kappa)$; this is a compact set, so since the sets $h(H_\alpha)$ are all separated from each other by at least c, only finitely many of them can intersect $B(0, N+LA+\kappa)$. Hence Σ is a finite collection. If ξ does not lie in $\overline{H_\alpha^A}$ for any of the $\overline{H_\alpha^A}$ in Σ, let U_1 be a neighborhood of ξ that is disjoint from the sets $\overline{H_\alpha^A}$ in Σ. Otherwise, let U_1 be a neighborhood of ξ sufficiently small that if $U_1 \cap \overline{H_\alpha^A} \neq \emptyset$ (where $\overline{H_\alpha^A} \in \Sigma$) then $h(U_1 \cap \overline{H_\alpha^A})$ is contained in $V_{\frac{\epsilon}{2}}$ (we can choose such a neighborhood since $h|_{H_\alpha^A}$ extends continuously to the boundary).

Let U_2 be a neighborhood of ξ in $\overline{X}$ such that $h(U_2)$ does not intersect $B(0,N)$ in Y. If r is a geodesic ray from 0 to ξ then by the previous theorem we can choose a point p_r on $r \cap U_1 \cap U_2$ such that $h(p_r) \in V_{\frac{\epsilon}{2}}$. The map h is not assumed to be continuous, but since h is (L,κ)-quasi-Lipschitz, if we choose p_r far enough out on r then we can find a neighborhood W_r of p_r in $U_1 \cap U_2$ that is mapped by h into $V_{\frac{\epsilon}{2}}$: concretely, choose p_r far enough along on r so that the ball in Y around $h(p_r)$ of radius $1+\kappa$ is contained in $V_{\frac{\epsilon}{2}}$; then the ball around p_r of radius $\frac{1}{L}$ is mapped by h into $V_{\frac{\epsilon}{2}}$. Let $W = \cup_r W_r$, where the union is taken over all geodesic rays from 0 to ξ (with only one point p_r being associated to each ray r). As before, let S be the set of all $x \in \overline{X}$ for which $[0,x]$ intersects W.

Claim: $U_1 \cap U_2 \cap S$ contains a neighborhood of ξ. U_1 and U_2 are open neighborhoods of ξ, so we need only show that there is a neighborhood of ξ contained in S; we will show that for any sequence (x_i) tending towards ξ, eventually all terms of the sequence are in S. Suppose instead that (x_i) has an infinite subsequence (which we will again call (x_i)) that is disjoint from S. By Ascoli's theorem, some subsequence of the geodesic rays $[0,x_i]$ converges to a path in $\overline{X}$, which is necessarily a geodesic ray from 0 to ξ. But then eventually the rays $[0,x_i]$ must intersect W, which gives a contradiction.

So to show that h extends continuously to ξ we need only show that $U = U_1 \cap U_2 \cap S$ is mapped by h into V_ϵ.

Let x be a point in U. If x lies in one of the sets $\overline{H_\alpha^A}$ in Σ then by construction $h(x) \in V_{\frac{\epsilon}{2}} \subset V_\epsilon$ and we are done. Otherwise, let s be a geodesic segment (or ray) from 0 to x. Let s' be the maximal subsegment of s that terminates at x, that lies completely in U, and that is disjoint from the sets $\overline{H_\alpha^A}$ in Σ. Note that if p is the initial endpoint of s' then $h(p) \in V_{\frac{\epsilon}{2}}$, since either p is in W or it is in the portion of one of the sets $\overline{H_\alpha^A}$ that is mapped into $V_{\frac{\epsilon}{2}}$. We will show that the diameter of $h(s')$ in the visual metric is less than $\frac{\epsilon}{2}$, so that $h(x)$ is in V_ϵ.

Let (χ_n) be the sequence of sets associated to s by Theorem 4.5. By accepting only those χ_n such that $\chi_n^A \cap s' \neq \emptyset$ we get a subsequence (χ_{n_k}) such that $\cup \chi_{n_k}^A$ covers s', and such that for all n_k, $h(\chi_{n_k}^A)$ does not intersect $B(0_Y, N)$ (this is true because if χ_{n_k} is an electric set H_α then $\overline{H_\alpha^A}$ is not in Σ, by construction; and if Ξ_{n_k} is a single point on s' then that point is in U_2, hence its image under h is not in $B(0_Y, N)$).

Now by Lemma 4.7,

$$\sum_{n_k} diam_{vis}(h(\chi^A_{n_k})) < \frac{\epsilon}{2}.$$

Hence $d_{vis}(h(p), h(x)) < \frac{\epsilon}{2}$, and $h(x) \in V_\epsilon$. □

5 The boundary of the electric space

If Δ is a proper, geodesic, δ-hyperbolic space and $\{S_\alpha\}$ is a collection of closed, disjoint, path-connected subsets of Δ, we have defined the electric space E_Δ to be the space obtained from Δ by collapsing each set S_α to a point; E_Δ inherits an electric distance from the distance function on Δ. In this section we will show that if $(\Delta, \{S_\alpha\})$ satisfies $(*_1) - (*_3)$ then the space E_Δ is hyperbolic in the sense of Gromov. As we have seen, a hyperbolic space can be equipped with a boundary at infinity in a natural way; we will show that the boundary of E_Δ is homeomorphic to the set $\partial_\infty \Delta_{res}$ of residual points of $\partial_\infty \Delta$ (recall that a point in $\partial_\infty \Delta$ is called residual if it does not lie on the boundary of any set S_α). The homeomorphism from $\partial_\infty \Delta_{res}$ to $\partial_\infty E_\Delta$ is obtained as the extension to $\partial_\infty \Delta_{res}$ of the natural projection $\Pi_\Delta : \Delta \to E_\Delta$.

If $h : X \to Y$ is as in Theorems C and D then the induced map $h_E : E_X \to E_Y$ on the electric spaces is a quasi-isometry, so by Theorem 3.10 it extends to an *injective* map $h_E : \partial_\infty E_X \to \partial_\infty E_Y$. If $\Pi_X : X \cup \partial_\infty X_{res} \to \overline{E}_X$ and $\Pi_Y : Y \cup \partial_\infty Y_{res} \to \overline{E}_Y$ are the projection maps then we have that $h|_{\partial_\infty X_{res}} = \Pi_Y^{-1} \circ h_E \circ \Pi_X|_{\partial_\infty X_{res}}$. So the injectivity of h_E on $\partial_\infty E_X$ tells us that $h|_{\partial_\infty X_{res}}$ is injective; in fact, it tells us more: we will see that if $\xi \in \partial_\infty X_{res}$ then for all $\eta \in \partial_\infty X$, $h(\xi) \neq h(\eta)$. This is the statement of Theorem D.

5.1 Hyperbolicity of the electric space

Suppose $(\Delta, \{S_\alpha\})$ satisfies $(*_1) - (*_3)$. We will begin by proving the following narrowness property for electric quasigeodesic "triangles" in Δ:

Proposition 5.1 *Let $(\Delta, \{S_\alpha\})$ satisfy $(*_1) - (*_3)$. Let a, b and c be points in Δ and let p, q and r be electric (k, μ)-quasigeodesics between a and b, b and c, and a and c, respectively. Then every point on p lies within electric distance $4\delta + 2D(\delta, k, \mu)$ of some point of q or r.*

Proof:

Let u be a point on p. By Proposition 4.3 there is a point u' on $[a, b]$ within electric distance D of u. Since Δ is δ-hyperbolic, there is a point v' on either $[b, c]$ or $[a, c]$ within hyperbolic distance 4δ of u'. Again by Proposition 4.3, there is a point v on either q or r within hyperbolic distance D of v'. Since electric distance is less than or equal to hyperbolic distance, we have $d_{el}(u, v) \leq 4\delta + 2D$. □

E_Δ is not necessarily a geodesic space, so the above proposition is as close as we can get to an analog of Proposition 3.2. However, it is sufficient to show that E_Δ is hyperbolic in the sense of Gromov. Let 0 denote the point in E_Δ that corresponds to the basepoint 0 in Δ. We will denote an electric Gromov product on both Δ and E_Δ by

$$(a|b)_{el} = \frac{1}{2}[d_{el}(a,0) + d_{el}(b,0) - d_{el}(a,b)].$$

A consequence of the previous proposition is the following, which will enable us to prove hyperbolicity:

Proposition 5.2 *Let x and y be points in E_Δ, and let p be a (1,1)-quasigeodesic whose endpoints are x and y. Then there exists a constant S that depends only on the hyperbolicity constant of Δ, such that*

$$d_{el}(p,0) - S \leq (x|y)_{el} \leq d_{el}(p,0) + S.$$

The proof of the above is analogous to the proof when the space in question is geodesic and p is a geodesic rather than a quasigeodesic (see e.g. [17]). This gives us

Theorem 5.3 *Let $(\Delta, \{S_\alpha\})$ satisfy $(*_1) - (*_3)$. The space E_Δ is Gromov-hyperbolic.*

Proof:

Let a, b and c be points in E_Δ, and let p, q and r be $(1,1)$-quasigeodesics between a and b, b and c, and a and c, respectively. Let $u \in p$ be such that $d(0,u) = d(0,p)$. By Proposition 5.1 there is a point v in q or r (say q) such that $d_{el}(u,v) \leq 4\delta + 2D(\delta,1,1)$. We have

$$(a|b)_{el} \geq d(p,0) - S = d(u,0) - S \geq d(v,0) - (4\delta + 2D) - S \geq d(q,0) - (4\delta + 2D + S)$$

$$\geq (b|c)_{el} - (4\delta + 2D + 2S).$$

Thus $(a|b)_{el} \geq \min\{(b|c)_{el}, (a|c)_{el}\} - (4\delta + 2D + 2S)$. $\Box$

5.2 The boundary of the electric space

We can now prove that the boundary at infinity of E_Δ is naturally homeomorphic to the space of residual points in $\partial_\infty \Delta$. We will need the following simple lemma:

Lemma 5.4 *Let (v_n) be a sequence in Δ that converges to a residual point ξ in $\partial_\infty \Delta$. Then $d_{el}(0, v_n) \to \infty$ as $n \to \infty$.*

Proof:

Let $r_n : [0, \infty] \to \Delta$ be the map that sends the interval $[0, d(0, v_n)]$ isometrically onto the geodesic segment $[0, v_n]$, and sends the interval $[d(0, v_n), \infty]$ onto v_n. By Ascoli's theorem, a subsequence of (r_n), which we will again call (r_n), converges to a geodesic ray r from 0 to ξ, in the topology of uniform convergence on compact sets.

Suppose that $d_{el}(0, v_n)$ does not tend to infinity as $n \to \infty$; then there is a constant $B < \infty$ and a subsequence of (v_n), which we will again call (v_n), such that $d_{el}(0, v_n) \leq B$ for all n.

Claim: there is a constant β such that if w is a point in the image of any ray r_i then $d_{el}(0, w) < \beta$.

This follows from the fact the d_{el} is roughly monotonic on any geodesic segment. More precisely, let (χ_k) be the sequence of sets associated to r_i in Theorem 4.5, and let n be such that v_i (the endpoint of r_i) is contained in χ_n^A. Then $d_{el}(0, \chi_n) \leq B + A$, so since $d_{el}(0, \chi_n) \geq a \cdot n - b$, we have that $n \leq \frac{B+A+b}{a}$. The point w, being before v_i as we travel along r_i away from 0, is contained in χ_m^A for some $m \leq n$. By Theorem 4.5 we have $d_{el}(0, \chi_m) \leq a'(\frac{B+A+b}{a})$, and $d_{el}(0, w) \leq a'(\frac{B+A+b}{a}) + A$.

So if we assume that $d_{el}(0, v_n)$ does not tend to infinity as $n \to \infty$ then we must have a sequence of rays (r_n) converging to a ray r from 0 to ξ, and for all n we have that every point on r_n is within a uniformly bounded electric distance of 0. But then every point on r must be within a bounded electric distance of 0; this is impossible by Theorem 4.5, since ξ is a residual point in $\partial_\infty \Delta$. $\square$

We are now able to prove

Theorem 5.5 *Let $(\Delta, \{S_\alpha\})$ satisfy $(*_1) - (*_3)$. Then a sequence (x_n) in Δ converges to a point in $\partial_\infty \Delta_{res}$ if and only if the sequence $(\Pi_\Delta(x_n))$ converges to a point in $\partial_\infty E_\Delta$, so that Π_Δ extends to a homeomorphism of $\partial_\infty \Delta_{res}$ with $\partial_\infty E_\Delta$.*

Proof:

First let (x_n) be a sequence of points in Δ that tend toward a point ξ in $\partial_\infty \Delta_{res}$. Since (x_n) is a convergent sequence, we have that $lim_{m,n \to \infty}(x_m | x_n) = \infty$. We wish to show that the sequence $(\Pi_\Delta(x_n))$ converges to a unique point in $\partial_\infty E_\Delta$.

Suppose instead that (x_n) contains two subsequences (y_i) and (z_j) such that the sequences $(\Pi_\Delta(y_i))$ and $(\Pi_\Delta(z_j))$ do not converge to the same point in $\partial_\infty E_\Delta$, or do not even tend towards the boundary. Then for some constant M we have

$$\liminf_{i,j \to \infty} (\Pi_\Delta(y_i) | \Pi_\Delta(z_j))_{el} = M < \infty.$$

For every i and j let e_{ij} be an electric $(1,1)$-quasigeodesic from y_i to z_j. Then we have

$$\liminf_{i,j \to \infty} d_{el}(0, e_{ij}) \leq M + S,$$

where S is the constant from Proposition 5.2. So there exist infinitely many pairs (i, j), with $i, j \to \infty$, for which there exists a point u_{ij} on e_{ij} such that $d_{el}(0, u_{ij}) \leq M + S + 1$.

By Proposition 4.3, for every such point u_{ij} there is a point v_{ij} on $[y_i, z_j]$ such that $d(u_{ij}, v_{ij}) \leq D(\delta, 1, 1)$; hence $d_{el}(0, v_{ij}) \leq M + S + 1 + D$.

The sequences (y_i) and (z_j) both tend towards ξ, so (v_{ij}) also tends towards ξ as $i, j \to \infty$ (see Figure 11). But this gives a contradiction, by the previous lemma. Thus, such subsequences (y_i)

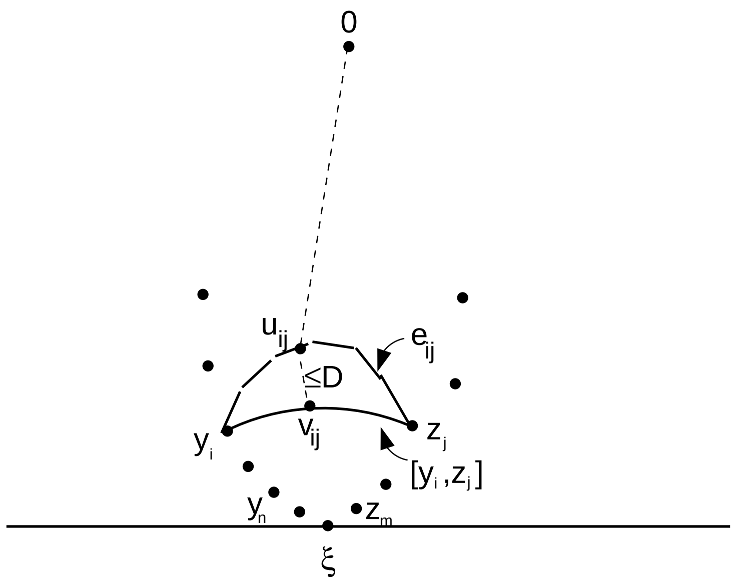

Figure 11: *As $i, j \to \infty$, the point u_{ij} stays within bounded electric distance of* 0*; hence so does v_{ij}. This is impossible since $v_{ij} \to \xi$, a residual point.*

and (z_j) cannot exist. Hence the map Π_Δ extends continuously to ξ.

We must now show that if (x_n) is a sequence in Δ such that $(\Pi_\Delta(x_n))$ converges to a point $\eta \in \partial_\infty E_\Delta$ then (x_n) converges to a unique point ξ, and that ξ is in $\partial_\infty \Delta_{res}$. Since the sequence $(\Pi_\Delta(x_n))$ converges to a boundary point, we have that

$$\lim_{m,n\to\infty} (\Pi_\Delta(x_m)|\Pi_\Delta(x_n))_{el} = \infty.$$

Suppose that (x_n) does not converge to a unique point $\xi \in \partial_\infty \Delta_{res}$. $\overline{\Delta}$ is compact, so a subsequence of (x_n) must converge in $\overline{\Delta}$. We must consider three possibilities:

(1) A subsequence (y_i) of (x_n) converges to a point ρ within Δ.

(2) A subsequence (y_i) of (x_n) converges to a boundary point ρ that is not a residual point (so that it lies in the boundary at infinity of some electric set S_α).

(3) There exist subsequences (y_i) and (z_j) that converge to two distinct points ρ and σ of $\partial_\infty \Delta_{res}$.

Case (1) cannot occur, since if $\rho \in \Delta$ then Π_Δ is defined and continuous at ρ and maps ρ to a point in E_Δ, not to η in $\partial_\infty E_\Delta$.

Suppose case (2) occurs. We will show that $\liminf_{i,j\to\infty}(\Pi_\Delta(y_i)|\Pi_\Delta(y_j))_{el} < M < \infty$ for some M, which will give a contradiction since (y_i) is a subsequence of (x_n). In fact, we will show that there exists a constant $M < \infty$ such that for every $i \in \mathbf{Z}_+$, there are infinitely many $j \in \mathbf{Z}_+$ such that $(\Pi_\Delta(y_i)|\Pi_\Delta(y_j))_{el} < M$.

Choose some point w in S_α , and let r be a geodesic ray from w to ρ. S_α is q-quasiconvex, so since $\xi \in \partial_\infty S_\alpha$, by Proposition 3.7 every point on r is within $q + 8\delta$ of a point on S_α. Fix i, and for each j, let s_j be a geodesic segment from y_i to y_j. By Ascoli's theorem, a subsequence (s_{j_k}) of (s_j) converges to a geodesic ray $s : [0,\infty) \to \Delta$ from y_i to ρ. By Proposition 3.5*(b)*, there exists a real number $T \geq 0$ such that $d(s(t), r) \leq 8\delta$ for all $t \geq T$. So there is some $N \in \mathbf{Z}_+$ such that for all indices $j_k > N$ there is a point z_{j_k} on s_{j_k} that is within $8\delta + 1$ of r, and hence within $16\delta + 1 + q$ of S_α .

Let e_{j_k} be an electric $(1,1)$-quasigeodesic from y_i to y_{j_k}. By proposition 4.3 there is a point z'_{j_k} in $EN_0(e_{j_k})$ that is within a (non-electric) distance $D(\delta, 1, 1)$ of z_{j_k} (see Figure 12).

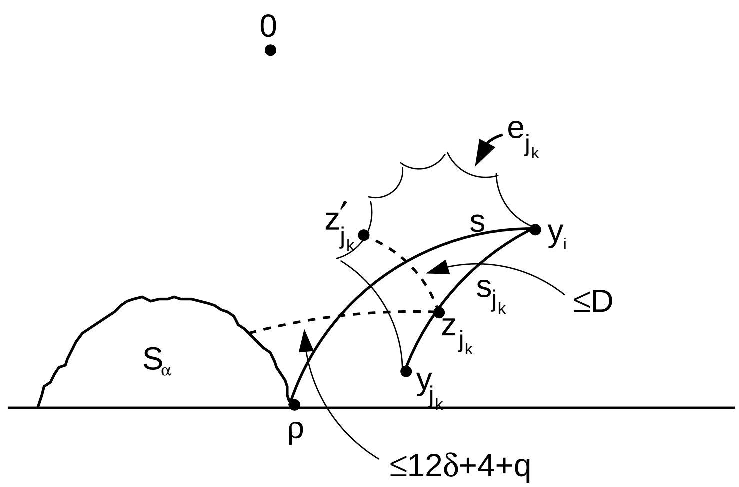

Figure 12: *There is an upper bound on hence also on* $(y_i|y_{j_k})_{el}$.

Now by Proposition 5.2 we have

$$(\Pi_\Delta(y_i)|\Pi_\Delta(y_{j_k}))_{el} \leq d_{el}(e_{j_k}, 0) + S$$

$$\leq d_{el}(z'_{j_k}, 0) + S \leq d_{el}(z'_{j_k}, z_{j_k}) + d_{el}(z_{j_k}, S_\alpha) + d_{el}(S_\alpha, 0) + S$$

$$\leq D + 16\delta + 1 + q + d_{el}(S_\alpha, 0) + S.$$

$d_{el}(S_\alpha, 0)$ is some finite number, so $(\Pi_\Delta(y_i)|\Pi_\Delta(y_{j_k}))_{el}$ is bounded above by a constant that is independent of i and j_k. This gives a contradiction.

Suppose case (3) occurs. Since $\rho \neq \sigma$ we have $\limsup_{i,j\to\infty}(y_i|z_j) < M < \infty$ for some M. We will show that $\limsup_{i,j\to\infty}(\Pi_\Delta(y_i)|\Pi_\Delta(z_j))_{el}$ is finite; this will give a contradiction since (y_i) and (z_j) are subsequences of (x_n). Let s_{ij} be a geodesic segment from y_i to z_j; by Proposition 3.4, $d(0, s_{ij}) \leq (y_i|z_j) + 4\delta$.

Let i and j be sufficiently large that $(y_i|z_j) < M + 1$, so that $d(0, s_{ij}) < M + 1 + 4\delta$; let a_{ij} be a point on s_{ij} such that $d(0, a_{ij}) < M + 1 + 4\delta$. If e_{ij} be an electric $(1, 1)$-quasigeodesic from y_i to z_j then there is a point a'_{ij} in $EN_0(e_{ij})$ that is within hyperbolic distance D of a_{ij}. We have

$$(\Pi_\Delta(y_i)|\Pi_\Delta(z_j))_{el} \leq d_{el}(e_{ij}, 0) + S \leq d_{el}(a'_{ij}, 0) + S$$

$$\leq d_{el}(a'_{ij}, a_{ij}) + d_{el}(a_{ij}, 0) + S \leq d(a'_{ij}, a_{ij}) + d(a_{ij}, 0) + S \leq D + M + 1 + 4\delta + S.$$

This again gives a contradiction.

Thus cases (1)-(3) cannot occur; so if the sequence $(\Pi_\Delta(x_n))$ converges to a point $\eta \in \partial_\infty E_\Delta$ then (x_n) converges to a unique point ξ that lies in $\partial_\infty \Delta_{res}$. $\square$

5.3 The identifications of the map $h : \partial_\infty X \to \partial_\infty Y$

We can now prove Theorem D, which can be written as:

Theorem D *Let $h : \partial_\infty X \to \partial_\infty Y$ be the map constructed in Theorem C. If ξ and ρ are points in $\partial_\infty X$ such that $h(\xi) = h(\rho)$ then ξ and ρ are not residual points.*

Proof:

Let $h(\xi) = h(\rho)$, and suppose that ξ is a residual point. Write $\eta = h(\xi) = h(\rho)$. Let (x_n) and (y_n) be sequences in X that converge to ξ and ρ, respectively. The sequences $(h(x_n))$ and $(h(y_n))$ both converge to η. We will show first that η must lie in $\partial_\infty Y_{res}$.

Since $\xi \in \partial_\infty X_{res}$, by the preceding theorem we have that the sequence $(\Pi_X(x_n))$ converges to a point $\Pi_X(\xi)$ in $\partial_\infty E_X$. The map $h_E : E_X \to E_Y$ induced by h extends continuously to a map $h_E : \partial_\infty E_X \to \partial_\infty E_Y$ by proposition 3.10, so the sequence $h_E(\Pi_X(x_n))$ converges to a point $h_E(\Pi_X(\xi))$ in $\partial_\infty E_Y$. For all $x \in X \cup \partial_\infty X_{res}$, we have $h_E \circ \Pi_X = \Pi_Y \circ h$. So for all n, $h_E(\Pi_X(x_n)) = \Pi_Y(h(x_n))$. So $\Pi_Y(h(x_n))$ converges to a point in $\partial_\infty E_Y$, and by the previous theorem the limit η of the sequence $(h(x_n))$ must be in $\partial_\infty Y_{res}$.

Now since $\eta \in \partial_\infty Y_{res}$ we have that the sequences $(\Pi_Y(h(x_n)))$ and $(\Pi_Y(h(y_n)))$ both converge to a point $\Pi_Y(\eta)$ in $\partial_\infty E_Y$.

Since $h_E : E_X \to E_Y$ is a quasi-isometry, it has a quasi-inverse, that is, a quasi-isometry $g_E : h_E(E_X) \to E_X$ such that for some constant A, $d_{el}(a, g_E \circ h_E(a)) \leq A$ for all $a \in E_X$. Since

g_E is a quasi-isometry it extends continuously to a map from the boundary of $h_E(E_X)$ to $\partial_\infty E_X$. We have then that the sequences $(g_E(\Pi_Y(h(x_n))))$ and $(g_E(\Pi_Y(h(y_n))))$ both converge to the point $g_E(\Pi_Y(\eta))$ in $\partial_\infty E_X$.

Now since $\Pi_Y(h(x_n)) = h_E(\Pi_X(x_n))$, we have for all n

$$d(\Pi_X(x_n), g_E(\Pi_Y(h(x_n)))) = d(\Pi_X(x_n), g_E \circ h_E(\Pi_X(x_n)) \leq A,$$

and likewise

$$d(\Pi_X(y_n), g_E(\Pi_Y(h(y_n)))) \leq A.$$

So the sequences $(\Pi_X(x_n))$ and $(\Pi_X(y_n))$ also converge to $g_E(\Pi_Y(\eta))$ by Proposition 3.5*(a)*.

But then by the previous theorem the sequences (x_n) and (y_n) must converge to the same point in $\partial_\infty X_{res}$. This is a contradiction, since $\xi \neq \rho$. $\square$

In the Kleinian groups problem, we know that the semiconjugacy $\tilde{g} : \hat{\mathbf{C}} \to \hat{\mathbf{C}}$ from the action of Γ_0 to the action of Γ identifies points on the closure of each component Ω_α of the domain of discontinuity of Γ_0 according to the ending lamination $\tilde{\lambda}_\alpha$; we wish to show that these are the only identifications that occur. Since the sets $\overline{\Omega}_\alpha$ can intersect in their boundaries, one obstacle to proving this is the possibility that the intersection of two sets $\overline{\Omega}_\alpha$ and $\overline{\Omega}_\beta$ might contain a point that is on the closure of a leaf of $\tilde{\lambda}_\alpha$ and also on a leaf of $\tilde{\lambda}_\beta$; then all the points on these two leaves will be identified by $\tilde{g}$. Points on different sets $\overline{\Omega}_\alpha$ and $\overline{\Omega}_\beta$ would clearly also be identified if there were a chain of components of the domain of discontinuity $\overline{\Omega}_0, ..., \overline{\Omega}_m$, with $\overline{\Omega}_0 = \overline{\Omega}_\alpha$ and $\overline{\Omega}_m = \overline{\Omega}_\beta$, and leaves $l_0, ..., l_m$ of the ending laminations of $\overline{\Omega}_0, ..., \overline{\Omega}_m$, respectively, such that for each i, one endpoint of l_i is an endpoint of l_{i-1} and the other endpoint of l_i is an endpoint of l_{i+1}, as in Figure 13 (we will show later in this section that in fact this cannot occur). Our next

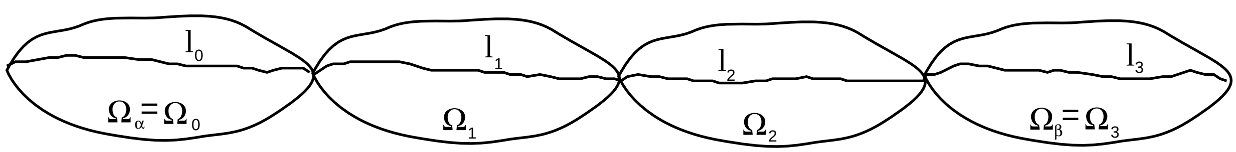

Figure 13: *If this configuration were to occur then the map $\tilde{g}$ would identify points in $\hat{\mathbf{C}}$ that do not lie on the same set $\overline{\Omega}_\alpha$.*

theorem shows that this is the only obstacle to proving Theorem B. The only part of the proof that is specific to the setting of the Kleinian groups problem is the fact that the map $\tilde{g} : \mathbf{H^3} \to \mathbf{H^3}$ of Theorem A descends to a map of electric spaces that is a bi-Lipschitz homeomorphism, not just a quasi-isometry. In fact, if $h : X \to Y$ is any map satisfying the conditions of Theorem C such that the map $h_E : E_X \to E_Y$ of electric spaces is a bi-Lipschitz homeomorphism, we can say the following:

Theorem E *Let X, Y, $\{H_\alpha\}$ and $h : X \to Y$ be as in Theorem C, and suppose that the induced map $h_E : E_X \to E_Y$ on electric spaces is a bi-Lipschitz homeomorphism. If ξ and ρ are points in $\partial_\infty X$ such that $h(\xi) = h(\rho)$ then there is a finite chain of electric sets $\{H_0, ...H_m\}$ such that*

(1) $\xi \in \partial_\infty H_0$ and $\rho \in \partial_\infty H_m$.

(2) For $i = 0, ..., m-1$ there is a point $p_i \in \partial_\infty H_i \cap \partial_\infty H_{i+1}$ such that $h(p_i) = h(\xi) = h(\rho)$.

Proof:

Write $\eta = h(\xi) = h(\rho)$. Since $h(\xi) = h(\rho)$, by Theorem *D* there exist electric sets H_α and H_β such that $\xi \in \partial_\infty H_\alpha$ and $\rho \in \partial_\infty H_\beta$. Choose a point $u \in h(H_\alpha)$ and a point $v \in h(H_\beta)$; the rays $[u, \eta)$ and $[v, \eta)$ terminate at the same point on $\partial_\infty Y$, so by Proposition 3.5*(b)* there are sequences (x'_n) on $[u, \eta)$ and (y'_n) on $[v, \eta)$ that tend towards η and such that $d(x'_n, y'_n) \leq 8\delta$. The sets $h(H_\alpha)$ and $h(H_\beta)$ are q-quasiconvex, so every point on $[u, \eta)$ or $[v, \eta)$ is within $q + 8\delta$ of a point on $h(H_\alpha)$ or $h(H_\beta)$, respectively, by Proposition 3.7; hence there are sequences (x_n) and (y_n) on $h(H_\alpha)$ and $h(H_\beta)$, respectively, that tend towards η and such that $d(x_n, y_n) \leq q + 16\delta$. Let s_n be the geodesic segment from x_n to y_n. Let a_n be a point in H_α that is mapped by h to x_n, and b_n a point in H_β that is mapped to y_n; we will construct a finite chain of points $\{a_{n0}, b_{n0}, a_{n1}, b_{n1}, ..., a_{nM}, b_{nM}\}$ from a_n to b_n whose behavior we can control.

Set $a_{n0} = a_n$. If $H_\alpha \neq H_\beta$, let s_{n1} be the component of $s_n \cap h(\mathcal{C})$ that begins at the last point in $s_n \cap h(H_\alpha)$, traveling towards y_n (recall that $\mathcal{C}$ is the complement in X of the electric sets). The map h restricted to $\mathcal{C}$ is a homeomorphism, so we may pull s_{n1} to obtain a path $h^{-1}(s_{n1})$ in $\mathcal{C}$ whose initial endpoint lies in H_α and whose terminal endpoint lies in some other electric set H_{n1}; let b_{n0} denote the initial endpoint and a_{n1} the terminal endpoint of $h^{-1}(s_{n1})$ (see Figure 14). If $H_{n1} \neq H_\beta$, let s_{n2} be the component of $s_n \cap h(\mathcal{C})$ that begins at the last point in $s_n \cap h(H_{n1}$,

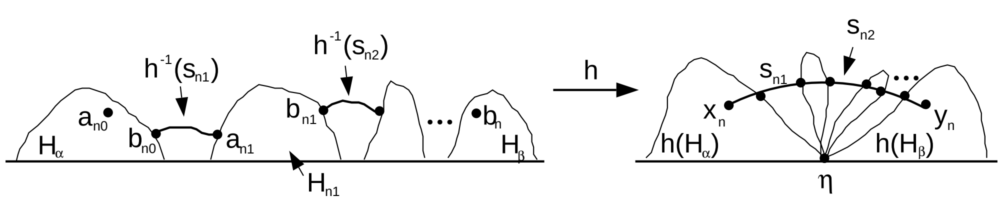

Figure 14: *The construction of the points $a_{n0}, b_{n0}, a_{n1}, b_{n1}, ..., a_{nk}, b_{nk}$.*

traveling towards y_n. The path s_{n2} pulls back to a path $h^{-1}(s_{n2})$ whose initial endpoint lies in H_{n1} and whose terminal endpoint lies in some new electric set H_{n2}; let b_{n1} denote the initial endpoint and a_{n2} the terminal endpoint of $h^{-1}(x_{n2})$. Continue constructing the points a_{ni} and b_{ni} in this manner; the process with terminate when we have reached a point a_{nk} that lies on H_β; set $b_{nk} = b_n$.

Since there is a uniform upper bound on the lengths of the paths s_n, there is (trivially) a uniform upper bound on the total lengths of the paths s_{ni} for each n; each such path must have length at least c (the minimum separation between any two electric sets), so it follows that there is a uniform upper bound on the number of paths s_{ni}, for each n. So since the map $h_E : E_X \to E_Y$ is a bi-Lipschitz homeomorphism, we have the following:

(1) There is some constant M such that for all n, the above construction terminates after producing at most M points a_{ni} and M points b_{ni}.

(2) There is an upper bound B (independent of n and i) on $d(b_{ni}, a_{n(i+1)})$.

We wish to produce collections $\{a_{n0}, b_{n0}, a_{n1}, b_{n1}, ..., a_{nk}, b_{nk}\}$ for each n that all have the same number of elements; so if $k < M$, set $a_{n(k+1)} = b_{n(k+1)} = a_{n(k+2)} = ... = a_{nM} = b_{nM} = b_n$.

Choose some subsequence $\{n_k\}$ of $\mathbf{Z}^+$ such that for all i between 0 and M, the sequences $(a_{n_k i})$ and $(b_{n_k i})$ converge in $\overline{X}$ to points p_i^a and p_i^b, respectively. By our construction we have $h(p_i^a) = h(p_i^b) = \eta$ for all $i = 0, 1, ..., M$, so that in particular the points p_i^a and p_i^b all lie in $\partial_\infty X$ and, by Theorem D, are not residual points.

For each i and each n_k, we have that $d(b_{n_k i}, a_{n_k(i+1)}) < B$; hence by Proposition 3.5(a) the sequences $(b_{n_k i})$ and $(a_{n_k(i+1)})$ converge to the same point in $\partial_\infty X$ as $n_k \to \infty$, that is, $p_{i-1}^b = p_i^a$. Set $p_i = p_i^a$; we will be done if we can show that for $i - 0, ..., M$, p_i and p_{i+1} lie in the boundary of the same electric set, since we already know by construction that p_0 lies in $\partial_\infty H_\alpha$ and p_M lies in H_β.

Fix i; we know that $p_i = \lim_{n_k \to \infty} a_{n_k i}$ and $p_{i+1} = \lim_{n_k \to \infty} b_{n_k i}$. For each n_k, $a_{n_k i}$ and $b_{n_k i}$ lie on the same electric set $H_{n_k} i$. There are two possibilities for the sequence $(H_{n_k} i)$: either infinitely many of the sets $H_{n_k i}$ are the same, or there is some infinite subsequence of $(H_{n_k i})$ whose elements are all distinct. Suppose first that infinitely many of the sets $H_{n_k i}$ are the same; call this electric set H_i. Then p_i and p_{i+1} both lie on H_i, and we are done. Otherwise there is some subsequence of $(H_{n_k i})$, which we will again call $(H_{n_k i}$, whose elements are all distinct. Since there is a minimum separation between any two electric sets, for every $N < \infty$ there is an upper bound on the number of electric sets that intersect $B(0, N)$ in X; hence as $n_k \to \infty$, $d(0, H_{n_k}) \to \infty$. The electric sets are uniformly quasiconvex, so by Proposition 3.9 we have $diam_{vis}(H_{n_k}) \to 0$ as $n_k \to \infty$. Thus $d_{vis}(a_{n_k i}, b_{n_k i}) \to 0$ as $n \to \infty$; so in this case the sequences $(a_{n_k i})$ and $(b_{n_k i})$ converge to the same point in ∂X, that is, $p_i = p_{i+1}$. So since $p_i = p_{i+1}$ is not a residual point, we certainly have again that p_i and p_{i+1} lie in the boundary of the same electric set. $\Box$

This enables us to prove

Theorem B *Let $\tilde{g} : \hat{\mathbf{C}} \to \hat{\mathbf{C}}$ be the map constructed in Theorem A. If x and y are in $\hat{\mathbf{C}}$ then $\tilde{g}(x) = \tilde{g}(y)$ iff for some Ω_α, x and y lie on the closure of some leaf or complementary component of $\tilde{\lambda}_\alpha$.*

Proof:

If the domain of discontinuity Ω of Γ_0 is empty or has two components, we have already seen that the statement of the theorem holds. For the remainder of the proof, we assume that Ω has infinitely many components (the only remaining possibility, by Lemma 2.4). In this case we have seen that the map $\tilde{g} : (\mathbf{H^3}, \{H_\alpha\}) \to (\mathbf{H^3}, \{\tilde{g}(H_\alpha)\})$ satisfies the conditions of Theorems C, D and E (recall that the sets H_α are the components of the complement of the convex hull of Γ_0, and each set H_α meets the boundary $\hat{\mathbf{C}}$ of $\mathbf{H^3}$ in the closure of a component Ω_α of Ω).

We wish to show that is $x \in \overline{\Omega}_\alpha$ and $y \in \overline{\Omega}_\beta$ where Ω_α and Ω_β are distinct components of Ω, then $\tilde{g}(x) \neq \tilde{g}(y)$. In view of Theorem E and the discussion immediately preceding it, it will suffice to show that if $\xi \in \overline{\Omega}_\alpha \cap \overline{\Omega}_\beta$ then ξ is not an endpoint of any leaf of the ending laminations $\tilde{\lambda}_\alpha$ and $\tilde{\lambda}_\beta$.

So let $\xi \in \overline{\Omega}_\alpha \cap \overline{\Omega}_\beta$; we will show that ξ is not an endpoint of a leaf of $\tilde{\lambda}_\alpha$. Let $\Gamma_\alpha \subset \Gamma_0$ and $\Gamma_\beta \subset \Gamma_0$ be the stabilizer subgroups of Ω_α and Ω_β, respectively, so that the boundary circle of Ω_α is the limit set of Γ_α; likewise for Ω_β. Write $\Gamma_{\alpha\beta} = \Gamma_\alpha \cap \Gamma_\beta$. By work of Susskind (see [29]), the intersection $\Lambda_{\Gamma_\alpha} \cap \Lambda_{\Gamma_\beta}$ of the limit sets of Γ_α and Γ_β is the limit set $\Lambda_{\Gamma_{\alpha\beta}}$ of $\Gamma_{\alpha\beta}$, and $\Gamma_{\alpha\beta}$ is geometrically finite.

Equip Ω_α with a Poincaré metric and realize the leaves of $\tilde{\lambda_\alpha}$ as geodesics. Let $CH(\Lambda_{\Gamma_{\alpha\beta}})$ denote the convex hull in Ω_α of $\Lambda_{\Gamma_{\alpha\beta}}$. Since $\Gamma_{\alpha\beta}$ is geometrically finite, $CH(\Lambda_{\Gamma_{\alpha\beta}})/\Gamma_{\alpha\beta}$ is a compact surface. We will show that the quotient of $CH(\Lambda_{\Gamma_{\alpha\beta}})$ by Γ_α is the same as its quotient by $\Gamma_{\alpha\beta}$, so that $CH(\Lambda_{\Gamma_{\alpha\beta}})/\Gamma_\alpha$ is a compact subsurface of $S_e = \Omega_\alpha/\Gamma_\alpha$.

Let $\gamma \in \Gamma_\alpha$. If $\gamma \in \Gamma_{\alpha\beta}$ then $\gamma(CH(\Lambda_{\Gamma_{\alpha\beta}})) = CH(\Lambda_{\Gamma_{\alpha\beta}})$. If γ is not in $\Gamma_{\alpha\beta}$ then we claim that $\gamma(CH(\Lambda_{\Gamma_{\alpha\beta}})) \cap CH(\Lambda_{\Gamma_{\alpha\beta}}) = \emptyset$. If $\gamma(CH(\Lambda_{\Gamma_{\alpha\beta}})) \cap CH(\Lambda_{\Gamma_{\alpha\beta}})$ is nonempty then there must be points a and b in $\Lambda_{\Gamma_{\alpha\beta}}$ and c and d in $\gamma(\Lambda_{\Gamma_{\alpha\beta}})$ such that c and d separate a and b on the boundary circle of Ω_α. But then since a and b lie on the boundary of Ω_β and c and d lie on the boundary of $\gamma(\overline{\Omega}_\beta)$, the sets $\overline{\Omega}_\beta$ and $\gamma(\overline{\Omega}_\beta)$ must intersect (see Figure 15); this is impossible since γ is not in Γ_β, the stabilizer of Ω_β.

Thus the quotient of $CH(\Lambda_{\Gamma_{\alpha\beta}})$ by Γ_α is the same as its quotient by $\Gamma_{\alpha\beta}$, so that $CH(\Lambda_{\Gamma_{\alpha\beta}})/\Gamma_\alpha$ is a compact subsurface of S_e with geodesic boundary. Hence the boundary of $CH(\Lambda_{\Gamma_{\alpha\beta}})/\Gamma_{\alpha\beta}$ in S_e consists of a finite collection of simple closed geodesics.

Now suppose $\xi \in \overline{\Omega}_\alpha \cap \overline{\Omega}_\beta$ and ξ lies on the closure of a leaf of $\tilde{\lambda}_\alpha$; then ξ is the endpoint of a geodesic half-leaf $\tilde{l}$ of $\tilde{\lambda_\alpha}$. Let l denote the projection of $\tilde{l}$ to S_e; l is a half-leaf of λ_e.

Since ξ is in $\Lambda_{\Gamma_{\alpha\beta}}$, either:

(1) $\tilde{l}$ is asymptotic to a boundary component of $CH(\Lambda_{\Gamma_{\alpha\beta}})$; or

(2) $\tilde{l}$ is eventually contained in $CH(\Lambda_{\Gamma_{\alpha\beta}})$. (see Figure 16).

In either case, if c is a simple closed curve in the boundary of $CH(\Lambda_{\Gamma_{\alpha\beta}})/\Gamma_\alpha$, c intersects l transversely at most once; so since l is dense in λ_e, in fact λ_e cannot ever intersect c transversely. But then either c is contained in λ_e, which is impossible since λ_e contains no closed curves, or c is disjoint from λ_e, which is also impossible since λ_e is maximal. So we have arrived at a contradiction. □

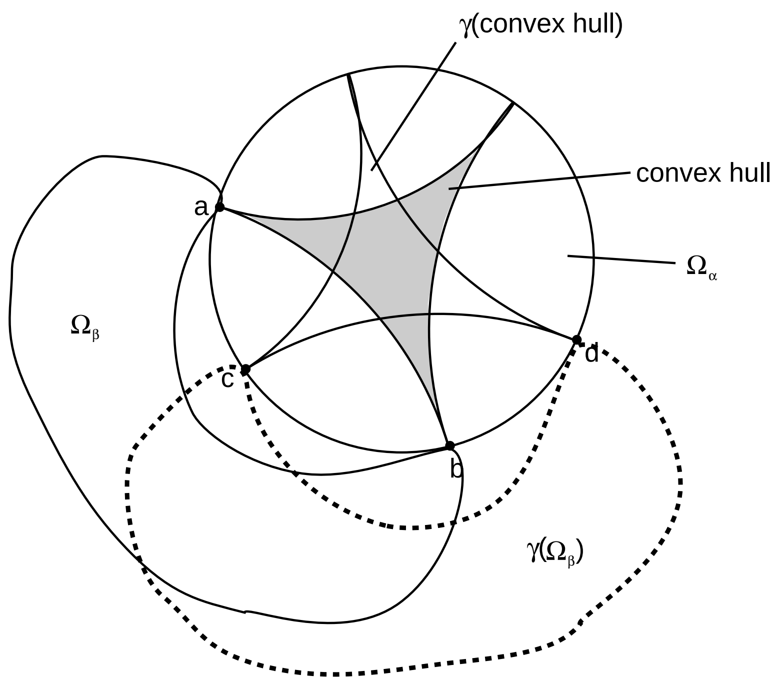

Figure 15: *The shaded region denotes the convex hull of $\Lambda_{\Gamma_{\alpha\beta}}$. If the points c and d separate a and b then Ω_β must intersect $\gamma(\Omega_\beta)$.*

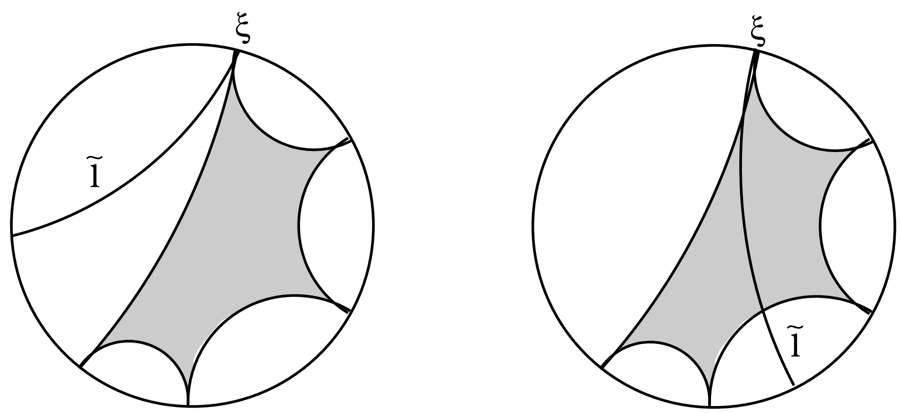

Figure 16: *The two possibilities for $\tilde{l}$, if $\xi \in \overline{\Omega}_\alpha \cap \overline{\Omega}_\beta$.*

References

[1] W. Abikoff. *Kleinian groups – geometrically finite and geometrically perverse.* Geometry of Group Representations, Comtemporary Math., **74**. Amer. Math. Soc., Providence, RI (1988), 11-50.

[2] W. Abikoff. *On boundaries of Teichmüller spaces and on Kleinian groups III.* Acta. Math., **134** (1975), 211-237.

[3] L. Bers. *On boundaries of Teichmüller spaces and on Kleinian groups I.* Ann. of Math., **91** (1970), 570-600.

[4] F. Bonahon. *Bouts des variétès hyperboliques de dimension 3.* Ann. of Math. (2) **124** (1986), 71-158.

[5] R. D. Canary. *Ends of hyperbolic 3-manifolds.* J. Amer. Math. Soc. **6** (1993), 1-35.

[6] R. D. Canary. *A covering theorem for hyperbolic 3-manifolds and its applications.* Preprint.

[7] R. D. Canary and E. Taylor. *Kleinian groups with small limit sets.* Duke Math. J. **73** (1994), 371-381.

[8] J. Cannon. *The theory of negatively curved spaces and groups.* Ergodic Theory, Symbolic Dynamics, and Hyperbolic Spaces (Trieste, 1989). Oxford Univ. Press, London and New York (1991), 315-369.

[9] A. J. Casson. *Automorphisms of surfaces after Nielsen and Thurston.* Notes by S. A. Bleiler. U. T. Austin, 1982.

[10] A. J. Casson and S. A. Bleiler. *Automorphisms of surfaces after Nielsen and Thurston.* Cambridge Univ. Press, London and New York, 1988.

[11] M. Coornaert, T. Delzant and A. Papadopoulos. *Géométrie et theorie des groupes: les groupes hyperboliques de Gromov.* Springer-Verlag, Berlin and New York, 1990.

[12] J. Cannon and W. Thurston. *Group invariant peano curves.* Preprint (1989).

[13] A. Douady. *Descriptions of compact sets in* $\mathbf{C}$. Topological Methods in Modern Mathematics (L. Goldberg and A. Phillips, eds.). Publish or Perish, Houston, 1993.

[14] D. B. A. Epstein and A. Marden. *Convex hulls in hyperbolic space, a theorem of Sullivan, and measured pleated surfaces.* Analytical and Geometric Aspects of Hyperbolic Space. London Math. Soc. Lecture Note Ser., no. 111, Cambridge Univ. Press, Cambridge and New York, 1987, 113-254.

[15] B. Farb. *Relatively hyperbolic and automatic groups with applications to negatively curved manifolds.* Ph.D. thesis, Princeton Univ., 1994.

[16] W. J. Floyd. *Group completions and limit sets of Kleinian groups.* Invent. Math. **57** (1980), 205-218.

[17] E. Ghys and P. de la Harpe (eds). *Sur les groupes hyperboliques d'aprés Mikhael Gromov.* Birkhaüser, Basel, 1990.

[18] M. Gromov. *Hyperbolic Groups.* Essays in Group Theory (S. M. Gersten, ed.). Math. Sci. Res. Inst. Publ. no 8, Springer-Verlag, Berlin and New York (1987).

[19] J. Lott. L^2*-cohomology of geometrically infinite hyperbolic 3-manifolds.* Preprint.

[20] B. Maskit. *On boundaries of Teichmüller spaces and on Kleinian groups II.* Ann. of Math., **91** (1970), 607-639.

[21] B. Maskit. *On extended quasifuchsian groups.* Ann. Acad. Scient. Fenn., Series A. I. Math., **15** (1990), 53-64.

[22] C. McMullen. *Riemann surfaces and the geometrization of 3-manifolds.* Bull. Amer. Math. Soc **27** (1992) no. 2, 207-216.

[23] Y. Minsky. *Teichmüller geodesics and ends of hyperbolic 3-manifolds.* Topology **32** (1993), 625-647.

[24] Y. Minsky. *On Rigidity, Limit Sets, and End Invariants of Hyperbolic 3-Manifolds.* J. Amer. Math. Soc. **7** (1994), 539-588.

[25] M. Mitra. *Cannon-Thurston maps for graphs of hyperbolic groups.* Preprint.

[26] J. Morgan. *On Thurston's uniformization theorem for three-dimensional manifolds.* The Smith Conjecture. Academic Press (1984), 37-125.

[27] G. D. Mostow. *Quasiconformal mappings in n-space and the rigidity of hyperbolic space forms.* Inst. Hautes Études Sci. Publ. Math. **34** (1968), 53-104.

[28] G. D. Mostow. *Strong rigidity of locally symmetric spaces.* Ann. of Math. Stud., no. 78, Princeton Univ. Press, Princeton, NJ, 1973.

[29] P. Susskind. *Kleinian groups with intersecting limit sets.* J. d'Anal. Math. **52** (1989), 26-38.

[30] W. Thurston. *The geometry and topology of 3-manifolds.* Princeton University lecture notes (1982).

[31] W. Thurston. *Hyperbolic structures on 3-manifolds, II: surface groups and manifolds which fiber over the circle.* Preprint.

Department of Mathematics, University of Michigan, East Hall, Ann Arbor, MI 48109-1109; klarreic@math.lsa.umich.edu