STABILITY OF HYPERBOLIC GROUPS ACTING ON THEIR BOUNDARIES

KATHRYN MANN, JASON FOX MANNING, AND THEODORE WEISMAN

Abstract. A hyperbolic group Γ acts by homeomorphisms on its Gromov boundary. We use a dynamical coding of boundary points to show that such actions are *topologically stable* in the dynamical sense: any nearby action is semi-conjugate to (and an extension of) the standard boundary action.

1. Introduction

A discrete, hyperbolic group Γ, viewed as a (coarse) metric space, admits a natural compactification by its Gromov boundary, denoted $\partial\Gamma$. This boundary is a compact metrizable space, and if Γ is not virtually cyclic, it has no isolated points. The action of Γ on itself by left-multiplication extends naturally to an action on $\partial\Gamma$ by homeomorphisms, with rich global dynamics. In fact, boundary actions of hyperbolic groups are entirely characterized by their dynamics, following work of Bowditch: a group action by homeomorphisms on a perfect compact metrizable space X is conjugate to the action of a hyperbolic group on its boundary if and only if the action has convergence group dynamics and each point in X is a conical limit point [Bow98, Tuk98]. (Equivalently, the induced action on the space of distinct triples in X is properly discontinuous and cocompact.)

This paper considers another dynamical question with a long history, namely stability of these actions under perturbation. The general study of stability of actions of groups on boundaries, in an algebraic rather than dynamical context, dates back to Mostow and Furstenburg. On the dynamical side, Sullivan [Sul85] showed that the actions of convex cocompact Kleinian groups on their limit sets are stable under C^1–small perturbation; his techniques were generalized to prove stability under "Lipschitz-small" perturbations for a much broader class of group actions, including those of hyperbolic groups on their boundaries, in [KKL].

Here we consider the more general question of C^0–small perturbations and stability in the sense of topological dynamics. Recall that a representation $\rho_0 \in \mathrm{Hom}(\Gamma, \mathrm{Homeo}(X))$ is said to be a *topological factor* of another such representation ρ if there exists a continuous, surjective map $h : X \to X$ such that $h \circ \rho(\gamma) = \rho_0(\gamma) \circ h$ for all $\gamma \in \Gamma$. In this case, ρ is said to be an *extension* of ρ_0, and the map h is called a *semi-conjugacy* of the two actions. An action ρ is said to be *topologically* or *C^0–stable* if every nearby action is an extension of it, or in other words, nearby actions encode the same dynamical information as ρ.

We prove the following.

Theorem 1.1. *Let Γ be a hyperbolic group. For any neighborhood $\mathcal{U}$ of the identity in the space of continuous self-maps of $\partial\Gamma$, there exists a neighborhood $\mathcal{V}$ of the standard boundary action in $\mathrm{Hom}(\Gamma, \mathrm{Homeo}(\partial\Gamma))$ such that any $\rho \in \mathcal{V}$ is an extension of the standard boundary action, via a semi-conjugacy in $\mathcal{U}$.*

In particular, Theorem 1.1 implies that actions C^0–close to the standard boundary action cannot have larger kernel than the standard action. This kernel is finite as long as Γ is not virtually cyclic.

C^0–stability of boundary actions was previously proved for the special case of fundamental groups of compact Riemannian manifolds of negative curvature in [BM19], and for hyperbolic groups with sphere boundary in [MM21]. The strategy in both was to translate perturbations of actions into nice maps between foliated spaces, translating the dynamical problem into a geometric one. Here, we use a different approach, inspired by the proof of a stability property for relative Anosov representations given in [Wei22]. This approach more closely follows the idea of *dynamical coding* of boundary points originally employed by Sullivan, and involves the construction of an automaton that outputs quasi-geodesic strings. While the coding we construct is tailored to proving stability of the action, the use of automata in the study of hyperbolic groups has a long history dating back to Cannon [Can84]; see [ECH$^+$92, Ch.3] for an introduction.

1.1. **Necessity of semi-conjugacy.** One cannot hope to improve the semi-conjugacy in Theorem 1.1 to a genuine conjugacy without stronger restrictions on the perturbation. One can easily build examples of arbitrarily small perturbations of the action of a free group on its boundary which are not conjugate to the original action, as in the following example.

Example 1.2 (Perturbation of action of F_2)**.** Let F_2 be the free group on the letters a and b. The boundary ∂F_2 is a Cantor set. We start by modifying the action of a in a small clopen neighborhood N of its attracting fixed point x_+, as follows. Pick a fundamental domain D for the standard action, so that the sets $a^k(D)$, $k = 0, 1, 2\ldots$ partition $N - \{x_+\}$ into countably many disjoint, clopen sets. Let $X \subset N$ be a proper clopen neighborhood of x_+, and let $x' \in N - X$. Let $\{D_k\}_{k=0}^{\infty}$ be a partition of $N - (X \cup \{x'\})$ into countably many clopen subsets of decreasing diameter accumulating to x', and set $D_{-1} = a^{-1}D$. Define a modified action of a on N as follows. On $\partial F_2 - (N \cup a^{-1}D)$, the action is unchanged. For each $k \geq -1$ define the restriction of a to D_k to be a homeomorphism to D_{k+1}. On X, define a to be the identity. This gives a homeomorphism of the Cantor set which is not conjugate to the original action of a, since it has uncountable fixed point set. One may extend this action of a to an action of F_2 by inserting the standard action of b (or a similar modification if desired).

Examples of non-conjugate perturbations of actions of Kleinian groups with sphere boundary are given in [BM19].

1.2. **Relatively hyperbolic groups.** It should be possible to use our methods to prove a relative version of Theorem 1.1, where Γ is replaced by a relatively hyperbolic group and $\partial\Gamma$ is replaced by the Bowditch boundary. In fact much of [Wei22] (adapted in this paper) takes place in this setting.

However, the action of a relatively hyperbolic group on its Bowditch boundary is not stable in general. That is, if we replace "hyperbolic" with "relatively hyperbolic" and "boundary" with "Bowditch boundary" in Theorem 1.1, the statement is no longer true. In fact in cases where the boundary has a C^1 structure, the statement can even fail for C^1–perturbations. As an example, consider the action of the fundamental group of a finite volume non-compact hyperbolic 3-manifold M on the ideal boundary of $\mathbb{H}^3$ (which is equivariantly identified with the Bowditch boundary of $\pi_1 M$). By Thurston's hyperbolic Dehn filling theorem, there are arbitrarily small C^1–deformations of the action of $\pi_1 M$ on $\partial\mathbb{H}^3$ which have infinite kernel—and therefore cannot be semi-conjugate to the original action, where the kernel is trivial. Thus any version of Theorem 1.1 in the relative case must in some way restrict the allowable deformations in $\mathrm{Hom}(\Gamma, \mathrm{Homeo}(\partial\Gamma))$.

Outline. Section 2 describes the process for coding boundary points. In Section 3 we show that coding sequences give quasi-geodesics, establish a technical result describing the relationship between two codings of the same point, and discuss how conjugacy changes codings. We also make some remarks on the relationship between our work and Bowditch's *annulus systems*. Section 4 uses the results of the previous sections to prove Theorem 1.1.

Acknowledgments. K.M. was partially supported by NSF grant DMS 1844516 and a Sloan fellowship. J.M. was partially supported by Simons Collaboration Grant 524176. T.W. was partially supported by NSF grant DMS 1937215.

2. Set-up: coding boundary points

We assume familiarity with the basics of hyperbolic groups and their boundaries. The reader may consult [Gro87, 8.2] or [BH99, III.H] for a general reference. Let Γ be a hyperbolic group, with fixed finite, symmetric, generating set $\mathcal{S}$. Theorem 1.1 is trivially true for two ended (virtually $\mathbb{Z}$) groups, so for the remainder of the paper we assume that Γ is not virtually cyclic. We denote by $\partial\Gamma$ the Gromov boundary of Γ, or equivalently the Gromov boundary of the Cayley graph of Γ with respect to $\mathcal{S}$.

Fix any metric (for instance, a visual metric) d_∂ on $\partial\Gamma$. The open ϵ–ball around a point $p \in \partial\Gamma$ with respect to this metric is denoted by $B_\epsilon(p)$, and the open ϵ–neighborhood of a set $K \subset \partial\Gamma$ is denoted by $N_\epsilon(K)$.

As is well known, the action of Γ on the set of pairs of distinct points in $\partial\Gamma$ is cocompact, hence we have the following.

Lemma 2.1. *There is some $D > 0$ so that for every pair a, b of distinct points in $\partial\Gamma$, there is a g so that $d_\partial(ga, gb) \geq D$.*

We fix such a constant D for the rest of the paper.

Our first goal is to "code" boundary points, i.e. to associate to each boundary point a collection of infinite paths in a certain automaton. In Section 3 we will see that each of these paths tracks a geodesic ray limiting to the boundary point. We wish to define the coding in a way that uses only the dynamics of the action of Γ on its boundary, so that the coding sequences will still contain meaningful information after the action of Γ is perturbed.

The inspiration for this comes from Sullivan [Sul85], who codes points using sequences of elements that contract subsets of the boundary a uniform amount. Sullivan's "uniform contraction" is not C^0–stable, so we instead follow the modification of this approach given in [Wei22], and use the topological dynamics of the action to construct nested sequences of sets that capture the idea of uniform contraction.

We recall the following.

Definition 2.2. Let G act by homeomorphisms on on a metric space X. A point $x \in X$ is a *conical limit point* for the action if there exists a sequence $\{g_n\}_{n\in\mathbb{N}}$ of elements of G and distinct points $a, b \in X$ such that $g_n x \to a$ and $g_n y \to b$ for all $y \neq x$, uniformly on compact sets in $X - \{x\}$.

When Γ is a hyperbolic group acting on its boundary $\partial\Gamma$, every point in $\partial\Gamma$ is a conical limit point. The first step in our construction is to use this property to set up a pair of good covers of $\partial\Gamma$. Then we use these covers to create a finite-state automaton that accepts words in (a finite index subgroup of) Γ which code points. To produce the covers, we use the following general lemma. Here and in what follows, metric notions like diameter always refer to the fixed metric d_∂ on $\partial\Gamma$.

Lemma 2.3 (Expanded neighborhoods)**.** *For any positive $\epsilon < \frac{D}{4}$ and any $z \in \partial\Gamma$, there is an $\alpha_z \in \Gamma$ and a pair of open neighborhoods $\hat{V}_z \subset W_z$ of z so that*

(1) $\mathrm{diam}(W_z) \leq \epsilon$*;*
(2) $\mathrm{diam}(\alpha_z^{-1} W_z) > 3\epsilon$*; and*
(3) $\overline{N_{1.5\epsilon}(\alpha_z^{-1}\hat{V}_z)} \subset \alpha_z^{-1} W_z$*.*

Proof. We choose some $\epsilon < \frac{D}{4}$ where D is the constant from Lemma 2.1.

Let $z \in \partial\Gamma$. Since z is a conical limit point, we can find distinct points a, b and a sequence of group elements $\{g_i\}_{i\in\mathbb{N}}$ so that $g_i z \to b$ and $g_i x \to a$ uniformly away from z. Up to post-composing all g_i with a fixed element g as in Lemma 2.1 if necessary, we may assume $d_\partial(a, b) \geq D$. Also, since $\partial\Gamma$ is perfect, there is some point $a' \neq a$ with $d_\partial(a, a') = \epsilon' < \epsilon$.

Let $W_z = B_{\epsilon/2}(z)$, so Property (1) is satisfied. Let K_z be the complement of W_z in $\partial\Gamma$. The set K_z is compact and does not contain z, so for i sufficiently large, we have $g_i K_z \subset B_{\epsilon'}(a)$ and $g_i z \in B_\epsilon(b)$.

Fixing some such i, set $\alpha_z = g_i^{-1}$, and let $\hat{V}_z = \alpha_z(B_\epsilon(b))$. Note that $B_\epsilon(a)$ contains $\alpha_z^{-1} K_z = \partial\Gamma - \alpha_z^{-1} W_z$. Let $\hat{V}_z = \alpha_z(B_\epsilon(b))$. Since $B_\epsilon(a)$ is disjoint from $B_\epsilon(b)$, we have

$$\hat{V}_z = \alpha_z(B_\epsilon(b)) \quad \subset \quad \partial\Gamma - \alpha_z(K_z) = W_z.$$

The set $\alpha_z^{-1} W_z = \partial\Gamma - \alpha_z^{-1} K_z$ contains both b and a', so $\mathrm{diam}(\alpha_z^{-1} W_z) \geq d_\partial(b, a') \geq D - \epsilon > 3\epsilon$, establishing Property (2).

Finally, since $d_\partial(B_\epsilon(b), \alpha_z^{-1} K_z) \geq D - 2\epsilon > 2\epsilon$, we have

$$\overline{N_\epsilon(\alpha_z^{-1}\hat{V}_z)} \subset \overline{B_{2\epsilon}(b)} \subset (\partial\Gamma - \alpha_z^{-1} K_z) = \alpha_z^{-1} W_z,$$

establishing Property (3). □

Definition 2.4 (Fixing ϵ)**.** Here we will fix a scale ϵ for the rest of the paper. In order to do so we first fix, as in the statement of Theorem 1.1, a neighborhood $\mathcal{U}$ of the identity in the space of continuous self-maps of $\partial\Gamma$. Now we fix a scale ϵ so that all of the following hold.

(1) $0 < \epsilon < D/4$, where D is the constant furnished by Lemma 2.1.
(2) $\mathcal{U}$ contains all maps f such that $d(x, f(x)) \leq \epsilon$ for all $x \in \partial\Gamma$.

Definition 2.5 (Fixing a pair of covers)**.** Since $\epsilon < D/4$, we can apply Lemma 2.3. For each $z \in \partial\Gamma$ choose a pair of open neighborhoods $\hat{V}_z \subset W_z$ of z as in the conclusion of Lemma 2.3. Let $I \subset \partial\Gamma$ be a finite collection so that the sets $\{\hat{V}_z\}_{z\in I}$ cover $\partial\Gamma$. Parts (1) and (2) of Lemma 2.3 imply that the W_z satisfy the following conditions.

(C1) $\mathrm{diam}(W_z) \leq \epsilon$;
(C2) $\mathrm{diam}(\alpha_z^{-1} W_z) > 3\epsilon$

Our sets $\hat{V}_z$ could have the inconvenient property that for some pair y, z, the intersection of *closures* $\overline{\hat{V}}_z \cap \alpha_y^{-1} \overline{\hat{V}}_y$ is non-empty, even when the corresponding intersection $\hat{V}_z \cap \alpha_y^{-1} \hat{V}_y$ is empty, meaning that this intersection may not persist under small perturbations of α_y. However, since there are only finitely many such pairs to consider, we may modify the sets $\hat{V}_z$ slightly to change each such unstable intersection to a stable one. Precisely, we replace each set $\hat{V}_z$ by a larger set V_z so that no unstable intersections occur, and we take the enlargements small enough so that $V_z \subset W_z$ still holds, and there are no further intersections of these sets and their images under the α_z. For a sufficiently small modification, condition (3) of Lemma 2.3 will still hold if 1.5ϵ is replaced by ϵ. In summary, the sets V_z have the following properties:

(C3) $\overline{N_\epsilon(\alpha_z^{-1} V_z)} \subset \alpha_z^{-1} W_z$.
(C4) For any pair $y, z \in I$, the intersection $\overline{V}_z \cap \alpha_y^{-1} \overline{V}_y$ is either empty or has nonempty interior.

For the rest of the paper we fix the scale ϵ, the indexing set I and pair of covers $V_z \subset W_z$ from Definition 2.5. For each $z \in I$ we also fix the group element α_z from Lemma 2.3.

Lemma 2.6. *For each $z, y \in I$, if $\overline{V}_z$ meets $\alpha_y^{-1} \overline{V}_y$, then $\overline{W_z}$ is contained in $\alpha_y^{-1}(W_y)$.*

Proof. Suppose $\overline{V}_z \cap \alpha_y^{-1} \overline{V}_y$ is non-empty. Property (C3) of Definition 2.5 implies that $\overline{V}_z \subset W_z$, so we deduce that $W_z \cap \alpha_y^{-1} \overline{V}_y$ is nonempty. Property (C1) of Definition 2.5 implies $\mathrm{diam}(W_z) \leq \epsilon$. Thus $\overline{W}_z$ is contained in $\overline{N_\epsilon(a_y^{-1} V_y)}$, which is contained in $\alpha_y^{-1} W_y$ by Property (C3) of Definition 2.5. $\square$

The index set I can be given the structure of a directed graph as follows.

Definition 2.7 (The associated automaton)**.** We let $\mathcal{G}$ be the graph with vertex set I, with an edge from y to z if and only if $\alpha_y^{-1}(\overline{V}_y) \cap \overline{V}_z \neq \emptyset$ (see Figure 1). A *$\mathcal{G}$–coding* is an infinite sequence $\{z(k)\}_{k \in \mathbb{N}}$ of points of I so that there is an edge from $z(k)$ to $z(k+1)$ for each k. If

$$p \in \bigcap_{k=0}^{\infty} \alpha_{z(1)} \cdots \alpha_{z(k)} \overline{W}_{z(k+1)}$$

we say that $\{z(k)\}_{k \in \mathbb{N}}$ is a *$\mathcal{G}$–coding of p*. If $\mathcal{G}$ is understood, we may omit it, and speak of a *coding of p*.

Remark 2.8. Lemma 2.6 implies that for any directed path $\{z(k)\}_{k \in \mathbb{N}}$, we have $\alpha_{z(k)} \overline{W}_{z(k+1)} \subset W_{z(k)}$. Thus, the intersection

$$\bigcap_{k=0}^{\infty} \alpha_{z(1)} \cdots \alpha_{z(k)} \overline{W}_{z(k+1)}$$

is an intersection of nested closed sets, and so by compactness of $\partial\Gamma$, it is always nonempty. In particular every $\mathcal{G}$–coding is a $\mathcal{G}$–coding of at least one point. We will see below that it is a $\mathcal{G}$–coding for *only* one point.

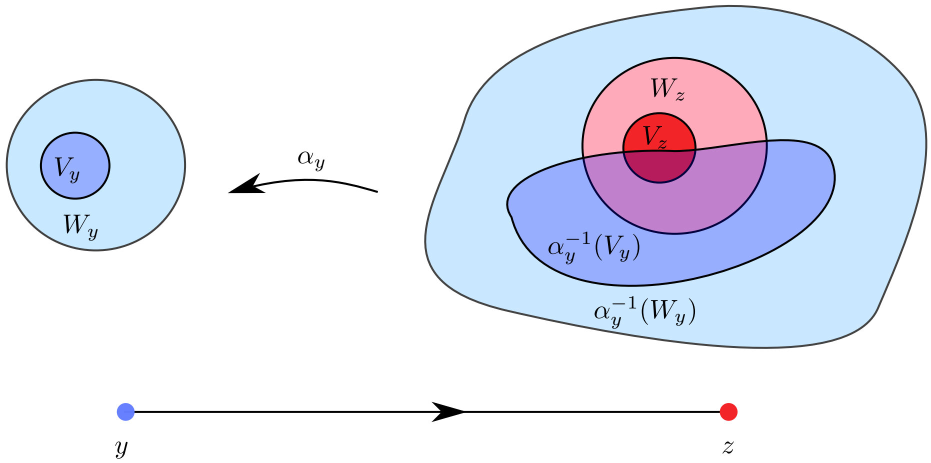

FIGURE 1. There is an edge from y to z when $\overline{\alpha_y^{-1}(V_y)}$ meets $\overline{V}_z$.

Lemma 2.9. *Every $p \in \partial\Gamma$ has a coding.*

Proof. Let $p \in \partial\Gamma$ be given. Take $z(1) \in I$ so that $p \in \overline{V}_{z(1)}$. Suppose that $z(1), \dots, z(k)$ have been defined, and let $g_k = \alpha_{z(1)} \cdots \alpha_{z(k)}$. Then there is some $z(k+1)$ so that $g_k^{-1} p \in \overline{V}_{z(k+1)}$. Inductively we have $g_{k-1}^{-1}(p) \in \overline{V}_{z(k)}$, so $g_k^{-1} p = \alpha_{z(k)}^{-1} g_{k-1}^{-1}(p)$ is a point in the intersection of $\alpha_{z(k)}^{-1} \overline{V}_{z(k)}$ with $\overline{V}_{z(k+1)}$. In particular there is an edge joining $z(k)$ to $z(k+1)$. From the construction we have

$$p \in \bigcap_{k=0}^{\infty} g_k \overline{V}_{z(k+1)} \subset \bigcap_{k=0}^{\infty} g_k \overline{W}_{z(k+1)},$$

so the sequence $\{z(k)\}_{k\in\mathbb{N}}$ is a $\mathcal{G}$–coding of p. □

Lemma 2.10 (Bounded backtracking property)**.** *Let $\{z(k)\}_{k\in\mathbb{N}}$ be a $\mathcal{G}$–coding. For any k define*

$$U_k = \alpha_{z(1)} \cdots \alpha_{z(k-1)} W_{z(k)}.$$

Then U_{k+1} is a proper subset of U_k for any $k \geq 1$. Moreover, in the sequence $\{g_k = \alpha_{z(1)} \cdots \alpha_{z(k)}\}_{k\in\mathbb{N}}$, no element g_k is repeated more than $\#I$ times.

Proof. Fix $k \geq 1$. By the definition of the graph $\mathcal{G}$, we have

$$\alpha_{z(k)}^{-1}(\overline{V}_{z(k)}) \cap \overline{V}_{z(k+1)} \neq \emptyset.$$

By Lemma 2.6, we have $W_{z(k+1)} \subset \alpha_{z(k)}^{-1}(W_{z(k)})$. By Properties (C1) and (C2) of Definition 2.5, the set $W_{z(k+1)}$ has diameter at most ϵ, whereas the set $\alpha_{z(k)}^{-1} W_{z(k)}$ has diameter at least 3ϵ. In particular the inclusion $W_{z(k+1)} \subset \alpha_{z(k)}^{-1}(W_{z(k)})$ is proper. Multiplying on the left by $\alpha_{z(1)} \cdots \alpha_{z(k)}$ then gives a proper inclusion $U_{k+1} \subset U_k$.

To see the last assertion, suppose $\#\{k \mid g_k = g\} > \#I$ for some g. Then there must be distinct k, k' so that $g_k = g_{k'} = g$ and $W_{z(k)} = W_{z(k')}$. For these indices, $U_k = U_{k'} = gW_{z(k)}$, a contradiction to proper nesting. □

3. Properties of $\mathcal{G}$–codings

In this section we establish key properties of $\mathcal{G}$–codings to be used in the proof of the main theorem. Our first goal is to show that a sequence in Γ defined by a coding lies a uniformly bounded Hausdorff distance from a geodesic ray based at identity.

3.1. **Codings and quasi-geodesics.** We begin with two general lemmas on hyperbolic spaces.

Lemma 3.1. *Let $\epsilon_0 < D/4$. There is a $c_0 > 0$ so that for any $p \in \partial\Gamma$, there are points q, q' so that $d_\partial(p,q) > \epsilon_0$ and $d_\partial(p,q') > \epsilon_0$, and $d_\partial(q,q') \geq c_0$.*

Proof. We recall that $D \leq \operatorname{diam}(\partial\Gamma)$ is the bound from Lemma 2.1 so that any pair of points in $\partial\Gamma$ can be translated to a pair whose distance is at least D. In particular, there are points $N, S \in \partial\Gamma$ with $d_\partial(N,S) \geq D$. Since $\partial\Gamma$ is perfect, there are points $N' \in B_{\epsilon_0}(N) - \{N\}$ and $S' \in B_{\epsilon_0}(S) - \{S\}$. Let $c_0 = \min\{d_\partial(N,N'), d_\partial(S,S')\}$.

If p is further than ϵ_0 from both N and S, we may take $q = N$, $q' = S$. Otherwise we may suppose after relabeling that $d_\partial(p,N) \leq \epsilon_0$. But then

$$d_\partial(p,S) \geq D - \epsilon_0 \geq 3\epsilon_0$$

and

$$d_\partial(p,S') \geq D - 2\epsilon_0 \geq 2\epsilon_0,$$

so we may take $q = S$ and $q' = S'$. □

Lemma 3.2. *Given $\delta \geq 0$, there is a constant T depending only on δ so that the following holds. Let (X, d_X) be a δ–hyperbolic metric space. Let $p, q, r \in \partial X$, and for each pair of distinct points $x, y \in \{p,q,r\}$, let (x,y) be a bi-infinite geodesic joining x to y.*

Then for each $x, y \in \{p,q,r\}$ distinct, there are points $c_{xy} \in (x,y)$ such that

(1) $\operatorname{diam}(\{c_{pq}, c_{qr}, c_{pr}\}) \leq T$;
(2) If w lies in the sub-ray $[c_{pq}, p) \subset (p,q)$, then

$$d(w,(q,r)) \geq d(w, c_{pq}) - T$$

and similar statements hold with p, q, r permuted.

Proof. We may assume that for any x, y distinct in $\{p,q,r\}$, we have $(x,y) = (y,x)$. Let c be any point which lies within 2δ of all three geodesics, and for each $\{x,y\} \subset \{p,q,r\}$, choose a closest point $c_{yx} = c_{xy} \in (x,y)$ to c. These points clearly satisfy (1).

The point c_{xy} cuts (x,y) into rays $[c_{xy}, x)$ and $[c_{xy}, y)$. For each $z \in \{p,q,r\}$, choose a geodesic ray $[c,z)$ from c to z. Let Y be the tripod which is the union of these three rays, and let Δ be the union of the geodesics $(p,q), (q,r)$ and (p,r). We define a map $\pi\colon \Delta \to Y$ so that $\pi(\{c_{pq}, c_{qr}, c_{rp}\}) = z$, and so that π sends the ray $[c_{xy}, y)$ isometrically to the ray $[c,y)$. It is straightforward to see that $d_X(\pi(w), w) \leq 4\delta$ for any $w \in \Delta$, so π is a $(1, 8\delta)$–quasi-isometry.

Let $\tilde{Y}$ be an abstract infinite tripod, the wedge of three rays. We claim the obvious map from $\tilde{Y}$ to Y, isometric on each ray, is a $(1, 8\delta)$–quasi-isometric embedding. Indeed, the embedding is clearly distance non-increasing, and it is enough to consider points on distinct rays. Up to relabeling, we may assume that these two rays are $[c,p)$ and $[c,q)$. We consider unit-speed parameterizations τ_p and τ_q of

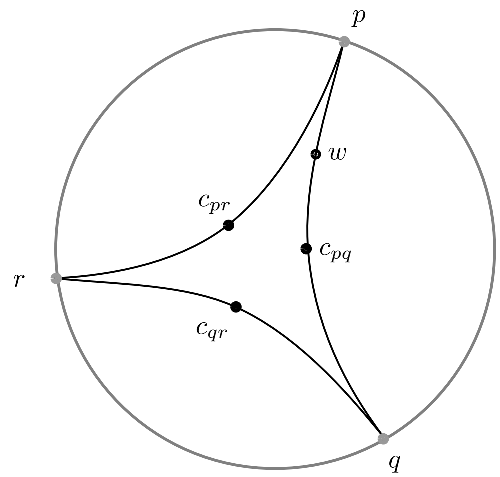

FIGURE 2. Lemma 3.2.

these rays, and unit-speed parameterizations α_p, α_q of the rays $[c_{pq}, p)$ and $[c_{pq}, q)$. Then, for any $s, t \geq 0$, we have $d_X(\alpha_p(s), \tau_p(s)) \leq 4\delta$ and $d_X(\alpha_q(t), \tau_q(t)) \leq 4\delta$. Since (p, q) is geodesic, $d_X(\tau_p(s), \tau_q(t)) \geq s + t - 8\delta$.

Combining the maps from the last two paragraphs, we see that the triangle Δ is $(1, 16\delta)$–quasi-isometric to the abstract tripod $\tilde{Y}$. Since (2) holds in the tripod with $T = 0$ (taking all the c_{xy} to be the central point of the tripod), it holds in Δ with $T = 32\delta$. □

Definition 3.3. The points c_{pq}, c_{qr}, c_{pr} from Lemma 3.2 will be referred to as *central points* of the ideal triangle with vertices p, q, r.

The next lemma is inspired by Proposition 5.11 from [Wei22].

Lemma 3.4. *Let X denote the Cayley graph of G with generating set $\mathcal{S}$. Fix a $\mathcal{G}$–coding $\{z(k)\}_{k\in\mathbb{N}}$. For ease of notation, let α_k denote $\alpha_{z(k)}$, and let W_k denote $W_{z(k)}$. Then the set $\{g_k = \alpha_1 \cdots \alpha_k\}_{k\in\mathbb{N}}$ is (uniformly) finite Hausdorff distance from a geodesic ray based at the identity in X. If $\{z(k)\}_{k\in\mathbb{N}}$ is a $\mathcal{G}$–coding for p, then this geodesic ray tends to p.*

Proof. Let $\{z(k)\}_{k\in\mathbb{N}}$ be a $\mathcal{G}$–coding for p and let $g_k = \alpha_1 \cdots \alpha_k$. The distances $d_X(g_k, g_{k+1})$ are uniformly bounded, so $\{g_k\}_{k\in\mathbb{N}}$ is uniformly finite Hausdorff distance from the image of some path in X. Lemma 2.10 implies this path must be a proper ray. So it suffices to show that the sequence $\{g_k\}_{k\in\mathbb{N}}$ lies within a uniformly bounded neighborhood of a geodesic ray based at the identity in X, tending towards p.

The general strategy of proof is as follows: we first show that $\{g_k\}$ is contained in a uniformly bounded neighborhood of a sub-ray of a geodesic between p and some point $q \in \partial\Gamma$, and that this sub-ray also passes close to the identity vertex in X. By varying the point q and running the same argument, we show that this sub-ray does not extend too far in the direction of q and conclude that it is close to a ray based at the identity vertex (see Figure 3 for a schematic).

Let c_1 be a positive lower bound for the distances $d_\partial(\partial\Gamma - \alpha_x^{-1}W_x, \overline{W}_y)$ as (x, y) ranges over the directed edges of $\mathcal{G}$. Let c_0 be the constant from Lemma 3.1 (applied

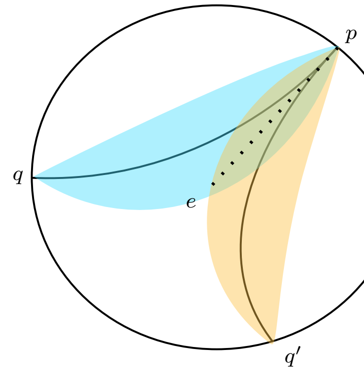

FIGURE 3. If $\{g_k\}$ lies in a bounded neighborhood of both pq and pq', then it must be near a ray to p based at e

with $\epsilon_0 = \epsilon$) and let $c = \min\{c_0, c_1, \epsilon\}$. Take $C > 0$ large enough (depending only on c) so that for any pair of points x, y with $d_\partial(x, y) \geq c$, any geodesic joining x to y in X comes within C of the identity vertex e. The existence of such a C is immediate if d_∂ is a visual metric based at e; since the boundary is compact, such a C will exist for any metric.

By Lemma 3.1, there are points q, q' so that $d_\partial(p, q) > \epsilon$, $d_\partial(p, q') > \epsilon$, and $d_\partial(q, q') > c_0$. In particular, any geodesic joining a pair of distinct points in $\{p, q, q'\}$ meets the ball of radius C around the identity.

For each $k \in \mathbb{N}$, let $\overline{U}_k = g_{k-1}\overline{W}_k$. This gives a family of nested closed neighborhoods of p, each of which has diameter $\leq \epsilon$. Since $d_\partial(p, q)$ and $d_\partial(p, q')$ are strictly larger than ϵ, neither q nor q' is contained in any of the sets $\overline{U}_k$. In particular,

$$q, q' \notin \overline{U}_k \subset \overline{W}_1. \tag{1}$$

Now let z be either q or q'. For any k we have $g_k^{-1}(p) \in g_k^{-1}\overline{U}_{k+1} = \overline{W}_{k+1}$ and, by (1), $g_k^{-1}(z) \notin g_k^{-1}\overline{U}_k = \alpha_k^{-1}(\overline{W}_k)$. Thus $d_\partial(g_k^{-1}(p), g_k^{-1}(z)) \geq c_1 \geq c$, which by our choice of C means that any geodesic from $g_k^{-1}(p)$ to $g_k^{-1}(z)$ passes within distance C of e; in other words, if we fix a geodesic γ from z to p, we have $d_X(e, g_k^{-1}\gamma) \leq C$ for each k. Multiplying by g_k we see that $d_X(g_k, \gamma) \leq C$ for each k. In particular the sequence $\{g_k\}_{k\in\mathbb{N}}$ is contained in the closed C–neighborhood of γ.

Now fix bi-infinite geodesics (q, q'), (p, q) and (p, q'), and consider central points $c_{pq}, c_{pq'}$ and $c_{qq'}$ for the ideal triangle with vertices p, q, q', given by Lemma 3.2. We claim that these central points lie uniformly close to e. To see this, for each $\{x, y\} \subset \{p, q, q'\}$, we let w_{xy} be a point lying on the geodesic from x to y such that w_{xy} is within distance C of the identity. If $\{x, y, z\} = \{p, q, r\}$, the geodesic segment $[c_{xy}, w_{xy}]$ lies in either the sub-ray $[c_{xy}, x)$ or $[c_{xy}, y)$; up to relabeling, we assume it lies in $[c_{xy}, x)$. Then, by Lemma 3.2, we have

$$2C \geq d_X(w_{xy}, w_{yz}) \geq d_X(w_{xy}, (y, z)) \geq d_X(w_{xy}, c_{xy}) - T.$$

This gives the estimate $d_X(c_{xy}, e) < 3C + T$, as desired.

Using this bound, we conclude that the rays $[c_{pq}, p) \subset (q, p)$ and $[c_{pq'}, p) \subset (q', p)$ are both Hausdorff distance at most $3C + T + 2\delta$ from any geodesic ray $[e, p)$ from e to p.

Let $z \in \{g_k\}_{k\in\mathbb{N}}$. Then z lies within C of both (q,p) and (q',p). If z lies within C of $[c_{pq},p) \subset (q,p)$, then it lies in the $4C+T+2\delta$–neighborhood of $[e,p)$. Otherwise, z lies within distance C of a point w on $[c_{pq},q)$. Then by Lemma 3.2 again, we have

$$d_X(z,c_{pq}) \le d_X(w,c_{pq}) + C \le d_X(w,(p,q')) + C + T \le d_X(z,(p,q')) + 2C + T.$$

Since $d_X(c_{pq},e) \le 3C+T$ and $d_X(z,(p,q')) \le C$, we have $d_X(z,e) \le 6C+2T$. In either case z lies in the $6C+2T+2\delta$–neighborhood of the ray $[e,p)$.

□

Corollary 3.5. *For any $\mathcal{G}$–coding $\{z(k)\}_{k\in\mathbb{N}}$, the intersection*

$$\bigcap_{k=0}^{\infty} \alpha_{z(1)} \cdots \alpha_{z(k)} \overline{W}_{z(k+1)}$$

is a singleton.

Proof. By Lemma 3.4, $\{z(k)\}_{k\in\mathbb{N}}$ determines a set $\{g_k\}_{k\in\mathbb{N}}$ which is finite Hausdorff distance from a geodesic ray with a unique endpoint p. If q is a point in the intersection, then $\{z(k)\}_{k\in\mathbb{N}}$ codes q, and therefore by Lemma 3.4, $q = p$. □

Combining Lemmas 3.4 and 2.9 we have the following.

Corollary 3.6. *The set $\{\alpha_z\}_{z\in I}$ generates a finite index subgroup of Γ.*

3.2. **Uniform contraction.** This section contains the primary technical application of our work above. The following is an immediate consequence of Lemma 3.4.

Lemma 3.7. *There is a finite set $F \subset G$ so that the following holds. Let $p \in \partial\Gamma$, and let $\{z(k)\}_{k\in\mathbb{N}}, \{y(k)\}_{k\in\mathbb{N}}$ be two $\mathcal{G}$–codings of p. For $k > 0$ let $g_k = \alpha_{z(1)} \cdots \alpha_{z(k)}$ and $h_k = \alpha_{y(1)} \cdots \alpha_{y(k)}$. For any $k > 0$ there is some $n(k)$ so that*

$$h_k = g_{n(k)} f_k \text{ for some } f_k \in F. \tag{2}$$

Symmetrically, there exists $m(k) \in \mathbb{N}$ such that $g_k = h_{m(k)} f'_k$ for some $f'_k \in F$, but it is the equality (2) that we will make use of in the next step of the proof.

The following lemma is the technical heart of our main theorem.

Lemma 3.8 (Uniform Contraction)**.** *Using the notation from Lemma 3.7, there exists a uniform $N > 0$ so that for any such pair of such sequences $\{g_k\}_{k\in\mathbb{N}}$ and $\{h_k\}_{k\in\mathbb{N}}$ and any $k > 0$, we have*

$$g_{n(k)+N} \overline{W}_{z(n(k)+N+1)} \subset h_k W_{y(k+1)}. \tag{3}$$

We emphasize that by "uniform" we mean that the constant N is independent of the point being coded and of the two chosen codings of the point.

Proof. The proof is by contradiction. We therefore assume we have a sequence of natural numbers N tending to ∞ so that for each N in the sequence there is a point $p^N \in \partial\Gamma$ with two distinct codings $\{z^N(k)\}_{k\in\mathbb{N}}$ and $\{y^N(k)\}_{k\in\mathbb{N}}$, so that the statement (3) fails for this pair of codings and some $k = k_N \in \mathbb{N}$.

Before continuing the proof, we clarify notation: Our sequences of codings $\{z^N(k)\}_{k\in\mathbb{N}}$ and $\{y^N(k)\}_{k\in\mathbb{N}}$ are fixed and indexed by N. As in Lemma 3.7, for

$g_k^{(N)} = \alpha_{z^N(1)} \cdots \alpha_{z^N(k)}$ and $h_k^{(N)} = \alpha_{y^N(1)} \cdots \alpha_{y^N(k)}$ we have indices $n^N(k) \in \mathbb{N}$ and elements $f_k^{(N)}$ of the finite set F such that for each N, k there is an equality

$$h_k^{(N)} = g_{n^N(k)}^{(N)} f_k^{(N)}. \tag{4}$$

To decrease notation slightly, we shorten $n^N(k_N)$ to n_N. The failure of the nesting condition (3) can then be expressed as

$$g_{n_N+N}^{(N)} \overline{W}_{z^N(n_N+N+1)} \not\subset h_{k_N}^{(N)} W_{y^N(k_N+1)}. \tag{5}$$

We immediately pass to a subsequence so that

(f) $\qquad f_{k_N}^{(N)}$ is constant, equal to $f \in F$.

We further refine our subsequence so the following three conditions are satisfied.

(z) The sets $W_{z^N(n_N+N+1)}$ are constant, equal to some W_z.

(y) The sets $W_{y^N(k_N+1)}$ are constant, equal to some W_y.

($*$) The sets $\alpha_{y^N(k_N+1)} W_{y^N(k_N+2)}$ are constant, equal to some W.

This is possible because there are only finitely many possibilities for $W_{z^N(n_N+N+1)}$ and $W_{y^N(k_N+1)}$ (being elements of the cover) and also for $\alpha_{y^N(k_N+1)} W_{y^N(k_N+2)}$, being a translate of an element of the cover by one of finitely many elements. A key property we will use at the end of the proof is that

$$\overline{W} \subset W_y. \tag{6}$$

For these values of N, we can multiply each side of (5) on the left by $(g_{n_N}^{(N)})^{-1}$ and use (4) together with condition (f) to obtain

$$\alpha_{z^N(n_N+1)} \cdots \alpha_{z^N(n_N+N)} \overline{W}_z \not\subset f W_y. \tag{7}$$

For each N, consider the sub-coding from $\{z^N(k)\}_{k\in\mathbb{N}}$ formed by terms $n_N + 1$ through infinity. We denote this coding by $\{\gamma^N(k)\}_{k\in\mathbb{N}}$. In other words, we define $\gamma^N(k) = z^N(n_N + k)$. After passing to a subsequence $\{N(j)\}_{j\in\mathbb{N}}$ one final time, we may obtain a sequence of codings $\{\gamma^{N(j)}(k)\}_{k\in\mathbb{N}}$ so that for all $l \geq j$, the initial segment $\{\gamma^{N(l)}(1), \ldots, \gamma^{N(l)}(j)\}$ is independent of l, so equal to the first j terms of $\{\gamma^{N(j)}(k)\}_{k\in\mathbb{N}}$. That is, the codings $\{\gamma^{N(j)}(k)\}_{k\in\mathbb{N}}$ converge to a coding $\{\gamma^\infty(k)\}_{k\in\mathbb{N}}$, which codes a unique point $p^\infty \in \partial\Gamma$.

For our subsequence $N(j)$, the non-containment in (7) takes the slightly simpler form

$$\alpha_{\gamma^{N(j)}(1)} \cdots \alpha_{\gamma^{N(j)}(N(j))} \overline{W}_z \not\subset f W_y,$$

which we can rewrite as

$$\left(\alpha_{\gamma^\infty(1)} \cdots \alpha_{\gamma^\infty(j)}\right) \left(\alpha_{\gamma^{N(j)}(j+1)} \cdots \alpha_{\gamma^{N(j)}(N(j))}\right) \overline{W}_z \not\subset f W_y.$$

Since $\{\gamma^{N(j)}(k)\}_{k\in\mathbb{N}}$ is a coding,

$$\left(\alpha_{\gamma^\infty(1)} \cdots \alpha_{\gamma^\infty(j)}\right) \left(\alpha_{\gamma^{N(j)}(j+1)} \cdots \alpha_{\gamma^{N(j)}(N(j))}\right) \overline{W}_z \subset \alpha_{\gamma^\infty(1)} \cdots \alpha_{\gamma^\infty(j-1)} W_{\gamma^\infty(j)},$$

so we must therefore have

$$\alpha_{\gamma^\infty(1)} \cdots \alpha_{\gamma^\infty(j-1)} \overline{W}_{\gamma^\infty(j)} \not\subset f W_y. \tag{8}$$

Since $\{\gamma^\infty(k)\}_{k\in\mathbb{N}}$ is a coding for p^∞, the sets on the left hand side of (8) give a nested basis of closed neighborhoods of p^∞ and we must have

$$p^\infty \notin f W_y. \tag{9}$$

On the other hand, since $\{z^N(k)\}_{k\in\mathbb{N}}$ is a $\mathcal{G}$–coding of p^N, for every j we have

$$p^{N(j)} \in \alpha_{z^{N(j)}(1)} \cdots \alpha_{z^{N(j)}(n_{N(j)}+j-1)} W_{z^{N(j)}(n_{N(j)}+j)},$$

or equivalently (multiplying both sides on the left by $(g^{(N(j))}_{n_{N(j)}})^{-1}$):

$$\begin{aligned}(g^{(N(j))}_{n_{N(j)}})^{-1} p^{N(j)} &\in \alpha_{z^{N(j)}(n_{N(j)}+1)} \cdots \alpha_{z^{N(j)}(n_{N(j)}+j-1)} W_{z^{N(j)}(n_{N(j)}+j)} \\ &= \alpha_{\gamma^\infty(1)} \cdots \alpha_{\gamma^\infty(j-1)} W_{\gamma^\infty(j)}.\end{aligned}$$

Again, since $\{\gamma^\infty(k)\}_{k\in\mathbb{N}}$ codes p^∞, this last sequence of sets gives a nested neighborhood basis for p^∞ and we must have

$$\lim_{j\to\infty} (g^{(N(j))}_{n_{N(j)}})^{-1} p^{N(j)} = p^\infty.$$

We also know that

$$(g^{(N)}_{n_N})^{-1} p^N \in (g^{(N)}_{n_N})^{-1} h^{(N)}_{k_N+1} W_{y^N(k_N+2)} \tag{10}$$

for any N, since $h^{(N)}_{k_N+1} W_{y^N(k_N+2)}$ is a neighborhood of p^N. By our assumptions (f) and ($*$) on our chosen subsequence, the right-hand side of (10) is always equal to a constant fW. But we have just seen that a subsequence of the left-hand side converges to p_∞, so we must have $p^\infty \in f\overline{W}$. Because of (6) this implies

$$p^\infty \in fW_y,$$

contradicting (9).

□

In our application of the uniform contraction principle above, we will use the following key observation:

Remark 3.9 (A finite list of nesting conditions suffices)**.** The conditions appearing in (3) (as we vary over all possible points and codings) appear to be infinite in number, but multiplying both sides of (3) on the left by $g^{-1}_{n(k)}$ gives an equivalent condition of the form

$$\alpha_{z(n(k)+1)} \cdots \alpha_{z(n(k)+N)} \overline{W}_{z(n(k)+N+1)} \subset f_k W_{y(k+1)}. \tag{11}$$

Since N is bounded, there are only finitely many possible conditions of this form.

3.3. **Conjugating by generators.** For each generator $s \in \mathcal{S}$, consider the action ρ_s of Γ on $\partial\Gamma$ given by conjugating the standard action by s. The sets sW_z and elements $s\alpha_z s^{-1}$ (for $z \in I$) give a graph $\mathcal{G}^s$ that is naturally isomorphic to $\mathcal{G}$ and codes points for the conjugated action. Under the natural isomorphism, a coding of sp in $\mathcal{G}^s$ corresponds to a coding of p in $\mathcal{G}$. Furthermore, if $\{z(k)\}_{k\in\mathbb{N}}$ is a $\mathcal{G}^s$–coding of sp, the corresponding path $\{sg_k s^{-1} = s\alpha_{z(1)}s^{-1} \cdots s\alpha_{z(k)}s^{-1}\}_{k\in\mathbb{N}}$ in Γ is a (uniformly, depending only on s) bounded distance from a geodesic ray in Γ tending to sp, since it is the image of the path $\{g_k = \alpha_{z(1)} \cdots \alpha_{z(k)}\}_{k\in\mathbb{N}}$ under conjugacy by s.

From the above uniformity we obtain the following analogue of Lemma 3.7.

Lemma 3.10. *For each $s \in S$ there is a finite set $F_s \subset G$ so that the following holds, for any $p \in \partial\Gamma$. Let $\{z(k)\}_{k\in\mathbb{N}}$ be a $\mathcal{G}$–coding of $sp \in \partial\Gamma$, and let $\{y(k)\}_{k\in\mathbb{N}}$ be a $\mathcal{G}$–coding of p (equivalently $\{y(k)\}_{k\in\mathbb{N}}$ is a $\mathcal{G}^s$–coding of sp). For $k > 0$ let*

$$g_k = \alpha_{z(1)} \cdots \alpha_{z(k)} \text{ and } h_k = s\alpha_{y(1)}s^{-1} \cdots s\alpha_{y(k)}s^{-1}.$$

For any $k > 0$ there is some $n(k)$ so that

$$h_k = g_{n(k)} f_k \text{ for some } f_k \in F_s.$$

Analogous to Lemma 3.8 we have the following.

Corollary 3.11. *Using the notation from Lemma 3.10, there exists a uniform $N_s > 0$ so that for any pair of such sequences $\{g_k\}_{k\in\mathbb{N}}$ and $\{h_k\}_{k\in\mathbb{N}}$ and any $k > 0$, we have*

$$g_{n(k)+N_s}\overline{W}_{z(n(k)+N_s+1)} \subset h_k s W_{y(k+1)}. \tag{12}$$

The proof follows that of Lemma 3.8 almost verbatim. The main change is that the constant sets W_y and W from that proof are defined slightly differently as $W_y = sW_{y^N(k_N+1)}$ and $W = s\alpha_{y^N(k_N+1)}W_{y^N(k_N+2)}$.

Remark 3.12. As in Remark 3.9, the equations of (12) reduce to only finitely many conditions; the fact that $h_k = g_{n(k)}f_k$ for some $f_k \in F_s$ means that each condition is equivalent to one of the form

$$\alpha_{z(n(k)+1)} \dots \alpha_{z(n(k)+N_s)}\overline{W}_z \subset fsW_y \tag{13}$$

for some $f \in F_s$, and W_z, W_y in our finite set, and $z(n(k)+1), \dots z(n(k)+N_s)$ a path of length N_s in $\mathcal{G}$.

3.4. **Relationship with annulus systems.** We conclude this section by indicating how Bowditch's framework of *annulus systems* on spaces with a convergence group action can give an alternative strategy towards the proof of the uniform contraction lemma (Lemma 3.8). In fact, one could use Bowditch's framework to prove a stronger version of Lemma 3.4, showing that for any $\mathcal{G}$–coding $\{z(k)\}_{k\in\mathbb{N}}$, the map $\mathbb{N} \to \Gamma$ given by $k \mapsto \alpha_{z(1)} \cdots \alpha_{z(k)}$ is actually a (uniform) quasi-geodesic embedding. However, since we do not need this stronger statement anywhere in the paper, and the setup is somewhat involved, we only provide a sketch of the argument for the purpose of describing the relationship.

The general setting is as follows. In [Bow98], Bowditch showed that whenever Γ acts on a perfect compact metrizable space Z as uniform convergence group, then it is possible to completely recover a word-hyperbolic metric on Γ from the data of the convergence group action. To do so, Bowditch defines a notion of a *system of annuli* on Z, and relates the topological behavior of such a system to a Γ–invariant *cross-ratio* (which in turn defines a hyperbolic metric on the space of triples in Z).

Definition 3.13. An *annulus* in $\partial\Gamma$ is an ordered pair $A = (A^-, A^+)$ of disjoint closed subsets of $\partial\Gamma$, such that $A_- \cup A_+ \neq \partial\Gamma$. A *symmetric system of annuli* is a collection $\mathcal{A}$ of annuli such that if $A = (A^-, A^+)$ is in $\mathcal{A}$, then $-A = (A^+, A^-)$ is also in $\mathcal{A}$.

A sequence of annuli $A_1, \dots, A_n$ is said to be *nested* if A_i^- contains $\partial\Gamma - A_{i+1}^+$ for all $1 \leq i < n$.

The terminology is inspired by the example $\partial\Gamma \simeq S^2$, where an "annulus" consisting of a pair of disjoint closed disks determines an annulus (in the usual sense) in S^2. Any system of annuli $\mathcal{A}$ defines a four-point cross-ratio on $\partial\Gamma$, denoted $(\cdot,\cdot;\cdot,\cdot)_{\mathcal{A}}$: for $x, y, z, w \in \partial\Gamma$, the cross-ratio $(x, y; z, w)_{\mathcal{A}}$ is the maximum length of a nested sequence of annuli $A_1, \dots, A_n$ in $\mathcal{A}$ such that A_1^+ contains $\{x, y\}$ and A_n^- contains $\{z, w\}$.

Work of Bowditch (see [Bow98], Proposition 4.8 and Sections 6 and 7) shows that whenever $\mathcal{A}$ is a symmetric annulus system satisfying certain hypotheses, then the cross-ratio $(\cdot,\cdot;\cdot,\cdot)_{\mathcal{A}}$ is within bounded additive and multiplicative error of a *standard* cross-ratio $(\cdot,\cdot;\cdot,\cdot)$ on $\partial\Gamma$. This cross-ratio is defined by realizing $\partial\Gamma$ as the Gromov boundary of a hyperbolic metric space X' (which is Γ–equivariantly quasi-isometric to the Cayley graph X of Γ) and defining $(a,b;c,d)$ to be the minimum distance between a geodesic in X' joining a to b and a geodesic joining c to d.

We can use our sets W_y for $y \in I$ to define a suitable annulus system: for each $y \in I$, take $A_y = (A_y^+, A_y^-)$, with

$$A_y^+ = \partial\Gamma - W_y, \qquad A_y^- = \bigcup_{y\to z} \alpha_y \overline{W_z}.$$

Then take $\mathcal{A} = \{\gamma A_y \mid \gamma \in \Gamma, y \in I\} \cup \{-(\gamma A_y) \mid \gamma \in \Gamma, y \in I\}$. One can then check that this system of annuli meets all of Bowditch's conditions, and therefore the induced cross-ratio $(\cdot,\cdot;\cdot,\cdot)_{\mathcal{A}}$ approximates $(\cdot,\cdot;\cdot,\cdot)$.

Now let $\{z(k)\}_{k\in\mathbb{N}}$ be a $\mathcal{G}$–coding, and for each $n \in \mathbb{N}$, let q_n, q_n' be a pair of distinct points in $U_{z(n+1)}$, joined by a geodesic passing within some uniformly bounded distance of the identity. Let $g_n = \alpha_{z(1)} \cdots \alpha_{z(n)}$. For any q_1, q_1' lying outside of $U_{z(1)}$, the cross-ratio $(g_n q_n, g_n q_n'; q_1, q_1')$ must be at least n. It follows that (up to uniform additive and multiplicative error) the distance from g_n to the identity in Γ is at least n.

4. Proof of Theorem 1.1

We can now prove the main theorem. Recall our standing assumptions: Γ is a hyperbolic group with fixed generating set $\mathcal{S}$. We have fixed a metric d_∂ on $\partial\Gamma$, and the constant D from Lemma 2.1. Moreover in Definition 2.4 we fixed a neighborhood $\mathcal{U}$ of the identity in the space of continuous self-maps of $\partial\Gamma$, and a constant $\epsilon < D/4$ small enough so that $\mathcal{U}$ contains all maps f such that $d(x, f(x)) \le \epsilon$ for all $x \in \partial\Gamma$.

We have also fixed covers $\{W_z \supset V_z\}_{z\in I}$ in Definition 2.5 that define a coding of boundary points as in Section 2, so that the results of Section 3 follow. Recall also that the sets W_z have diameter bounded by ϵ; we will use this property later.

Our goal is to specify a neighborhood $\mathcal{V}$ of the standard boundary action in $\mathrm{Hom}(\Gamma, \mathrm{Homeo}(\partial\Gamma))$ such that every action in $\mathcal{V}$ is an extension of the standard boundary action via a semi-conjugacy in $\mathcal{U}$.

Definition 4.1. We say that $\rho \in \mathrm{Hom}(\Gamma, \mathrm{Homeo}(\partial\Gamma))$ *has the same combinatorics as the standard boundary action* if the following hold for every $y, z \in I$.

(1) $\overline{V}_z \cap (\alpha_y^{-1}\overline{V}_y) \neq \emptyset$ iff $\overline{V}_z \cap (\rho(\alpha_y)^{-1}\overline{V}_y) \neq \emptyset$, and
(2) if $\overline{V}_z \cap (\alpha_y^{-1}\overline{V}_y) \neq \emptyset$, then $\overline{W_z} \subset \rho(\alpha_y)^{-1}(W_y)$.

Note that "having the same combinatorics" is an open condition, because of Property (C4) from Definition 2.5 of our covers.

Definition 4.2. Suppose ρ has the same combinatorics as the standard boundary action. We say that an infinite path $\{z(k)\}_{k\in\mathbb{N}}$ in $\mathcal{G}$ is a $(\mathcal{G},\rho)$*–coding of* p if

$$p \in \bigcap_{k=0}^{\infty} \rho(\alpha_{z(1)}) \cdots \rho(\alpha_{z(k)}) \overline{W}_{z(k+1)}.$$

Since Lemma 2.9 only used the intersection pattern of the sets V_i and their images under the action of Γ, its proof applies verbatim to show the following.

Lemma 4.3. *If $\rho\colon \Gamma \to \mathrm{Homeo}(\partial\Gamma)$ has the same combinatorics as the standard boundary action, then every $p \in \partial\Gamma$ has a $(\mathcal{G},\rho)$–coding.*

Corollary 3.5 stated that for the standard action, each $\mathcal{G}$–coding determined a *unique* point. But the proof of that corollary used strongly that the action on $\partial\Gamma$ was induced by the isometric action on the Cayley graph. Indeed there may be a nondegenerate closed subset of $\partial\Gamma$ all of whose points share the same $(\mathcal{G},\rho)$–coding.

Definition 4.4 (The neigborhood $\mathcal{V}$)**.** We now describe a neighborhood of the standard boundary action in $\mathrm{Hom}(\Gamma, \mathrm{Homeo}(\partial\Gamma))$ that will satisfy the requirements of the theorem. In the following description F is the finite set from Lemma 3.7 and N is the constant from the Uniform Contraction Lemma 3.8. Similarly for $s \in S$, the set F_s is the finite set from Lemma 3.10, and the constant N_s is the constant from Corollary 3.11. We let $\mathcal{V}$ be some neighborhood sufficiently small that all the following hold.

(V1) If $\rho \in \mathcal{V}$ then ρ has the same combinatorics as the standard boundary action, in the sense of Definition 4.1.

(V2) If $\rho \in \mathcal{V}$, $y, z \in I$, and $z(1)\dots z(N)$ is a length-N path in $\mathcal{G}$ so that $\alpha_{z(1)}\cdots\alpha_{z(N)}\overline{W}_z \subset fW_y$ for some $f \in F$, then

$$\rho(\alpha_{z(1)})\cdots\rho(\alpha_{z(N)})\overline{W}_z \subset \rho(f)W_y$$

(V3) If $\rho \in \mathcal{V}$, $s \in \mathcal{S}$, $y, z \in I$, and $z(1)\dots z(N_s)$ is a length-N_s path in $\mathcal{G}$ so that $\alpha_{z(1)}\dots\alpha_{z(N_s)}\overline{W}_z \subset fsW_y$ for some $f \in F_s$, then

$$\rho(\alpha_{z(1)})\dots\rho(\alpha_{z(N_s)})\overline{W}_z \subset \rho(fs)W_y.$$

$\boxed{\text{From here on we fix a representation } \rho \text{ from the neighborhood } \mathcal{V} \text{ just defined.}}$

Our next goal is to define a function $\Phi\colon\ \partial\Gamma \to 2^{\partial\Gamma}$ associating each point $p \in \partial\Gamma$ that is $\mathcal{G}$–coded by a sequence $\{z(k)\}_{k\in\mathbb{N}}$ to the closed set

$$\bigcap_{k=0}^{\infty} \rho(\alpha_{z(1)})\cdots\rho(\alpha_{z(k)})\overline{W}_{z(k+1)}.$$

The sets $\Phi(p)$ will be the fibers of our semi-conjugacy. To show Φ is well defined, we use the following.

Lemma 4.5. *If $\{z(k)\}_{k\in\mathbb{N}}$ and $\{y(k)\}_{k\in\mathbb{N}}$ are two distinct $\mathcal{G}$–codings of p, then*

$$\bigcap_{k=0}^{\infty} \rho(\alpha_{z(1)})\cdots\rho(\alpha_{z(k)})\overline{W}_{z(k+1)} = \bigcap_{k=0}^{\infty} \rho(\alpha_{y(1)})\cdots\rho(\alpha_{y(k)})\overline{W}_{y(k+1)}$$

Proof. It suffices to show that for any finite k, we can find some n such that

$$\rho(\alpha_{z(1)})\cdots\rho(\alpha_{z(n)})\overline{W}_{z(n+1)} \subseteq \rho(\alpha_{y(1)})\cdots\rho(\alpha_{y(k)})\overline{W}_{y(k+1)}. \tag{14}$$

Given k, we choose $n(k)$ as in Lemma 3.7, so that if $g_k = \alpha_{z(1)}\cdots\alpha_{z(k)}$ and $h_k = \alpha_{y(1)}\cdots\alpha_{y(k)}$, then $h_k = g_{n(k)}f$ for some $f \in F$.

Then we choose N as in Lemma 3.8, so that

$$\alpha_{z(1)}\cdots\alpha_{z(n(k)+N)}\overline{W}_{z(n(k)+N+1)} \subseteq \alpha_{y(1)}\cdots\alpha_{y(k)}W_{y(k+1)}.$$

We claim that the containment above still holds even after we replace each α_k with its perturbed image $\rho(\alpha_k)$, i.e. that

$$\rho(\alpha_{z(1)}) \cdots \rho(\alpha_{z(n(k)+N)})\overline{W}_{z(n(k)+N+1)} \subseteq \rho(\alpha_{y(1)}) \cdots \rho(\alpha_{y(k)})W_{y(k+1)}.$$

In other words, (14) is satisfied with $n = n(k)+N$. To prove the claim, multiply each side of the above inclusion by $\rho(g_{n(k)})^{-1}$ to obtain one of the conditions assumed in Item (V2) of Definition 4.4. $\square$

Thus, the following gives a well-defined map from $\partial\Gamma$ to the space of closed subsets of $\partial\Gamma$.

Definition 4.6. Let

$$\Phi(p) := \bigcap_{k=0}^{\infty} \rho(\alpha_{z(1)}) \cdots \rho(\alpha_{z(k)})\overline{W}_{z(k+1)},$$

where $\{z(k)\}_{k\in\mathbb{N}}$ is any coding of p.

Lemma 4.7 (Equivariance)**.** *For any $p \in \partial\Gamma$ and $g \in \Gamma$, we have*

$$\Phi(gp) = \rho(g)(\Phi(p)).$$

This proof is where we use the fact that the nesting conditions for conjugates by generators from Item (V3) also hold under our perturbation.

Proof of Lemma 4.7. It suffices to prove the statement for an element in the finite generating set $\mathcal{S}$, then apply iteratively. Let $s \in \mathcal{S}$. Let $\{z(k)\}_{k\in\mathbb{N}}$ be any $\mathcal{G}$–coding of sp and let $\{y(k)\}_{k\in\mathbb{N}}$ be a $\mathcal{G}$–coding of p. The left-hand side of the claimed equality is, by definition, $\Phi(sp) = \bigcap_{k=0}^{\infty} \rho(\alpha_{z(1)}) \cdots \rho(\alpha_{z(k)})\overline{W}_{z(k+1)}$.

The right-hand side is

$$\begin{aligned}\rho(s)(\Phi(p)) &= \rho(s) \bigcap_{k=0}^{\infty} \rho(\alpha_{y(1)}) \cdots \rho(\alpha_{y(k)})\overline{W}_{y(k+1)} \\ &= \bigcap_{k=0}^{\infty} \rho(s\alpha_{y(1)}s^{-1}) \cdots \rho(s\alpha_{y(k)}s^{-1})\rho(s)\overline{W}_{y(k+1)}.\end{aligned}$$

Since $\{y(k)\}_{k\in\mathbb{N}}$ gives a $\mathcal{G}^s$–coding of sp we may apply Corollary 3.11 to find some N_s such that for all k,

$$g_{n(k)+N_s}\overline{W}_{z(n(k)+N_s+1)} \subset h_k s W_{y(k+1)},$$

where g_k, h_k, and $n(k)$ are as in Lemma 3.10. Explicitly, this means that for any k,

$$\alpha_{z(1)} \dots \alpha_{z(n(k)+N_s)}\overline{W}_{z(n(k)+N_s+1)} \subset s\alpha_{y(1)}s^{-1} \cdots s\alpha_{y(k)}s^{-1}sW_{y(k+1)},$$

which gives, after multiplying both sides on the left by $(\alpha_{z(1)} \dots \alpha_{z(n(k))})^{-1}$, one of the finitely many containments

$$\alpha_{z(n(k)+1)} \dots \alpha_{z(n(k)+N_s)}\overline{W}_z \subset fsW_y$$

with $f \in F_s$ as in (13).

Our assumption (V3) implies that this containment is still satisfied after perturbation, i.e.

$$\rho(\alpha_{z(n(k)+1)}) \dots \rho(\alpha_{z(n(k)+N_s)})\overline{W}_z \subset \rho(fs)W_y.$$

After multiplying on the left by $\rho(\alpha_{z(1)}\dots\alpha_{z(n(k))})$ we obtain

$$\rho(\alpha_{z(1)}\cdots\alpha_{z(n(k)+N_s)})\overline{W}_{z(n(k)+N_s+1)} \subset \rho(s\alpha_{y(1)}s^{-1}\cdots s\alpha_{y(k)}s^{-1})\rho(s)W_{y(k+1)}$$

which shows that $\Phi(sp) \subset \rho(s)(\Phi(p))$. Applying the same argument using s^{-1} we also see that

$$\rho(s)(\Phi(p)) = \rho(s)\Phi(s^{-1}sp) \subset \rho(s)\rho(s^{-1})\Phi(sp) = \Phi(sp).$$

$\square$

Combined with Lemma 2.9 which states that every point has a $(\mathcal{G},\rho)$ coding, the next lemma shows that the sets $\Phi(p)$ partition $\partial\Gamma$ as p ranges over $\partial\Gamma$.

Lemma 4.8. *If $p \neq q$, then $\Phi(p)\cap\Phi(q) = \emptyset$.*

Proof. First consider the case where $d_\partial(p,q) > D$. Then for any coding $\{z(k)\}_{k\in\mathbb{N}}$ of p and $\{y(k)\}_{k\in\mathbb{N}}$ of q, respectively, we have $W_{z(1)}\cap W_{y(1)} = \emptyset$. Since $\Phi(p) \subset W_{z(1)}$ and $\Phi(q) \subset W_{y(1)}$, this proves the lemma in this case. Lemma 2.1 (that any pair $a \neq b$ can be taken to a pair separated by distance D by some group element) and 4.7 (equivariance) reduce the general case to this one. $\square$

In summary, Lemma 4.3, Lemma 4.5, and Lemma 4.8 together imply that the sets $\Phi(p)$ give a partition of $\partial\Gamma$, indexed by the points in $\partial\Gamma$. Thus this partition defines a surjection $\phi : \partial\Gamma \to \partial\Gamma$, determined by the condition

$$\phi(x) = p \iff x \in \Phi(p).$$

Lemma 4.7 implies the function ϕ is *ρ–equivariant* in the sense that for every $g \in \Gamma$ and $x \in \partial\Gamma$,

$$g\phi(x) = \phi(\rho(g)x).$$

Lemma 4.9. *For every $x \in \partial\Gamma$, we have $d_\partial(x,\phi(x)) \leq \epsilon$.*

Proof. Let $p = \phi(x)$, equivalently $x \in \Phi(p)$. For any $\mathcal{G}$–coding $\{z(k)\}_{k\in\mathbb{N}}$ of p,

$$\Phi(p) = \bigcap_{k=0}^{\infty} \rho(\alpha_{z(1)})\cdots\rho(\alpha_{z(k)})\overline{W}_{z(k+1)}.$$

In particular p and $\Phi(p)$ are both contained in $\overline{W}_{z(1)}$, which has diameter at most ϵ. Since $x \in \Phi(p)$, $d_\partial(x,p) \leq \epsilon$. $\square$

We have now verified all of the conditions needed for ϕ to be a semi-conjugacy in the specified neighborhood $\mathcal{U}$ of the identity, except for the fact that ϕ is continuous. This last condition is implied by the properties already established, as follows.

Lemma 4.10. *Let $\rho : \Gamma \to \mathrm{Homeo}(\partial\Gamma)$ be an action of Γ on its Gromov boundary. For any $\epsilon_0 < D/4$, if $\phi : \partial\Gamma \to \partial\Gamma$ is a ρ–equivariant surjection satisfying $d_\partial(x,\phi(x)) < \epsilon_0$ for all $x \in \partial\Gamma$, then ϕ is continuous.*

Proof. We proceed by contradiction, and suppose that for a sequence $\{x_n\}_{n\in\mathbb{N}}$ in $\partial\Gamma$, x_n converges to x, but $\phi(x_n) = p_n$ does not converge to $\phi(x) = p$. Taking a subsequence, we may assume that p_n converges to $q \neq p$. Since ϕ is ρ–equivariant, we may use Lemma 2.1 to assume that $d_\partial(p,q) > D$.

For sufficiently large n we have $d_\partial(x_n,x) < \epsilon_0$. Then by the triangle inequality we have $d_\partial(p_n,p) \leq d_\partial(p_n,x_n) + d_\partial(x_n,x) + d_\partial(x,p) < 3\epsilon_0$ and thus $d_\partial(p_n,q) > D - 3\epsilon_0 > D/4$, contradicting the fact that $p_n \to q$. $\square$

This concludes the proof of Theorem 1.1.

References

[BH99] M. R. Bridson and A. Haefliger. *Metric Spaces of Non–Positive Curvature*, volume 319 of *Grundlehren der mathematischen Wissenschaften*. Springer–Verlag, Berlin, 1999.

[BM19] J. Bowden and K. Mann. C^0 stability of boundary actions and inequivalent Anosov flows. Preprint, arXiv:1909.02324, 2019. [math.DS].

[Bow98] B. H. Bowditch. A topological characterisation of hyperbolic groups. *J. Amer. Math. Soc.*, 11(3):643–667, 1998.

[Can84] J. W. Cannon. The combinatorial structure of cocompact discrete hyperbolic groups. *Geom. Dedicata*, 16(2):123–148, 1984.

[ECH+92] D. B. A. Epstein, J. W. Cannon, D. F. Holt, S. V. F. Levy, M. S. Paterson, and W. P. Thurston. *Word processing in groups*. Jones and Bartlett Publishers, Boston, MA, 1992.

[Gro87] M. Gromov. Hyperbolic groups. In *Essays in group theory*, volume 8 of *Math. Sci. Res. Inst. Publ.*, pages 75–263. Springer, New York, 1987.

[KKL] M. Kapovich, S. Kim, and J. Lee. Structural stability of meandering-hyperbolic group actions. Preprint, arXiv:1904.06921. [math.GR].

[MM21] K. Mann and J. F. Manning. Stability for hyperbolic groups acting on boundary spheres. Preprint, arXiv:2104.01269, 2021.

[Sul85] D. Sullivan. Quasiconformal homeomorphisms and dynamics. II. Structural stability implies hyperbolicity for Kleinian groups. *Acta Math.*, 155(3-4):243–260, 1985.

[Tuk98] P. Tukia. Conical limit points and uniform convergence groups. *J. Reine Angew. Math.*, 501:71–98, 1998.

[Wei22] T. Weisman. An extended definition of Anosov representation for relatively hyperbolic groups. Preprint, arXiv:2205.07183, 2022. [math.GR].

Department of Mathematics, Cornell University, Ithaca, NY 14853, USA
Email address: `k.mann@cornell.edu`

Department of Mathematics, Cornell University, Ithaca, NY 14853, USA
Email address: `jfmanning@cornell.edu`

Department of Mathematics, The University of Texas at Austin, 2515 Speedway, PMA 8.100, Austin, TX 78712, USA
Email address: `weisman@math.utexas.edu`