

TOPOLOGICAL STABILITY OF RELATIVELY HYPERBOLIC GROUPS ACTING ON THEIR BOUNDARIES

KATHRYN MANN, JASON FOX MANNING, AND THEODORE WEISMAN

ABSTRACT. We prove a topological stability result for the actions of hyperbolic groups on their Bowditch boundaries. More precisely, we show that a sufficiently small perturbation of the standard boundary action, if assumed on each parabolic subgroup to be a perturbation by semi-conjugacy, is in fact always globally semi-conjugate to the standard action. This proves a relative version of the main result of [MMW22]. The assumption of control on the perturbation of parabolics is necessary.

1. INTRODUCTION

The study of stability of boundary actions in the presence of some hyperbolicity has a long history. Sullivan [Sul85] proved that the action of a Kleinian group on its limit set is stable in the sense of C^1 dynamics, meaning that C^1 -close actions remain conjugate. This was generalized by Kapovich–Kim–Lee [KKL] to the broader class of what they call “meandering hyperbolic actions” under Lipschitz–close perturbations, examples that include the boundary actions of hyperbolic groups and of uniform lattices. More recent work has treated the more general question of *topological* or C^0 stability. An action ρ_0 of a group Γ on a topological space X is said to be C^0 *stable* if any action ρ sufficiently close to ρ_0 in $\text{Hom}(\Gamma, \text{Homeo}(X))$ has ρ_0 as a *topological factor*, i.e. there exists a continuous, surjective map $h : X \rightarrow X$ such that $h\rho(\gamma) = \rho_0(\gamma)h$ for all $\gamma \in \Gamma$. The map h is called a *semi-conjugacy*. This is sometimes strengthened to require that h can be taken arbitrarily close to the identity map $X \rightarrow X$ by controlling how close ρ is to ρ_0 .

Topological stability of hyperbolic groups acting on their boundaries was shown in [BM22, MM23, MMW22]; see also [Gro87]. The approach from [BM22] was very recently adapted (with much additional work) to establish topological stability for actions of uniform lattices on Furstenberg boundaries in [CINS23]. The current paper treats the case of boundary actions of *relatively hyperbolic* groups. We prove the following.

Theorem 1.1 (Stability for relatively hyperbolic groups). *Let Γ be hyperbolic relative to \mathcal{P} and ρ_0 the natural action of Γ on the Bowditch boundary $\partial(\Gamma, \mathcal{P})$.*

For any neighborhood \mathcal{V} of the identity in $C(\partial(\Gamma, \mathcal{P}))$, there exists a neighborhood \mathcal{U} of ρ_0 in $\text{Hom}(\Gamma, \text{Homeo}(\partial(\Gamma, \mathcal{P})))$ and a neighborhood \mathcal{V}' of the identity in $C(\partial(\Gamma, \mathcal{P}))$ such that, if $\rho \in \mathcal{U}$ is an action whose restriction to each $P \in \mathcal{P}$ is an extension of ρ_0 via a semi-conjugacy in \mathcal{V}' , then ρ is an extension of ρ_0 by a semi-conjugacy in \mathcal{V} .

Examples. Control on the restriction of ρ to each subgroup $P \in \mathcal{P}$ is necessary. In fact it is necessary even when Γ is a lattice in a Lie group G , and the deformed action ρ is induced by deforming the inclusion representation in the character variety.

Example 1.2. Consider a finite-area hyperbolic surface with a single cusp, isometric to a quotient \mathbb{H}^2/Γ for a discrete subgroup $\Gamma \subset \mathrm{PSL}(2, \mathbb{R})$. Then Γ is a relatively hyperbolic group, relative to the fundamental group of the cusp, and the action of this group on its Bowditch boundary is the action of Γ on $\mathbb{RP}^1 = \partial\mathbb{H}^2$.

- (i) Let $\rho : \Gamma \rightarrow \mathrm{PSL}(2, \mathbb{R})$ be a small deformation of the inclusion representation $\rho_0 : \Gamma \hookrightarrow \mathrm{PSL}(2, \mathbb{R})$ for which a generator of the cusp group acts by an elliptic transformation. As elliptic transformations have no fixed points in $\partial\mathbb{H}^2$, this deformed action cannot be semi-conjugate to the original boundary action.
- (ii) Now let $\rho : \Gamma \rightarrow \mathrm{PSL}(2, \mathbb{R})$ be a deformation of ρ_0 where the cusp group instead acts by a loxodromic transformation, which fixes a pair of points in $\partial\mathbb{H}^2$ close to the parabolic fixed point for the original cusp group action. The map $\partial\mathbb{H}^2 \rightarrow \partial\mathbb{H}^2$ which collapses the small arc A joining this pair of fixed points is a semi-conjugacy for the ρ -action of the cusp group, which is close to the identity. Additionally collapsing all of the arcs in the $\rho(\Gamma)$ -orbit of A yields a semi-conjugacy for the ρ -action of the whole group Γ . This is an example of the situation described by Theorem 1.1.

It is actually possible to generalize both of the examples above to higher dimensions.

Example 1.3. Suppose that M is a finite-volume non-compact hyperbolic n -manifold, isometric to \mathbb{H}^n/Γ for a discrete group $\Gamma \subset \mathrm{Isom}(\mathbb{H}^n) \simeq \mathrm{PO}(n, 1)$; as in the previous example, Γ is then relatively hyperbolic, relative to its collection of cusp subgroups, and the action on its Bowditch boundary is the induced action on $\partial\mathbb{H}^n \simeq S^{n-1}$.

- (i) When $n = 3$, Thurston's hyperbolic Dehn filling theorem implies that there are arbitrarily small deformations of the inclusion $\rho_0 : \Gamma \hookrightarrow \mathrm{PO}(n, 1)$ for which the action on \mathbb{H}^3 (and hence the induced action on $\partial\mathbb{H}^3$) has infinite kernel. These deformations cannot be semi-conjugate to the original action, where the kernel is trivial.
- (ii) For each $n \geq 2$, there are examples of (non-uniform) lattices $\Gamma \subset \mathrm{PO}(n, 1)$ which have deformations in $\mathrm{Hom}(\Gamma, \mathrm{PGL}(n+1, \mathbb{R}))$ which are discrete, faithful, and preserve a convex open subset $\Omega \subset \mathbb{RP}^n$, arbitrarily close (but not equivalent) to the projective model for \mathbb{H}^n ; in some of these examples, the cusp groups preserve small k -simplices embedded in $\partial\Omega$, for some $1 \leq k < n$ (see [BDL18], [BM20], [Bob19]). The induced action of Γ on $\partial\Omega \simeq S^{n-1}$ then cannot be conjugate to the original action of Γ on $\partial\mathbb{H}^n$. However, it follows from [Wei23, Section 5] that (for sufficiently small deformations) the map collapsing all of the k -simplices is a semi-conjugacy to the standard action of Γ on its Bowditch boundary.

We also note that, in the broad context of C^0 deformations covered by Theorem 1.1, even if the restriction of ρ to each peripheral subgroup is actually conjugate to the restriction of ρ_0 , the ρ -action of the full group Γ does *not* need to be conjugate to ρ_0 . This can occur even when Γ is hyperbolic and $\mathcal{P} = \emptyset$; for examples see [BM22, Section 4] and [MMW22, Example 1.4].

Outline. The broad strategy of this work follows that of [MMW22], but much additional technical work is needed in the presence of parabolic elements. We adapt

and use tools and ideas from [Wei22] and the geometry of relatively hyperbolic groups. A reader looking for a gentler introduction to rigidity of boundary actions may wish to first read the simpler proof in [MMW22].

In Section 2 we review necessary material on relatively hyperbolic groups and set the stage for the proof. In Section 3 we define an automaton which codes boundary points, adapted to the setting of the desired stability theorem for relatively hyperbolic groups. Essential properties of this automaton are proved in Section 4

Section 5 is the technical heart of the paper. The key proposition (Proposition 5.2) is a “uniform nesting” condition for sequences of nested sets furnished by the automaton, which can be translated into a stable condition under perturbation that allows us to define the desired semi-conjugacy between the standard boundary action and a sufficiently small perturbation. Section 6 uses all the previous work to define this semi-conjugacy and conclude the proof.

Acknowledgments. KM was partially supported by NSF grant DMS 1844516 and a Sloan fellowship. JM was partially supported by the Simons Foundation, grant number 942496. TW was partially supported by NSF grant DMS-2202770.

2. SET-UP

We assume the reader has basic familiarity with the theory of relatively hyperbolic groups; general background can be found in [Bow12], or [GM08, Section 2] and references therein. In this section we set notation and recall the essential properties that we will use.

Let Γ be a relatively hyperbolic group, relative to a finite collection \mathcal{P} of infinite subgroups. Fix a finite, symmetric generating set \mathcal{S} for Γ . We let $\text{Cay}(\Gamma) = \text{Cay}(\Gamma, \mathcal{S})$ denote the Cayley graph of Γ with respect to the generating set \mathcal{S} , and we let d_Γ denote the metric on $\text{Cay}(\Gamma)$ induced by this generating set. We also assume the generating set \mathcal{S} is compatible with \mathcal{P} , i.e. that for every parabolic subgroup $P \in \mathcal{P}$, the intersection $P \cap \mathcal{S}$ is a generating set for P . If Γ is hyperbolic (i.e. $\mathcal{P} = \emptyset$) then Theorem 1.1 follows from [MMW22], so we assume for the duration of this work that $\mathcal{P} \neq \emptyset$. We also assume that (Γ, \mathcal{P}) is non-elementary (meaning that $\mathcal{P} \neq \{\Gamma\}$ and Γ is not finite or virtually cyclic), since in these cases the Bowditch boundary $\partial(\Gamma, \mathcal{P})$ contains at most two points and the theorem is trivial.

The Bowditch boundary $\partial(\Gamma, \mathcal{P})$ can be identified with the Gromov boundary of a hyperbolic space $X = X(\Gamma, \mathcal{P}, \mathcal{S})$ which was defined in [GM08] and called a *cusped space* for the pair (Γ, \mathcal{P}) . The space X is a locally finite metric graph, with each edge having length 1. We use d_X to refer to the metric on X . The group Γ acts properly by isometries on X , and the Cayley graph $\text{Cay}(\Gamma, \mathcal{S})$ embeds properly and Γ -equivariantly (though not quasi-isometrically) as a subgraph of X .

We will need to use some results on the geometry of horoballs in X and geodesics near horoballs, especially in Section 5. To this end, define the *depth* D_X of a vertex of X to be its distance from the Cayley graph in X and extend linearly across edges to get a continuous function $D_X : X \rightarrow [0, \infty)$. Thus, the Cayley graph in X is the set $D_X^{-1}(0)$; a *horoball* is the smallest full subgraph containing the closure of a component of $D_X^{-1}(0, \infty)$. For an integer $k > 0$, we refer to a component of $D_X^{-1}[k, \infty)$ as a *k-horoball*. If \mathcal{H} is a horoball, we say that a geodesic $\gamma : [a, b] \rightarrow \mathcal{H}$

is *regular* if there are $a \leq A < B \leq b$ with $B - A \leq 3$ and

$$\frac{d}{dt}D_X \circ \gamma(t) = \begin{cases} 1 & t < A \\ 0 & A < t < B \\ -1 & t > B \end{cases}.$$

A geodesic in a horoball is *vertical* if it is regular and either $a = A = B$ or $A = B = b$. We will say a geodesic in X is *regular* if every intersection with a horoball is regular. Whenever necessary, we will assume our geodesics are regular. We also fix an integer $\delta \geq 1$ so that the cusped space X is δ -hyperbolic (in the sense that all geodesic triangles are δ -thin).

The following result was proved in [GM08].

Lemma 2.1. [GM08, Lemmas 3.10, 3.26] *For any $k \geq \delta + 1$, the k -horoballs are convex. Moreover, any geodesic in a horoball is Hausdorff distance at most 4 from a regular unit speed geodesic $\gamma: [a, b] \rightarrow X$ with the same endpoints.*

We also need one more lemma about the geometry of geodesics which pass through horoballs.

Lemma 2.2. *Let \mathcal{H} be a k -horoball of X for some $k \geq 0$, let $\sigma: [a, b] \rightarrow X$ be a regular unit speed geodesic, and let $[a', b']$ be a connected component of $\sigma^{-1}(\mathcal{H})$. Assume that $b' - a' \geq (4\delta + 3)$. Then for $t \in [a, a']$ we have*

$$(a' - t) - \delta \leq d_X(\sigma(t), \mathcal{H}) \leq a' - t.$$

Proof. Let x be a closest point in \mathcal{H} to $\sigma(t)$, and let $y = \sigma(a')$. Any geodesic from $\sigma(t)$ to x can be extended by a vertical path to a point x' on the $k + 2\delta$ -horoball \mathcal{H}' nested inside \mathcal{H} . Call this extended geodesic τ . Now consider a geodesic triangle two of whose sides are τ and $\sigma(t, a' + 2\delta)$. (Regularity of σ implies that $\sigma(a' + 2\delta)$ is on the boundary of \mathcal{H}' .) The third side of the triangle has endpoints in \mathcal{H}' , which is convex by the first part of Lemma 2.1. In particular the distance from y to this third side is at least 2δ , so there is a point y' on $\sigma(t, a' + 2\delta)$ which is within δ of y . Either y' lies between $\sigma(t)$ and x , or on $\tau \cap \mathcal{H}$ at depth at most δ . In either case, we deduce that $a' - t \leq d_X(\sigma(t), x) + \delta$, as desired. \square

The Gromov boundary of the cusped space X is the Bowditch boundary $\partial(\Gamma, \mathcal{P})$ of the pair (Γ, \mathcal{P}) (see [Bow12], [GM08]). We fix a metric d_∂ on $\partial(\Gamma, \mathcal{P})$. Metric notions such as diameter, ϵ -neighborhoods are always with respect to this metric. We write $B_r(x)$ for the (open) ball about x of radius r , and $N_r(Y)$ for the open r -neighborhood of a subset Y .

When we need to distinguish the natural action of Γ on $\partial(\Gamma, \mathcal{P})$ from another action of this group on the space, we will use the notation $\rho_0(g)(x)$ for the action of $g \in \Gamma$ on $x \in \partial(\Gamma, \mathcal{P})$. However, in the first part of this work we use *only* this action, and so for convenience shorten this to gx .

Each subgroup in \mathcal{P} acts on $\partial(\Gamma, \mathcal{P})$ with a unique fixed point. We denote the finite set of fixed points of groups in \mathcal{P} by Π , and for $p \in \Pi$, we let $\Gamma_p \in \mathcal{P}$ denote the subgroup fixing p . Thus,

$$\mathcal{P} = \{\Gamma_p : p \in \Pi\}.$$

Any Γ -translate in $\partial(\Gamma, \mathcal{P})$ of a point in Π is called a *parabolic point*.

We now recall some important properties of the action of Γ on $\partial(\Gamma, \mathcal{P})$, which we use to set up the next construction. First, the group Γ acts on $\partial(\Gamma, \mathcal{P})$ as a

convergence group (see [Tuk94]), meaning that the induced action on the space of distinct triples in $\partial(\Gamma, \mathcal{P})$ is properly discontinuous.

Secondly, Γ acts cocompactly on pairs of distinct points in the Bowditch boundary $\partial(\Gamma, \mathcal{P})$. Thus, we can set the following separation constant:

Definition 2.3. We fix a constant $D > 0$ such that for any pair of distinct points $x, y \in \partial(\Gamma, \mathcal{P})$, there is some $g \in \Gamma$ such that $d_\partial(gx, gy) > D$.

Finally, the action of Γ on $\partial(\Gamma, \mathcal{P})$ is *geometrically finite* in a dynamical sense, meaning that every non-parabolic point in $\partial(\Gamma, \mathcal{P})$ is a *conical limit point*, and every parabolic point is *bounded*.

We recall that a point $z \in \partial(\Gamma, \mathcal{P})$ is a conical limit point if there are distinct points $a, b \in \partial(\Gamma, \mathcal{P})$ and a sequence $(g_i)_{i \in \mathbb{N}}$ so that $g_i z \rightarrow b$ and $g_i x \rightarrow a$ uniformly on compacts in $\partial(\Gamma, \mathcal{P}) - \{z\}$. A parabolic point p is *bounded* if its stabilizer $\Gamma_p \subset \Gamma$ acts with compact quotient on the space $\partial(\Gamma, \mathcal{P}) - \{p\}$. Thus, we define the following.

Definition 2.4. For each parabolic point $p \in \Pi$, we fix a compact set $K_p \subset \partial(\Gamma, \mathcal{P}) - \{p\}$ such that $\Gamma_p \cdot K_p = \partial(\Gamma, \mathcal{P}) - \{p\}$.

For each p , since K_p is compact, the distance $d_\partial(x, p)$ is bounded below by a positive constant. Further, since we assume Γ is non-elementary, $\partial(\Gamma, \mathcal{P})$ is uncountable. Since Γ_p is countable, K_p must have positive diameter. So, we can also make the following definition.

Definition 2.5. We fix a constant $D_\Pi > 0$ so that for each $p \in \Pi$, we have $D_\Pi < \text{diam}(K_p)$ and $D_\Pi < d_\partial(K_p, p)$.

3. AUTOMATON

In this section we give a version of the construction in Section 2 of [MMW22] (see also Sections 5 and 6 of [Wei22]). The former paper shows that every point in the Gromov boundary of a hyperbolic group has an “expanded neighborhood” which is well-adapted to the construction of an automaton which codes points. In the setting of this paper, conical points can be coded in essentially the manner of [MMW22]. However, parabolic fixed points need separate treatment, for which we adapt the approach in [Wei22].

We start with the conical case, adapting Lemma 2.3 in [MMW22] as follows.

Lemma 3.1 (Expanded neighborhoods). *For any positive $\epsilon < \frac{D}{5}$ and any conical limit point $z \in \partial(\Gamma, \mathcal{P})$, there exists $\alpha_z \in \Gamma$ and a pair of open neighborhoods $V(z) \subset W(z)$ of z so that*

- (1) $\text{diam}(W(z)) \leq \epsilon$;
- (2) $\text{diam}(\alpha_z^{-1}W(z)) > 4\epsilon$; and
- (3) $\overline{N_{2\epsilon}(\alpha_z^{-1}V(z))} \subset \alpha_z^{-1}W(z)$.

Proof. We choose some $\epsilon < \frac{D}{5}$ where D is the constant from Definition 2.3.

Let $z \in \partial(\Gamma, \mathcal{P})$ be a conical limit point. Since z is conical, we can find distinct points a, b and a sequence of group elements $(g_i)_{i \in \mathbb{N}}$ so that $g_i z \rightarrow b$ and $g_i x \rightarrow a$ uniformly for all $x \neq z$. Up to post-composing all g_i with a fixed element g as in Definition 2.3 if necessary, we may assume $d_\partial(a, b) \geq D$. Also, since $\partial(\Gamma, \mathcal{P})$ 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 be the complement of $W(z)$ in $\partial(\Gamma, \mathcal{P})$. The set K is compact and does not contain z , so for i sufficiently large, we have $g_i K \subset B_{\epsilon'}(a)$ and $g_i z \in B_\epsilon(b)$.

Fixing some such i , set $\alpha_z = g_i^{-1}$, and let $V(z) = \alpha_z(B_\epsilon(b))$. Note that $B_\epsilon(a)$ contains $\alpha_z^{-1}K = \partial(\Gamma, \mathcal{P}) - \alpha_z^{-1}W(z)$. Since $B_\epsilon(a)$ is disjoint from $B_\epsilon(b)$, we have

$$V(z) = \alpha_z(B_\epsilon(b)) \subset \partial(\Gamma, \mathcal{P}) - \alpha_z(K) = W_z.$$

The set $\alpha_z^{-1}W(z) = \partial(\Gamma, \mathcal{P}) - \alpha_z^{-1}K$ contains both b and a' , so $\text{diam}(\alpha_z^{-1}W(z)) \geq d_\partial(b, a') \geq D - \epsilon > 4\epsilon$, establishing Property (2).

Finally, since $d_\partial(b, \alpha_z^{-1}K) \geq D - \epsilon > 4\epsilon$, we have

$$\overline{N_{2\epsilon}(\alpha_z^{-1}V(z))} \subset \overline{B_{3\epsilon}(b)} \subset (\partial(\Gamma, \mathcal{P}) - \alpha_z^{-1}K) = \alpha_z^{-1}W(z),$$

establishing Property (3). \square

To treat parabolic points, we follow an argument given in [Wei22]; compare the lemma below to [Wei22, Lemma 6.7].

Lemma 3.2 (Parabolic points). *For each point $p \in \Pi$, any $\epsilon < D_\Pi/5$, and each $q = gp$, there exist open sets $\hat{V}(p) \subset \hat{W}(p) \subset \partial(\Gamma, \mathcal{P})$, neighborhoods $V(q) \subset W(q)$ of q and a finite set $F_q \subset g\Gamma_p$ such that:*

- (1) $\text{diam}(W(q)) \leq \epsilon$;
- (2) $\text{diam}(\hat{W}(p)) > 4\epsilon$;
- (3) $\overline{N_{2\epsilon}(\hat{V}(p))} \subset \hat{W}(p)$;
- (4) $\overline{N_\epsilon(\hat{W}(p))}$ does not contain p ;
- (5) The set $W(q)$ is equal to $\{q\} \cup \bigcup_{\alpha \in g\Gamma_p - F_q} \alpha \hat{W}(p)$;
- (6) The set $V(q)$ is equal to $\{q\} \cup \bigcup_{\alpha \in g\Gamma_p - F_q} \alpha \hat{V}(p)$.

Proof. Let $q \in \partial(\Gamma, \mathcal{P})$ be a parabolic point with $q = gp$ for $p \in \Pi$. Let K_p be the compact set from Definition 2.4, so that $\Gamma_p \cdot K_p = \partial(\Gamma, \mathcal{P}) - \{p\}$, $\text{diam}(K_p) > D_\Pi$, and $d_\partial(x, p) > D_\Pi$ for every $x \in K_p$. We let $\hat{V}(p) = N_\epsilon(K_p)$, and we let $\hat{W}(p) = N_{4\epsilon}(K_p)$.

Conditions (2) and (3) above are immediate. Further, since $N_\epsilon(\hat{W}(p)) \subset N_{5\epsilon}(K_p)$ and we assume $\epsilon < D_\Pi/5$, we know that the closed neighborhood $\overline{N_\epsilon(\hat{W}(p))}$ does not contain p , so Condition (4) holds as well.

The stabilizer of a parabolic point $p \in \partial(\Gamma, \mathcal{P})$ acts properly discontinuously on $\partial(\Gamma, \mathcal{P}) - \{p\}$. So, for any neighborhood U of p in $\partial(\Gamma, \mathcal{P})$, there is a finite set F (depending on U) so that for any $\alpha \in \Gamma_p - F$, we have $\alpha \overline{\hat{W}(p)} \subset U$. Thus, by taking $U = g^{-1}B_{\epsilon/2}(q)$, we see that for $F_q = gF$, we have $\alpha \overline{\hat{W}(p)} \subset B_{\epsilon/2}(q)$ for all $\alpha \in g\Gamma_p - F_q$.

We now define $W(q)$ and $V(q)$ exactly as stated in Conditions (5) and (6). Our choice of F_q ensures that $W(q) \subset B_{\epsilon/2}(q)$, so we know Condition (1) holds.

It remains to verify that $W(q)$ and $V(q)$ are open neighborhoods of q . For this, it is enough to show that $W(q)$ and $V(q)$ each contain an open neighborhood of q . However, since $g\Gamma_p \cdot K_p = \partial(\Gamma, \mathcal{P}) - \{p\}$, and $\hat{V}(p)$ and $\hat{W}(p)$ both contain K_p , we know that $V(q)$ and $W(q)$ both contain the complement of the set $\bigcup_{\beta \in F_q} \beta K_p$. This set is a finite union of compact subsets not containing q , so its complement contains an open neighborhood of q . \square

3.1. Fixing ε and the geometric automaton. The construction of the automaton depends on a pair of nested finite open covers of $\partial(\Gamma, \mathcal{P})$, coming from the sets $V(z)$ and $W(z)$ provided by Lemma 3.1 and Lemma 3.2. The covers we construct via these lemmas depend on the target neighborhood of the identity \mathcal{V} needed in Theorem 1.1. Thus, we suppose now that we have been given some $\mathcal{V} \subset C(\partial(\Gamma, \mathcal{P}))$, and make the following definition.

Definition 3.3. We fix a constant $\varepsilon > 0$ so that ε is smaller than $\min(D/5, D_\Pi/5)$, and additionally such that every continuous self-map $\partial(\Gamma, \mathcal{P}) \rightarrow \partial(\Gamma, \mathcal{P})$ ε -close to the identity lies in \mathcal{V} .

For each conical limit point $z \in \partial(\Gamma, \mathcal{P})$, choose a pair of open neighborhoods $V(z)^* \subset W(z)$ of z and an element $\alpha(z)$ as in the statement of Lemma 3.1, for our chosen ε . Later, we will modify the $V(z)^*$ slightly, hence the provisional $*$ in the notation, rather than simply calling the set $V(z)$ as in the Lemma statement.

For each parabolic point $p \in \Pi$ choose open sets $\hat{V}(p)^* \subset \hat{W}(p)$ as in Lemma 3.2, again using the fixed ε . In addition, for each parabolic point $q \in \partial(\Gamma, \mathcal{P})$ with $q = gp$ for $p \in \Pi$, choose neighborhoods $V(q)^* \subset W(q)$ of q and a finite set $F_q \subset g\Gamma_p$ as in the same lemma.

Let $Z \subset \partial(\Gamma, \mathcal{P})$ be a finite collection so that the sets $\{V(z)^*\}_{z \in Z}$ cover $\partial(\Gamma, \mathcal{P})$. We streamline notation as follows.

Notation 3.4. For each conical limit point $z \in Z$, we denote $\alpha_z^{-1}V(z)^*$ by $\hat{V}(z)^*$ and $\alpha_z^{-1}W(z)$ by $\hat{W}(z)$. These sets will play an important role.

If $q = gp$ for $p \in \Pi$, we will define $\hat{V}(q)^* := \hat{V}(p)^*$.

We will set up an automaton with vertex set Z , and edges determined by the combinatorics of the intersections these sets and labeled by certain elements of Γ . For this we use the following definition.

Definition 3.5. For each point $z \in Z$, let $L(z) \subset \Gamma$ be defined as follows.

- If z is a conical limit point, then $L(z)$ is the singleton $\{\alpha_z\}$.
- If z is a parabolic point $q = gp$ for $p \in \Pi$, then we define $L(z) := g\Gamma_p - F_q$, where F_q is the set given by Lemma 3.2.

Since we ultimately want to use this automaton to prove stability properties of the action of Γ on $\partial(\Gamma, \mathcal{P})$, we need to make sure that this intersection pattern is stable. More precisely, we want to modify the sets $V(z)^*$ and $\hat{V}(z)^*$ slightly to sets $V(z)$ and $\hat{V}(z)$ which satisfy $\overline{\hat{V}(z)} \cap \overline{V(y)} = \emptyset$ if and only if $\hat{V}(z) \cap V(y) = \emptyset$.

For this we use the following lemma.

Lemma 3.6. *For any $z \in Z$ and any $\eta > 0$, there exists $r > 0$ so that the set*

$$\bigcup_{\alpha \in L(z)} \alpha N_r(\hat{V}(z)^*)$$

is contained in the η -neighborhood of $V(z)^$.*

Proof. When z is a conical limit point, then $L(z)$ is a singleton $\{\alpha_z\}$ and $V(z)^* = \alpha_z \hat{V}(z)^*$. Thus, $\alpha_z N_r(\hat{V}(z)^*) \subset N_\eta(\alpha_z \hat{V}(z)^*)$ holds for all sufficiently small r because Γ acts by homeomorphisms. Now assume that z is parabolic point $q = gp$ for $p \in \Pi$, and $L(z) = g\Gamma_p - F_q$.

Part 4 of Lemma 3.2 ensures that we can choose r small enough so that the closure of $N_r(\hat{V}(p)^*)$ does not contain p . Then, for a finite subset $E_\eta \subset g\Gamma_p$, every $\alpha \in g\Gamma_p - E_\eta$ satisfies $\alpha N_r(\hat{V}(p)) \subset B_\eta(q)$. As $V(q)$ contains q , this ensures that $\bigcup_{\alpha \in g\Gamma_p - E_\eta} N_r(\hat{V}(p)^*)$ is contained in $N_\eta(V(q)^*)$.

Since the set $E_\eta - F_q$ is finite, we can choose r small enough so that $\alpha N_r(\hat{V}(p)^*)$ is contained in $N_\eta(\alpha \hat{V}(p)^*)$ for every $\alpha \in E_\eta - F_q$. Since $V(q)$ contains the union of sets $\alpha \hat{V}(p)^*$ for $\alpha \in g\Gamma_p - F_q$, this guarantees that the desired inclusion holds. \square

The following proposition collects the key properties of the combinatorics of an open cover that we will use to build the automaton.

Proposition 3.7. *There exist sets $V(z) \supset V(z)^*$ and $\hat{V}(z) \supset \hat{V}(z)^*$ so that for each $z \in Z \setminus \Pi$ we have $\hat{V}(z) = \alpha_z^{-1}V(z)$, and for every $z \in Z$, we have:*

- (C1) $\text{diam}(W(z)) < \varepsilon$;
- (C2) $\text{diam}(\hat{W}(z)) > 4\varepsilon$;
- (C3) $\overline{N_{2\varepsilon}(\hat{V}(z))} \subset \hat{W}(z)$
- (C4) *The set $W(z)$ is equal to $\{z\} \cup \bigcup_{\alpha \in L(z)} \alpha \hat{W}(z)$;*
- (C5) *The set $V(z)$ is equal to $\{z\} \cup \bigcup_{\alpha \in L(z)} \alpha \hat{V}(z)$ and $\overline{V(z)} \subset W(z)$.*
- (C6) *For any pair $y, z \in Z$, we have $\overline{\hat{V}(z)} \cap \overline{V(y)} = \emptyset$ iff $\hat{V}(z) \cap V(y) = \emptyset$.*

Note that $\{V(z)\}_{z \in Z}$ is still a cover of $\partial(\Gamma, \mathcal{P})$, and the index set Z , elements α_z , and sets $W(z), \hat{W}(z)$ and $L(z)$ remain unchanged.

Proof. That the sets $V(z)^*$ and $\hat{V}(z)^*$ already chosen satisfy conditions (C1)-(C5) nearly follows from Lemma 3.1 and 3.2. For the last part of (C5) in the case of parabolic z we must also use Lemma 3.6.

We can replace each of the sets $\hat{V}(z)^*$ with a slightly larger set $\hat{V}(z)$, and define $V(z) = \alpha_z(\hat{V}(z))$ for conical z , and via the expression in (C5) for parabolic z . Lemma 3.6 tells us that we can do this so that $V(z), \hat{V}(z)$ are respectively contained in arbitrarily small neighborhoods of $V(z)^*$ and $\hat{V}(z)^*$, and property (C6) above holds for $V(z), \hat{V}(z)$.

By choosing $V(z)$ and $\hat{V}(z)$ sufficiently close to $V(z)^*, \hat{V}(z)^*$, we can ensure that all of the properties (C1)-(C5) still hold after we replace $V(z)^*, \hat{V}(z)^*$ with $V(z), \hat{V}(z)$. (Note that Condition (C3) is open, since it is about the *closure* of $N_{2\varepsilon}(\hat{V}(z))$.) In particular the condition is preserved when $\hat{V}(z)$ is enlarged slightly. \square

For the rest of the paper, we fix the index set Z , the open cover $V(z)$, as well as open sets $\hat{V}(z), W(z), \hat{W}(z)$, and the sets $L(z)$ for each $z \in Z$, and assume that these sets satisfy properties (C1)-(C6) above.

Definition 3.8 (The automaton). Let \mathcal{G} be the directed graph with vertex set Z . For $(z, y) \in Z \times Z$, if $\hat{V}(z) \cap V(y) = \emptyset$ there are no directed edges from z to y , and if $\hat{V}(z) \cap V(y) \neq \emptyset$, then for each $\alpha \in L(z)$ we put a directed edge from z to y labeled by α . Thus, the set of edges from z to y is either empty or in bijective correspondence with $L(z)$. See Figure 1 and Figure 2.

Note that there are infinitely many outgoing edges from $z \in Z$ if and only if z is parabolic. We have constructed our automaton so that it satisfies the following key property:

Proposition 3.9. *If there is an edge from z to y in \mathcal{G} labeled by a group element $\alpha \in L(z)$, then we have the inclusions*

$$(E1) \quad \alpha \overline{N}_\varepsilon(W(y)) \subsetneq \alpha \hat{W}(z) \subset W(z).$$

Proof. If there is an edge from z to y , then $\hat{V}(z) \cap V(y)$ is nonempty. By (C5), we have $V(y) \subset W(y)$, so $W(y) \cap \hat{V}(z)$ is nonempty. By (C3) we have $\overline{N}_{2\varepsilon}(\hat{V}(z)) \subset \hat{W}(z)$. Since $\text{diam}(W(y)) < \varepsilon$ by (C1), we have $\overline{N}_\varepsilon(W(y)) \subset \hat{W}(z)$. Moreover, since the diameter of $W(y)$ is at most ε , and the diameter of $\hat{W}(z)$ is at least 4ε by (C2), this inclusion is proper. This proves the left-hand inclusion above.

Finally, property (C4) implies that if $\alpha \in L(z)$, then $\alpha \hat{W}(z) \subset W(z)$, which gives us the right-hand inclusion as well. \square




FIGURE 1. If z is conical, and \hat{V}_z meets V_y , then there is an edge from z to y labeled α_z . The group element α_z may or may not fix the point z .

We note the following consequence of Proposition 3.9 for future use.

Remark 3.10. Consider a parabolic vertex $z = gp \in Z$, where $p \in \Pi$. If there is an edge from z to another vertex $y \in Z$, by (E1) we have we have $\alpha \overline{N}_\varepsilon(W(y)) \subsetneq \alpha \hat{W}(z) \subset W(z)$ for all $\alpha \in L(z)$, hence, for all but *finitely many* $\alpha \in g\Gamma_p$. It follows that the closed neighborhood $\overline{N}_\varepsilon(W(y))$ does not contain p .

As a consequence, for any given $\eta > 0$ and any parabolic vertex $z \in Z$, if there is an edge from z to y , then the inclusion $\alpha N_\varepsilon(W(y)) \subset B_\eta(z)$ holds for all but finitely many $\alpha \in L(z)$. In particular, by choosing η sufficiently small, we can ensure that for all but finitely many exceptional $\alpha \in L(z)$, we have

$$N_\eta(\alpha N_\varepsilon(W(y))) \subset W(z).$$

Further, since the edge inclusion condition (E1) still holds for the finitely many exceptional α , there is some $\varepsilon_z > 0$ so that for *every* $\alpha \in L(z)$, we have

$$(1) \quad N_{\varepsilon_z}(\alpha(N_\varepsilon W(y))) \subset W(z).$$




FIGURE 2. If z is parabolic, and \hat{V}_z meets V_y , then for each α_i in $L(z)$ there is an edge from z to y labeled α_i . In general α_i may or may not fix z , meaning that $W_z \cap \alpha_i W_z$ may or may not be empty.

4. PROPERTIES OF THE ρ_0 -AUTOMATON

In this section we explain how to associate points of $\partial(\Gamma, \mathcal{P})$ to edge paths in \mathcal{G} , called “codings.” While points of $\partial(\Gamma, \mathcal{P})$ may have more than one coding, we also show that any two codings of the same point are geometrically related.

Notation 4.1. We will often need to work with both finite-length and infinite-length edge paths in the graph \mathcal{G} ; some of our results will apply to all edge paths, while others may apply only to infinite paths or only to finite paths.

Whenever we refer to an edge path (or any other sequence) which may be either finite or infinite, we will let I denote an index set for the sequence, which can be equal to either $\mathbb{N} = \{1, \dots\}$ or $\{1, \dots, n\}$ for some n , depending on context.

Definition 4.2. A *strict conical \mathcal{G} -coding* is an infinite edge path in \mathcal{G} . A *strict parabolic \mathcal{G} -coding* is a finite edge path terminating in a parabolic point. Note that it is possible that this edge path has length zero, in which case the coding is just a single parabolic point $z \in Z$.

We use the notation $\iota(e)$, $\tau(e)$, and $\mathbf{Lab}(e)$ to denote (respectively) the initial vertex, terminal vertex, and label of an edge e in a directed graph \mathcal{G} . If $\mathbf{e} = (e_k)_{k \in I}$ is an edge path in \mathcal{G} , we call $(\tau(e_k))_{k \in I}$ the *terminal vertex sequence* and $(\mathbf{Lab}(e_k))_{k \in I}$ the *label sequence* for \mathbf{e} . We also define the *initial vertex sequence* for \mathbf{e} in the same way, except that for convenience we index this sequence starting from zero (so its k th term is $\iota(e_{k+1})$).

Definition 4.3. For a strict conical coding $\mathbf{e} = (e_k)_{k \in \mathbb{N}}$ with label sequence $(\alpha_k)_{k \in \mathbb{N}}$, and terminal vertex sequence $(z_k)_{k \in \mathbb{N}} = (\tau(e_k))_{k \in \mathbb{N}}$, if

$$\zeta \in \bigcap_{k=0}^{\infty} \alpha_1 \cdots \alpha_k \overline{W(z_k)}$$

we say that \mathbf{e} is a *strict \mathcal{G} -coding of ζ* . If \mathcal{G} is understood, we may omit it, and speak of a *strict coding of ζ* .

Similarly, a strict parabolic coding with label sequence $(\alpha_k)_{k \in \{1, \dots, n\}}$ is a *strict \mathcal{G} -coding of the parabolic point ζ* if

$$\zeta = \alpha_1 \cdots \alpha_n q,$$

where q is the terminal point of the last edge. If \mathcal{G} is understood, we may speak simply of a sequence that (strictly) *codes ζ* .

Ultimately we want to use \mathcal{G} -codings of points in $\partial(\Gamma, \mathcal{P})$ to understand perturbations of the Γ -action on $\partial(\Gamma, \mathcal{P})$, so it is useful to introduce a formalism which also allows Γ to act on the set of codings. In [KKL], this kind of idea is referred to as “Sullivan’s trick.”

Definition 4.4. A *generalized \mathcal{G} -coding* is a pair (g_0, \mathbf{e}) , where g_0 is any element of Γ , and \mathbf{e} is a strict \mathcal{G} -coding. The generalized coding (g_0, \mathbf{e}) is *conical* if \mathbf{e} is conical, and *parabolic* if \mathbf{e} is parabolic.

If \mathbf{e} codes a point $\xi \in \partial(\Gamma, \mathcal{P})$, then we say that (g_0, \mathbf{e}) is a \mathcal{G} -coding of the point $g_0\xi$. We refer to the element g_0 as the *initial point* of the coding (g_0, \mathbf{e}) . Slightly abusing terminology, we also refer to the terminal vertex sequence of \mathbf{e} as the *terminal vertex sequence* for (g_0, \mathbf{e}) , and similarly for label sequences and initial vertex sequences.

Note that any strict coding can be viewed as a generalized coding by taking the initial point g_0 to be the identity. Occasionally we will refer to generalized \mathcal{G} -codings as “ \mathcal{G} -codings,” or even just “codings” if \mathcal{G} is understood from context.

Lemma 4.5. *Every point in $\partial(\Gamma, \mathcal{P})$ has a strict \mathcal{G} -coding.*

Proof. Fix a point $\zeta \in \partial(\Gamma, \mathcal{P})$. Since the $V(z)$ cover $\partial(\Gamma, \mathcal{P})$, we can choose z_0 so that $\zeta \in V(z_0)$. If $\zeta = z_0$ and z_0 is parabolic, we stop. The length-zero edge path consisting of the vertex z_0 is a strict parabolic coding of ζ .

Otherwise, we inductively define a strict \mathcal{G} -coding as follows. Assume that we have already defined a (possibly empty) edge path e_1, \dots, e_n starting at the point z_0 chosen in the last paragraph, and with $\tau(e_k) = z_k$ for each $k > 0$. Set $\alpha_k = \mathbf{Lab}(e_k)$ and make the inductive hypothesis that

$$\zeta_n := \alpha_n^{-1} \cdots \alpha_1^{-1} \zeta \in V(z_n).$$

This implies that

$$\zeta \in \alpha_1 \cdots \alpha_n W(z_n).$$

There are now two possibilities. If z_n is not a parabolic point, then we let α_{n+1} be the only element in the singleton set $L(z_n)$. Then define $\zeta_{n+1} = \alpha_{n+1}^{-1} \zeta_n$, which lies in $\alpha_{n+1}^{-1} V(z_n) = \hat{V}(z_n)$. Since the sets $V(z)$ cover, we can find some $z_{n+1} \in Z$ so that $\zeta_{n+1} \in V(z_{n+1})$. Then there is an edge e_{n+1} between z_n and z_{n+1} because $\hat{V}(z_n)$ and $V(z_{n+1})$ have nonempty intersection.

If instead z_n is a parabolic point, then we may have that $\zeta = \alpha_1 \cdots \alpha_n z_n$, in which case we are done and have found a strict parabolic coding of ζ . Otherwise,

by property (C5) of our open covers, there is some $\alpha_{n+1} \in L(z_n)$ so that $\alpha_{n+1}^{-1}(\zeta_n) \in \hat{V}(z_n)$. Define $\zeta_{n+1} = \alpha_{n+1}^{-1}\zeta_n$, as above, pick some z_{n+1} so that $\zeta_{n+1} \in V(z_{n+1})$. There is an edge e_{n+1} between z_n and z_{n+1} because $\hat{V}(z_n)$ and $V(z_{n+1})$ have nonempty intersection.

Thus, if the inductive procedure terminates at a finite stage, we have produced a parabolic coding of ζ . Otherwise, by construction we produce an infinite edge path with labels α_k such that $\zeta \in \bigcap_n \alpha_1 \cdots \alpha_n \overline{W}(z_n)$, as desired. \square

It turns out that conical \mathcal{G} -codings actually code (unique) *conical* points in $\partial(\Gamma, \mathcal{P})$, although this is not obvious from the lemma above. We will prove this fact later in Corollary 4.11. We will also eventually show that conical codings determine sequences in Γ that are well-defined up to some bounded error (Lemma 4.13).

To prove these two facts, we first show that codings define sequences of elements that stay close to geodesic rays in the cusped space X . We use the following basic fact about hyperbolic metric spaces. The statement is a rephrasing of Lemma 3.2 of [MMW22], and we refer the reader there for a proof.

Lemma 4.6 (See [MMW22] Lemma 3.2). *Let X be a proper δ -hyperbolic metric space, and fix a metric d_∂ on the Gromov boundary ∂X and a basepoint $x_0 \in X$. For any $\epsilon_0 > 0$ and any $R_1 > 0$, there exists a constant R_2 satisfying the following.*

Let z_1, z_2, z_3 be three points in ∂X , and for each i, j distinct in $\{1, 2, 3\}$ let $[z_i, z_j]$ be a geodesic joining z_i to z_j . If the distance between z_i and z_j is at least ϵ_0 for each distinct pair z_i, z_j , then the intersection

$$N_{R_1}([z_1, z_2]) \cap N_{R_1}([z_1, z_3])$$

is contained in the R_2 -neighborhood of a geodesic ray from x_0 to z_1 .

Now, using the strategy of Lemma 3.4 in [MMW22], we show the following.

Lemma 4.7. *There exists a uniform constant $R > 0$ so that, for any strict coding \mathbf{e} of any point $\zeta \in \partial(\Gamma, \mathcal{P})$, if $(\alpha_k)_{k \in I}$ is the associated label sequence, then the sequence*

$$g_k := \alpha_1 \cdots \alpha_k$$

lies in the R -neighborhood of any geodesic ray in the cusped space X based at the identity in $\text{Cay}(\Gamma) \subset X$ and with endpoint $\zeta \in \partial X = \partial(\Gamma, \mathcal{P})$.

Proof. Let $\epsilon_0 > 0$ be small enough so that $\epsilon_0 < \epsilon$, and so that for every $z \in \partial(\Gamma, \mathcal{P})$, there exist points $z_+, z_- \in \partial(\Gamma, \mathcal{P})$ so that

$$d_\partial(z, z_\pm) > \epsilon, \quad d_\partial(z_+, z_-) > \epsilon_0.$$

That such an ϵ_0 exists follows from an easy geometric argument using the fact that $\epsilon < D/5$ (recall D is the constant from Definition 2.3); see Lemma 3.1 in [MMW22] for a proof. We also choose a constant $R_1 > 0$ so that, for any pair of points $a, b \in \partial(\Gamma, \mathcal{P})$ such that $d_\partial(a, b) \geq \epsilon$, any geodesic in X joining a to b passes within an R_1 -neighborhood of the identity.

Now fix $\zeta \in \partial(\Gamma, \mathcal{P})$, and let \mathbf{e} be a strict \mathcal{G} -coding for ζ , with label sequence $(\alpha_k)_{k \in I}$ and terminal vertex sequence $(z_k)_{k \in I}$. We define z_0 to be the initial vertex of the coding, and let $g_k := \alpha_1 \cdots \alpha_k$. For all $k \in I \cup \{0\}$, we write W_k for $W(z_k)$. Since \mathbf{e} is a coding for ζ we know $\zeta \in W_0$.

Choose points $\zeta_\pm \in \partial(\Gamma, \mathcal{P})$ so that $d_\partial(\zeta, \zeta_-) > \epsilon$, $d_\partial(\zeta, \zeta_+) > \epsilon$, and $d_\partial(\zeta_-, \zeta_+) > \epsilon_0$. Condition (C1) implies that the diameter of W_0 is less than ϵ , so we know that

ζ_+ and ζ_- both lie in $\partial(\Gamma, \mathcal{P}) - W_0$. Let $[\zeta_-, \zeta]$ be a geodesic in X from ζ_- to ζ and $[\zeta_+, \zeta]$ a geodesic from ζ_+ to ζ .

The edge inclusion condition (E1), together with the fact that \mathbf{e} is a strict coding for ζ , implies that $g_k^{-1}\zeta \in W_k$ and $g_k N_\varepsilon(W_k) \subset W_0$ for every k . Thus $g_k^{-1}(\partial(\Gamma, \mathcal{P}) - W_0)$ is a subset of $\partial(\Gamma, \mathcal{P}) - N_\varepsilon(W_k)$, so $g_k^{-1}\zeta_-$ lies in $\partial(\Gamma, \mathcal{P}) - N_\varepsilon(W_k)$ and therefore $d_\partial(g_k^{-1}\zeta_-, g_k^{-1}\zeta) \geq \varepsilon$.

Thus, by our choice of R_1 , the geodesic $g_k^{-1}[\zeta_-, \zeta]$ enters an R_1 -neighborhood of the identity. Equivalently, g_k lies in the R_1 -neighborhood of the geodesic $[\zeta_-, \zeta]$. The same argument also shows that g_k lies in an R_1 -neighborhood of the geodesic $[\zeta_+, \zeta]$. Now apply Lemma 4.6 with $z_1 = \zeta$, $z_2 = \zeta_-$, and $z_3 = \zeta_+$ to see that there is a constant R_2 (independent of ζ and \mathbf{e}) so that g_k lies in the R_2 -neighborhood of some geodesic ray from the identity to ζ . Setting $R = R_2 + 2\delta$, we conclude that all g_k lie in the R -neighborhood of any such geodesic ray. \square

Although the previous lemma only applies directly to strict codings, we can use it to obtain a statement for generalized codings as well. Here and in what follows, we use the notation $|\alpha|_X$ for $d_X(\alpha, id)$, where id is the image of the identity element of Γ in $\text{Cay}(\Gamma) \subset X$.

Corollary 4.8. *Let $R > 0$ be the constant from Lemma 4.7, and let (g_0, \mathbf{e}) be a generalized \mathcal{G} -coding of a point $\zeta \in \partial(\Gamma, \mathcal{P})$, with label sequence $(\alpha_k)_{k \in I}$. Then the sequence*

$$g_k := g_0 \cdot \alpha_1 \cdots \alpha_k$$

lies in the $(R + |g_0|_X + 2\delta)$ -neighborhood of any geodesic ray in X from id to ζ .

Proof. Lemma 4.7 immediately implies that the sequence $(g_k)_{k \in I \cup \{0\}}$ lies within an R -neighborhood of a geodesic $[g_0, \zeta]$ in X from g_0 to ζ . However, this geodesic lies within distance $|g_0|_X + 2\delta$ of any geodesic $[id, \zeta]$ in X from id to ζ . To see this, simply make a long quadrilateral with one side a geodesic segment from id to g_0 , two sides long subsegments of the rays from id and g_0 to ζ , and the fourth side a short path between these rays far from id . That quadrilaterals are 2δ -slim now implies that $[g_0, \zeta]$ lies within the $R + 2\delta$ neighborhood of $[id, \zeta]$, and vice versa. \square

The sequence $(g_k)_{k \in I \cup \{0\}}$ from Lemma 4.7 and Corollary 4.8 will make many appearances, so we make the following definition.

Definition 4.9. If (g_0, \mathbf{e}) is a generalized coding with label sequence $(\alpha_k)_{k \in I}$, we call the sequence $(g_k)_{k \in I \cup \{0\}}$ defined by $g_k := g_0 \cdot \alpha_1 \cdots \alpha_k$ the *quasi-geodesic sequence associated to (g_0, \mathbf{e})* .

The terminology “quasi-geodesic sequence” comes from the fact that $(g_k)_{k \in I \cup \{0\}}$ lies bounded distance from a geodesic in the cusped space X , and behaves like a quasi-geodesic in the *relative Cayley graph* for Γ (i.e. the Cayley graph for Γ defined with respect to the infinite generating set $\mathcal{S} \cup \bigcup_{p \in \Pi} \Gamma_p$). However, note that $(g_k)_{k \in I \cup \{0\}}$ may not be a quasi-geodesic in X , since the distances $d_X(g_k, g_{k+1})$ may be arbitrarily large.

At this point, we have only shown (via Lemma 4.7) that the associated quasi-geodesic sequence $(g_k)_{k \in I \cup \{0\}}$ stays in a uniform neighborhood of a geodesic ray in X . To see that the sequence actually follows the ray to infinity, we will use an argument from Lemma 2.11 in [MMW22]:

Lemma 4.10 (Bounded backtracking I). *Let (g_0, \mathbf{e}) be a generalized coding with terminal vertex sequence $(z_k)_{k \in I}$ and associated quasi-geodesic sequence $(g_k)_{k \in I \cup \{0\}}$. For every k , there is a proper inclusion*

$$g_{k+1}W(z_{k+1}) \subsetneq g_kW(z_k).$$

Moreover, no element in $(g_k)_{k \in I \cup \{0\}}$ is repeated more than $\#Z$ times.

Proof. The first statement follows directly from Proposition 3.9. Given this, we now prove the second statement. Suppose for a contradiction that for some $g \in \Gamma$ we have $\#\{k : g_k = g\} > \#Z$. Then there are distinct $k, k' \in \mathbb{N}$ such that $g_k = g_{k'} = g$ and $z_k = z_{k'}$. But then $g_kW(z_k) = g_{k'}W(z_{k'})$, which contradicts the proper inclusion already established. \square

Corollary 4.11. *Every conical \mathcal{G} -coding codes a unique conical limit point.*

Proof. The statement for generalized codings follows immediately from the statement for strict codings, so fix a strict conical coding \mathbf{e} . That *some* point ζ is coded by \mathbf{e} follows immediately from Definition 4.3. Lemma 4.7 and Lemma 4.10 imply that the point ζ is uniquely determined, because the associated quasi-geodesic sequence $(g_k)_{k \in \mathbb{N}}$ tends to infinity in Γ and stays in a uniform neighborhood of a geodesic ray in X with ideal endpoint ζ . But this is just another way of saying that ζ is a conical limit point (see e.g. [BH20, Prop. A.2] for the equivalence). \square

Corollary 4.12 (Bounded backtracking II). *Let R be the constant from Lemma 4.7. Suppose (g_0, \mathbf{e}) is a generalized conical coding with associated quasi-geodesic sequence $(g_k)_{k \in I \cup \{0\}}$. For any $m \in \mathbb{N}$ and any $n > m$, we have $d_X(id, g_n) > d_X(id, g_m) - (3R + 2|g_0|_X + 6\delta)$.*

Proof. Let ζ be the point coded by (g_0, \mathbf{e}) , and let σ be a ray in X from the identity to ζ . By Corollary 4.8, we have $d(g_k, \sigma) < R + |g_0|_X + 2\delta$ for all k . Fix indices $m < n$, and let x_m be a point along σ which is within distance $R + |g_0|_X + 2\delta$ of g_m .

Consider the strict conical coding \mathbf{e}' in \mathcal{G} given by the edge path $(e_{k+m})_{k \in \mathbb{N}}$. The label sequence for this coding is the tail of the label sequence for \mathbf{e} , so the associated quasi-geodesic is the sequence $(g'_k)_{k \in \mathbb{N}}$, where $g'_k := g_m^{-1}g_{k+m}$. It follows that \mathbf{e}' codes the point $g_m^{-1}\zeta$, hence (by Lemma 4.7) the sequence g'_k stays within distance R of a geodesic ray σ' from id to $g_m^{-1}\zeta$.

Let σ'' be the sub-ray of σ from x_m to ζ , so that $g_m^{-1}\sigma''$ is a ray from $g_m^{-1}x_m$ to $g_m^{-1}\zeta$. Since $d_X(x_m, g_m) = d_X(g_m^{-1}x_m, id) < R + |g_0|_X + 2\delta$, the rays σ' and $g_m^{-1}\sigma''$ have Hausdorff distance bounded by $R + |g_0|_X + 4\delta$ (see the argument in the proof of Corollary 4.8).

Now, let x'_n be a point on the ray σ' so that $g'_{n-m} = g_m^{-1}g_n$ lies within distance R of x'_n . There is a point x''_n on σ'' so that $d_X(g_m^{-1}x''_n, x'_n) < R + |g_0|_X + 4\delta$, meaning $d_X(g_m^{-1}g_n, g_m^{-1}x''_n) = d_X(g_n, x''_n) \leq 2R + |g_0|_X + 4\delta$. Using the fact that σ'' is a sub-ray of σ based at x_m , we have

$$d_X(id, g_n) + d_X(g_n, x''_n) \geq d_X(id, x''_n) = d_X(id, x_m) + d_X(x_m, x''_n) \geq d_X(id, x_m)$$

We have seen that

$$d_X(id, x_m) \geq d_X(id, g_m) - (R + |g_0|_X + 2\delta)$$

and

$$d_X(g_n, x''_n) \leq 2R + |g_0|_X + 4\delta.$$

Combining the above inequalities gives the desired bound $d_X(id, g_n) > d_X(id, g_m) - (3R + 2|g_0|_X + 6\delta)$. \square

As a consequence of Corollary 4.12, whenever some edge label α_i in a generalized coding (g_0, \mathbf{e}) satisfies $|\alpha_i|_X > 3R + 2|g_0|_X + 4\delta$, the associated quasi-geodesic sequence makes positive progress along the ray it tracks. This will be important in the following section.

The final lemma of this section shows that conical codings are “unique up to bounded distance,” as follows.

Lemma 4.13. *For any $g_0, h_0 \in \Gamma$, there exists a constant $D_0 > 0$ satisfying the following. Suppose that (g_0, \mathbf{e}) , (h_0, \mathbf{f}) are two generalized codings of a common conical point ζ . Then the Hausdorff distance between the sets $\{g_k : k \geq 0\}$ and $\{h_k : k \geq 0\}$ (with respect to the metric d_X) is at most D_0 .*

Before proving the lemma, we fix notation for some more data related to the automaton \mathcal{G} which will appear in both the proof below and in several arguments in the following section.

Definition 4.14. For each parabolic vertex q of our automaton, we choose an element $t_q \in \Gamma$ so that $t_q^{-1}q \in \Pi$; we make this choice so that $|t_q|_X$ is minimized.

Then define the quantity C by

$$C = 2\delta + 6 + \max(\{|t_q|_X : q \in Z \text{ parabolic}\} \cup \{|\alpha_z|_X : z \in Z \text{ conical}\}).$$

Recall that when $z \in Z$ is conical, α_z is the unique element in the label set $L(z)$.

Proof of Lemma 4.13. Fix a conical point $\zeta \in \partial(\Gamma, \mathcal{P})$ and a geodesic ray σ in X from id to ζ , and let g_0, h_0 be arbitrary elements of Γ . Consider a pair of generalized codings (g_0, \mathbf{e}) and (h_0, \mathbf{f}) of ζ . Letting R be the constant from Lemma 4.7, and defining

$$R_0 = R + \max\{|g_0|_X, |h_0|_X\} + 6\delta,$$

Corollary 4.8 implies that the associated quasi-geodesic sequences for both (g_0, \mathbf{e}) and (h_0, \mathbf{f}) lie in the set

$$N_{R_0}(\sigma) \cap \text{Cay}(\Gamma).$$

Let $(g_k)_{k \in I \cup \{0\}}$ be the quasi-geodesic sequence associated to (g_0, \mathbf{e}) , and let $(e_k)_{k \in \mathbb{N}}$ and $(\alpha_k)_{k \in \mathbb{N}}$ be the sequence of edges and labels for \mathbf{e} .

Suppose that for some particular $k > 0$, we have $|\alpha_k|_X \geq C$. Then $\iota(e_k)$ is parabolic, equal to $t_k p_k$ for some $p_k \in \Pi$ and $t_k \in \Gamma$ chosen in Definition 4.14. We thus have $\alpha_k = t_k a_k$ for $a_k \in \Gamma_{p_k}$, and the group elements $g_{k-1} t_k$ and $g_k = g_{k-1} t_k a_k$ lie on the boundary of a common horoball \mathcal{H} in X . We let τ_k be a regular \mathcal{H} -geodesic joining $g_{k-1} t_k$ to g_k .

Let τ'_k be the subsegment of τ_k contained in the $\delta + 1$ -horoball nested inside \mathcal{H} . Our assumption on $|\alpha_k|_X$ ensures that τ'_k is non-empty. Lemma 2.1 says that τ'_k is an X -geodesic, and the remaining subsegments of τ_k are vertical, so they are also X -geodesics. The endpoints $g_{k-1} t_k$ and g_k of τ_k are distance at most $R_0 + C$ from points s_k and s_{k+1} on σ . We thus obtain a geodesic hexagon, one of whose sides is a part of σ , with the opposite side equal to τ'_k . Any point of $\tau_k - \tau'_k$ is at most $R_0 + C + \delta + 1$ from either s_k or s_{k+1} . And any point of τ'_k is within 4δ of one of the other five sides of the hexagon, hence within $R_0 + C + 5\delta + 1$ of a point of σ .

For each index k such that $|\alpha_k|_X \geq C$, we fix a path τ_k as above. Consider the set

$$(2) \quad Y_{\mathbf{e}} = \{id\} \cup \{g_j : j \geq 0\} \cup \bigcup \{\tau_k : |\alpha_k|_X \geq C\}.$$

Since each g_j for $j \geq 0$ is contained in an R_0 -neighborhood of σ , and each segment τ_k is contained within an $(R_0 + C + 5\delta + 1)$ -neighborhood of σ , the whole set $Y_{\mathbf{e}}$ is also contained within an $(R_0 + C + 5\delta + 1)$ -neighborhood of σ . We also know that the set $Y_{\mathbf{e}}$ is $(C + |g_0|_X)$ -coarsely connected, since $d_X(id, g_0) = |g_0|_X$ by definition, and for each $k > 0$ either $d_X(g_{k-1}, g_k) = |\alpha_k|_X < C$ or there is a path τ_k in $Y_{\mathbf{e}}$ with one endpoint equal to g_k and the other within C of g_{k-1} .

Lemma 4.10 implies that the set of points $\{g_j : j \geq 0\}$ has unbounded diameter in X , which means that there are points of X arbitrarily far along σ that lie within distance $R_0 + C + 5\delta + 1$ of $Y_{\mathbf{e}}$. This means that $Y_{\mathbf{e}}$ is actually within Hausdorff distance R' of σ , for a constant R' depending only on R_0, C, δ , and $|g_0|_X$.

The coding (h_0, \mathbf{f}) has an associated set $Y_{\mathbf{f}}$ defined analogously to the way $Y_{\mathbf{e}}$ was defined above in (2). The same argument shows that the Hausdorff distance from σ to $Y_{\mathbf{f}}$ is at most R'' , for a constant R'' depending only on R_0, C, δ , and $|h_0|_X$. Thus the Hausdorff distance between $Y_{\mathbf{e}}$ and $Y_{\mathbf{f}}$ is at most $R' + R''$.

Now consider some h_k in the associated quasi-geodesic sequence for (h_0, \mathbf{f}) . We wish to show that h_k lies uniformly close to some point g_j in the associated sequence for (g_0, \mathbf{e}) . We know that there is some point $p \in Y_{\mathbf{e}}$ so that $d_X(p, h_k) \leq R' + R''$. If $p = g_j$ for some j there is nothing left to show. If $p = id$, then $d_X(g_0, h_k) \leq d_X(g_0, id) + d_X(id, h_k) \leq R' + R'' + |g_0|_X$. Finally, if $p \in \tau_j$ for some $j \geq 0$, we note that p has depth at most $R' + R''$, so it is distance at most $R' + R'' + 2$ from an endpoint of τ_j , and hence at most $R' + R'' + C + 2$ from either g_{j-1} or g_j . In each case, we have shown that h_k lies in the $(R' + R'' + |g_0|_X + C + 2)$ -neighborhood of the associated quasi-geodesic sequence for (g_0, \mathbf{e}) ; we can then argue symmetrically to obtain the desired uniform bound on Hausdorff distance. \square

5. UNIFORM NESTING

In this section, we prove an analog of [MMW22, Lemma 3.8] (there called the Uniform Contraction Lemma). More care is needed here because of the presence of parabolic elements.

We have seen that any two generalized codings of the same conical point have infinitely many nearby pairs of points along their quasi-geodesic sequences g_k and h_k . The next condition says, roughly, that there are two possibilities: either these pairs of points are eventually uniformly spaced along the quasi-geodesic, or there is an infinite sequence of large parabolic jumps between them. In either case, one can control the behavior of a sequence of nested sets determined by one coding in terms of the other coding, which is what we will need to prove that codings determine a well-defined semiconjugacy between the standard action and a small perturbation.

Definition 5.1. Let (g_0, \mathbf{e}) and (h_0, \mathbf{f}) be generalized codings of the same conical point in $\partial(\Gamma, \mathcal{P})$. Let g_k, h_k be their associated quasi-geodesic sequences, and let $W^{\mathbf{e}}(k) = W(\tau(e_k))$ and $W^{\mathbf{f}}(k) = W(\tau(f_k))$. This pair of codings has the c -uniform nesting property if for any $\varepsilon' < c$, there are constants $D_1 > 0$ and $D_2 = D_2(\varepsilon') > 0$ so that at least one of the following conditions holds.

- (1) (*Uniform nesting with short words*) There exist $N, M \in \mathbb{N}$ such that for every term in the sequence of indices $n_k := k + M$, we have $|\mathbf{Lab}(e_{n_k})|_X \leq$

D_2 , there exists m_k such that $d_X(g_{n_k}, h_{m_k}) \leq D_1$ and

$$g_{n_k+N} \overline{W^e(n_k+N)} \subset h_{m_k} W^f(m_k).$$

- (2) (*Uniform nesting with long parabolics*) There are infinite sequences of indices n_k, m_k such that for every $k \in \mathbb{N}$, $d_X(g_{n_k}, h_{m_k}) \leq D_1$ and
- $\tau(e_{n_k})$ is a parabolic point;
 - $g_{n_k} \overline{B_{3\varepsilon'}(\tau(e_{n_k}))} \subset h_{m_k} W^f(m_k)$;
 - $\mathbf{Lab}(e_{n_k+1}) N_\varepsilon(W^e(n_k+1)) \subset B_{\varepsilon'}(\tau(e_{n_k}))$, where ε is as in Definition 3.3.

Our main goal in this section is to prove:

Proposition 5.2. *Given $g_0, h_0 \in G$, there exists $c > 0$ such that if (g_0, \mathbf{e}) and (h_0, \mathbf{f}) are generalized codings of the same conical point $\zeta \in \partial(\Gamma, \mathcal{P})$, then (g_0, \mathbf{e}) and (h_0, \mathbf{f}) have the c -uniform nesting property. Furthermore, the constant D_1 depends only on g_0, h_0 , and the constants N and D_2 depend only on g_0, h_0 and the choice of $\varepsilon' < c$, and not on ζ, \mathbf{e} , or \mathbf{f} .*

As a first step towards the proof, the following lemma describes the behavior of codings that involve edges labeled by long words in Γ . Since codings are close to geodesics, these long words correspond to parabolics – geometrically, the geodesic has a long segment through the horosphere based at this parabolic point. The lemma makes precise the notion that two codings of the same point, being close to the same geodesic, have long segments in common horospheres.

Lemma 5.3 (Large jumps come from common parabolics). *Given $g_0, h_0 \in \Gamma$, there are constants $D_1, J > 0$ so the following holds. Let (g_0, \mathbf{e}) and (h_0, \mathbf{f}) be two generalized codings of the same conical point ζ , let $\mathbf{Lab}(e_k) = \alpha_k$, $\mathbf{Lab}(f_k) = \beta_k$, and let $(g_k)_{k \in \mathbb{N} \cup \{0\}}, (h_k)_{k \in \mathbb{N} \cup \{0\}}$ be the respective associated quasi-geodesic sequences. Then for each n with $|\alpha_n|_X > J$ there exists $m = m(n)$ such that*

- (1) $d_X(g_{n-1}, h_{m-1}) < D_1$ and $d_X(g_n, h_m) < D_1$;
- (2) There exists $p \in \Pi$ and $g, h \in G$ such that $\iota(e_n) = gp$ and $\iota(f_m) = hp$;
- (3) $g_{n-1}g, h_{m-1}h, g_n$ and h_m are all in the same coset of Γ_p .

The rough idea of the proof is as follows. We consider a regular geodesic ray σ in X from id to ζ , close to the quasi-geodesic sequences associated to both generalized codings. We first prove that, if the distance in X between two consecutive points g_{n-1}, g_n is large, then a geodesic segment between g_{n-1} and g_n must spend a large amount of its lifetime in some horoball \mathcal{H} (this will follow from Lemma 5.4 below). Using the fact that the quasi-geodesic sequence tracks σ , we prove that σ must also spend some large amount of time in the same horoball \mathcal{H} . If this time is long enough, we can even conclude that σ spends some large time in a horoball \mathcal{H}' nested deeply inside of \mathcal{H} .

Then, using the fact that the quasi-geodesic sequence associated to \mathbf{f} also tracks σ , we can show that there are points h_{m-1}, h_m in this sequence on “either side” of \mathcal{H}' . Since these points are far apart, a regular geodesic joining h_{m-1} to h_m must spend a long amount of time in \mathcal{H} , which essentially proves the lemma.

To make the above reasoning precise, we first prove two general lemmas about the geometry of the cusped space X . The first lemma gives us a way to estimate the length of time a geodesic in X spends inside a horoball \mathcal{H} when one of the endpoints of the geodesic lies on the boundary of \mathcal{H} .

Lemma 5.4. *Suppose that $\alpha \in G$ lies in a coset $g\Gamma_p$ for some $g \in \Gamma$ and $p \in \Pi$, and let $\mathcal{H} \subset X$ be the horoball in X based at $g\Gamma_p$. Let τ be a regular geodesic in X from id to α . Then $\tau \cap \mathcal{H}$ contains a subsegment with length at least $|\alpha|_X - (|g|_X + 12\delta)$.*

Moreover, if $\mathcal{H}' \neq \mathcal{H}$ is any other horoball in X , then any component of $\tau \cap \mathcal{H}'$ has length at most $|g|_X + 12\delta$.

Proof. Note that the elements g and α both lie on the boundary of the horoball \mathcal{H} . Let $\check{\mathcal{H}}$ be the 3δ -horoball nested inside \mathcal{H} . (See Section 2 for definitions.) Let a be the point on $\partial\check{\mathcal{H}}$ which is connected by a vertical path to g , and let b be the point on $\partial\check{\mathcal{H}}$ which is connected by a vertical path to α .

Consider a geodesic pentagon whose vertices are (in cyclic order) id , g , a , b , and α , so that the side connecting id to α is τ . The sides $[g, a]$ and $[\alpha, b]$ are vertical, and we may suppose that the side $[a, b]$ is regular.

Every point on τ is within 3δ of some point on one of the other sides of this pentagon. At most the initial subsegment of length $|g|_X + 6\delta$ and the final subsegment of length 6δ can be close to some other side than $[a, b]$.

We deduce that a subsegment τ' of length at least $|\alpha|_X - (|g|_X + 12\delta)$ is within 3δ of $[a, b]$, and hence within 3δ of $\check{\mathcal{H}}$. In particular τ' is completely contained in \mathcal{H} , establishing the first part of the lemma.

The last sentence follows since horoballs have disjoint interior. \square

The next lemma is essentially a consequence of the quasi-convexity of horoballs in X .

Lemma 5.5 (Nearby geodesics enter common horoballs). *Let $\mathcal{H} \subset X$ be a k -horoball for some $k \geq 0$, and let τ_1, τ_2 be two regular geodesic segments in X such that the endpoints of τ_1 lie within distance L of the endpoints of τ_2 . If $\tau_1 \cap \mathcal{H}$ contains a segment with length at least $T \geq 4L + 8\delta + 3$, then $\tau_2 \cap \mathcal{H}$ contains a segment with length at least $T - (4L + 8\delta)$.*

Proof. Let a_1, b_1 and a_2, b_2 be the endpoints of τ_1, τ_2 respectively, so that $d_X(a_1, a_2) \leq L$ and $d_X(b_1, b_2) \leq L$. We let $\check{\mathcal{H}}$ be the $(k + 2\delta)$ -horoball nested inside of the k -horoball \mathcal{H} . Then, since τ_1 is regular, and its intersection with \mathcal{H} has length at least $T > 4\delta + 3$, it also intersects $\check{\mathcal{H}}$, and in fact the intersection contains a segment with length at least $T - 4\delta$. We let x_1, y_1 be the endpoints of such a segment, so that a_1, x_1, y_1, b_1 are arranged on τ_1 in that order.

The length of τ_1 is

$$\begin{aligned} d_X(a_1, x_1) + d_X(x_1, y_1) + d_X(y_1, b_1) &= d_X(a_1, x_1) + (T - 4\delta) + d_X(y_1, b_1) \\ &\geq d_X(a_1, x_1) + d_X(y_1, b_1) + 4L + 4\delta + 3 \end{aligned}$$

The geodesic τ_2 has endpoints at most L from those of τ_1 , so its length is at least $d_X(a_1, x_1) + d_X(y_1, b_1) + 2L + 4\delta + 3$. In particular, there are points x_2 and y_2 of τ_2 satisfying

$$\begin{aligned} d_X(a_2, x_2) &= d_X(a_1, x_1) + L + 2\delta, \\ d_X(y_2, b_2) &= d_X(y_1, b_1) + L + 2\delta. \end{aligned}$$

Moreover, the points a_2, x_2, y_2, b_2 must lie on the segment τ_2 in that order, since $2(L + 2\delta) < 2L + 4\delta + 3$.

Consider a geodesic quadrilateral with opposite sides τ_1, τ_2 . The subsegment $[x_2, y_2] \subset \tau_2$ is contained in a 2δ -neighborhood of the other three sides of this

quadrilateral. In fact, $[x_2, y_2]$ must be contained in the 2δ -neighborhood of $[x_1, y_1]$, meaning it is contained in the 2δ -neighborhood of $\check{\mathcal{H}}$ and therefore in \mathcal{H} . Thus, this segment in $\tau_2 \cap \mathcal{H}$ has length at least

$$|\tau_2| - 2L - 4\delta - d_X(a_1, x_1) - d_X(y_1, b_1).$$

Then, since $|\tau_2| \geq |\tau_1| - 2L$, by our choice of x_1, y_1 we have

$$|\tau_1| \geq d_X(a_1, x_1) + (T - 4\delta) + d_X(y_1, b_1),$$

so we obtain the desired bound. \square

Proof of Lemma 5.3. Let C be the constant from Definition 4.14, which we recall is (partly) determined by fixing a group element $t_q \in \Gamma$ for each parabolic vertex $q \in Z$, satisfying $t_q^{-1}q \in \Pi$. Letting R be the constant from Lemma 4.7, we set

$$R' = R + \max\{|g_0|_X, |h_0|_X\} + 2\delta,$$

and then let

$$J > 11R' + 2C + 50\delta.$$

Consider two conical codings (g_0, \mathbf{e}) and (h_0, \mathbf{f}) of a point $\zeta \in \partial(\Gamma, \mathcal{P})$, and let $(g_n)_{n \in \mathbb{N} \cup \{0\}}$ and $(h_m)_{m \in \mathbb{N} \cup \{0\}}$ be the associated quasi-geodesic sequences. Fix some $n \in \mathbb{N}$ such that $|\alpha_n|_X > J$. Since $J > C$ the edge labeled by α_n has its initial vertex at a parabolic point, meaning α_n is in some coset $t_q\Gamma_p$ with $q = \iota(e_n) = t_qp$ and $p \in \Pi$. In particular $\alpha_n = t_qa$ where $a \in \Gamma_p$. Let \mathcal{H} be the horoball associated to the coset $g_n\Gamma_p = g_{n-1}t_q\Gamma_p$, and let τ be a regular geodesic from g_{n-1} to g_n . Now, $g_{n-1}^{-1}\tau$ is a regular geodesic from id to $\alpha_n = g_{n-1}^{-1}g_n$, so by Lemma 5.4, it contains a segment in $g_{n-1}^{-1}\mathcal{H}$ of length at least $J - (C + 12\delta)$. Thus, $\tau \cap \mathcal{H}$ contains a segment of length at least $J - (C + 12\delta)$.

Let σ be a regular geodesic from id to ζ . For a pair of points $a, b \in \sigma$, let $[a, b]$ denote the subsegment of σ with endpoints a, b . By Corollary 4.8, we can find a sequence $(x_i)_{i \in \mathbb{N}}$ on σ , so that $d_X(g_i, x_i) \leq R'$ for all i . Then Lemma 5.5 implies that the intersection $[x_{n-1}, x_n] \cap \mathcal{H}$ contains a segment with length at least $J - C - 4R' - 20\delta$.

We now wish to find consecutive h_{m-1}, h_m in the quasi-geodesic sequence associated to (h_0, \mathbf{f}) so that a regular geodesic between h_{m-1} and h_m spends all but a uniformly bounded amount of its length inside the horoball \mathcal{H} . For this, we let \mathcal{H}' be the R' -horoball nested inside of \mathcal{H} . Since $R' > 2\delta$ by definition, Lemma 2.1 implies that \mathcal{H}' is convex, so the intersection $\sigma \cap \mathcal{H}'$ is a geodesic segment, which must have length at least $J - C - 6R' - 24\delta$ (In particular the length of this segment is greater than $4R' + 8\delta + 3$, so we will be able to apply Lemma 5.5 to it.) Let σ_1 be the connected component of $\sigma \setminus \mathcal{H}'$ containing id and let σ_2 be the unbounded connected component.

Since every element in our sequence h_m lies in Γ , no h_m can lie in an R' -neighborhood of the R' -horoball \mathcal{H}' , and thus each h_m lies within distance R' of exactly one of σ_1 or σ_2 . We let m be the first index so that h_m is within distance R' of σ_2 , so h_{m-1} lies distance at most R' from σ_1 . Let y_{m-1}, y_m be points on σ which are within distance R' of h_{m-1} and h_m respectively. These points must lie on either side of the intersection $\mathcal{H}' \cap \sigma$, so the segment $[y_{m-1}, y_m] \cap \mathcal{H}'$ is a geodesic segment with length at least $J - C - 6R' - 24\delta$. Then, if τ' is a regular geodesic between h_{m-1} and h_m , Lemma 5.5 tells us that $\tau' \cap \mathcal{H}'$ contains a segment of length at least $J - C - 10R' - 32\delta > C$.

In particular, we have $d_X(h_{m-1}, h_m) = |\beta_m|_X = |\tau'| > C$, so the initial vertex of the edge labeled by β_m is a parabolic point $q' = t_{q'}p' \in Z$ for some $p' \in \Pi$. Lemma 5.4 implies that the horoball based on $t_{q'}p'$ is the only horoball which τ' can intersect in a segment of length longer than $C + 12\delta$. Since τ' meets \mathcal{H} in a segment longer than this, we must have $q = q'$ and $p' = p$, establishing the last two items of the lemma.

It remains to bound $d_X(g_{n-1}, h_{m-1})$ and $d_X(g_n, h_m)$. We first show that τ and τ' cross \mathcal{H} in the same direction. That is, we show:

Claim 1. g_{n-1} is in the R' -neighborhood of the initial segment σ_1 of σ , and g_n is in the R' -neighborhood of the final ray σ_2 .

Proof. The claim essentially follows from the fact that the distance $d_X(g_{n-1}, g_n)$ is large and from the bounded backtracking property of the quasi-geodesic sequence $(g_n)_{n \in \mathbb{N}}$ (Corollary 4.12). To be specific, we consider the points x_{n-1}, x_n on σ which are within distance R' of g_{n-1}, g_n respectively. Recall that the regular geodesic τ between g_{n-1} and g_n intersects \mathcal{H} in a segment of length at least $J - C - 12\delta$, so $\tau \cap \mathcal{H}'$ contains a segment of length at least $J - C - 2R' - 12\delta$. Then Lemma 5.5 implies that $[x_{n-1}, x_n] \cap \mathcal{H}'$ contains a segment of length at least $J - C - 6R' - 20\delta > 0$. So, either $x_{n-1} \in \sigma_1$ and $x_n \in \sigma_2$, or vice versa.

However, if $x_{n-1} \in \sigma_2$ and $x_n \in \sigma_1$, then

$$d_X(id, x_{n-1}) = d_X(id, x_n) + d_X(x_{n-1}, x_n),$$

implying that

$$d_X(id, x_n) \leq d_X(id, x_{n-1}) - (J - C - 6R' - 24\delta).$$

Since $d_X(x_{n-1}, g_{n-1}) \leq R'$ and $d_X(x_n, g_n) \leq R'$, this would imply $|g_n|_X \leq |g_{n-1}|_X - (J - C - 8R' - 24\delta) < |g_{n-1}|_X - (3R' + 2\delta)$. Since $R' \geq R + |g_0|_X + 2\delta$ by definition, this contradicts Corollary 4.12. This proves the claim. \square

We now consider the points x_{n-1}, y_{m-1} on σ_1 within distance R' of g_{n-1}, h_{m-1} . Let z_1, z_2 be the endpoints of σ_1 and σ_2 on $\partial\mathcal{H}'$, respectively.

Since $d_X(g_{n-1}, \mathcal{H}') \leq C + R'$ and $d_X(h_{m-1}, \mathcal{H}') \leq C + R'$, we have $d_X(x_{n-1}, \mathcal{H}') \leq 2R' + C$ and $d_X(y_{m-1}, \mathcal{H}') \leq 2R' + C$. Then, we apply Lemma 2.2 to the subsegment $[x_{n-1}, z_2]$ and the R' -horoball \mathcal{H}' to deduce that $d_X(x_{n-1}, \partial\mathcal{H}')$ differs by at most δ from $d_X(x_{n-1}, z_1)$, and similarly for y_{m-1} . In particular, $d_X(x_{n-1}, y_{m-1}) \leq 2(2R' + C + \delta)$, hence $d_X(g_{n-1}, h_{m-1}) \leq 6R' + 2C + 4\delta$.

A nearly identical argument applied to the points x_n, y_m (with the roles of σ_1, z_1 interchanged with the roles of σ_2, z_2) shows that $d_X(g_n, h_m)$ is also at most $6R' + 2C + 4\delta$, meaning we can set $D_1 = 6R' + 2C + 4\delta$. \square

Proof of Proposition 5.2. Let (g_0, \mathbf{e}) and (h_0, \mathbf{f}) be generalized codings of the same conical point ζ , with $\mathbf{Lab}(e_n) = \alpha_n$, $\mathbf{Lab}(f_m) = \beta_m$ and let $(g_n)_{n \in \mathbb{N} \cup \{0\}}$ and $(h_m)_{m \in \mathbb{N} \cup \{0\}}$ be the associated quasigeodesic sequences.

First, we set the constants c and D_1 . (The numbers D_2, N depends on ε' so they will be set later.) Choose D_1 and J so that they are at least as large as the corresponding constants in Lemma 5.3, and so that D_1 is also at least the Hausdorff distance bound in Lemma 4.13. Note that these only depend on g_0 and h_0 , not \mathbf{e}, \mathbf{f} or ζ .

By Lemma 5.3, whenever $|\alpha_n|_X > J$, there exists a parabolic point $p \in \Pi$, an index m , and elements $g, h \in \Gamma$ such that $d(g_{n-1}, h_{m-1}) < D_1$, $\iota(e_n) = gp$ and $\iota(f_m) = hp$, and the elements $g_{n-1}g$ and $h_{m-1}h$ lie in the same coset of Γ_p .

Thus, $g\Gamma_p = (g_{n-1}^{-1}h_{m-1})h\Gamma_p$, and $|g_{n-1}^{-1}h_{m-1}|_X < D_1$, so $g_{n-1}^{-1}h_{m-1}$ lies in a finite set $F_1 := \{f \in \Gamma : |f|_X < D_1\}$. Restating the above, we have $gp = fhp$, for $f \in F_1$. Since $hp \in W(hp)$, we have $fhp = gp \in fW(hp)$. Consider all possible triples (z, y, f) such that $z, y \in Z$, $f \in F_1$, and $z \in fW(y)$. There are finitely many such, so we may choose some $c > 0$ such that $\overline{B}_{3c}(z) \subset fW(y)$ holds for every such triple.

Now fix any $\varepsilon' < c$, and choose $D'_2 = D'_2(\varepsilon') > J$ large enough so that, for each edge from z to a vertex x where $z = gp$ is parabolic, and each $\alpha \in L(z)$, if $|\alpha|_X > D'_2$ then $\alpha\overline{N}_\varepsilon(W(x)) \subset B_{\varepsilon'}(z)$. We know such a D'_2 exists because only finitely many elements of the coset $g\Gamma_p$ fail this contraction condition: as we observed at the beginning of Remark 3.10, since the edge inclusion condition

$$\alpha\overline{N}_\varepsilon(W(x)) \subset \alpha\hat{W}(z) \subset W(z)$$

holds for all but finitely many $\alpha \in g\Gamma_p$, the closed neighborhood $\overline{N}_\varepsilon(W(x))$ cannot contain p and therefore all but finitely many $\alpha \in g\Gamma_p$ take $\overline{N}_\varepsilon(W(x))$ into an arbitrarily small neighborhood of z . We define $D_2 = D_2(\varepsilon') := D'_2(\varepsilon') + 2D_1$. Again, this depends on h_0, g_0 and ε' , but not the strict codings \mathbf{e}, \mathbf{f} or the point ζ .

We now can prove that the generalized codings satisfy uniform nesting. As a first case, suppose $|\alpha_n|_X > D'_2$ for infinitely many n . We will show that, in this case, we have *uniform nesting with long parabolics*. Choose an infinite sequence of indices n_k so that $|\alpha_{n_k+1}|_X > D'_2$ is always satisfied. Lemma 5.3 then provides a sequence of indices m_k with $d_X(g_{n_k}, h_{m_k}) < D_1$ and $d_X(g_{n_k+1}, h_{m_k+1}) < D_1$, which satisfies all of the requirements of “uniform nesting with long parabolics” by our choice of constants D'_2, ε' . This finishes the proof in this case.

On the other hand, if $|\beta_m|_X > D_2 = D'_2 + 2D_1$ is satisfied for infinitely many m , then we can use Lemma 5.3 again to find an infinite subsequence $(\alpha_{n_k})_{k \in \mathbb{N}}$ so that $|\alpha_{n_k}|_X > D'_2$ for every k , and we are in the previous case.

If neither of the first two cases hold, then we know that both $|\alpha_n|_X$ and $|\beta_m|_X$ are bounded by D_2 for all but finitely many n, m , respectively. For this case, we will create new generalized codings of ζ by “shifting the indices:” we replace g_0 and h_0 with terms further along in the associated quasi-geodesic sequences, and truncating the first terms of the sequence so that all labels are bounded by D_2 . This will put us in a position to consider a sub-automaton only including edges with short labels, and then apply Lemma A.3 from the Appendix to conclude the proof.

In more detail: first, choose M large enough so that $|\alpha_n|_X \leq D_2$ for all $n \geq M$. Lemma 4.13 tells us that for each $n \in \mathbb{N}$, there is some index $m(n)$ so that $d_X(g_n, h_{m(n)}) < D_1$. Lemma 4.10 implies that $m(n)$ tends to infinity as n tends to infinity, so increasing M if necessary, we can ensure that $|\beta_k|_X \leq D_2$ for every $k \geq m(M)$.

Let \mathbf{e}' be the sub-path of \mathbf{e} starting with the edge e_M and let \mathbf{f}' be sub-path of \mathbf{f} starting with the edge $f_{m(M)}$. Consider the generalized codings (g_{M-1}, \mathbf{e}') and $(h_{m(M)-1}, \mathbf{f}')$, and let $g'_n = g_{n+M-1}$ and $h'_m = h_{m+m(M)-1}$ be their associated quasi-geodesic sequences. By construction, both of these codings are generalized codings of ζ , and their associated quasi-geodesic sequences are tails of the associated

quasi-geodesic sequences for (g_0, \mathbf{e}) and (h_0, \mathbf{f}) , respectively. Further, the label sequences of \mathbf{e}' and \mathbf{f}' consist of elements whose length in X is bounded by D_2 .

Let \mathcal{F} be the subgraph of \mathcal{G} obtained by deleting all labels of length more than D_2 , and then deleting all edges with empty label set. Then \mathcal{F} is a *finitary point coder* in the sense defined in Appendix A, and (g_{M-1}, \mathbf{e}') and $(h_{m(M)-1}, \mathbf{f}')$ are generalized \mathcal{F} -codings.

Let F be the finite subset of Γ consisting of elements of length at most D_1 , and apply this to Lemma A.3. We conclude there is a constant $N > 0$, so that whenever $d_X(g'_n, h'_m) < D_1$, we have $g'_{n+N} \overline{W(z_{n+N})} \subset h'_m W(y_m)$, where $z_k = \tau(e'_k)$ and $y_k = \tau(f'_k)$. Translating this statement, for each m_k such that $d_X(g_{n_k}, h_{m_k}) \leq D_1$ and $n_k > M$, we have

$$g'_{n_k+N} \overline{W^{\mathbf{e}}(n_k+N)} \subset h_{m_k} W^{\mathbf{f}}(m_k).$$

The constant N in the above depends only on the point coder \mathcal{F} , which in turn depends only on D_2 and the original automaton \mathcal{G} , so this completes the proof. \square

Remark 5.6. The containment

$$g_{n_k+N} \overline{W^{\mathbf{e}}(n_k+N)} \subset h_{m_k} W^{\mathbf{f}}(m_k)$$

from *uniform nesting with short words* is equivalent to

$$g_{n_k} \alpha_{n_k+1} \dots \alpha_{n_k+N} \overline{W^{\mathbf{e}}(n_k+N)} \subset h_{m_k} W^{\mathbf{f}}(m_k)$$

or, multiplying on the left by $g_{n_k}^{-1}$,

$$(\dagger) \quad \alpha_{n_k+1} \dots \alpha_{n_k+N} \overline{W^{\mathbf{e}}(n_k+N)} \subset g_{n_k}^{-1} h_{m_k} W^{\mathbf{f}}(m_k)$$

Since the assumptions of uniform nesting with short words stipulate that $|\alpha_{n_k}| \leq D_2$, and $|g_{n_k}^{-1} h_{m_k}|_X < R$, and since there are only finitely many sets of the form $W(z)$, the inclusions given by (\dagger) are only finite in number. So these inclusions correspond to finitely many open conditions, meaning they are stable under small perturbation. This is essential to our argument in Lemma 6.13.

Similarly, the containments $g_{n_k} \overline{B_{3e'}(\iota(e_{n_k+1}))} \subset h_{m_k} W^{\mathbf{e}}(m_k+1)$ from the parabolic nesting conditions are only finite in number.

6. PROOF OF MAIN THEOREM

Thus far we have suppressed notation for the action of Γ on its boundary, simply writing $g(\zeta)$ or $g\zeta$ for the image of ζ under g . We will now need to consider other actions of Γ on this space, so we reintroduce the following notation.

Notation 6.1. As in the introduction, let $\rho_0 : \Gamma \rightarrow \text{Homeo}(\partial(\Gamma, \mathcal{P}))$ denote the standard action of Γ on its Bowditch boundary. Thus, what was previously written $gW(z_k)$ now becomes $\rho_0(g)W(z_k)$, for example.

Definition 6.2. Given $\mathcal{V} \subset C(\partial(\Gamma, \mathcal{P}))$, we define $\mathcal{R}_{\mathcal{V}} \subset \text{Hom}(\Gamma, \text{Homeo}(\partial(\Gamma, \mathcal{P})))$ to be the set of representations ρ such that, for each parabolic point $p \in \mathcal{P}$ with stabilizer $P = \Gamma_p$, there exists $\phi_p \in \mathcal{V}$ such that $\rho|_P$ is an extension of $\rho_0|_P$ via ϕ_p .

Remark 6.3. The map ϕ_p defining a semiconjugacy for the action of P as in Definition 6.2 is a priori not determined by the representation $\rho \in \mathcal{R}_{\mathcal{V}}$. To get around this, for the rest of this section, whenever we fix $\rho \in \mathcal{R}_{\mathcal{V}}$ for some $\mathcal{V} \subset C(\partial(\Gamma, \mathcal{P}))$, then for each $p \in \Pi$ we implicitly choose a semi-conjugacy $\phi_p \in \mathcal{V}$ which extends the restriction of ρ_0 to $P = \Gamma_p$.

Notation 6.4. Assuming that some action $\rho \in \mathcal{R}_{\mathcal{V}}$ has been fixed, then for each parabolic point $z = \rho_0(g)p$, where $p \in \Pi$, we let $\phi_z := \rho_0(g)\phi_p\rho(g)^{-1}$.

Observe the definition of ϕ_z depends only on z and ϕ_p , and not on the choice of $g \in \Gamma$ such that $z = \rho_0(g)p$. To see this, suppose $\rho_0(g)p = \rho_0(h)p$. Then

$$\rho_0(g)\phi_p\rho(g)^{-1} = \rho_0(h)\rho_0(h^{-1}g)\phi_p\rho(h^{-1}g)^{-1}\rho(h^{-1}).$$

Since $h^{-1}g \in \Gamma_p$, $\rho_0(h^{-1}g)\phi_p = \phi_p\rho(h^{-1}g)$, which shows

$$\rho_0(g)\phi_p\rho(g)^{-1} = \rho_0(h)\phi_p\rho(h)^{-1}$$

as desired.

In addition, note that for fixed z , if ρ approaches ρ_0 , and ϕ_p tends to the identity, then ϕ_z tends to the identity as well. We also record the following.

Lemma 6.5. *If $z = \rho_0(g)p$, then for any $x \in \partial(\Gamma, \mathcal{P})$ and any $\alpha \in g\Gamma_p$, we have*

$$(3) \quad \phi_z(\rho(\alpha)x) = \rho_0(\alpha)\phi_p(x).$$

Proof. Let x be given, and suppose that $z = \rho_0(g)p$ and $\alpha = gh$ for $h \in \Gamma_p$. Then

$$\begin{aligned} \phi_z(\rho(\alpha)x) &= \phi_z(\rho(gh)x) \\ &= \rho_0(g)\phi_p(\rho(g)^{-1}\rho(gh)x) \\ &= \rho_0(g)\phi_p(\rho(h)x) \\ &= \rho_0(gh)\phi_p(x) = \rho_0(\alpha)\phi_p(x). \end{aligned}$$

□

6.1. Specifying the neighborhoods \mathcal{U} and \mathcal{V}' .

Definition 6.6. Given a subset $\mathcal{V} \subset C(\partial(\Gamma, \mathcal{P}))$ and a representation $\rho \in \mathcal{R}_{\mathcal{V}}$, we define families of sets $\{V_\rho(z)\}_{z \in Z}$ and $\{\hat{V}_\rho(z)\}_{z \in Z}$ as follows.

- If z is a conical limit point, then we define

$$V_\rho(z) = V(z), \quad \hat{V}_\rho(z) = \rho(\alpha_z^{-1})V_\rho(z),$$

where α_z is the unique element in the label set $L(z)$.

- If $z = \rho_0(g)p$ for some $g \in \Gamma$, $p \in \Pi$, we define

$$V_\rho(z) = \phi_z^{-1}(V(z)),$$

and define

$$\hat{V}_\rho(z) = \phi_p^{-1}(\hat{V}(z)).$$

Definition 6.7. Suppose that $\mathcal{V} \subset C(\partial(\Gamma, \mathcal{P}))$ and $\rho \in \mathcal{R}_{\mathcal{V}}$. We say that ρ has the same combinatorics as the standard action ρ_0 if the following conditions hold:

- The collection $\{V_\rho(z)\}_{z \in Z}$ is an open covering of $\partial(\Gamma, \mathcal{P})$.
- For every $z \in Z$, we have $V_\rho(z) \subset W(z)$.
- For any $y, z \in Z$, we have $\hat{V}_\rho(y) \cap V_\rho(z) = \emptyset$ if and only if $\hat{V}(y) \cap V(z) = \emptyset$.
- If there is an edge from z to y in \mathcal{G} labeled by α , then $\rho(\alpha)(W(y)) \subset W(z)$.

Lemma 6.8 (Same combinatorics is relatively open). *For any sufficiently small neighborhood \mathcal{V}' of the identity in $C(\partial(\Gamma, \mathcal{P}))$ and any sufficiently small neighborhood \mathcal{U} of ρ_0 in $\text{Homeo}(\Gamma, \partial(\Gamma, \mathcal{P}))$, each $\rho \in \mathcal{U} \cap \mathcal{R}_{\mathcal{V}'}$ has the same combinatorics as ρ_0 .*

Proof. First, we ensure that Item (i) in the definition holds. Let r be a Lebesgue number for the original open covering $\{V(z)\}_{z \in Z}$, so that every set of diameter at most r is contained in some $V(z)$. We may choose \mathcal{V}' small enough so that $d_{\partial}(\phi_z(x), x) < r/2$ for every parabolic point $z \in Z$ and $x \in \partial(\Gamma, \mathcal{P})$.

Then, for any $x \in \partial(\Gamma, \mathcal{P})$, the $(r/2)$ -ball about x is contained in some $V(z)$, where $z \in \partial(\Gamma, \mathcal{P})$ is either conical or parabolic. If z is conical then $V_{\rho}(z) = V(z)$ and thus $x \in V_{\rho}(z)$. Otherwise, if z is parabolic, then $\phi_z(x)$ lies in $B_{r/2}(x) \subset V(z)$, hence $x \in \phi_z^{-1}(V(z)) = V_{\rho}(z)$.

Item (ii) can be arranged because $\overline{V(z)} \subset W(z)$ for each z (see Condition (C5)). In particular, there is some minimum distance from any $V(z)$ to $\partial(\Gamma, \mathcal{P}) - W(z)$, and we can choose \mathcal{V}' small enough so that no ϕ_z moves any point more than that distance.

For Item (iii), we argue similarly. Property (C6) from Proposition 3.7 implies there is a minimum distance from any $V(z)$ to any disjoint $\hat{V}(y)$. Shrinking \mathcal{V}' if necessary, we may assume that no ϕ_z moves any point more than half that distance. By choosing sufficiently small \mathcal{U} , we can ensure that the distance between $\rho(\alpha_z)$ and $\rho_0(\alpha_z)$ is also at most half that distance. This will ensure that $V(z) \cap \hat{V}(y) = \emptyset$ implies $V_{\rho}(z) \cap \hat{V}_{\rho}(y) = \emptyset$. For the converse, observe that there is some positive radius r so that each nonempty intersection $V(z) \cap \hat{V}(y)$ contains a ball of radius r . The sets $V_{\rho}(z)$ and $\hat{V}_{\rho}(y)$ each contain the preimage of this ball by a continuous map arbitrarily close to the identity, so if \mathcal{V}' and \mathcal{U} are small enough then these preimages have nonempty intersection.

We now turn to condition (iv). By condition (E1) in Proposition 3.9 we have $\rho_0(\alpha)\overline{N_{\varepsilon}(W(y))} \subsetneq W(z)$ for every edge from z to y labeled by α . For conical vertices z , provided ρ is a small enough perturbation of ρ_0 , each of the containments $\rho(\alpha)(W(y)) \subset W(z)$ will hold, since there are only finitely many such edges and labels. For parabolic vertices, we argue as follows.

Consider the constants ϵ_z from Remark 3.10, and fix $\epsilon_{\min} \leq \min\{\epsilon_z : z \in Z \text{ parabolic}\}$. The inclusion (1) from Remark 3.10 says if z is connected to y by an edge labeled α , we have

$$(4) \quad N_{\epsilon_{\min}}(\rho_0(\alpha)(N_{\varepsilon}W(y))) \subset W(z).$$

We choose our neighborhood \mathcal{V}' small enough so that, for each of the finitely many parabolic vertices $z \in Z$, the semi-conjugacy ϕ_z satisfies $d_{\partial}(x, \phi_z(x)) < \min\{\epsilon_{\min}, \varepsilon\}$. This ensures that for any subset $A \subset \partial(\Gamma, \mathcal{P})$ we have $\phi_z(A) \subset N_{\varepsilon}(A)$ and $\phi_z^{-1}(A) \subset N_{\epsilon_{\min}}(A)$. In particular, we have $\phi_z W(y) \subset N_{\varepsilon}(W(y))$, so our choice of ϵ_{\min} and the containment from (4) above ensures that

$$\rho(\alpha)W(y) \subset \phi_z^{-1}\phi_z\rho(\alpha)W(y) \subset N_{\epsilon_{\min}}(\rho_0(\alpha)\phi_z W(y)) \subset W(z)$$

as desired. \square

Our next step is to define (\mathcal{G}, ρ) -codings, which provide a modified notion of a \mathcal{G} -coding which is compatible with the perturbed action ρ instead of the standard action ρ_0 . A (\mathcal{G}, ρ_0) -coding is the same thing as a \mathcal{G} -coding, but since our convention in this section is to make the standard action explicit, we will only refer to these as (\mathcal{G}, ρ_0) -codings from this point forward.

Definition 6.9. Let $\mathcal{V} \subset C(\partial(\Gamma, \mathcal{P}))$, and suppose that $\rho \in \mathcal{R}_{\mathcal{V}}$ has the same combinatorics as ρ_0 . If $\mathbf{e} = (e_k)_{k \in \mathbb{N}}$ is an infinite edge path in \mathcal{G} with $\mathbf{Lab}(e_k) = \alpha_k$,

we say \mathbf{e} is a (\mathcal{G}, ρ) -coding for ζ if

$$\zeta \in \bigcap_{k=1}^{\infty} \rho(\alpha_1) \cdots \rho(\alpha_k) W(z_k).$$

If $e_1 \dots e_n$ is a finite edge path giving a (\mathcal{G}, ρ_0) -coding for a parabolic point $z \in \partial(\Gamma, \mathcal{P})$, we say that it is a (\mathcal{G}, ρ) -coding for ζ if $\zeta \in \phi_z^{-1}(z)$, or equivalently if $\zeta \in \rho(\alpha_1 \cdots \alpha_n) \phi_{z_n}^{-1}(z_n)$, where $z_n = \tau(e_n)$.

Lemma 6.10. *Let $\mathcal{V} \subset C(\partial(\Gamma, \mathcal{P}))$, and suppose that $\rho \in \mathcal{R}_{\mathcal{V}}$ has the same combinatorics as ρ_0 . Then every point in $\partial(\Gamma, \mathcal{P})$ has a (\mathcal{G}, ρ) -coding.*

Proof. The proof follows the same strategy of Lemma 4.5. The first part of the inductive procedure goes through verbatim, replacing the sets $V(z)$ with $V_{\rho}(z)$, and elements α_i with $\rho(\alpha_i)$. The only modification required occurs at the inductive step when $z_n = \tau(e_n)$ is a parabolic rather than conical limit point. In this case we simply need to pay attention to the semi-conjugacies ϕ_{z_n} .

In detail, adopting the notation and setting from the proof of Lemma 4.5, suppose we are given $\zeta \in \partial(\Gamma, \mathcal{P})$, and assume we have found a partial coding so that $\zeta \in \rho(\alpha_1 \cdots \alpha_n) V_{\rho}(z_n)$ where z_n is a parabolic point. If $\zeta \in \rho(\alpha_1 \cdots \alpha_n) \phi_{z_n}^{-1}(z_n)$, then we have found a parabolic (\mathcal{G}, ρ) -coding for ζ and are done. Otherwise, $\zeta \notin \rho(\alpha_1 \cdots \alpha_n) \phi_{z_n}^{-1}(z_n)$. Let $\zeta_n := \rho(\alpha_1 \cdots \alpha_n)^{-1} \zeta$, so that $\zeta_n \in V_{\rho}(z_n) - \phi_{z_n}^{-1}(z_n)$. Then $\phi_{z_n}(\zeta_n) \in V(z_n) - \{z_n\}$ by the definition of $V_{\rho}(z_n)$. So, by property (C5) of the ρ_0 -automaton, there is some $\alpha_{n+1} \in L(z_n)$ so that $\rho_0(\alpha_{n+1}^{-1}) \phi_{z_n}(\zeta_n) \in \hat{V}(z_n)$.

We can then apply Lemma 6.5 (with $x = \rho(\alpha_{n+1}^{-1}) \zeta_n$) to see that

$$\rho_0(\alpha_{n+1}^{-1}) \phi_{z_n}(\zeta_n) = \phi_p(\rho(\alpha_{n+1}^{-1}) \zeta_n).$$

We conclude that $\phi_p(\rho(\alpha_{n+1}^{-1}) \zeta_n) \in \hat{V}(z_n)$. So, setting $\zeta_{n+1} := \rho(\alpha_{n+1}^{-1}) \zeta_n$, we have

$$\zeta_{n+1} \in \phi_p^{-1}(\hat{V}(z_n)) = \hat{V}_{\rho}(z_n).$$

Since the $V_{\rho}(z)$ sets still cover $\partial(\Gamma, \mathcal{P})$, there is some z_{n+1} with $\zeta_{n+1} \in V_{\rho}(z_{n+1})$. Since ρ has the same combinatorics as ρ_0 , and ζ_n lies in $\hat{V}_{\rho}(z_n) \cap V_{\rho}(z_{n+1})$ there is an edge in \mathcal{G} from z_n to z_{n+1} , which completes the inductive step. \square

Convention 6.11 (Choosing constants). *For each pair s, s' in $\mathcal{S} \cup \{id\}$, Proposition 5.2 gives a constant $c(s, s')$ such that any generalized codings of a common point of the form (s, \mathbf{e}) and (s', \mathbf{f}) satisfy c -uniform nesting. Fix $\varepsilon' < \min\{c(s, s') : s, s' \in \mathcal{S} \cup \{id\}\}$.*

Proposition 5.2 now states that, for each pair (s, s') , for this fixed ε' there exist constants $N(s, s'), D_1(s, s')$ and $D_2(s, s')$ so that, for every pair of generalized codings (s, \mathbf{e}) and (s', \mathbf{f}) of a common conical point, uniform nesting is satisfied with the constants $N(s, s'), D_1(s, s')$ and $D_2(s, s')$ and ε' . Fix N, D_1 and D_2 greater than the maximum of all such respective constants ranging over all pairs $s, s' \in \mathcal{S} \cup \{id\}$.

We note that these constants also implicitly depended on ε , which was specified in Definition 3.3 using our target neighborhood \mathcal{V} , and will reappear in Lemma 6.19.

Lemma 6.12. *There exists an open neighborhood \mathcal{U} of ρ_0 in $\text{Hom}(\Gamma, \partial(\Gamma, \mathcal{P}))$ and an open neighborhood \mathcal{V}' of the identity in $C(\partial(\Gamma, \mathcal{P}))$ such that for any $\rho \in \mathcal{U} \cap \mathcal{R}_{\mathcal{V}'}$, the following hold.*

(V1) *The representation ρ has the same combinatorics as ρ_0 .*

- (V2) For any $g \in \Gamma$ satisfying $|g|_X \leq D_1 + D_2 N$, and any $y, z \in Z$, if $\rho_0(g)\overline{W(z)} \subset W(y)$, then $\rho(g)\overline{W(z)} \subset W(y)$.
- (V3) For any $q, y \in Z$ with q parabolic, and any $g \in \Gamma$ satisfying $|g|_X \leq D_1$, if $\rho_0(g)\overline{B_{3\varepsilon'}(q)} \subset W(y)$, then $\rho(g)\overline{B_{3\varepsilon'}(q)} \subset W(y)$.
- (V4) For every parabolic vertex $q \in Z$ and every edge in \mathcal{G} from q to y labeled by α , if $\rho_0(\alpha)N_\varepsilon(W(y)) \subset B_{\varepsilon'}(q)$, then $\rho(\alpha)W(y) \subset B_{3\varepsilon'}(q)$.

Proof. Lemma 6.8 shows that having the same combinatorics is a relatively open condition, thus we can find $\mathcal{U}, \mathcal{V}'$ so that Item (V1) holds. Both (V2) and (V3) correspond to open conditions on finitely many elements in Γ . Thus, we need only demonstrate that by further shrinking \mathcal{V}' and \mathcal{U} if needed, we may satisfy (V4).

As in the proof of Lemma 5.3, for each parabolic vertex $q \in Z$, we choose some $t_q \in \Gamma$ so that $q = \rho_0(t_q)p$ for some $p \in \Pi$. We choose our neighborhood \mathcal{U} of ρ_0 small enough so that for any $\rho \in \mathcal{U}$ and every parabolic vertex $q \in Z$, we have

$$(5) \quad \rho(t_q)\rho_0(t_q)^{-1}B_{2\varepsilon'}(q) \subset B_{3\varepsilon'}(q).$$

Choose \mathcal{V}' sufficiently small so that for any $\phi \in \mathcal{V}'$ and every $y \in Z$, we have

$$\phi(W(y)) \subset N_\varepsilon(W(y))$$

and for every parabolic vertex $q \in Z$, we have

$$\phi^{-1}(\rho_0(t_q)^{-1}B_{\varepsilon'}(q)) \subset \rho_0(t_q)^{-1}B_{2\varepsilon'}(q).$$

Let $\rho \in \mathcal{U} \cap \mathcal{R}_{\mathcal{V}'}$. Fix a parabolic vertex $q = t_q p \in Z$. Suppose that for some edge e from q to y labeled by $\alpha \in L(q) \subset t_q \Gamma_p$, we have $\rho_0(\alpha)N_\varepsilon(W(y)) \subset B_{\varepsilon'}(q)$. We may write $\alpha = t_q \alpha'$ for some $\alpha' \in \Gamma_p$, so $\rho_0(\alpha')N_\varepsilon(W(y)) \subset \rho_0(t_q)^{-1}B_{\varepsilon'}(q)$.

From this it follows that

$$\begin{aligned} \rho(\alpha')W(y) &\subset \rho(\alpha')\phi_p^{-1}N_\varepsilon(W(y)) \\ &= \phi_p^{-1}(\rho_0(\alpha')N_\varepsilon(W(y))) \\ &\subset \phi_p^{-1}(\rho_0(t_q)^{-1}B_{\varepsilon'}(q)) \\ &\subset \rho_0(t_q)^{-1}B_{2\varepsilon'}(q). \end{aligned}$$

Putting this together with (5), we conclude that

$$\rho(\alpha)W(y) = \rho(t_q)\rho(\alpha')W(y) \subset \rho(t_q)\rho_0(t_q)^{-1}B_{2\varepsilon'}(q) \subset B_{3\varepsilon'}(q)$$

as desired. \square

6.2. Defining the semi-conjugacy. For the rest of the paper, we fix open sets $\mathcal{U} \subset \text{Hom}(\Gamma, \text{Homeo}(\partial(\Gamma, \mathcal{P})))$ and $\mathcal{V}' \subset C(\partial(\Gamma, \mathcal{P}))$ satisfying the conditions of Lemma 6.12. We let $\rho : \Gamma \rightarrow \text{Homeo}(\partial(\Gamma, \mathcal{P}))$ be a representation in $\mathcal{U} \cap \mathcal{R}_{\mathcal{V}'}$. We recall from Remark 6.3 that this means that for each $P \in \mathcal{P}$, we fix a semi-conjugacy $\phi_P \in \mathcal{V}'$, such that $\rho|_P$ is an extension of $\rho_0|_P$ by ϕ_P . In turn this determines semi-conjugacies ϕ_z extending $\rho|_{\Gamma_z}$ for each parabolic $z \in \partial(\Gamma, \mathcal{P})$ with stabilizer Γ_z .

Our goal is to show that ρ is semi-conjugate to ρ_0 via some map ϕ . We now set about constructing a map ϕ , and will then show that it is indeed a well-defined semi-conjugacy, and in the neighborhood \mathcal{V} of the identity that was fixed in Section 3.1. The general strategy is to use codings to assign a well-defined closed subset $\Phi(\zeta) \subset \partial(\Gamma, \mathcal{P})$ to each $\zeta \in \partial(\Gamma, \mathcal{P})$, then specify that ϕ send $\Phi(\zeta)$ to ζ . To ensure this gives a well-defined continuous map, we need the following lemma.

Lemma 6.13. *Let ζ be a conical limit point in $\partial(\Gamma, \mathcal{P})$, and $s \in \mathcal{S} \cup \{id\}$. Suppose that $\mathbf{e} = (e_k)_{k \in \mathbb{N}}$ is a strict (\mathcal{G}, ρ_0) -coding of $\rho_0(s)\zeta$, and $\mathbf{f} = (f_k)_{k \in \mathbb{N}}$ is a strict (\mathcal{G}, ρ_0) -coding of ζ ; equivalently (s, \mathbf{f}) is a generalized (\mathcal{G}, ρ_0) -coding of $\rho_0(s)\zeta$. Then, for $\alpha_k = \mathbf{Lab}(e_k)$ and $\beta_k = \mathbf{Lab}(f_k)$, we have*

$$\bigcap_{k=1}^{\infty} \rho(\alpha_1) \cdots \rho(\alpha_k) \overline{W(\tau(e_k))} = \rho(s) \bigcap_{k=1}^{\infty} \rho(\beta_1) \cdots \rho(\beta_k) \overline{W(\tau(f_k))}.$$

Proof. For each $k \in \mathbb{N}$, write $W^{\mathbf{e}}(k) = W(\tau(e_k))$ and $W^{\mathbf{f}}(k) = W(\tau(f_k))$. Let g_k and h_k be the associated quasi-geodesic sequences to the codings $\mathbf{e} = (id, \mathbf{e})$ and (s, \mathbf{f}) , respectively (so $h_0 = s$). Then $h'_k := s^{-1}h_k = \beta_1 \cdots \beta_k$ is the associated quasi-geodesic sequence to the strict coding \mathbf{f} .

We will prove the inclusion

$$(6) \quad \bigcap_{k=1}^{\infty} \rho(g_k) \overline{W^{\mathbf{e}}(k)} \subseteq \rho(s) \bigcap_{k=1}^{\infty} \rho(h'_k) \overline{W^{\mathbf{f}}(k)}.$$

Since \mathcal{S} is symmetric, the other inclusion then follows immediately.

Since both intersections in (6) are given by decreasing sequences of sets, it suffices to show that we can find arbitrarily large pairs of indices L, R so that

$$\rho(g_L) \overline{W^{\mathbf{e}}(L)} \subseteq \rho(s) \rho(h'_R) \overline{W^{\mathbf{f}}(R)}.$$

Since both (id, \mathbf{e}) and (s, \mathbf{f}) are generalized (\mathcal{G}, ρ_0) -codings of the point $\rho_0(s)\zeta$, Proposition 5.2 and our choice of constants imply that (id, \mathbf{e}) and (s, \mathbf{f}) have the uniform nesting property, with respect to the constants chosen in Convention 6.11. Thus, there exist sequences n_k, m_k satisfying the properties in Definition 5.1 for the corresponding pair of quasi-geodesic sequences. Let L be one such choice of n_k , and let R denote m_k (for the same k). Thus, $|h_R^{-1}g_L|_X \leq D_1$.

If *uniform nesting with short words* holds, we have:

$$(7) \quad \rho_0(g_{L+N}) \overline{W^{\mathbf{e}}(L+N)} \subset \rho_0(h_R) W^{\mathbf{f}}(R)$$

and $|\alpha_{L+k}|_X \leq D_2$ for all $k \in \mathbb{N}$.

Let $g_{(L+1, L+N)} = \alpha_{L+1} \cdots \alpha_{L+N}$, so that

$$g_{L+N} = g_L \cdot g_{(L+1, L+N)}.$$

Thus,

$$\rho_0((h_R)^{-1}g_L) \rho_0(g_{(L+1, L+N)}) \overline{W^{\mathbf{e}}(L+N)} \subset W^{\mathbf{f}}(R).$$

Since each α_{L+k} has length at most D_2 , we know $|g_{(L+1, L+N)}|_X \leq ND_2$. By definition, $|h_R^{-1}g_L|_X \leq D_1$, and therefore $|(h_R)^{-1}g_L g_{(L+1, L+N)}|_X \leq D_1 + ND_2$.

Thus, by condition (V2) of Lemma 6.12 and our choice of \mathcal{U} , we have

$$\rho(g_{L+N}) \overline{W^{\mathbf{e}}(L+N)} \subset \rho(h_R) W^{\mathbf{f}}(R) = \rho(s) \rho(h'_R) W^{\mathbf{f}}(R)$$

which is what we needed to show.

If instead we have uniform nesting with long parabolics, then $z_L = \tau(e_L) = \iota(e_{L+1})$ is parabolic, and we have

$$(8) \quad \rho_0(g_L) \overline{B_{3\varepsilon'}(z_L)} \subset \rho_0(h_R) W^{\mathbf{f}}(R),$$

$$(9) \quad \rho_0(\alpha_{L+1}) N_{\varepsilon}(W^{\mathbf{e}}(L+1)) \subset B_{\varepsilon'}(z_L).$$

Since $|(h_R)^{-1}g_L|_X \leq D_1$, (8), together with condition (V3) on the neighborhood \mathcal{U} , implies that

$$\rho(g_L)\overline{B_{3\varepsilon'}(z_L)} \subset \rho(h_R)W^f(R).$$

And, our assumption (V4), together with (9), implies that

$$\rho(\alpha_{L+1})W^e(L+1) \subset B_{3\varepsilon'}(z_L).$$

We can then conclude that

$$\rho(g_{L+1})W^e(L+1) \subset \rho(h_R)W^f(R).$$

This means that $\rho(g_{L+2})\overline{W^e(L+2)} \subset \rho(h_R)W^f(R) = \rho(s)\rho(h'_R)W^f(R)$, and we are done. \square

Definition 6.14. Define a map Φ from $\partial(\Gamma, \mathcal{P})$ to the space of closed subsets of $\partial(\Gamma, \mathcal{P})$ as follows:

- If $p \in \partial(\Gamma, \mathcal{P})$ is a conical limit point, choose a strict (\mathcal{G}, ρ_0) -coding of p with terminal vertex sequence $(z_k)_{k \in \mathbb{N}}$ and label sequence $(\alpha_k)_{k \in \mathbb{N}}$, and define

$$\Phi(p) = \bigcap_{k=1}^{\infty} \rho(\alpha_1 \cdots \alpha_k) \overline{W(z_k)}.$$

- If $q \in \partial(\Gamma, \mathcal{P})$ is a parabolic point, then we choose some $g \in \Gamma$ so that $q = \rho_0(g)p$ for a point $p \in \Pi$. Then define

$$\Phi(q) = \rho(g)\phi_p^{-1}(p).$$

Taking $s = id$ in Lemma 6.13 shows that the map Φ is well-defined on conical limit points. This lemma also implies that $\Phi(\rho_0(s)z) = \rho(s)\Phi(z)$ for any $s \in \mathcal{S}$ and conical z , which means that Φ is equivariant on conical limit points.

To show Φ is well-defined on parabolic points, suppose a parabolic point q satisfies $q = \rho_0(h)p_i = \rho_0(g)p_j$ for $p_i, p_j \in \Pi$. Let $P_i = \Gamma_{p_i}$ and $P_j = \Gamma_{p_j}$. Then $p_i = p_j$ and $h \in gP_i$, so $\rho(g)\phi_{P_i}^{-1}(p_i) = \rho(h)\phi_{P_j}^{-1}(p_j)$. The same reasoning shows that in fact Φ is equivariant on parabolic points: for any $g \in \Gamma$ and parabolic $q \in \partial(\Gamma, \mathcal{P})$, we have $\rho(g)\Phi(q) = \Phi(\rho_0(g)q)$. We record this fact for future use.

Lemma 6.15. *The map Φ is equivariant, in the sense that for any $g \in \Gamma$ and $\zeta \in \partial(\Gamma, \mathcal{P})$, we have*

$$\Phi(\rho_0(g)\zeta) = \rho(g)\Phi(\zeta).$$

Our next goal is to show that the sets $\Phi(\zeta)$ partition $\partial(\Gamma, \mathcal{P})$. Towards this, we first observe that the sets $W(z)$ for $z \in Z$ can be used to “approximate” the map Φ , in the following sense.

Lemma 6.16. *Let $\zeta \in \partial(\Gamma, \mathcal{P})$, and let \mathbf{e} be a strict (\mathcal{G}, ρ_0) -coding for ζ whose first vertex is $z_0 \in Z$. Then $\Phi(\zeta) \subset W(z_0)$.*

Proof. If ζ is a conical point, this is immediate from the definition of Φ and the fact that ρ has the same combinatorics as ρ_0 . Otherwise, we consider a finite (\mathcal{G}, ρ_0) -coding \mathbf{e} for the parabolic point ζ , with label sequence $\alpha_1, \dots, \alpha_n$ and initial vertex sequence z_0, \dots, z_n . Then z_n is parabolic and $\zeta = \rho_0(\alpha_1 \cdots \alpha_n)z_n$, so by equivariance we have $\Phi(\zeta) = \rho(\alpha_1 \cdots \alpha_n)\Phi(z_n)$. We know $z_n \in V(z_n)$ by Property (C5) of our original automaton, so $\Phi(z_n) = \phi_{z_n}^{-1}(z_n) \subset V_\rho(z_n)$. Then, because ρ has the same combinatorics as ρ_0 , we know $\Phi(z_n) \subset W(z_n)$ (from part (ii) of the definition) and therefore $\Phi(\zeta) \subset W(z_0)$ (from part (iv)). \square

The endgame of the proof of Theorem 1.1 is identical to that in the case without parabolics, given at the end of Section 4 in [MMW22]. For convenience, we repeat it below.

Lemma 6.17. *For any distinct $a, b \in \partial(\Gamma, \mathcal{P})$, the sets $\Phi(a), \Phi(b)$ are disjoint.*

Proof. Given $a \neq b \in \partial(\Gamma, \mathcal{P})$, let $g \in \Gamma$ be such that $d_\partial(\rho_0(g)a, \rho_0(g)b) > D$, where D is the constant from Definition 2.3. Let z_0 and y_0 be the initial vertices of strict \mathcal{G} -codings for $\rho_0(g)a$ and $\rho_0(g)b$, respectively. Then $\rho_0(g)a \in W(z_0)$ and $\rho_0(g)b \in W(y_0)$. Since each set $W(z)$ has diameter strictly less than $D/4$, and $d_\partial(\rho_0(g)a, \rho_0(g)b) > D$, we conclude that $W(z_0) \cap W(y_0) = \emptyset$.

Now Lemma 6.16 states that $\Phi(\rho_0(g)a) = \rho(g)\Phi(a) \subset W(z_0)$, and similarly $\rho(g)\Phi(b) \subset W(y_0)$. Since $W(z_0) \cap W(y_0) = \emptyset$, we also have $\Phi(a) \cap \Phi(b) = \emptyset$. \square

The fact that every point in $\partial(\Gamma, \mathcal{P})$ has a strict (\mathcal{G}, ρ) -coding (Lemma 6.10) ensures that the union of the sets $\Phi(z)$ for $z \in \partial(\Gamma, \mathcal{P})$ is all of $\partial(\Gamma, \mathcal{P})$ and Lemma 6.17 ensures that this union is in fact a partition. So, the following definition makes sense and gives a surjective map.

Definition 6.18. Define $\phi : \partial(\Gamma, \mathcal{P}) \rightarrow \partial(\Gamma, \mathcal{P})$ by taking $\phi(x) = y$ if $x \in \Phi(y)$.

We know that ϕ is equivariant by Lemma 6.15, so to prove that it is a semi-conjugacy lying in our chosen neighborhood $\mathcal{V} \subset C(\partial(\Gamma, \mathcal{P}))$, we just need to prove that it is a continuous map which is ε -close to the identity (see Definition 3.3).

Lemma 6.19. *For every $\zeta \in \partial(\Gamma, \mathcal{P})$, we have $d_\partial(\zeta, \phi(\zeta)) < \varepsilon$.*

Proof. Fix $\zeta \in \partial(\Gamma, \mathcal{P})$ and let $\xi = \phi(\zeta)$, so $\zeta \in \Phi(\xi)$. Then by Lemma 6.16, if e is a $(\mathcal{G}\rho_0)$ -coding for ξ with initial vertex z_0 , we have both $\xi \in W(z_0)$ and $\Phi(\xi) \subset W(z_0)$, hence $\zeta \in W(z_0)$. Since $\text{diam}(W(z_0)) < \varepsilon$ by property (C1), the result follows. \square

Lemma 6.20. *The map ϕ is continuous.*

Proof. Fix $\zeta \in \partial(\Gamma, \mathcal{P})$. By the previous lemma we know that $d_\partial(\zeta, \phi(\zeta)) < \varepsilon$. To show continuity at ζ we will use equivariance and the convergence action of Γ on $\partial(\Gamma, \mathcal{P})$. Suppose $\zeta_n \rightarrow \zeta$, but along some subsequence we have $\phi(\zeta_n) \rightarrow \xi \neq \phi(\zeta)$. By equivariance of ϕ , we may assume without loss of generality (after applying some $g \in \Gamma$) that $d_\partial(\xi, \phi(\zeta)) > D$, where D is the constant from Definition 2.3. On the other hand, by the triangle inequality we have

$$d_\partial(\phi(\zeta_n), \phi(\zeta)) \leq d_\partial(\phi(\zeta_n), \zeta_n) + d_\partial(\zeta_n, \zeta) + d_\partial(\zeta, \phi(\zeta)) < 3\varepsilon$$

provided that n is sufficiently large. Then $d_\partial(\xi, \phi(\zeta_n)) > D - 3\varepsilon > \varepsilon$, which gives a contradiction. \square

This completes the demonstration that ϕ is a semiconjugacy satisfying the properties of Theorem 1.1, and concludes the proof.

APPENDIX A. UNIFORM NESTING FOR AUTOMATA

The purpose of this appendix is to prove a *uniform nesting property* for finite-state automata which “code points” with respect to a Γ -action on some Hausdorff space M . Special cases of this result were originally stated as Lemma 3.8 and Corollary 3.11 in [MMW22]. We use essentially the same proof as in that paper to show the general statement (Lemma A.3 below).

We start with some general set-up. Let G be any group acting by homeomorphisms on a Hausdorff space M . As in Section 4 of this paper, for any edge e in a directed graph, we let $\iota(e)$ and $\tau(e)$ respectively denote the initial and terminal vertices of e .

Definition A.1. A *finitary point coder* \mathcal{Q} for the action $G \curvearrowright M$ consists of

- (1) A finite collection $\mathcal{W}(\mathcal{Q})$ of open sets of M .
- (2) A finite collection $F(\mathcal{Q})$ of elements of G .
- (3) A finite directed graph, with each vertex v labeled by an open set $W(v) \in \mathcal{W}(\mathcal{Q})$, and each edge e labeled by an element $\mathbf{Lab}(e) \in F(\mathcal{Q})$, satisfying the following conditions:
 - (a) Whenever there is an edge from z_1 to z_2 labeled by α , there is an inclusion $\alpha \overline{W}(z_2) \subset W(z_1)$.
 - (b) Whenever $\mathbf{e} = (e_k)_{k \in \mathbb{N}}$ is an infinite edge path with terminal vertex sequence $\tau(e_k) = z_k$ and edge labels $\alpha_k = \mathbf{Lab}(e_k)$, the sequence of sets $(\alpha_1 \cdots \alpha_n \overline{W}(z_n))_{n \in \mathbb{N}}$ is a system of neighborhoods for a point $p \in M$. Such an edge path is called a *strict \mathcal{Q} -coding* of the point p .

If $\mathbf{e} = (e_k)_{k \in \mathbb{N}}$ is a strict \mathcal{Q} -coding of p with labels $\mathbf{Lab}(e_k) = \alpha_k$ and terminal vertices $v_k = \tau(e_k)$, it follows immediately that the intersection

$$\bigcap_{n \in \mathbb{N}} \alpha_1 \cdots \alpha_n \overline{W}(v_n)$$

is equal to $\{p\}$. If M is metrizable and every set $W \in \mathcal{W}(\mathcal{Q})$ has compact closure (which is the case for all of our applications), then the converse also holds: if $\mathbf{e} = (e_k)_{k \in \mathbb{N}}$ is an edge path with label sequence $(\alpha_k)_{k \in \mathbb{N}}$, and $\bigcap_{n \in \mathbb{N}} \alpha_1 \cdots \alpha_n \overline{W}(\tau(e_n)) = \{p\}$, then \mathbf{e} is a \mathcal{Q} -coding of p .

Definition A.2. Let \mathcal{Q} be a finitary point coder. A *generalized \mathcal{Q} -coding* is a pair (g_0, \mathbf{e}) , where $g_0 \in G$ and \mathbf{e} is a strict \mathcal{Q} -coding. If \mathbf{e} is a strict coding of $p \in M$, then we say that the generalized coding (g_0, \mathbf{e}) codes the point $g_0 p \in M$.

The label sequence and initial/terminal vertex sequences of a generalized coding (g_0, \mathbf{e}) are defined to be the same as the corresponding sequences for the strict coding \mathbf{e} . If $(\alpha_k)_{k \in \mathbb{N}}$ is the label sequence for a generalized coding (g_0, \mathbf{e}) , then the *path sequence* associated to the coding is the sequence $(g_k)_{k \in \mathbb{N} \cup \{0\}}$ in G defined by

$$g_k := g_0 \cdot \alpha_1 \cdots \alpha_k.$$

We prove the following:

Lemma A.3 (Uniform nesting for finitary point coders). *For any finite subset $F \subset G$ and any finitary point coders $\mathcal{Q}, \mathcal{Q}'$, there is a number $N = N(\mathcal{Q}, \mathcal{Q}', F)$ satisfying the following. Suppose that (g_0, \mathbf{c}) is a generalized \mathcal{Q}' -coding of p with path sequence $(g_k)_{k \in \mathbb{N} \cup \{0\}}$ and terminal vertex sequence $(z_k)_{k \in \mathbb{N}}$, and (h_0, \mathbf{d}) is a generalized \mathcal{Q} -coding of p with path sequence $(h_k)_{k \in \mathbb{N} \cup \{0\}}$ and terminal vertex sequence $(y_k)_{k \in \mathbb{N}}$. Then, for any indices $m, n \in \mathbb{N}$ satisfying $g_n^{-1} h_m \in F$, we have*

$$(10) \quad g_{n+N} \overline{W}(z_{n+N}) \subset h_m W(y_m).$$

Proof. Fix a finite set $F \subset G$ and finitary point coders $\mathcal{Q}, \mathcal{Q}'$. The proof is by contradiction, so 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 $p^N \in M$, a generalized \mathcal{Q}' -coding $(g_0^{(N)}, \mathbf{c}^{(N)})$ of $p^{(N)}$ with path sequence $(g_k^{(N)})_{k \in \mathbb{N} \cup \{0\}}$ and

terminal vertex sequence $(z_k^{(N)})_{k \in \mathbb{N}}$, a generalized \mathcal{Q} -coding $(h_0^{(N)}, \mathbf{d}^{(N)})$ of $p^{(N)}$ with path sequence $(h_k^{(N)})_{k \in \mathbb{N} \cup \{0\}}$ and terminal vertex sequence $(y_k^{(N)})_{k \in \mathbb{N}}$, and indices $m_N, n_N \in \mathbb{N}$ so that $(g_{n_N}^{(N)})^{-1} h_{m_N}^{(N)} \in F$, but the inclusion (10) fails. That is, for each N , we have

$$(11) \quad g_{n_N+N}^{(N)} \overline{W(z_{n_N+N}^{(N)})} \not\subset h_{m_N}^{(N)} W(y_{m_N}^{(N)}).$$

We immediately pass to a subsequence so that

$$(f) \quad (g_{n_N}^{(N)})^{-1} h_{m_N}^{(N)} \text{ is constant, equal to } f \in F.$$

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

$$(z) \quad \text{The sets } W(z_{n_N+N}^{(N)}) \text{ are constant, equal to some } W \in \mathcal{W}(\mathcal{Q}').$$

$$(y) \quad \text{The sets } W(y_{m_N}^{(N)}) \text{ are constant, equal to some } U \in \mathcal{W}(\mathcal{Q}).$$

$$(*) \quad \text{The sets } (h_{m_N}^{(N)})^{-1} h_{m_N+1}^{(N)} W(y_{m_N+1}^{(N)}) \text{ are constant, equal to some } U'.$$

The first two are possible because the sets $\mathcal{W}(\mathcal{Q}')$ and $\mathcal{W}(\mathcal{Q})$ are finite; the third also uses the fact that each $\mathbf{d}^{(N)}$ is a (strict) \mathcal{Q} -coding, so the elements $(h_{m_N}^{(N)})^{-1} h_{m_N+1}^{(N)}$ lie in the finite set $F(\mathcal{Q})$. A key property we will use at the end of the proof is that

$$(12) \quad \overline{U'} \subset U.$$

The fact that $\mathbf{c}^{(N)}$ is a \mathcal{Q}' -coding implies that for each $k \geq 1$, we have $g_k^{(N)} = g_{k-1}^{(N)} \alpha$ for some α chosen from the finite set $F(\mathcal{Q}')$. For N in our subsequence, we can multiply each side of (11) on the left by $(g_{n_N}^{(N)})^{-1}$ and apply condition (f) to obtain

$$(13) \quad \alpha_1^{(N)} \cdots \alpha_N^{(N)} \overline{W} \not\subset fU,$$

where $\alpha_k^{(N)} := (g_{n_N+k-1}^{(N)})^{-1} g_{n_N+k}^{(N)}$ is the label of an edge $e_k^{(N)}$ in \mathcal{Q}' .

For each N we consider the infinite edge path $\gamma^{(N)}$ in \mathcal{Q}' given by

$$\gamma^{(N)} := (e_1^{(N)}, e_2^{(N)}, \dots)$$

We note that $\gamma^{(N)}$ is a strict \mathcal{Q}' -coding, coding the point $(g_{n_N}^{(N)})^{-1} p^N$.

Passing to a subsequence $\{N(j)\}_{j \in \mathbb{N}}$ one final time, we obtain a sequence of strict \mathcal{Q}' -codings $\{\gamma^{(N(j))}\}_{j \in \mathbb{N}}$ so that for all $l \geq j$, the initial subsegment of length j of the \mathcal{Q}' -coding $\gamma^{N(l)}$ is independent of l . In particular this subsequence of codings converges to a strict \mathcal{Q}' -coding

$$\gamma^\infty = (e_1^\infty, e_2^\infty, \dots)$$

with edge labels $\alpha_k^\infty := \mathbf{Lab}(e_k^\infty)$. By property (3b) of \mathcal{Q}' , this coding determines a unique point $p^\infty \in M$.

For our subsequence $N(j)$, the non-containment in (13) takes the form

$$\alpha_1^{(N(j))} \cdots \alpha_{N(j)}^{(N(j))} \overline{W} \not\subset fU.$$

We may assume that $N(j) > j$ for all j , so we can rewrite this as

$$(\alpha_1^\infty \cdots \alpha_j^\infty) \left(\alpha_{j+1}^{(N(j))} \cdots \alpha_{N(j)}^{(N(j))} \right) \overline{W} \not\subset fU.$$

By the nesting property of codings we have

$$(\alpha_1^\infty \cdots \alpha_j^\infty) \left(\alpha_{j+1}^{(N(j))} \cdots \alpha_{N(j)}^{(N(j))} \right) \overline{W} \subset (\alpha_1^\infty \cdots \alpha_j^\infty) W_j^\infty,$$

so we must therefore have

$$(14) \quad (\alpha_1^\infty \cdots \alpha_j^\infty) W_j^\infty \not\subset fU.$$

Since γ^∞ is a coding for p^∞ , the sets on the left hand side of (14) give a nested basis of neighborhoods of p^∞ and so we conclude

$$(15) \quad p^\infty \notin fU.$$

On the other hand, for each N , $\gamma^{(N)}$ is a coding for $(g_{n_N}^{(N)})^{-1}p^N$. Thus for each j we have

$$\begin{aligned} (g_{n_{N(j)}}^{(N(j))})^{-1}p^{N(j)} &\in \alpha_1^{(N(j))} \cdots \alpha_j^{(N(j))} W_j^{(N(j))} \\ &= \alpha_1^\infty \cdots \alpha_j^\infty W_j^\infty. \end{aligned}$$

As before this last sequence of sets is a nested neighborhood basis for p^∞ and thus

$$\lim_{j \rightarrow \infty} (g_{n_{N(j)}}^{(N(j))})^{-1}p^{N(j)} = p^\infty.$$

We also know that, for any N ,

$$(16) \quad (g_{n_N}^{(N)})^{-1}p^N \in (g_{n_N}^{(N)})^{-1}h_{m_N+1}^{(N)} \overline{W(y_{m_N+1}^{(N)})},$$

since $\mathbf{d}^{(N)}$ also codes p^N . By our assumptions (f) and (*) on our chosen subsequence, the right-hand side of (16) is always equal to a constant fU' . But we have just seen that a subsequence of the left-hand side converges to p^∞ , so we must have $p^\infty \in f\overline{U'}$. Because of (12) this implies

$$p^\infty \in fU,$$

contradicting (15). □

REFERENCES

- [BDL18] S. A. Ballas, J. Danciger, and G.-S. Lee. Convex projective structures on nonhyperbolic three-manifolds. *Geom. Topol.*, 22(3):1593–1646, 2018.
- [BH20] B. Burns Healy and G. C. Hruska. Cusped spaces and quasi-isometries of relatively hyperbolic groups. *arXiv e-prints*, page arXiv:2010.09876, October 2020.
- [BM20] S. Ballas and L. Marquis. Properly convex bending of hyperbolic manifolds. *Groups Geom. Dyn.*, 14(2):653–688, 2020.
- [BM22] J. Bowden and K. Mann. C^0 stability of boundary actions and inequivalent Anosov flows. *Ann. Sci. Éc. Norm. Supér. (4)*, 55(4):1003–1046, 2022.
- [Bob19] M. D. Bobb. Convex projective manifolds with a cusp of any non-diagonalizable type. *J. Lond. Math. Soc. (2)*, 100(1):183–202, 2019.
- [Bow12] B. H. Bowditch. Relatively hyperbolic groups. *Internat. J. Algebra Comput.*, 22(3):1250016, 66, 2012.
- [CINS23] C. Connell, M. Islam, T. Nguyen, and R. Spatzier. Boundary actions of lattices and C^0 local semi-rigidity. Preprint, arXiv:2303.00543, 2023.
- [GM08] D. Groves and J. F. Manning. Dehn filling in relatively hyperbolic groups. *Israel J. Math.*, 168:317–429, 2008.
- [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].
- [MM23] K. Mann and J. F. Manning. Stability for hyperbolic groups acting on boundary spheres. *Forum Math. Sigma*, 11:25, 2023. Id/No e83.
- [MMW22] K. Mann, J. F. Manning, and T. Weisman. Stability of hyperbolic groups acting on their boundaries. Preprint, arXiv:2206.14914, 2022.
- [Sul85] D. Sullivan. Quasiconformal homeomorphisms and dynamics. II. Structural stability implies hyperbolicity for Kleinian groups. *Acta Math.*, 155(3-4):243–260, 1985.

- [Tuk94] P. Tukia. Convergence groups and Gromov's metric hyperbolic spaces. *New Zealand J. Math.*, 23(2):157–187, 1994.
- [Wei22] T. Weisman. An extended definition of Anosov representation for relatively hyperbolic groups. Preprint, arXiv:2205.07183, 2022. [math.GR].
- [Wei23] T. Weisman. Examples of extended geometrically finite representations. *arXiv e-prints*, page arXiv:2311.18653, November 2023.

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, UNIVERSITY OF MICHIGAN, ANN ARBOR, MI 48109, USA
Email address: `tjwei@umich.edu`